April 2015, Vol. 53, No. 4

IEEE Communications MAGAZINE

www.comsoc.org

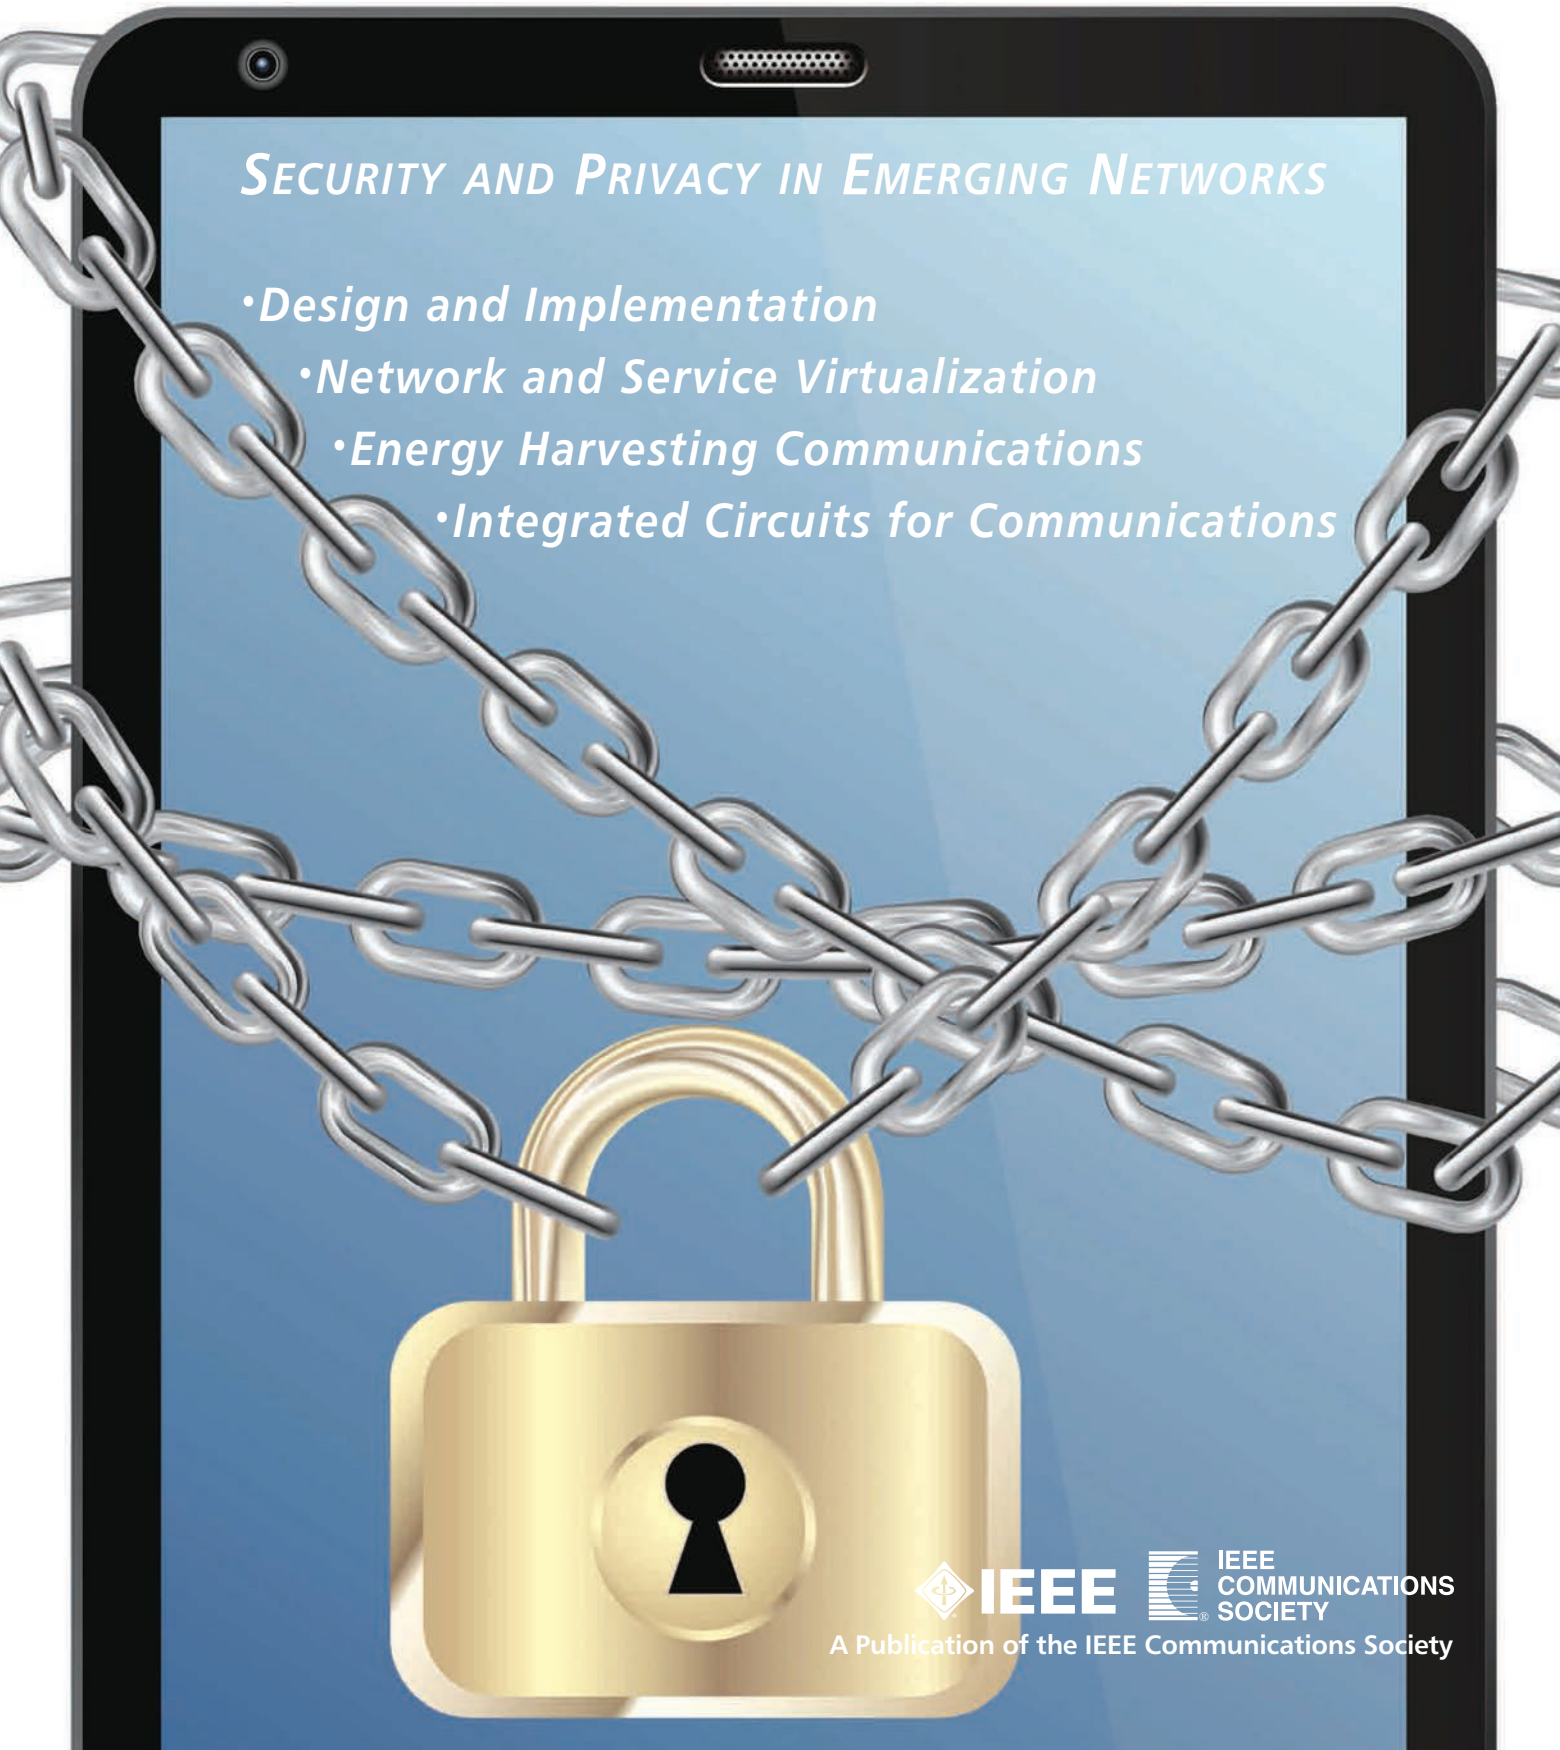

SECURITY AND PRIVACY IN EMERGING NETWORKS

- Design and Implementation
- Network and Service Virtualization
- Energy Harvesting Communications
- Integrated Circuits for Communications

IEEE COMMUNICATIONS SOCIETY
A Publication of the IEEE Communications Society

IEEE Communications MAGAZINE

THANKS OUR CORPORATE SUPPORTERS

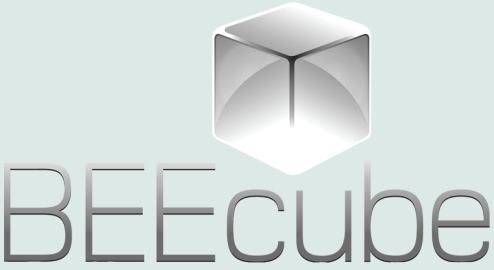

April 2015, Vol. 53, No. 4

IEEE Communications MAGAZINE

www.comsoc.org

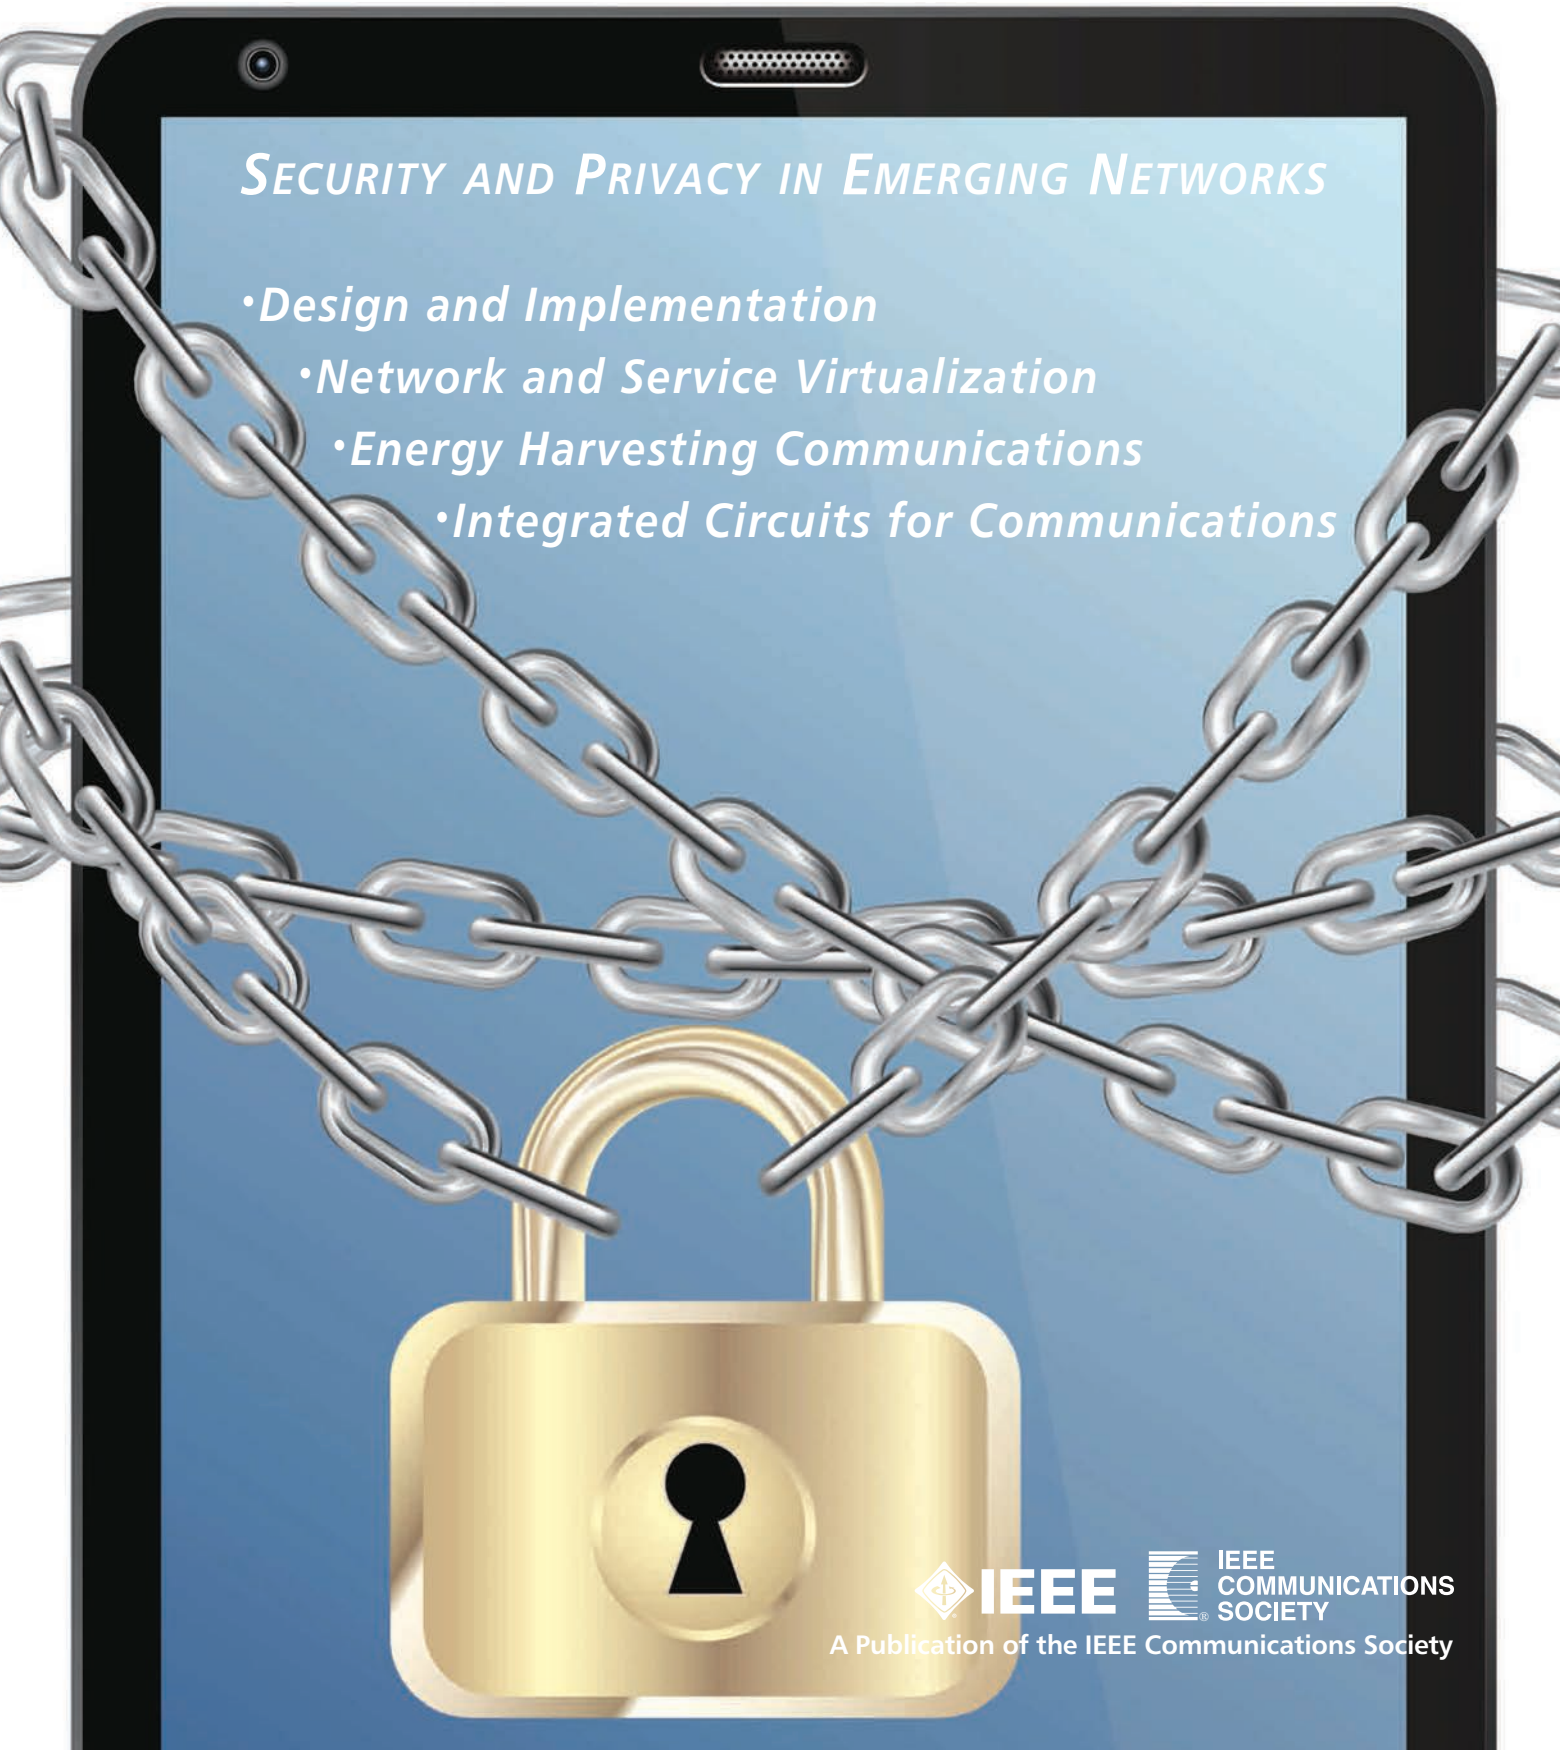

SECURITY AND PRIVACY IN EMERGING NETWORKS

- Design and Implementation
- Network and Service Virtualization
- Energy Harvesting Communications
- Integrated Circuits for Communications

IEEE COMMUNICATIONS SOCIETY
A Publication of the IEEE Communications Society

The Aha! moment.

We'll help you feel it.

It takes more than silicon to push the limits of DDR memory. It also takes gray matter. The stuff inside your head. A brain capable of genuine insight. If you're a DDR design engineer, we can give you expert advice from some of the brightest minds in the measurement world. And our end-to-end solutions range from simulation software to advanced hardware. Working together, they can help you determine precisely where your memory challenges are and how to overcome them.

HARDWARE + SOFTWARE + PEOPLE = DDR INSIGHTS

Keysight W2211BP Advanced Design System electronic design automation software
W2351EP ADS DDR4 Compliance Test Bench

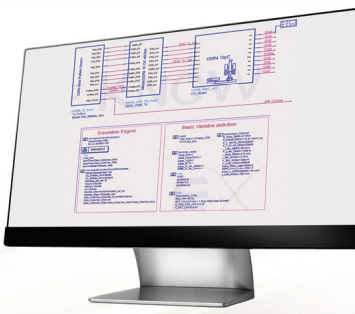

Order our complimentary 2015 DDR memory resource DVD at
www.keysight.com/find/HSD-insight

USA: 800 829 4444 CAN: 877 894 4414

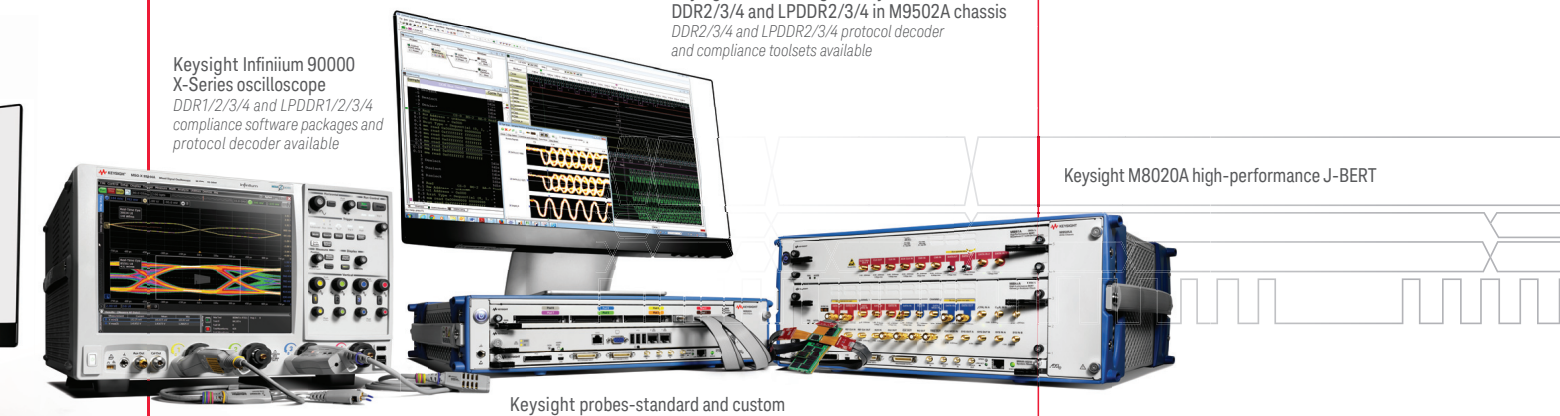

Keysight Infiniium 90000 X-Series oscilloscope
DDR1/2/3/4 and LPDDR1/2/3/4 compliance software packages and protocol decoder available

Keysight U4154B logic analyzer module for DDR2/3/4 and LPDDR2/3/4 in M9502A chassis
DDR2/3/4 and LPDDR2/3/4 protocol decoder and compliance toolsets available

Keysight M8020A high-performance J-BERT

Keysight probes-standard and custom
Standard and custom DDR and LPDDR oscilloscope and logic analyzer BGA interposer solutions

HARDWARE + SOFTWARE

If you're an engineer on the leading edge of DDR memory design, chances are, you feel challenged to go faster, using less power and a smaller footprint. We can help. Keysight is the only test and measurement company that offers hardware and software solutions across all stages of DDR chip development. From simulation to debug, from validation to compliance, we've got you covered.

- *More than 4,000 electronic measurement tools*
- *Benchtop, modular and software solutions from simulation to compliance*
- *Full line of high-speed, high-density probes*

PEOPLE

Keysight engineers are leading the industry in the design of the next generation of DDR memory standards and solutions. This means that in the future, we can help you solve cutting-edge design challenges by sharing our expertise. It also means that we can rapidly integrate new DDR specs into our hardware and software. So they'll be fully functional the day you need them.

- *JEDEC and UFSA Board of Directors Chairman, JEDEC Digital Logic (JC40.5) and UFSA Compliance Committees*
- *Hundreds of applications engineers in 100 countries around the world*
- *Thousands of patents issued in Keysight's history*

KEYSIGHT TECHNOLOGIES

Unlocking Measurement Insights

IEEE Communications Magazine

APRIL 2015, Vol. 53, No. 4

www.comsoc.org/commag

Redefining RF and Microwave Instrumentation

with open software and modular hardware

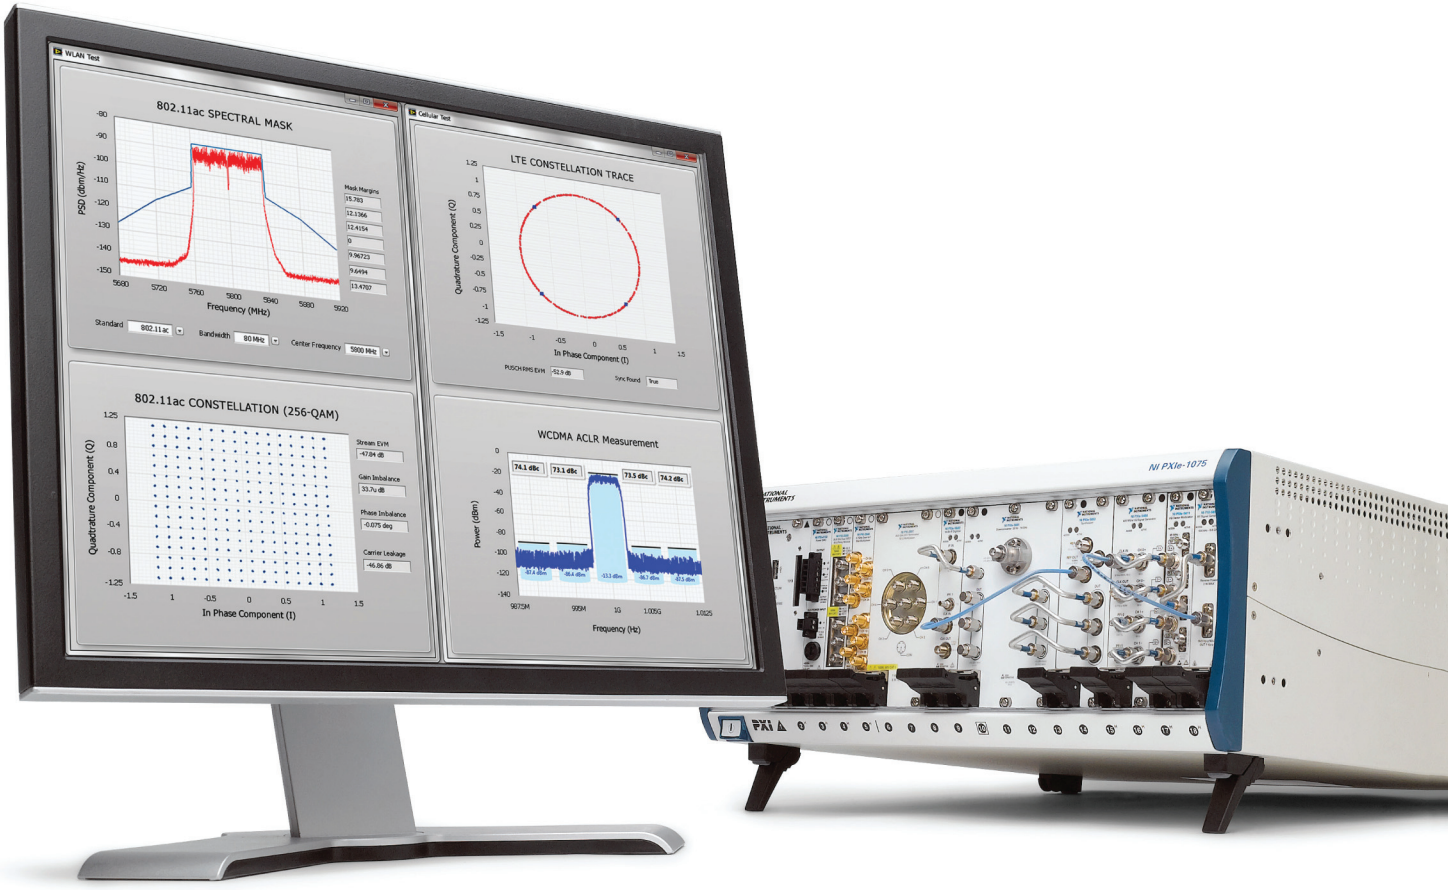

Achieve speed, accuracy, and flexibility in your RF and microwave test applications by combining National Instruments open software and modular hardware. Unlike rigid traditional instruments that quickly become obsolete by advancing technology, the system design software of NI LabVIEW coupled with NI PXI hardware puts the latest advances in PC buses, processors, and FPGAs at your fingertips.

WIRELESS TECHNOLOGIES

National Instruments supports a broad range of wireless standards including:

802.11a/b/g/n/ac	LTE
CDMA2000/EV-DO	GSM/EDGE
WCDMA/HSPA/HSPA+	Bluetooth

>> Learn more at ni.com/redefine

800 813 5078

NATIONAL INSTRUMENTS™

IEEE COMMUNICATIONS MAGAZINE (ISSN 0163-6804) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. Headquarters address: IEEE, 3 Park Avenue, 17th Floor, New York, NY 10016-5997, USA; tel: +1 (212) 705-8900; http://www.comsoc.org/commag. Responsibility for the contents rests upon authors of signed articles and not the IEEE or its members. Unless otherwise specified, the IEEE neither endorses nor sanctions any positions or actions espoused in *IEEE Communications Magazine.*

ANNUAL SUBSCRIPTION: $27 per year print subscription. $16 per year digital subscription. Non-member print subscription: $400. Single copy price is $25.

EDITORIAL CORRESPONDENCE: Address to: Editor-in-Chief, Osman S. Gebizlioglu, Huawei Technologies, 400 Crossing Blvd., 2nd Floor, Bridgewater, NJ 08807, USA; tel: +1 (908) 541-3591, e-mail: Osman.Gebizlioglu@huawei.com.

POSTMASTER: Send address changes to *IEEE Communications Magazine*, IEEE, 445 Hoes Lane, Piscataway, NJ 08855-1331. GST Registration No. 125634188. Printed in USA. Periodicals postage paid at New York, NY and at additional mailing offices. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40030962. Return undeliverable Canadian addresses to: Frontier, PO Box 1051, 1031 Helena Street, Fort Eire, ON L2A 6C7.

SUBSCRIPTIONS: Orders, address changes — IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08855-1331, USA; tel: +1 (732) 981-0060; e-mail: address.change@ieee.org.

ADVERTISING: Advertising is accepted at the discretion of the publisher. Address correspondence to: Advertising Manager, *IEEE Communications Magazine*, 3 Park Avenue, 17th Floor, New York, NY 10016.

SUBMISSIONS: The magazine welcomes tutorial or survey articles that span the breadth of communications. Submissions will normally be approximately 4500 words, with few mathematical formulas, accompanied by up to six figures and/or tables, with up to 10 carefully selected references. Electronic submissions are preferred, and should be sumitted through Manuscript Central: http://mc.manuscriptcentral.com/commag-ieee. Submission instructions can be found at the following: http://www.comsoc.org/commag/paper-submission-guidelines. For further information contact Zoran Zvonar, Associate Editor-in-Chief (zoran.zvonar@mediatek.com). All submissions will be peer reviewed.

The Brooklyn 5G Summit
Live Video Stream

sponsored by KEYSIGHT TECHNOLOGIES

9-10 April 2015

#BKLYN5G

register.ComSoc.org/BKLYN5G

Keynote Speakers:

Hossein Moiin, CTO & EVP, Nokia
James Truchard, CEO, National Instruments
Edward G. Tiedemann, SVP, Qualcomm
Gerhard P. Fettweis, TU Dresden
Seizo Onoe, CTO, NTT DoCoMo
Tom Keathley, SVP, AT&T

Topics Address:

- Channel Models: Key to 5G Air-Interface Technology
- Massive MIMO Technology for 5G and LTE-A below 6 GHz
- Spectrum for 5G
- Massive MIMO Technology for 5G
- RFIC Technology for Massive MIMO
- Safety, Exposure Assessment and Dosimetry from RF to mmWave

organized by:

Ted S. Rappaport
Professor and Founding Director of NYU WIRELESS

Amitabha Ghosh
Head, North America Radio Systems, Technology and Innovation Office, Nokia Networks

WATCH THIS EVENT LIVE OR ON-DEMAND

Find out how you can experience this $350 streaming event for free, register at ***register.ComSoc.org/BKLYN5G.***

NOKIA | NYU WIRELESS | NYU POLYTECHNIC SCHOOL OF ENGINEERING | IEEE COMMUNICATIONS SOCIETY

powered by IEEE.tv

IEEE FELLOWSHIP FOR COMSOC MEMBERS

This month's President's Page is devoted to the illustration of the IEEE and ComSoc process dealing with the evaluation of nominees to receive the IEEE Fellow grade. Within ComSoc, the assessment is performed by the Fellow Evaluation Committee, which is charged with the heavy, yet crucial task of analyzing from 70 to 100 nominations a year in order to rank and score them before passing them to the IEEE Fellow Committee.

I have the pleasure to introduce the Chair of the ComSoc Fellow Committee, Kin K. Leung, who will describe the process in detail.

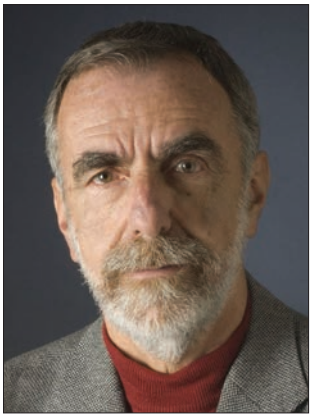

SERGIO BENEDETTO

KIN K. LEUNG

Kin received his B.S. degree from the Chinese University of Hong Kong in 1980, and his M.S. and Ph.D. degrees from the University of California, Los Angeles, in 1982 and 1985, respectively.

He joined AT&T Bell Labs in New Jersey in 1986 and worked at its successor companies, AT&T Labs and Bell Labs of Lucent Technologies, until 2004. Since then he has been the Tanaka Chair Professor in the Electrical and Electronic Engineering (EEE), and Computing Departments at Imperial College in London. He serves as the Head of Communications and the Signal Processing Group in the EEE Department at Imperial. His research interests focus on networking, protocols, optimization, and modeling issues of wireless broadband, sensor, and ad-hoc networks. He also works on multi-antenna and cross-layer designs for the physical layer of these networks.

He received the Distinguished Member of Technical Staff Award from AT&T Bell Labs in 1994, and was a co-recipient of the 1997 Lanchester Prize Honorable Mention Award. He was elected as an IEEE Fellow in 2001. He received the Royal Society Wolfson Research Merits Award from 2004 to 2009 and became a member of Academia Europaea in 2012. Along with his co-authors, he also received several best paper awards at major conferences. He has actively served on many conference committees. He serves as a member (2009-11) and the chairman (2012-15) of the IEEE Fellow Evaluation Committee for Communications Society. He was a guest editor for the *IEEE Journal on Selected Areas in Communications* (JSAC), *IEEE Wireless Communications* and the MONET journal, and as an editor for the JSAC: Wireless Series, *IEEE Transactions on Wireless Communications*, and *IEEE Transactions on Communications*. Currently he is an editor for the *ACM Computing Survey* and the *International Journal on Sensor Networks*.

IEEE FELLOW PROGRAM

The IEEE Fellow program has evolved through the years to become the program it is today. The current IEEE is the result of the merging of two professional organizations: the American Institute of Electrical Engineers (AIEE) and the Institute of Radio Engineers (IRE), in 1963. The grade of Fellow was first established in the AIEE constitution of 1912 for those engineers who had demonstrated outstanding proficiency and had achieved distinction in their profession. When the IRE established its Fellow grade in 1914, the requirements were modeled on those of the AIEE.

As it stands today, the grade of IEEE Fellow is conferred by the Board of Directors upon a member with an extraordinary record of accomplishments in any of the IEEE fields of interest. Election to the grade of IEEE Fellow is one of the highest honors that can be bestowed upon an individual by the Institute, and recognition of a member's technical, educational and leadership achievements is one of its major goals. The total number of selected Fellows in any one year does not exceed one-tenth of one percent of the total voting membership on record as of December 31st of the preceding year.

Nominations for the Fellow grade can be submitted once every year with the current deadline of March 1. At the time a nomination is submitted, the nominee must:

- Hold the IEEE Senior Member or IEEE Life Senior Member grade.
- Have accomplishments that have contributed importantly to the advancement or application of engineering, science and technology within the IEEE fields of interest, bringing the realization of significant value to society.
- Have been an IEEE member in good standing in any grade for a period of at least five years prior to January 1 of the year of elevation.

NOMINATION EVALUATION PROCESS

According to the nature and characteristics of contributions, each IEEE Fellow nominee is classified into one of the following four categories:

- Application engineer/practitioner
- Educator
- Research engineer or scientist
- Technical leader

Each nomination is also required to identify an IEEE Technical Society or Council for the purpose of carrying out the first stage of evaluation.

The assessment of the Fellow grade nominations consists of a two-stage evaluation process. The first stage of evaluation is performed by the IEEE Technical Society or Council, as identified on the nomination form. As for this society, ComSoc has an established committee known as the ComSoc Fellow Evaluation Committee (FEC) to carry out this first stage of evaluation of all nominations received by ComSoc each year. The FEC evaluation is extremely important because it provides an impartial and even-handed view of the merits of each nomination by the FEC members, who are leading experts familiar with the work and contributions by the nominees in the related fields. (The ComSoc FEC will be discussed further below.) Once the FEC evaluation is completed, their comments and scores are forwarded to the IEEE Fellow Committee for the second stage of evaluation.

The IEEE Fellow Committee has 51 members plus a chairperson. All committee members are IEEE Fellows and selected to represent the 10 IEEE Regions, and have expertise in the technical areas represented by IEEE Technical Societies or Councils. Given the very wide variety of engineering disciplines involved, it is possible that a member of the IEEE Fellow Committee may not have worked in the same area of expertise as a given nominee. As the outcome of the second-stage evaluation, the IEEE Fellow Committee recommends nominees for the Fellow grade elevation to the IEEE Board of Directors, according to the following criteria:

- Significant contributions as application engineer/practitioner, educator, research engineer/scientist, or technical leader.
- Evidence of technical accomplishments and realization of significant impact to society.
- Evaluation by the IEEE Technical Society or Council selected by the nominator.
- Confidential opinions of references and endorsers.
- Service to professional engineering societies.
- Total years in the profession.

Each nominee is rated numerically on the basis of this set of criteria.

The IEEE Board of Directors will act upon the Fellow grade recommendations from the IEEE Fellow Committee at the Board meeting during the third quarter of each year. The new IEEE Fellows are usually announced in late November or early December each year for the grade of Fellow for the selected members effective from the beginning of the following year.

COMSOC FELLOW EVALUATION COMMITTEE

At this point, the ComSoc Fellow Evaluation Committee (FEC) consists of nine members plus a chairperson, reporting to the ComSoc Vice President – Technical Activities, currently Khaled Letaief. Following is the roster of the current FEC members, the year in which they joined the committee, and their affiliations:

Kin Leung (2009, Chair), Imperial College, U.K.
Chengshan Xiao (2012), Missouri University of Science and Technology, U.S.A.
Constantinos Papadias (2013), Athens Information Technology, Greece.
Michele Zorzi (2013), University of Padova, Italy.
Dakshi Agrawal (2014), IBM Research, U.S.A.
Robert Heath (2014), University of Texas at Austin, U.S.A.
Sherman Shen (2014), University of Waterloo, Canada.
Mung Chiang (2015), Princeton University, U.S.A.
Chen-Nee Chuah (2015), University of California at Davis, U.S.A.
Ross Murch (2015), Hong Kong University of Science and Technology, China.

Our aim is to form the FEC by including leading experts to cover a wide variety of technical areas within the scope of interest of ComSoc, ranging from communication theory, communication technologies, wireless communications, networking and protocols, Internet to optical communications, satellite communications, and communication processor or hardware design, to name a few. We also aim to include members with diverse backgrounds and balanced representation from different geographic regions, academia versus industry, and gender. Each FEC member is an IEEE Fellow, and the term of service on the committee is three years. To ensure the highest degree of fairness and objectivity in our evaluation process, we request that our committee members not serve as a nominator, referee, or endorser for any Fellow nomination, while serving on the FEC. For the same reason, each committee member is also asked to refrain from assessing or expressing opinions on any nomination, if there is a potential conflict of interest between the committee member and the nominee, such as being colleagues in the same organization, research collaborators and coauthors of publications.

The first-stage evaluation of all Fellow nominations electing ComSoc as the technical society is carried out by the ComSoc FEC. The evaluation usually takes place from April to June, which unfortunately represents a very tight schedule, given the amount of work involved in addition to the regular job responsibilities of our committee members. IEEE requires that each nomination must be evaluated in detail by at least five FEC members. While all assessment work is performed through a special website, the FEC usually holds two long teleconference meetings in May so the merits of contributions and their impacts made by nominees can be discussed thoroughly among committee members. Even the teleconference meetings are not easy to schedule simply because our committee members are located across different areas of the world and many time zones are involved. In the end, we often held our teleconference calls starting at 6 am for committee members on the West Coast of the U.S. and correspondingly 10 pm for others located in Asia. One can readily see the dedication and professionalism of the FEC members.

Depending on the nature and characteristics of the contributions, each nominee is classified into one of the following four categories: application engineer/practitioner, educator, research engineer/scientist, and technical leader. According to the category of each nominee, the FEC's main focus is to identify the actual contributions and their impacts made by the nominee as well as the evidence provided to support the claim of contributions and impacts. For example, for a nominee in the category of

Class of	Number of All Nominations Received	Number of ComSoc Nominations Received	Number of ComSoc Nominees Elevated
2015	874	101	39
2014	852	86	28
2013	831	77	18
2012	799	83	33
2011	813	69	32

Table 1. Numbers of received and successful nominations.

research engineer/scientist, the committee naturally looks for the actual research contributions and their impacts. In this case, the contributions may be an introduction of a new research topic or framework, or invention of new techniques or technologies for solving a significant technical problem. The corresponding impacts may include influencing the direction of other researchers in the fields, widespread use of new technologies in engineering practices, adoption of the results in industrial standards, and so on. As another example, for the category of application engineer/ practitioner, the focus will be on the actual technical contributions and the associated practical impacts in the engineering fields. Since the FEC members come from both industry and academia in a wide variety of communications areas, our goal is to provide thorough, fair, and objective evaluations of all nominees regardless of whether their contributions are leaning more toward academic research, practical engineering techniques, engineering leadership, or education.

To assist the FEC evaluations, it would be helpful if each nominator could make the cited contributions of the nominee easy to identify and to make their importance and impact ready to understand and verify. Toward that goal, it is often helpful to include in the nomination sufficient background material and clearly state the outstand-

		Nominations Received					Nominations Received				
Class of	Affiliation	Application Engineer/ Practitioner	Educator	Technical Leader	Research Engineer/ Scientist	Total Nominations Received	Application Engineer/ Practitioner	Educator	Technical Leader	Research Engineer/ Scientist	Total Nominees Elevated
2011	Industry	2	0	3	8	**13**	1	0	0	8	**9**
2012	Industry	4	0	8	9	**21**	1	0	2	4	**7**
2013	Industry	4	0	8	8	**20**	1	0	1	7	**9**
2014	Industry	7	0	5	7	**19**	0	0	1	3	**4**
2015	Industry	6	0	9	12	**27**	2	0	2	5	**9**
2011	Education	1	0	1	46	**48**	0	0	0	22	**22**
2012	Education	0	3	4	51	**58**	0	1	0	24	**25**
2013	Education	1	4	3	42	**50**	0	2	0	7	**9**
2014	Education	0	2	2	58	**62**	0	0	0	23	**23**
2015	Education	3	3	3	55	**64**	0	0	1	27	**28**
2011	Government	1	0	1	1	**3**	0	0	0	0	**0**
2012	Government	0	0	3	0	**3**	0	0	1	0	**1**
2013	Government	1	0	1	3	**5**	0	0	0	0	**0**
2014	Government	0	0	1	3	**4**	0	0	0	1	**1**
2015	Government	0	0	5	2	**7**	0	0	0	1	**1**
2011	Other	1	0	1	3	**5**	1	0	0	1	**2**
2012	Other	0	0	1	0	**1**	0	0	0	0	**0**
2013	Other	0	0	1	1	**2**	0	0	0	0	**0**
2014	Other	0	0	0	1	**1**	0	0	0	0	**0**
2015	Other	0	0	1	2	**3**	0	0	0	1	**1**

Table 2. ComSoc nominations and elevations for different affiliations and categories.

ing contributions of the nominee and why they are important. It is also desirable to describe the state-of-the-art before the nominee made his/her contributions and highlight the significance of the contributions in this context. The nomination should also summarize how the nominee overcame the challenges involved. Finally, it is very important to illustrate the impact of the nominee's work and provide adequate and verifiable evidence to support the claim of contributions and impact.

After the thorough assessment and discussion in the teleconference meetings, a consensus on all nominations for evaluation by ComSoc can be achieved among the FEC members. To meet the requirement of only a very small percentage of members who can be elevated to the grade of Fellow each year, the FEC needs to use its assessment in order to rank all nominations submitted in the given year. The ranked list of nominees along with scores and descriptions of their key contributions and impacts represent the output of the evaluation from the FEC, which are then forwarded via the special website to the IEEE Fellow Committee for the second stage of evaluation. Although we have already optimized the evaluation process over the years in order to make it as efficient as possible, each FEC member is still required to perform a significant amount of work during the short time period of two months. It is particularly so because the committee aims to provide the most objective and thorough evaluation of every nomination, which takes time and collaboration among committee members. It is what the FEC is expected to do: to serve the ComSoc community by identifying the most deserving members for the highest honor to be given by the IEEE as an institute.

It is worth noting that the ComSoc FEC carries out only the first stage of evaluation, and the assessment results including the ranking and scoring of all ComSoc nominees serve as inputs to the second stage of evaluation to be performed by the IEEE Fellow Committee. It is likely that nominees ranked highly by the ComSoc FEC will be elevated to the Fellow grade at the end, but the final outcome of the second stage of evaluation does not strictly follow the ranked order of candidates from the FEC. This is so because the IEEE Fellow Committee uses its own set of evaluation criteria and additional information about the merits of a nominee may be revealed in the reference letters, which are not accessible by the ComSoc FEC in the first stage of evaluation.

Table 1 shows the total number of Fellow nominations, the number of those submitted to ComSoc, and the number of successful ComSoc nominees in the last several years. As shown, the number of nominations handled and evaluated by the ComSoc FEC ranged from about 70 to 100 each year. By averaging across the class years from 2011 to 2015, 36 percent of the ComSoc nominees were elevated to the Fellow grade.

The statistics can be further broken down in terms of the affiliations of the nominees, including industry, academia, government, and others, as well as the four categories: application engineer/practitioner, educator, technical leader, and research engineer/scientist. For ComSoc candidates, the numbers of received and successful nominations for various types of affiliations and categories are given in Table 2.

FURTHER INFORMATION

Further detailed information about the IEEE Fellow Program and the online nomination process can be found at http://www.ieee.org/membership_services/membership/fellows/index.html

Readers are also welcomed to forward any questions about the IEEE Fellow nomination process or other related questions to the IEEE Fellow staff at fellows@ieee.org.

BOOK REVIEWS

EDITED BY PIOTR CHOLDA

MOBILITY PROTOCOLS AND HANDOVER OPTIMIZATION: DESIGN, EVALUATION AND APPLICATION

BY ASHUTOSH DUTTA AND HENNING SCHULZRINNE

WILEY-IEEE PRESS, 2014, ISBN 978-0-470-74058-3, HARDCOVER, 476 PAGES

REVIEWER: AGNIESZKA CHODOREK

From the dawn of time, people have dreamed about the possibility of immediate communication. Ancient stories and fairy tales are full of mysterious objects, which enable one to talk with other people anytime and anywhere and/or to have real-time monitoring of distant places. Nowadays, crystal balls and magical mirrors, known from fairy tales, have found their communication equivalent in mobile technologies. Mobile terminals (e.g. smartphones, tablets) are associated with different types of wireless networks, which are constantly evolving to offer a suitable quality of services while ensuring mobility and high bandwidth. Efficient management of this complex system of mobile devices connected via wireless media is one of the most important challenges facing modern communications.

Mobility management is characterized by two basic aspects: location management and handoff management. Location management allows for the discovery of the binding point to establish a new session for the mobile user. Handoff management allows for the change of the network binding point of the mobile terminal when mobile users change their physical location. The book by Dutta and Schulzrinne is devoted to handoff management and is focused mainly on mobility in IP networks, starting from mobility in wireless technologies, such as 802.11, through mobility support for the IP protocol, to session mobility. As the reference point for discussion, the authors take switching between cells of cellular telephony, which has been optimized over many years and is now satisfactory for voice service.

Chapter 1 introduces the four types of mobility: terminal mobility, personal mobility, session mobility, and service mobility. Chapter 2 presents typical mobile solutions for multimedia transmission, ranging from the handoff solution used in cellular networks (from 1G to 4G) to problems of heterogeneous mobile environments. IP mobility support and other key mechanisms and protocols, which have an impact on the efficiency of the handoff, also are indicated. Chapter 3 is devoted to the analysis of the handoff delays imposed by each element of a mobile system. To determine whether systems (existing or newly designed) provide seamless mobility, specialized tools are needed. Chapter 4 presents one of these tools, i.e. the handoff event model based on a Petri net. This model can be used both for an analysis of handoff events (e.g. deadlocks) and in the analysis of a handoff performance. The three subsequent chapters describe handoff optimization carried out in different layers of the OSI Reference Model. Chapter 5 shows the exemplary optimization of Layer 2 (i.e. 802.11 wireless networks). Optimization of the mechanisms of layers higher than Layer 2 is presented in Chapter 6. Chapter 7 analyzes the impact of handover optimization, provided in different layers. The mentioned optimization approaches are based on the assumption that only one of the end systems (the sender or the receiver) is a mobile node. The next two chapters are devoted to the optimization of systems with multiple mobile nodes. Chapter 8 shows the problem of optimizing two communicating nodes that are simultaneously mobile, while Chapter 9 discusses handoff optimization in multicast systems. Chapter 10 analyzes the possibility of obtaining seamless handoffs in Layers 2 and 3 through cooperative roaming. A very strong point of the book is its presentation of handoff optimization techniques in given Chapter 11. This chapter, describing how to make the handoff validation using Petri net modeling, shows practical aspects of the theory introduced in Chapter 4.

An extensive bibliography allows the readers to broaden their knowledge of issues discussed in the book. The work offers important insights into mobility protocols and handover optimization. It is a recommended resource for graduate students, researchers, and IT professionals interested in the study of handoff management.

MILLIMETER WAVE WIRELESS COMMUNICATIONS

BY THEODORE S. RAPPAPORT, ROBERT W. HEATH JR., ROBERT C. DANIELS, AND JAMES N. MURDOCK

PRENTICE HALL, 2015, ISBN 978-0-13-217228-8, HARDCOVER, 680 PAGES

REVIEWER: LUIS M. CORREIA

Millimeter waves ('mm waves' in short) captured the attention of researchers on mobile and wireless communications ('M&W Comms.' in short) in the early 1990s in Europe, within a series of projects funded by the European Commission. Recently, with the emergence of the exploration of the 5th Generation (5G), this area of research gained a new momentum, since, in fact, mm waves present themselves as one of the key options for the transmission of high data rates (i.e. several Gbps) at short distances (i.e. a couple of hundred meters). Thus, this book from Ted Rappaport and his co-authors is very timely, indeed filling in a gap.

The book is structured into three parts: 'Prerequisites', 'Fundamentals', and 'Mm Wave Design and Applications'. This is a good approach to the topic. Besides an introduction (Chapter 1), the first part covers basic knowledge in M&W Communications (Chapter 2), enabling a non-specialized reader to gain a view of various matters, ranging from modulation and coding to equalization and synchronization, encompassing Multiple Input Multiple Output (MIMO), as well as system and hardware architectures. The second part addresses various aspects related to radio interface techniques supporting mm waves: wave propagation (Chapter 3), antennas (Chapter 4), radio-frequency devices and circuits (Chapter 5), as well as baseband circuits (Chapter 6) focusing on converters, both analog to digital and digital to analog. Finally, the third part offers a view into some of the specific problems of using mm waves, i.e. physical and higher-layer design (Chapters 7 and 8), and standardization (Chapter 9). Therefore, the book provides an almost complete view of the several matters related to the implementation of a radio interface at this frequency band. A List of Abbreviations and an Index help the readers to make the most of the information provided in the book. All chapters end with a summary, which enables an occasional reader to capture the essential results.

Basically, in each chapter the basic models are presented, followed by the application of mm waves to M&W Communications. This approach extends the readership to both specialized researchers (which may skip the basic aspects, well known to them) searching for detailed information, and laymen, e.g. graduate students looking for information in an area that may not be that familiar to them, but with enough background so the book is appropriately self-contained. In general, prior work in the area is adequately cited (there are over 800 references), but the early work in the 1990s, namely the one produced in Europe and Japan, should have captured more attention of the authors. The reader can really find some of the latest developments in this area, not only from the authors them-

selves, but from many other research groups as well, which is of great value.

This book does not have much competition, i.e. there are not many recent books in this area, and those that exist do not provide such a comprehensive view. The result of my personal browsing led to a list of books focusing on specific aspects, but not providing this broader perspective, and those that do are already from several years ago. Since mm waves have witnessed much work in recent years, this book provides a good update in the area.

The manuscript is also available as an e-book, which I find quite valuable, since these days many people (including myself) tend to read books on tablets. The authors/publisher also prepared a downloadable file with a color version of many figures, which may be helpful in some specific cases.

At the end of the Introduction the authors state, "It is our hope that you find this text to be a useful guide, enabling the creation of myriad new devices and applications that will soon be using the mm wave spectrum." I fully agree with them, i.e. the book is in fact a useful guide for those working, or intending to work, in mm waves. I am not sure if 5G will include mm waves as a technological option, but I do believe that the future of mobile and wireless communications does go across mm waves. Enjoy the reading.

CONFERENCE CALENDAR

Updated on the Communications Society's Web Site
www.comsoc.org/conferences

2015

APRIL

Brooklyn 5G — The Brooklyn 5G Summit, 3–4 Apr.
Brooklyn, NY
http://brooklyn5gsummit.com/

IEEE NETSOFT 2015 — IEEE Conference on Network Softwarization 2015, 13–17 Apr.
London, U.K.
http://sites.ieee.org/netsoft/

WTS 2015 — 2015 Wireless Telecommunications Symposium, 15–17 Apr.
New York, NY
http://www.csupomona.edu/~wtsi/index.html

FRUCT17 2015 — 17th Conference of Open Innovations Association, 20–24 Apr.
20–24 Apr.
Yaroslavi, Russia.
http://e-werest.org/cfp17

IEEE LANMAN 2015 — 21st IEEE Int'l. Workshop on Local and-Metropolitan Networks, 22–24 Apr.
Beijing, China.
http://www.ieee-lanman.org/

ICT 2015 — 2015 22nd Int'l. Conference on Telecommunications, 27–29 Apr.
Sydney, Australia.
http://www.engineersaustralia.org.au/ict2015-conference

IEEE INFOCOM 2015 — IEEE Int'l. Conference on Computer Communications, 26 Apr.–1 May
http://infocom2015.ieee-infocom.org/

MAY

IEEE CQR 2015 — IEEE Int'l. Workshop Technical Committee on Communications Quality and Reliability, 10–15 May
Charleston, SC.
http://www.ieee-cqr.org/

IEEE CTW 2015 — IEEE Communications Theory Workshop, 10–13 May
Dana Point, CA
http://www.ieee-ctw.org/

ONDM 2015 — 2015 Int'l. Conference on Optical Network Design and Modeling, 11–14 May
Pisa, Italy.
http://ondm2015.sssup.it/

IFIP/IEEE IM 2015 — Int'l. Symposium on Integrated Network Management, 11–15 May
Ottawa, Canada.
http://im2015.ieee-im.org/

IEEE BlackSeaCom 2015 — IEEE Int'l. Black Sea Conference on Communications and Networking, 18–21 May
Constanta, Romania
http://www.ieee-blackseacom.org/2015/index.html

IEEE 5G — 1st Int'l. 5G Summit, 26 May
Princeton, NJ
http://www.5gsummit.org/

UBI-HEALTHTECH 2015 — 2nd Int'l. Symposium on Future Information and Communication Technologies for Ubiquitous HealthCare, 8–12 May
Beijing, China.
http://www.ubi-health.org/

JUNE

IEEE ICC 2015 — 2015 IEEE Int'l Conference on Communications, 8–12 June
London, U.K.
http://icc2015.ieee-icc.org/

IEEE IWQOS 2015 — IEEE/ACM Int'l. Symposium on Quality and Service, 15–16 June
Portland, OR.
http://www.ieee-iwqos.org/

MED-HOC-NET 2015 — Mediterranean Ad Hoc Networking Workshop, 17–19 June
Vilamoura, Portugal.
http://www.ieee-iwqos.org/

IEEE SECON 2015 — IEEE Int'l. Conference on Sensing, Communication and Networking, 21–25 June
Seattle, WA.
http://secon2015.ieee-secon.org/

WMNC 2015 — 8th IFIP Wireless and Mobile Networking Conference, 23–25 June
Belgrade, Russia.
http://www.wmnc2015.com/

EUCNC 2015 — European Conference on Networks and Communications, 29 June–2 July
Paris, France.
http://www.eucnc.eu/?q=node/156

IEEE ICME 2015 — IEEE Int'l. Conference on Multimedia and Expo, 29 June–3 July
Torino, Italy.
http://www.icme2015.ieee-icme.org/

JULY

ICUFN 2015 — 7th Int'l. Conference on Ubiquitous and Future Networks, 7–10 July
Sapporo, Japan.
http://www.icufn.org/main/

AUGUST

IEEE PIMRC 2015 — 26th Annual IEEE Int'l. Symposium on Personal, Indoor, and Mobile Radio Communications — 30 Aug.–2 Sept.
Hong Kong.
http://pimrc2015.eee.hku.hk/index.html

SEPTEMBER

IEEE DySPAN 2015 — IEEE Symposium on Dynamic Spectrum Access Networks, 29 Sept.–2 Oct.
Stockholm, Sweden.
http://dyspan2015.ieee-dyspan.org/

–Communications Society portfolio events appear in bold colored print.

–Communications Society technically co-sponsored conferences appear in black italic print.

–Individuals with information about upcoming conferences, Calls for Papers, meeting announcements, and meeting reports should send this information to: IEEE Communications Society, 3 Park Avenue, 17th Floor, New York, NY 10016; e-mail: p.oneill@comsoc.org; fax: + (212) 705-8996. Items submitted for publication will be included on a space-available basis.

GCN GLOBAL COMMUNICATIONS NEWSLETTER

April 2015
ISSN 2374-1082

MEMBER RELATIONS

WICE: Promoting the Role of Women in Communications Engineering

Interview with Octavia Dobre, Chair of WICE Standing Committee

By Stefano Bregni, Vice-President for Member Relations, and Octavia A. Dobre, Chair of WICE Standing Committee

This is the eighth article in the series started in September 2014 and published monthly in the *Global Communications Newsletter*, which covers all areas of IEEE ComSoc Member Relations. In this series of articles, I introduced first the seven Member Relations Directors (Sister and Related Societies; Membership Programs Development; AP, NA, LA, EAME Regions; Marketing and Industry Relations) and then the Chairs of the Women in Communications Engineering (WICE) and IEEE Young Professionals (YP, formerly Graduates Of the Last Decade, GOLD) Committees. In each article, one by one they present their activities and plans.

Stefano Bregni

Octavia Dobre

In this issue I interview Octavia A. Dobre, Chair of the IEEE ComSoc Standing Committee on Women in Communications Engineering (WICE). Octavia is an Associate Professor with Memorial University, Canada. She is a Senior Editor with *IEEE Communications Letters*, and an Editor with *IEEE Transactions on Communications* and *IEEE Communications Surveys and Tutorials*. She also serves as Vice-Chair of the Signal Processing for Communications and Electronics Technical Committee and Vice-Chair Americas for the Technical Committee on Cognitive Networks.

It is my pleasure to interview Octavia and offer her the opportunity to outline her current activities and plans for ComSoc WICE.

Stefano: Hello Octavia! I am particularly glad to have the chance to interview you here and to offer you this opportunity to present the activities of the IEEE ComSoc Standing Committee on Women in Communications Engineering (WICE). At the beginning of my term as VP-MR, I indicated women among the five strategic directions to pursue in Member Relations to ensure innovation and balanced growth of ComSoc. Therefore, your Committee is key to this goal. Would you recall its history and mission?

Prof. Andrea Goldsmith receiving the WICE Outstanding Achievement Award (2014).

Prof. Sarah Kate Wilson receiving the WICE Outstanding Service Award (2014).

Octavia: In 2014, WICE became a Standing Committee of ComSoc, after being an Ad-Hoc Committee for three years. The dedication of its former chairs — Heather Yu (2011), Sarah Kate Wilson (2012), and Shalinee Kishore (2012) — made this possible. WICE's mission is to promote the visibility and roles of women communications engineers and to provide a venue for their professional growth. We hope that many ComSoc members will get involved and contribute to the mentoring and promotion of women communications engineers, as well as to increasing their membership.

Stefano: Due to a number of historical reasons, women are significantly under-represented in most engineering disciplines, including communications. What's more, for a long time there have been even fewer women in visible leadership positions. Therefore, our challenge in WICE is working to improve balance and to grant true equal opportunities to both genders, in ComSoc and in our professional world of communications engineers. How is the WICE Committee facing this challenge? What have been its main activities in 2014?

Octavia: A mandate of WICE is to promote women communications engineers. Certainly, this cannot be done without the support of professional women leaders, as well as the entire ComSoc community. WICE has been fortunate to have such support. However, there are also challenges. One of them is to have a larger number of dedicated volunteers, which would allow us to initiate additional activities. We have spread the word about WICE through Facebook, LinkedIn, and the WICE website, as well as speaking directly with attendees at Globecom and ICC, in order to attract more volunteers.

In 2014, we organized the Second Women's Workshop on Communications and Signal Processing, where we presented the Best Poster Presentation Award to a junior participant for the first time, and additionally presented the WICE Awards for the first time.

Stefano: Awards are well appreciated to acknowledge publicly outstanding contributions to the profession and ComSoc.

Octavia: To recognize the ComSoc members who bring outstanding contributions to the profession and WICE, starting in 2014 we will present three annual awards:

- Outstanding Achievements Award for exceptional technical work in the broad field of communications engineering.
- Outstanding Service Award for a distinguished record of service and excellent leadership within WICE.
- Mentorship Award for strong commitment to mentoring WICE members, significant positive impact on the mentees' education and career, and advancing communications engineering through mentees.

The first-time recipients of the WICE awards were Prof. Andrea Goldsmith (Stanford University) for the Outstanding Achievements Award; Prof. Sarah Kate Wilson (University of Santa Clara) for the Outstanding Service Award; and Dr. Larry Greenstein for

(Continued on Newsletter page 4)

CONFERENCE REPORT

Second IEEE ICCC International Workshop on Internet of Things (IOT 2014), Shanghai, China

By Antonio J. Jara, Kaoru Ota, Ruonan Zhang, and Wei Wang, ICCC IoT Workshop Publicity CoChairs

The IEEE/CIC International Conference on Communications in China (ICCC) is an international conference series newly incubatedby the IEEE Communications Society (ComSoc) in partnership with the China Institute of Communications (CIC) aiming at realistic globalization by extending ComSoc's reach to the fastest growing regions. ICCC is the flagship conference of the IEEE ComSoc in China that is held every year in the greater China region. Since its commencement in 2012, ICCC has grown steadily from a brand new conference with a strong vision, to a unique venue that brings together global researchers and practitioners in areas of communications.

ICCC 2014, the third edition of ICCC that follows the great success of ICCC 2012 in Beijing and ICCC 2013 in Xi'an, was held in Shanghai on October 13-15, 2014. Shanghai is the largest Chinese city by population and the largest city proper by population in the world. It is one of the four direct-controlled municipalities, with a population of more than 24 million as of 2013. It is a global financial center, and a transport hub with the world's busiest container port. Located in the Yangtze River Delta in East China, Shanghai sits at the mouth of the Yangtze in the middle portion of the Chinese coast.

ICCC 2014 featured four world-class plenary keynote speeches, eight technical symposia, five tutorials, four industrial and academic panels, and one workshop. The second IEEE ICCC international workshop on Internet of Things (IOT 2014) was the only workshop within ICCC 2014 and consisted of two sessions, sharing the same aim as IOT 2013: to provide a forum for authors to present early research results on Internet of Things (IOT) that advance the state of the art and practice in IOT, including theoretical principles, tools, applications, systems infrastructure, and test beds for IOT.

IOT has been the national strategy of China since 2009, and it maintains collaboration worldwide. The counterpart of IoT in the United States is cyber-physical systems (CPS) with initiatives such as SmartAmerica and with the leadership of the Industrial Internet Consortium (IIC) and IPSO Alliance to promote the use of Internet of Things, standardization, and market development. Finally, many other federal agencies have a common stake in the IoT, CPS, and Industrial Internet research and development. Finally, the European Union presents the strongest support worldwide for Internet of Things development through the European Research Cluster on the Internet of Things (IERC), IoT iniative

Shanghai Chenhuangmiao (Yuyuan Garden Bazar).

IOT 2014 Welcome Opening Speech by TPC Chair Prof. Qinghe Du.

IOT 2014 Invited Talk by Prof. Bijan Jabbari.

(IoT-i) with the IoT Forum development, IoT architecture (IoT-A) and a wide range of projects in order to address the key challenges of Internet of things in terms of cloud computing integration (OpenIoT, iCore), IPv6 support (IoT6), gateway integration (BUTLER), integration (SmartAction), and other key projects as part of the Framework Programme 7. Nowadays, the Internet of Things continues to be a priority for the European Union, with a special focus on the development of consolidated ecosystems that enable the exploitation for a wide range of markets, users, entrepreneurs, and consequently makes feasible the development of a collaborative value-chain for the Interne of Things. These ecosystems are expected to build a high impact with the development of Large Scale Pilots and Developments toward 2020.

IOT 2014 was a joint effort between the Internet of Things Emerging Technical Subcommittee within the IEEE Communications Society, including Prof. Latif Ladid and Prof. Antonio J. Jara, the founders of the IOT workshop series within ICCC, including Prof. Houbing Song, Prof. Qinghe Du, and Prof. Ruonan Zhang, and the broader IOT research community, including Prof. Bin Xia, Prof. Shengjie Zhao, Prof. Kaoru Ota, Prof. Xiaohua Tian, and Prof. Wei Wang.

The current IEEE positioning of the Internet of Things goes beyond with the development of the IEEE IoT World Forum, IEEE IoT Journal, IEEE P2413 Standard for the Internet of Things Architecture, and other activities from the IEEE IoT initiatie and the IEEE ComSoc Internet of Things Emerging Technical Subcommittee. The support of the community in the ICCC conference and the great numbers of attendees and contributions, demonstrate that Internet of Things is a top topic and enabler worldwide with a highly expected impact in our society during the coming years. Nowadays several challenges are pending in issues such as security, scalability, interoperability, and user-acceptability. We expect that the research community worldwide will continue to work on these topics within IEEE ICCC IOT 2015.

(Continued on Newsletter page 4)

TECHNOLOGY UPDATE

LTE-Assisted WiFi Direct Trial

By Sergey Andreev and Yevgeni Koucheryavy, Tampere University of Technology, Finland; Jiri Hosek, Brno University of Technology, Czech Republic; and Kerstin Johnsson, Intel Corporation, US

As the dust around 5G settles, one thing is clear: 5G will be a synergistic integration of numerous, diverse technologies (rather than one defining technology) that radically improves the performance of wireless networks. LTE-assisted WiFi Direct is considered a vital part of this 5G vision due to its ability to augment network capacity and enhance user performance without increased infrastructure cost. The use of WiFi Direct offloads the costly cellular bands, improves reuse (WiFi has a shorter range than LTE), and boosts user data rates (WiFi has a much larger bandwidth than LTE). These benefits are further enhanced by assistance from the LTE operator. LTE assistance automates WiFi Direct device discovery as well as connection establishment. It expands the number of potential WiFi Direct connections by providing secure access to strangers. It reduces battery and channel consumption by performing device proximity detection on the user's behalf. Finally, it provides service continuity to users by enabling a communication path via the LTE infrastructure if/when users move too far apart for successful WiFi Direct communication.

LTE-Assisted WiFi Direct on a 3GPP LTE Deployment

Motivated by the many potential benefits of LTE-assisted WiFi Direct and building on our extensive past research in this area, we committed to deploy this promising technology and comprehensively demonstrate its benefits in the summer of 2014. To this end, we have completed a full-scale trial of LTE-assisted WiFi Direct on a live 3GPP LTE deployment in Brno, Czech Republic. This unique trial builds upon 3GPP-compliant D2D technology, features our patented signaling protocols, and significantly extends our initial demo shown at Mobile World Congress in early 2014. This effort unites partners from Tampere U. of Technology (TUT), Brno U. of Technology (BUT), and Intel Labs US.

Our trial reveals that LTE-assisted WiFi Direct does, in fact, significantly improve network and user performance. This technology essentially creates large numbers of "small cells" that relieve cellular network congestion without the additional CAPEX/OPEX associated with deploying pico-/femto- cells. As long as the proper offloading criteria are set, LTE-assisted WiFi Direct overcomes the limitations of conventional WiFi, such as session continuity failures, excessive user contention, and cumbersome security and connection establishment procedures.

Group photo from the seminar on Network-Assisted D2D.

Our Continued Work on LTE-assisted WiFi Direct

Upon conclusion of the trial, we hosted a seminar at Brno University of Technology to discuss lessons learned, industry implications of the technology, and future research directions. This seminar brought together representatives from a variety of mobile network operators, vendors, and manufacturers including Intel Labs, AT&T, Nokia, T-Mobile, France Telecom, Honeywell International, Fraunhofer Institute (FOKUS), Brno University of Technology, Tampere University of Technology, and Vienna University of Technology.

The described trial was performed in laboratories of the SIX Research Centre; our current work on network-assisted proximate communications continues with support from the Academy of Finland. With this funding we are researching issues of security and privacy in the context of D2D, as well as its performance evaluation aspects (as part of a postdoctoral researcher grant by the first author). We are also supported by grants from the Internet of Things program of DIGILE (funded by Tekes). This funding supports research into D2D technology improvements specific to IoT usage scenarios. More details on the recent trial and the follow-up seminar are available at the following links:
http://winter-group.net/brno-trial/
http://wislab.cz/our-work/lte-assisted-wifi-direct

Please contact the authors (sergey.andreev@tut.fi, hosek@feec.vutbr.cz, kerstin.johnsson@intel.com, yk@cs.tut.fi) for further information about the research, demos, or trial.

WORKSHOP REPORT

Nets4Cars: 2014 Fall Workshop in Saint Petersburg, Russia

By Alexey Vinel, Sweden

The International Workshop on Communication Technologies for Vehicles (Nets4Cars) is a series of workshops that provides an international forum on the latest technologies and research in the field of intra-vehicle and inter-vehicle communications. The workshops are organized annually to present original research results in all areas related to physical layer, communication protocols and standards, mobility and traffic models, experimental and field operational testing, and performance analysis.

First launched by Tsutomu Tsuboi, Alexey Vinel, and Fei Liu in Saint Petersburg, Russia in 2009, Nets4Cars workshops have been held in Newcastle-upon-Tyne, UK (2010), Oberpfaffenhofen, Germany (2011), Vilnius, Lithuania (2012), Villeneuve d'Ascq, France (2013), and Offenburg, Germany (2014). The 2014 workshop, the seventh in the series, took place at Hotel New Peterhof, Saint Petersburg, Russia, on 6-8 October 2014, with the technical support of the V. A. Trapeznikov Institute of Control Sciences of the Russian Academy of Sciences, Russia, and Halmstad University, Sweden. The technical sponsors of the event were the IEEE Russia (Northwest) Section BT/CE/COM Joint Chapter and IFIP WG 6.8. Open call for papers resulted in 18 submissions. Each of the papers was assigned to the Technical Program Committee members or external reviewers, with each paper receiving at least three independent reviews. A total of 11 papers were accepted for publication, and all were presented at the workshop and are available in IEEE Xplore.

The Best Paper Award Ceremony: Tetiana Zinchenko (Volkswagen AG, Germany) and Alexey Vinel (Halmstad University, Sweden).

The keynote speakers were Antonio Bicchi (University of Pisa, Italy), Panos Papadimitratos (KTH, Sweden), and Evgeny Belyaev

(Continued on Newsletter page 4)

MEMBER RELATIONS/*Continued from page 1*

the Mentorship Award. These awards were presented at the WICE meeting held at IEEE Globecom 2014, in Austin, Texas.

Stefano: And what about Women's Workshops?

Octavia: In July 2014, WICE organized its Second Women's Workshop on Communications and Signal Processing at Princeton University. The event was a success, representing not only a unique opportunity for both junior and senior women researchers to interact about the new developments in their fields, but also an excellent possibility for mentoring and networking. There were 33 participants from the USA, Canada, Turkey, UK, UAE, and Sweden, among which 19 were junior attendees. The junior participants had poster presentations, while the senior participants gave invited talks. A best poster presentation award, sponsored by ComSoc, was presented for the first time. Different workshop attendees shared their personal stories, related to both career and family, which are available as interviews on the ComSoc Beats website (http://beats.comsoc.org/).

Stefano: You mentioned that you aim at achieving a wider participation of women in ComSoc activities by direct advertising and via the social media.

Octavia: Our members are informed about the WICE activities and events through the WICE website, Facebook, and LinkedIn, as well as through the mailing list. Last year we created the mailing list, which has more than 4,000 subscribers already. If you are interested in becoming a member of WICE, please subscribe to our mailing list. Information about that is provided on our website (committees.comsoc.org/WICE).

Stefano: This year we attempted to launch a new initiative, which attracted utmost interest by everyone: child care at ComSoc conferences. Regrettably, we experienced some difficulties which hindered our plans. Would you please summarize what is the current situation?

Octavia: We believe that providing child care at ComSoc conferences represents a valuable service for our members, both male and female. For this reason, we discussed with IEEE the possibility of launching the program through a pilot project at ICC 2015. Recently, IEEE advised that the attendees should contact corresponding services directly. Hopefully, in the future offering such a service to our ComSoc members will be possible.

Stefano: You are also working to promote highly-qualified participation of women among candidates for distinguished positions in ComSoc.

Octavia: Another initiative is to nominate outstanding WICE members for the Distinguished Lecturer Program. An announcement about nominations will be sent to our members via the mailing list. We also aim to increase the number of IEEE Senior and Fellow members from the WICE community, by nominating suitable candidates.

Stefano: What is your final call to ComSoc members?

Octavia: When different voices are represented in a group, everyone benefits. ComSoc recognizes that diversity is essential to the technical community and to the greater society. WICE welcomes participation from all ComSoc members. If you are interested in contributing, we would like to hear from you.

GCN GLOBAL COMMUNICATIONS NEWSLETTER

STEFANO BREGNI
Editor
Politecnico di Milano – Dept. of Electronics and Information
Piazza Leonardo da Vinci 32, 20133 MILANO MI, Italy
Tel: +39-02-2399.3503 – **Fax:** +39-02-2399.3413
Email: bregni@elet.polimi.it, s.bregni@ieee.org

IEEE COMMUNICATIONS SOCIETY

STEFANO BREGNI, VICE-PRESIDENT MEMBER RELATIONS
PEDRO AGUILERA, DIRECTOR OF LA REGION
MERRILY HARTMANN, DIRECTOR OF NA REGION
HANNA BOGUCKA, DIRECTOR OF EAME REGION
WANJIUN LIAO, DIRECTOR OF AP REGION
CURTIS SILLER, DIRECTOR OF SISTER AND RELATED SOCIETIES

REGIONAL CORRESPONDENTS WHO CONTRIBUTED TO THIS ISSUE

ALEXEY VINEL, FINLAND (ALEXEY.VINEL@GMAIL.COM)
EWELL TAN, SINGAPORE (EWELL.TAN@IEEE.ORG)

A publication of the IEEE Communications Society

www.comsoc.org/gcn
ISSN 2374-1082

IOT 2014/*Continued from page 2*

Following are the organizers of IOT 2014

General chairs: Houbing Song, West Virginia University & West Virginia Center of Excellence for Cyber-Physical Systems; Latif Ladid, ComSoc Internet of Things Subcommitte; Bin Xia, Shanghai JiaoTong University.

Technical Program Chairs: Shengjie Zhao, Tongji University; Qinghe Du, Xi'an Jiaotong University; Xiaohua Tian, Shanghai JiaoTong University

IOT 2014 received 18 papers and accepted eight papers. Papers were submitted from three continents: Asia, North America, and Europe. These papers covered the following topics:

- Theoretical Foundations of Cyber-Physical Systems/Internet of Things
- Modeling, Analysis and Synthesis Techniques
- Architectures for Cyber-Physical Systems/Internet of Things
- Building Blocks for Cyber-Physical Systems/Internet of Things
- Systems Abstractions, Services and OS Support
- Evaluation Approaches and Metrics
- Novel Cyber-Physical Systems/Internet of Things Applications
- Detailed Case Studies
- Security/Privacy

IOT 2014 also featured two invited talks by Prof. Bijan Jabbari, professor in the Department of Electrical and Computer Engineering at George Mason University, on the topic "Ad Hoc Networks: The Enabler for the IoT"; and Prof. Hengchang Liu, assistant professor at the School of Computer Science and Technology at the University of Science and Technology of China, on the topic "A Penalized Maximum Likelihood Approach for M-Year Precipitation Return Values Estimation with Lattice Spatial Data".

IOT 2015 will be held in Shenzhen, China in conjunction with ICCC 2015 (http://www.ieee-iccc.org/2015/).

NETS4CARS 2014/*Continued from page 3*

(Tampere University of Technology, Finland). The Best Paper Award was presented to Tetiana Zinchenko (Volkswagen AG, Germany) for her talk entitled "Reliability of Vehicle-to-Vehicle Communication at Urban Intersections".

We invite all the experts in the field of vehicular networking to join us in Sousse, Tunisia for Nets4Cars/Nets4Trains/Nets4Planes-2015 in May 2015, and in Munich, Germany for Nets4Cars-2015-Fall in October 2015. Visit: www.nets4cars.eu

IEEE ICC 2015 EXPLORES "SMART CITY & SMART WORLD" INNOVATIONS FROM JUNE 8–12 IN LONDON'S TECH CITY

The 2015 IEEE International Conference on Communications, the premier international venue dedicated to the worldwide advancement of wireless and wireline communications, will host its next annual event from 8-12 June 2015 at the ExCel London convention centre in London, UK. In addition to experiencing the numerous local iconic attractions like Canary Wharf, The Olympic Park, Westfield Stratford City, The O2 Arena and Museum of London, attendees will also have the ability to explore London's Tech City, the fastest growing technology cluster in Europe and the heart of the nation's rapidly accelerating digital economy.

Themed "Smart City & Smart World," IEEE ICC 2015 will be highlighted by more than 1,800 presentations, panels and forums as well as the exhibits and patronage of numerous technology industry leaders including Qualcomm, Huawei, Nokia, GENBAND, Keysight, InterDigital, P.I. Works, Airvana, SoliD, Imperial College London, DIGILE, Fore-Mont, Anritsu, Artech House, Springer, Cambridge University Press, Willey, EE, National Instruments and Three UK.

On Monday, June 8, the five-day international event will officially commence with the first of two full days of tutorials and workshops addressing topics like 5G Evolution and Candidate Technologies, Massive MIMO for 5G, Dedicated Short Range Vehicular Communications, Communication Architectures and Networking for Electric Vehicles in the Smart Grid, Wireless Physical Layer Security, Device-to-Device Communication for Cellular and Wireless Networks and 5G and Beyond: Enabling Technologies and Applications and Visible Light Communications. For instance, the session on Green Cloud Networks will detail measures for reducing cloud network power consumption and carbon footprints at a given power consumption level, while optimizing the use of renewable energy. In addition, the tutorial on Energy Harvesting and Energy Cooperation in Wireless Communications will focus on the new breed of "energy harvesting" wireless networks and the fundamentals for building next generation green and energy self-sufficient communication systems.

The conference will then proceed over the next three days with a comprehensive agenda exploring the entire communications spectrum ranging from mobile cloud computing and cooperative intelligent transport communications to social and Internet of Things (IoT) networking services and applications. This will be punctuated by the keynote addresses of industry experts like Dr. Paul E. Jacobs, Executive Chairman of Qualcomm Incorporated, who will discuss the "Mobile-Powered Future" and the way mobile device, network and service innovations are impacting consumer electronics, automotive, health care, robotics, and smart cities as well as creating intelligently-connected mobile ecosystems. Other notable authorities delivering keynotes at the event include:

- Professor H. Vincent Poor of Princeton University, who will speak on "Smart Grid: The Role of the Information Sciences" and the methods for improving the efficiency and lower the cost of power use and distribution, and enabling the effective integration of renewable energy sources and distributed storage into the grid
- Dr. Wen Tong, CTO, Huawei Wireless, who will focus on "5G to Embrace the Vertical Industries," the creation of new radio access capabilities needed to meet diverse requirements from different applications and the evolution of network virtualization and slicing technologies necessary for enabling single and unified network architectures
- Professor Alwyn Seeds of the University College London, who will address "Wireless over Fibre Systems: from MHz to THz" and how WoF system technologies and photonic techniques can enable ultra-high capacity wireless data transmission using signals at millimetre-wave and TeraHertz (THz) frequencies
- Professor Xiaohu You of Southeast University, who will discuss "5G Mobile Communications in China" and the nation's mobile communications strategies and key techniques for enabling 5G radio transmissions, networking and open architecture testbeds based on massive cooperative cloud radio capabilities
- Professor Lajos Hanzo of the University of Southampton, who will speak on "A Stroll with Shannon to Next-Generation Plaza: Large-Scale MIMOs, Single versus Multiple RF Chains and All That..." This will involve the pros and cons of coherent versus non-coherent large-scale MIMO systems and the benefits and disadvantages of their multi-functional antenna array based and spatial modulation aided manifestations

In addition, Tuesday through Thursday, will also be highlighted by a wide selection of industry panels and business forums detailing the latest advances in communications and policies. As an example, Bryn Jones, Chief Technology Officer at Three UK will chair the session titled "How far can we evolve Mobile Networks - What is next?" During this forum, the CTOs from Three, EE, O2, Qualcomm and Huawei will provide their views on the steps required to overcome data traffic challenges and offer higher capacities and more reliable data services, while ensuring a high Quality of Service.

Other milestones include the conference's technical program comprised of the presentation of more than 1250 original papers exploring topics within next-generation, mobile and wireless, ad hoc and sensor networking; signal processing and wireless communications: communications theory; optical networks and systems; communication QoS, reliability and modeling; communication software, services and multimedia applications; communications and information systems security; and cognitive radio and networks.

The last day of IEEE ICC 2015 on Friday, June 12 will then offer a second full day of tutorials and workshops addressing subjects like Network Coding: From Theory to Practice, Emerging Concepts and Technologies Toward 5G Wireless Networks, The Path Towards 5G – Essential Technologies, Protocols and Tools for Enabling 5G Mobile Communications, Game Theory for Future Wireless Networks: Challenges and Opportunities, Android Security, Massive Uncoordinated Access Protocols, Dependable Vehicular Communications (DVC) and Next Generation Green ICT.

For ongoing updates on IEEE ICC 2015, please visit http://www.ieee-icc.org/2015. All website visitors are also invited to network with colleagues and peers, share their professional experiences through the conference's Facebook, LinkedIn and Twitter pages.

SECURITY AND PRIVACY IN EMERGING NETWORKS: PART 1






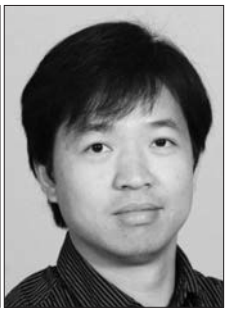

Mohsen Guizani *Daojing He* *Kui Ren* *Joel Rodrigues* *Sammy Chan* *Yan Zhang*

With the recent advancements in networking technologies, some new emerging networks are being implemented that have the potential to be deployed broadly and on a large scale in the near future. In the wired domain, these emerging networks include, for example, networks based on software-defined networking (SDN) and named data networking (NDN). In the wireless domain, they include mobile and wireless networks involving handheld computing devices, sensors and RFID devices, body area sensor networks, and participatory sensing networks, to name a few.

Although these kinds of networks have attracted much research effort, the security and privacy issues have not been studied well so far. Thus, there is an urgent need to protect these networks from various security and privacy threats. This could pave the way to implement these networks without major security obstacles. This Feature Topic aims to promote further research interest in security and privacy in emerging networks by providing a vehicle for researchers and practitioners to discuss research challenges and open issues, and disseminate their latest research results. We received an overwhelmingly large number of high-quality submissions (40 papers) out of which we accepted only the top 13 articles. We are lucky to get the permission of the Editor-in-Chief of *IEEE Communications Magazine* to divide the accepted papers into two parts. Part 1 will be composed of six manuscripts that deal with the theory of security and privacy threats of emerging networks, and the second part is composed of seven papers addressing the same issues with more specific applications of security and privacy. Part 2 is scheduled to appear in June 2015. We invite you to stay tuned and check that issue as it will complement the topics discussed in this first part.

The first article, by Yang *et al.*, "Safeguarding 5G Wireless Communication Networks Using Physical Layer Security," attempts to shed some light on the physical layer security related to fifth generation (5G) mobile and wireless networks. The authors examine some inherent vulnerabilities in 5G wireless networks and focus on three main technologies: heterogeneous networks, massive multiple-input multiple-output, and millimeter-wave. They identify possible opportunities and challenges in these technologies and warn security designers of the possible problems that could exist and must be tackled.

On the other hand, 5G mobile networks use densified small cell deployment with overlay coverage through coexisting heterogeneous networks (HetNets). This type of multi-tier architecture along with stringent latency requirements in 5G bring new challenges in security provisioning due to the potential frequent handovers and authentications. In the second article, the authors overview related studies and introduce SDN into 5G as a platform to enable efficient authentication handover and privacy protection. Thus, "Authentication Handover and Privacy Protection in 5G HetNet Using Software-Defined Networking" by X. Suan and X. Wang attempts to simplify authentication handover by sharing the user-dependent security context information among related access points. They demonstrate that SDN-enabled security solutions are highly efficient when using a centralized controlling capability.

The growth of software defined networks (SDNs) promises to dramatically simplify network management and enable innovation through network programmability. However, security is expected to remain the main impediment to SDNs' growth. This is due in part to the fact that security is not considered as part of the initial SDN design. The third article, "Securing the Software Defined Networks: Taxonomy, Requirements, and Open Issues" by A. Akhunzada *et al.*, discusses the state-of-the-art security solutions in order to overcome those challenges. The authors classify the existing security solutions based on SDN layers/interfaces, security measures, simulation environments, and security objectives. They then point out possible attacks and threats targeting SDNs with potential key security requirements. Finally, open issues and challenges of SDN security are presented that may be deemed appropriate for researchers to address in order to help SDNs achieve their potential goals.

Along the same line, Zhou *et al.* conceived a novel conceptual network security mechanism called the evolving defense

mechanism (EDM). In their contribution, "Evolving Defense Mechanism for Future Network Security," they show that EDM is an inspiration of a network configuration originating from a biological gene. They provide an overview of EDM and argue that it is able to avoid deficiencies of conventional network security approaches. They first discuss dynamic network configuration for preventing attacks and then sketch a way to implement EDM as an ideal framework based on SDN serving as an ecosystem and coexisting environments.

The next article, "Distributed Denial of Service Attacks in SDN with Cloud Computing," may help us make full use of SDN's advantages to defeat DDoS attacks in cloud computing environments. The authors, Q. Yan and R. Yu, first discuss the new trends and characteristics of DDoS attacks in cloud computing environments. Then they show that SDN brings new opportunities and special features in defeating DDoS attacks. They finally present a number of challenges that need to be addressed to mitigate DDoS attacks when SDN is combined with cloud computing.

In the final article of Part 1 of this Feature Topic, "De-Anonymizing and Countermeasures in Anonymous Communication Networks," M. Yang *et al.* classify and provide an overview of existing de-anonymizing techniques and propose countermeasures to mitigate those risks.

We are confident that this selection of high-quality articles will provide some research directions in the field. While most of the above articles discuss SDN security, there are plenty of issues that have been presented that will need more focus and attention to be developed for emerging networks. We strongly believe that all of us (from multiple disciplines) have to join our efforts, and must come together and strive hard to overcome technical roadblocks in order to bring the vision of emerging network security to reality.

The Guest Editors would like to thank the outgoing Editor-in-Chief (Sean Moore) and the incoming Editor-in-Chief (Osman Gebizlioglu) for their guidance, feedback, and encouragement along the way. We are very grateful to them for allowing us to schedule two issues of the Feature Topic due to the large number of submissions received from highly qualified researchers. We also thank the *IEEE Communications Magazine* Publications staff for their patience and hard work in making this issue a reality.

BIOGRAPHIES

MOHSEN GUIZANI [S'85, M'89, SM'99, F'09] (mguizani@ieee.org) is currently a professor and associate vice president of Graduate Studies at Qatar University. He previously served as Chair of the Computer Science Department at Western Michigan University, 2002–2006, and Chair of the Computer Science Department at the University of West Florida, 1999–2002. He received his B.S., M.S., and Ph.D. degrees in electrical and computer engineering all from Syracuse University, New York. His research interests include wireless communications and mobile computing, cloud computing, cyber security, and smart grid. He is the author of nine books and more than 400 publications in refereed journals and conferences He served as an IEEE Computer Society Distinguished Speaker from 2003 to 2005. He is a member of Computer Societies and ASEE, and a Senior Member of ACM.

DAOJING HE (hedaojinghit@gmail.com) received his B.Eng.(2007) and M.Eng. (2009) degrees from Harbin Institute of Technology, China, and his Ph.D. degree (2012) from Zhejiang University, China. He is currently a professor at the Software Engineering Institute, East China Normal University. His research interests include network and systems security. He is an Associate Editor or on the Editorial Boards of a number of international journals such as *IEEE Communications Magazine*.

KUI REN (kuiren@buffalo.edu) is an associate professor at the State University of New York at Buffalo. Hiss research interest spans cloud and outsourcing security, and wireless and wearable security. His research has been supported by NSF, DoE, AFRL, MSR, and Amazon. He was a recipient of an NSF CAREER Award in 2011 and a Sigma Xi/IIT Research Excellence Award in 2012. He is an Associate Editor for IEEE TMC, TIFS, TSG, and others. He is a Distinguished Lecturer of IEEE.

JOEL RODRIGUES [S'01, M'06, SM'06] (joeljr@ieee.org) is a professor in the Department of Informatics of the University of Beira Interior, Covilhã, Portugal, and a researcher at the Instituto de Telecomunicações, Portugal. He is the leader of the NetGNA Research Group (http://netgna.it.ubi.pt), Chair of the IEEE ComSoc TC on eHealth, Past Chair of the IEEE ComSoc TC on Communications Software, and a Steering Committee member of the IEEE Life Sciences Technical Community. He is the Editor-in-Chief of three international journals, and a co-author of over 400 papers, two books, and three patents. He is the recipient of several Outstanding Leadership and Outstanding Service Awards from the IEEE Communications Society and several best paper awards.

SAMMY CHAN [S'87, M'89] (eeschan@cityu.edu.hk) received his B.E. and M.Eng.Sc. degrees in electrical engineering from the University of Melbourne, Australia, in 1988 and 1990, respectively, and a Ph.D. degree in communication engineering from the Royal Melbourne Institute of Technology, Australia, in 1995. He is an associate professor in the Department of Electronic Engineering, City University of Hong Kong.

YAN ZHANG (yanzhang@simula.no) received a Ph.D. degree from Nanyang Technological University, Singapore. Since August 2006, he has been working with Simula Research Laboratory, Norway. He is currently head of the Department of Networks and an adjunct associate professor at the Department of Informatics, University of Oslo, Norway. He is a Regional Editor, Associate Editor, on the Editorial Board, or Guest Editor of a number of international journals. His recent research interests include wireless networks, cyber physical systems, and smart grid communications.

Safeguarding 5G Wireless Communication Networks Using Physical Layer Security

Nan Yang, Lifeng Wang, Giovanni Geraci, Maged Elkashlan, Jinhong Yuan, and Marco Di Renzo

ABSTRACT

The fifth generation (5G) network will serve as a key enabler in meeting the continuously increasing demands for future wireless applications, including an ultra-high data rate, an ultra-wide radio coverage, an ultra-large number of devices, and an ultra-low latency. This article examines security, a pivotal issue in the 5G network where wireless transmissions are inherently vulnerable to security breaches. Specifically, we focus on physical layer security, which safeguards data confidentiality by exploiting the intrinsic randomness of the communications medium and reaping the benefits offered by the disruptive technologies to 5G. Among various technologies, the three most promising ones are discussed: heterogenous networks, massive multiple-input multiple-output, and millimeter wave. On the basis of the key principles of each technology, we identify the rich opportunities and the outstanding challenges that security designers must tackle. Such an identification is expected to decisively advance the understanding of future physical layer security.

Nan Yang is with the Australian National University.

Lifeng Wang and M. Elkashlan are with Queen Mary University of London.

Giovanni Geraci is with the Singapore University of Technology and Design.

Jinhong Yuan is with the University of New South Wales.

Marco Di Renzo is with Paris-Saclay University, Laboratory of Signals and Systems (UMR-8506), CNRS - CentraleSupelec - University Paris-Sud XI, 91192 Gif-sur-Yvette (Paris), France.

INTRODUCTION

Mobile wireless communication has experienced an unprecedented growth in data traffic in recent years, spurred by the popularity of various intelligent devices, the demand for exuberant multimedia content, and the rapid increase in the number of base stations (BSs). In particular, global mobile data traffic in 2013 was nearly 18 times the size of the entire global Internet in 2000, and monthly global mobile data traffic by 2018 will surpass 15 exabytes [1]. While the mature third generation network and the currently deploying fourth generation (4G) network may accommodate the data traffic surge for the next few years, they will not be able to support a very large number of devices with a huge network traffic demand in 2020 and beyond [2]. Against this backdrop, a number of disruptive trends and technologies shaping the fifth generation (5G) network are emerging worldwide through research and development. For example, academia is researching robust and efficient wireless transmission technologies for the 5G era, such as the heterogeneous network (HetNet), massive multiple-input multiple-output (MIMO), and millimeter wave (mmWave). At the same time, the industry is undertaking 5G standardization. Given the ubiquitousness and necessity of 5G connections in the near future, an enormous amount of sensitive and confidential information, e.g. financial data, electronic media, medical records, and customer files, will be transmitted via wireless channels. Thus, providing an unrivalled security service is one of the top priorities in the design and implementation of the 5G network.

Despite the current efforts from academia and industry, the security paradigms protecting the confidentiality of wireless communication in the 5G network remain elusive. Indeed, how to secure wireless data transmission is one of the core problems that any 5G network designer can face. Differing from the traditional approach which protects data security through cryptographic techniques, physical layer security is identified as a promising strategy that provides secure wireless transmissions by smartly exploiting the imperfections of the communications medium. Using this strategy, 5G network designers can effectively degrade the quality of signal reception at unauthorized receivers and devices, and therefore prevent them from acquiring confidential information from the received signal. With careful planning and execution, physical layer security will protect the communication phase of the network while cryptography will protect the processed data after the communication phase. As such, they will form a well-integrated security solution that efficiently safeguards sensitive and confidential data for the 5G era.

Notably, physical layer security offers two major advantages compared to cryptography, making it particularly suitable for the 5G network. First, physical layer security techniques do not depend on computational complexity, which implies that the achieved level of security will not be compromised even if the unautho-

0163-6804/15/$25.00 © 2015 IEEE

rized smart devices in the 5G network have powerful computational capabilities. This is in contrast to computation-based cryptography, which is based on the premise that the unauthorized devices have insufficient computational capabilities for hard mathematical problems. Second, physical layer security techniques have a high scalability. In the 5G network, devices are always connected to the nodes with different powers and computation capabilities at the different levels of the hierarchical architecture. Also, devices always join in or leave the network at random time instants, due to the decentralized nature of the network. As a consequence, cryptographic key distribution and management become very challenging. To cope with this, physical layer security can be used to either provide direct secure data communication or facilitate the distribution of cryptographic keys in the 5G network.

Given the potential of physical layer security for the 5G era, the goal of this article is to identify the opportunities and challenges offered by the disruptive technologies enabling 5G for achieving a high security level at the physical layer. Among the various technologies, we focus on the three most promising ones, which we now describe in detail.

- **HetNet**: The HetNet creates a multi-tier topology where multiple nodes are deployed with dissimilar characteristics such as transmit powers, coverage areas, and radio access technologies. Obviously, it offers a rather provocative departure from the conventional single-tier wireless network and creates a new trend to reduce the cost per bit of future wireless connections. In such a trend, the full exploitation of the opportunities offered by the multi-tier topology, such as spatial modeling of nodes, association of mobile users, and direct connection between devices, is a core component in the design of physical layer security. This exploitation is discussed later.
- **Massive MIMO**: By deploying a very large number of antennas (e.g. a few hundred) at BSs to serve many tens of users simultaneously, massive MIMO reaps all the benefits offered by conventional MIMO, but on a much larger scale. To leverage the advantages of massive MIMO in physical layer security, some challenges need to be resolved during the design process, such as pilot contamination, power management, channel reciprocity, and eavesdropper-targeted signal processing. Motivated by this, we argue for physical layer security along with these challenges later.
- **mmWave**: As an innovative solution to meet the 5G's requirement, mmWave communication systems use a huge swath of spectrum, from 30 GHz to 300 GHz, to shift wireless transmissions away from the nearly fully occupied spectral band for current wireless networks. Notably, mmWave technologies have been standardized for short-range transmission, e.g. IEEE 802.11ad, as well as deployed for small cell backhaul, e.g. Siklu's Etherhaul 1200T. Since secure mmWave transmission is a completely new and promising research frontier, we advocate the potential of mmWave communication for physical layer security in a later Section.

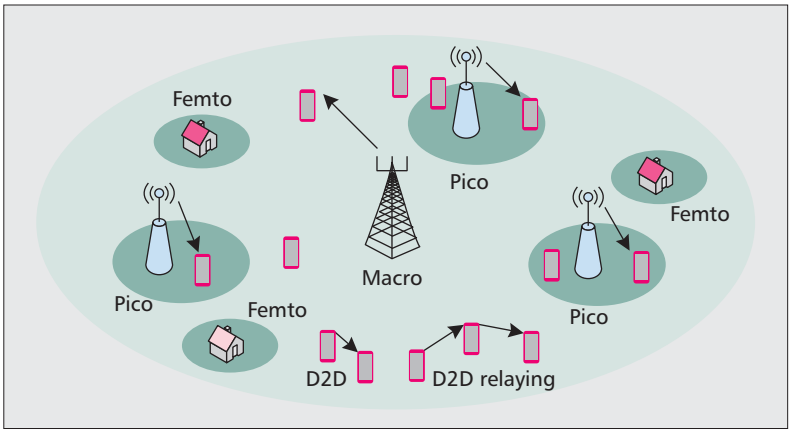

Figure 1. Heterogeneous network.

Physical Layer Security in Heterogeneous Networks

The HetNet is a promising network densification architecture in the 5G era. The aim of the HetNet is to provide a spectrum-efficient and energy-efficient solution that satisfies the dramatic growth in data demands of future wireless applications. In the HetNet, nodes with different transmit powers, coverage areas, and radio access technologies are deployed to form a multi-tier hierarchical architecture, as depicted in Fig. 1. Specifically, high-power nodes (HPNs) with large radio coverage areas are placed in *macro cell* tiers, while low-power nodes (LPNs) with small radio coverage areas are placed in *small cell* tiers. Small cells, such as pico cells and femto cells, are deployed under macro cell umbrellas to augment indoor coverage in highly populated buildings, and multi-tenant dwelling units, enterprises, and outdoor coverage in dense urban, suburban, or rural areas. In addition to the macro cell and small cell tiers that support HPN-to-device and LPN-to-device communications, respectively, the HetNet also involves a *device* tier that supports device-to-device (D2D) communications. The D2D communication allows geographically close devices to directly connect and interact with each other without using HPNs/LPNs, thus being a powerful enabler of low-latency and high-throughput data applications. Among multiple tiers, different radio access technologies such as WCDMA, LTE, WiMAX, and WLAN are adopted to provide various communication services. Therefore, the HetNet is a clearly different paradigm from conventional macrocell-only wireless networks. Obviously, the current physical layer security technologies for the direct single-user and multi-user transmissions [3, 4] and the relay-aided communications [5] in conventional networks cannot be readily applied in this paradigm. It follows that the compelling potential of the HetNet will trigger a new wave of innovation — in terms of spatial modeling, mobile association, and device connection — in securing multi-tier communications. These innovations are detailed in the following subsections.

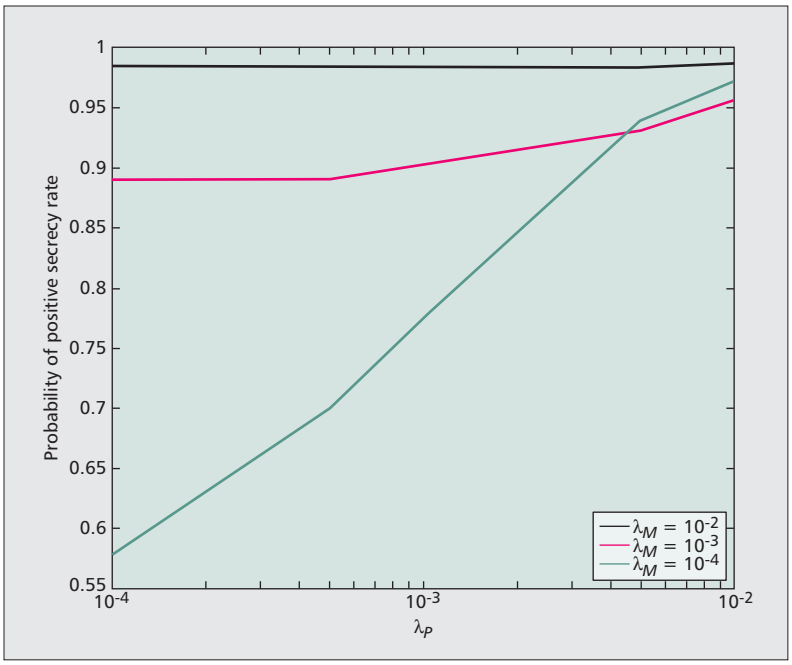

Figure 2. The probability of positive secrecy rate in a two-tier network where a macro cell tier is overlaid with a pico cell tier. The marco and pico cells are assumed to share the same frequency band. The locations of macro cell HPNs and pico cell LPNs follow independent homogeneous Poisson point processes (HPPPs) with densities λ_M and λ_P, respectively. The location of eavesdroppers also follows a HPPP with density $\lambda_E = 8 \times 10^{-4}$. User association is based on the maximal average received signal-to-noise ratio.

Spatial Modeling

In the HetNet, different spatial modeling of HPNs' and LPNs' locations raises a natural question: "*How does the spatial modeling of nodes' locations affect and guide the physical layer security design?*" The rationale behind this question is that a HPN's location is currently modeled as a point at the center of a hexagonal grid, while LPNs' locations can be modeled as a uniform distribution, in particular a Poisson point process in the two-dimensional plane [6]. Evidently, the deterministic model for HPNs' locations provides no randomness, whereas the Poisson model for LPNs' locations corresponds to complete randomness. This necessitates different mathematical tools to be employed to accommodate the nature and properties of the two models into the design of physical layer security.

Due to the deterministic nature of the model, the impact of HPNs' locations on physical layer security can be evaluated by applying system-level simulations to approximate the distributions of signal-to-interference-plus-noise ratios (SINRs) and corresponding quality of service (QoS) parameters. When small cells are deployed as add-ons, however, the amount of interference grows dramatically. Accordingly, the complexity of simulations substantially increases, making performance evaluation and optimization more complex and time-intensive. More recently, the HPNs' locations have also been modeled as a PPP to gauge the degree of randomness [7, 8]. Against this background, network security designers need to develop effective methods rooted in probability processes and order statistics to characterize the SINR distributions and the QoS parameters from the theoretical perspective. In particular, the generalized methods characterizing channel dynamics such as signal fading, spatial correlation, practical path loss, random channel errors, and mobility-induced channel variations are of paramount importance since they reduce the simulation burden to the minimum. The development of the generalized methods will allow for accurate capture of the multi-tier structure of the HetNet.

The effect of HPNs' and LPNs' locations on physical layer security can be examined based on the knowledge from the fields of graph theory and stochastic geometry. This motivates network security designers to develop effective mathematical tools from these two areas such that the SINR distributions and the QoS parameters impacted by HPNs and LPNs can be characterized. Of course, the characterization needs to be as general as possible, which enables the unequivocal establishment of secure connectivity and the accurate assessment of secrecy capacity in the HetNet with arbitrary transmit powers and densities of HPNs and LPNs. Moreover, new solutions to the SINR distributions are required if an appropriate level of practical correlation is introduced into the placement of HPNs and LPNs. Such solutions will force a substantial advancement over the current studies relying on the assumption of independent placement of HPNs/LPNs.

Figure 2 evaluates the impact of HPNs' and LPNs' densities on the secrecy performance via simulations. In this figure we show the probability of positive secrecy rate in a two-tier network where a macro cell tier is overlaid with a pico cell tier. It is evident that the probability of positive secrecy rate[1] increases as the density of pico cell LPNs increases. Moreover, it is observed that if the density of pico cell LPNs increases beyond a critical point, a higher density of macro cell HPNs does not improve the secrecy performance any further. Therefore, Fig. 2 provides a guide for network security designers to decide the best density for the implementation of HPNs and LPNs in the HetNet. Of course, as previously discussed, the development of effective mathematical tools will enable us to undertake the evaluation involved in this figure in a computationally efficient manner.

Mobile Association

Associating mobile users with HPNs and LPNs leads to a challenging and promising question: "*What is the optimal strategy for users to select HPNs/LPNs under security constraints?*" In traditional macro-cell-only cellular networks, it is typically assumed that mobile users select the strongest HPN to connect such that the best channel quality with the highest SINR is obtained. Accordingly, the physical layer security technologies in the open literature are designed based on this assumption. However, in the multi-tier HetNet, such a selection causes a load balancing problem. This is due to the fact that the HPNs with high transmit power and large coverage areas are often fully loaded or even "over" loaded, whereas the LPNs with low transmit

[1] Here, the probability of positive secrecy rate reveals the probability that the secrecy rate is higher than zero, where the secrecy rate can be characterized by the difference between the capacity of the main channel and the capacity of the eavesdropper's channel.

power and small coverage areas are often very lightly loaded [6]. Such an unbalanced load is detrimental to the ubiquitous applications of real-time services with stringent delay constraints and high power consumption, e.g. streaming video and gaming. As such, the unbalanced load should be addressed in the design of physical layer security.

In order to secure transmission and overcome the unbalanced load problem, new security-oriented mobile association policies are required to monitor and balance the instantaneous load of HPNs and LPNs. In designing these policies, the optimization of secrecy performance, e.g. the secrecy rate and the secrecy outage probability, should be prioritized. Under this prioritization, some intelligent mobile association policies can be developed such that the mobile users are wisely assigned to some HPNs or LPNs based on the achievable secrecy performance, the instantaneous load, and other factors such as the transmit power, coverage area, and density of HPNs/LPNs. Considering that such intelligent and optimal policies would impose a high computational complexity, some simple yet suboptimal mobile association policies are required to achieve a complexity-quality tradeoff. Aided by these suboptimal policies, the near-optimal secrecy performance is guaranteed with a lower computational cost. In addition, the cooperation among HPNs and LPNs offers a feasible way to enhance the secrecy performance. To explore this feasibility, network security designers should develop new cooperative strategies to allow neighboring HPNs/LPNs to exchange the secure data for mobile users, the instantaneous load of themselves, and other factors of the network with each other for achieving close-to-maximum secrecy performance.

Device Connection

The introduction of D2D communication triggers a pertinent question on the security issue: "*How to protect data confidentiality between connected devices against data leakage?*" Doubtlessly, maintaining data security is an essential task in D2D communications since the transmitted data between connected devices may be overheard by all of the surrounding devices. This task becomes more arduous particularly given the fact that the connected devices may not be able to handle complex signal processing algorithms as HPNs and LPNs do. One possible solution to tackle this task is *closed access* [9], where the intended device has a list of "trusted" devices. In closed access, the non-listed devices can only communicate with the intended device by getting authenticated in the macro cell or the micro cell tier. Hence, the establishment of closed access safeguards the data exchange between the intended device and the "trusted" devices against eavesdropping.

It is worth noting that closed access may not always be implemented, due to the lack of authentication in the macro cell or the micro cell tiers. In this case, referred to as *open access*, not only surrounding devices but geographically close HPNs and LPNs may act as potential eavesdroppers for the connected devices, meaning that they benefit from listening to the transmitted data and pose an acute threat to data security. To address security issues in open access, network designers need to construct new secure data exchange strategies that fully consider the physical characteristics of unintended devices and malicious HPNs/LPNs, e.g. ambiguous location, uncertain mobility, and unknown configuration. In addition, the potential attacks and threats induced by unintended devices and malicious HPNs/LPNs need to be carefully analyzed and incorporated into the construction.

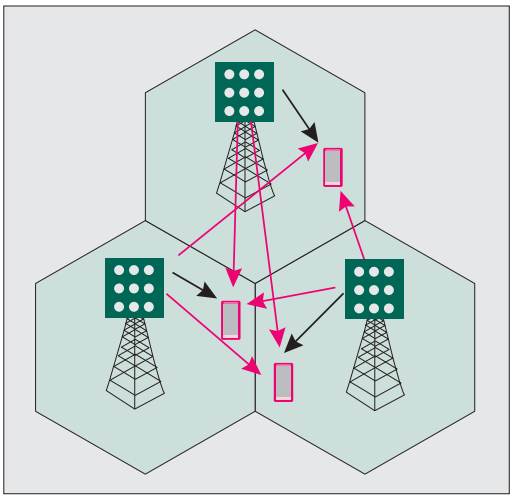

Figure 3. Cellular network with the deployment of massive MIMO.

Apart from direct D2D connections, another interesting paradigm in D2D communications is D2D relaying [10] where a device having better geometry to the transmitter device may act as a relay for the receiver device. The physical layer security design in this paradigm should exploit cooperative spatial diversity to maximize the secrecy performance. Despite the current relay-aided physical layer security techniques, such a design introduces new security problems to be solved. For example, the optimal selection of candidate relays need to be determined and the protection against untrusted relays needs to be investigated. Furthermore, if multiple devices are required for relaying the data between the connected devices, the impact of multi-hop coordination on the secrecy performance needs to be examined.

Physical Layer Security in Massive MIMO Systems

Massive MIMO systems are emerging as a new research field and have attracted significant interest from both scientists and industrialists. The benefits of the massive MIMO technique are realized by using very large antenna arrays (typically tens or even hundreds) at the transmitter and/or the receiver. In future cellular networks with massive MIMO, as depicted in Fig. 3, the number of antenna arrays at the BSs is much larger, e.g. 10 times, than the number of data streams served to all users in a cell [11]. Compared to the current counterpart, massive MIMO systems provide high power and spectrum effi-

Massive MIMO systems are emerging as a new research field and have attracted substantial interests from both scientists and industrialists. The benefits of the massive MIMO technique are realized by using very large antenna arrays (typically tens or even hundreds) at the transmitter and/or the receiver.

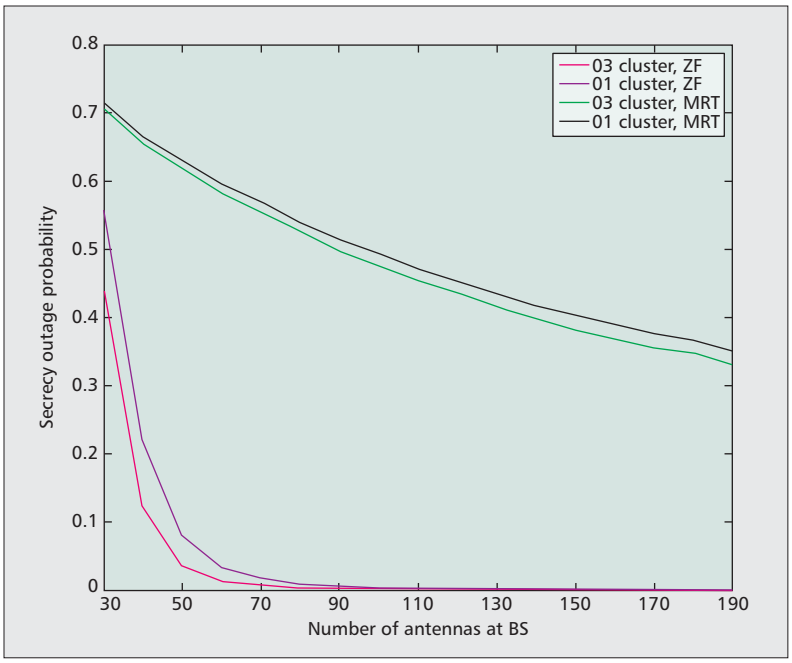

Figure 4. The secrecy outage probability of downlink with MRT and ZF precoding techniques at the BS. We consider three hexagonal cells without sectorization, where the radius of each cell is 350 meters. In each cell, seven single-antenna users are uniformly distributed and one of them is an eavesdropper. In this figure "01 cluster" means that the frequency reuse pattern is 1 while "03 cluster" means that the frequency reuse pattern is 3.

ciencies by exploiting the large arrays gain offered by low-complexity transmission designs. Moreover, random impairments such as small-scale fading and noise can be averaged out when a large number of antennas are deployed at the BS [12]. Furthermore, the interference, channel estimation errors, and hardware impairments [13] vanish when the number of antennas grows large, leaving only pilot contamination as the performance limit [14].

Given the fact that massive MIMO will serve as an essential enabling technology for the 5G wireless network, we next discuss the design of physical layer security based on the tremendous potential of massive MIMO systems. Needless to say, such a design opens a new and promising research avenue, extending current research efforts in conventional MIMO systems to a new area.

Low Power Consumption

In massive MIMO systems, the secrecy performance can be remarkably enhanced by adopting a reduced power consumption. The enhancement is attributed to two major causes. First, since the transmit power level is cut, the receive signal-to-noise ratios (SNRs) at the eavesdroppers are highly reduced. This leads to a significant decrease in the eavesdroppers' channel capacities. Second, given the transmit power and the expected secrecy rate at the transmitter, the secrecy outage probability can be arbitrarily small when the number of antennas grows unbounded.[2] In Fig. 4 we show the secrecy outage probability for a rate threshold of 2 bits/s/Hz in the downlink of a three-hexagonal-cell network. We consider two commonly-used precoding methods: maximal ratio transmission (MRT) and zero-forcing (ZF) [11]. It is seen from this figure that ZF outperforms MRT because the intra-cell interference can be cancelled through ZF. It is also seen that the secrecy outage probability profoundly declines when the number of antennas at the BS grows large. As such, the minimum power consumption achieving the target secrecy performance level needs to be determined. In this determination, the development of new and powerful mathematical tools, e.g. random matrix theory, will eliminate the burden of performance evaluation incurred by time-consuming simulations.

[2] Eavesdroppers are typically passive to hide their existence. As such, the secrecy outage is addressed as a principal concern of security.

Time Division Duplex Operation

Massive MIMO systems are recommended to operate in a time division duplex (TDD) mode [11], which is different from conventional MIMO systems that generally operate in a frequency division duplex (FDD) mode. This is due to the fact that the channel training overhead in the FDD mode scales linearly with the number of transmit antennas, which in turn imposes a severe limit on the number of antennas. In the TDD mode, the training burden is independent of the number of BS antennas and channel reciprocity is exploited. In TDD massive MIMO systems, eavesdroppers may experience particular difficulties for wiretapping, because downlink pilot signals from the BS to the users are not required in the TDD mode. Specifically, the BS with massive antenna arrays obtains the uplink channel state information (CSI) via uplink pilot signals from the users. It then obtains the downlink CSI relying on the reciprocity between the uplink and downlink. As such, it becomes difficult for eavesdroppers to know the CSI between themselves and the BS, as well as the CSI from other users to the BS. Therefore, how to design secure transmission under the assumption of imperfect (or no) CSI at the eavesdroppers is of practical importance in massive MIMO systems. Moreover, pilot contamination occurs in the TDD mode if the pilot signals employed in different cells are not orthogonal. As such, the effect of an inaccurate channel estimate caused by pilot contamination on the secrecy performance should be understood and counteracted. In addition, the TDD operation requires reciprocity calibration [11]. In practical systems, the hardware chains at the BS and users may not be reciprocal between the uplink and the downlink. This motivates the examination of the impact of improper calibration on the secrecy performance.

Artificial Noise

The deliberate deterioration of the eavesdropper's channel quality in massive MIMO systems is a fruitful avenue to explore. In conventional MIMO systems, the artificial noise (AN)-based transmission has been identified as an effective method to cause interference to the eavesdroppers and degrade their received signals. In massive MIMO systems, new challenges are opened for AN-based transmission. For example, transmitting AN signals in a spatial null space may not be practical since the computation complexity of the null space is extremely high for the

large-dimensional channel matrix. Moreover, random and independent AN is averaged out given the availability of a large number of antennas. Therefore, new AN-based transmission schemes need to be developed. Correspondingly, the optimal power allocation between information signals and AN signals needs to be determined and the achievable secrecy performance needs to be evaluated.

Antenna Correlation

Antenna correlation is a practical challenge underlying the deployment of massive MIMO systems. Specifically, a significant amount of correlation may exist between large antenna arrays, due to either the limited aperture of the antenna array or a lack of scattering. For the uplink transmission, for example, the antenna correlation is experienced in different diversity branches at the BS, due to the non-isotropic antennas with reduced separation. Although the impact of antenna correlation on the secrecy performance of conventional MIMO systems has been revealed, e.g. [15], very little detailed work has specifically been carried out to analyze the effect of antenna correlation on the secrecy performance of massive MIMO systems. The research efforts in this area are of enormous value since they enable us to decide how to compensate for antenna correlation in the uplink and downlink massive MIMO systems.

Confidential Broadcasting

In massive MIMO systems, each BS simultaneously communicates with a large number of users. One challenge to multiuser security is achieving confidential broadcasting in the downlink. In particular, each message needs to be kept confidential from all the users other than the intended one, i.e. each user can be treated as an eavesdropper for all messages other than its own. In order to preserve this confidentiality, a precoder needs to be associated with each data stream not only to limit the interference at other users, but also to limit the information leakage. Designing the optimal precoder often involves optimization problems that can only be solved numerically. More practical and near-optimal precoders are thus required. Therefore, it is pivotal to provide design guidelines and to quantify the optimal achievable secrecy performance of linear precoders that guarantee confidential broadcasting in massive MIMO systems.

Hardware Impairments

In contrast with conventional MIMO systems with ideal hardware, the inexpensive hardware components used by massive MIMO systems may give rise to hardware impairments [13]. Although hardware impairments deteriorate the legitimate receivers' channels, the impact of hardware impairments vanishes asymptotically when large-scale arrays are deployed. We note that the presence of hardware impairments also deteriorates the eavesdroppers' channels, which appears to be beneficial for security enhancement. It is therefore worth investigating physical layer security in massive MIMO systems with non-ideal hardware.

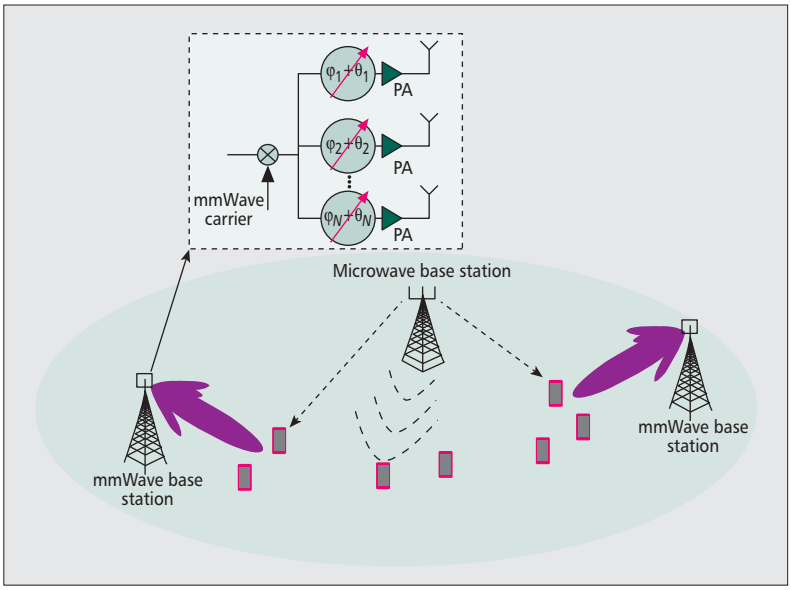

Figure 5. Deployment of mmWave BSs.

Physical Layer Security in Millimeter Wave Communication

Almost all mobile communication systems today restrict their operation to the spectrum in the range of 300 MHz–3GHz. Unfortunately, this spectral band has now become nearly fully occupied. In the 5G network, mmWave communication systems, operating in the frequency range of 30–300 GHz, have been recognized as a promising solution to remove the restriction and meet a thousand-fold capacity increase [16]. As depicted in Fig. 5, mmWave BSs can be deployed with microwave BSs to ensure reliable and fast data transmissions.

Although some efforts need to be made to render the GHz frequency bands available on mobile cellular networks, a series of research initiatives have been undertaken to explore the potential of mmWave communication technologies. Needless to say, security and privacy issues need to be addressed in the implementation of mmWave communication systems. We believe that the investigation of physical layer security in mmWave communication systems is a very promising and highly rewarding area, due to the following factors.

Large Bandwidth: Current maximum aggregated bandwidth in 4G LTE is 20+20 MHz by using carrier aggregation. However, mmWave communication systems provide GHz bandwidths. Therefore, the secrecy outage probability in the passive eavesdropping scenario is remarkably reduced if the transmitter sets a lower transmit secrecy rate in mmWave communications. Also, high secrecy throughput can be obtained with large mmWave bandwidths.

Short-Range Transmission: Compared to the current microwave communication systems, mmWave signals in the higher frequencies experience an increase in free-space path loss by several orders of magnitude. Therefore, only geographically neighboring eavesdroppers are

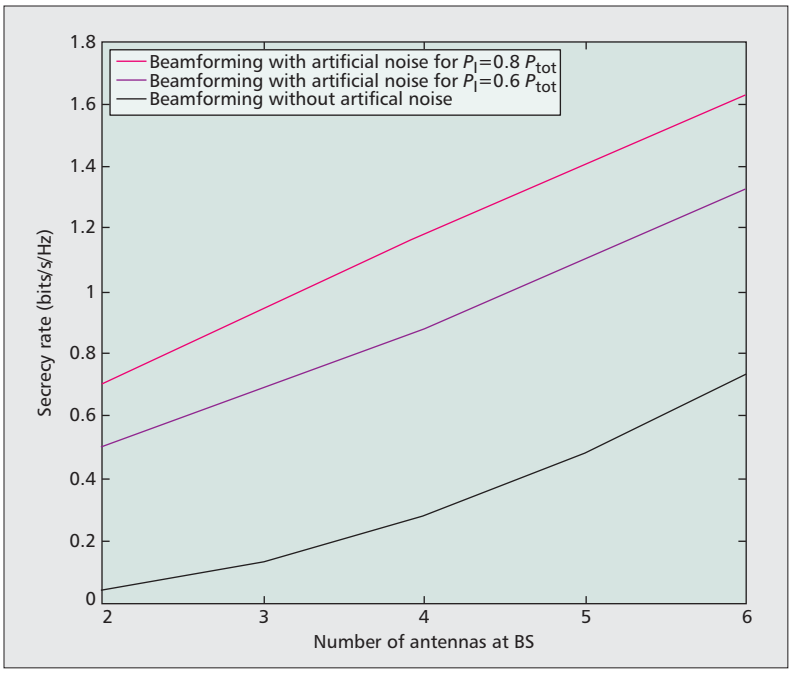

Figure 6. Secure mmWave downlink transmission with AN: An N-antenna mmWave BS transmits the confidential messages to a single-antenna user in the presence of a single-antenna eavesdropper. The BS uses analog beamforming with AN. The total transmit power is P_{tot} = 43 dBm, the power allocated to the information signal is P_I, and the power allocated to the AN signal is $P_{tot} - P_I$.

able to overhear the signals, whereas geographically remote users cannot capture the data transmission.

Directionality: In mmWave systems, highly directional communication with narrow beams is employed for suppressing the interference from neighbors. Therefore, the receive SNRs at the eavesdroppers may be extremely low such that the eavesdroppers are not able to recover information signals from the overheard messages.

Large Antenna Arrays: Large antenna arrays provide high beamforming gains to mitigate the propagation attenuation and save transmit power. In light of the array aperture constraint, current cellular systems[3] in the microwave frequency bands are expected to implement large antenna arrays in a two-dimensional (2-D) or three-dimensional (3-D) array structure. However, 2-D or 3-D arrays increase the coupling effects due to the increase in the number of adjacent antennas [17]. For a fixed array aperture, the shorter wavelengths at the mmWave frequencies enable the mmWave BSs to pack more antennas. Therefore, mmWave systems with large antenna arrays offer a wealth of opportunities at the physical layer security to secure mmWave communication.

Based on the aforementioned factors, the aim of physical layer security design in mmWave communication systems is to fully exploit the potentials of these factors. In this design, several challenging tasks need to be solved. First, the propagation characteristics at higher frequencies need to be precisely modeled. Indeed, an accurate and comprehensive quantification of the impact of path loss, blocking, penetration, and rain absorption on mmWave transmission enables network security designers to theoretically capture the properties of mmWave channels and address these properties in their design. Second, new secure transmission schemes need to be developed. It has been shown that beamforming is a key enabler of mmWave mobile broadband service [18]. Since digital beamforming with a large number of radio frequency (RF) chains incurs a very high implementation cost and power consumption, secure mmWave transmission needs to be designed based on analog beamforming and RF beamforming with a small number of RF chains. Against this background, the transmission of AN signals becomes promising in mmWave communication systems. With the aid of analog beamforming with phase shifters, the beam pattern of AN signals can be easily restricted to the orthogonal direction to the beam pattern of information signals. As depicted in Fig. 6, the secrecy rate is profoundly improved by incorporating AN signals into secure transmission. Moreover, the power allocated to the information signal plays a pivotal role in determining the secrecy performance. For example, beamforming with AN transmission with $P_I = 0.8P_{tot}$ yields a higher secrecy rate than with $P_I = 0.6P_{tot}$ in the system considered in Fig. 6. Motivated by this, an interesting question that needs to be explored is how to optimally allocate the transmit power between the information signal and AN signal in mmWave communication systems. Apart from the AN-based transmission, other secure schemes such as hybrid beamforming that mixes analog and digital signal processing techniques can be devised to secure mmWave transmission. In addition, a secure backhaul link between mmWave BSs in the mmWave cellular networks needs to be established.

[3] In current cellular systems, only a small number of antennas at the BS are used. For example, LTE allows for up to 8 antenna ports.

Conclusions

With the introduction of small cell deployments and D2D connections, the use of a very large number of antennas, and the exploration of the underutilized mmWave frequency spectrum, we believe that the 5G network is well positioned to meet the ever-increasing demand on data-centric applications over the next decade. The path to 5G is essentially irreversible, and will impose a profound impact on the design of physical layer security. In this article we have identified the scientific opportunities and discussed the technical challenges driven by the HetNet, massive MIMO, and mmWave communication. The novel solutions we have developed can take data confidentiality to a whole new level, inaugurating a brand new security paradigm that is truly worthy of the 5G designation.

References

[1] Cisco, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast update, 2013–2018, Cisco White Paper, Feb. 2014.
[2] W. Roh *et al.*, "Millimeter-Wave Beamforming as an Enabling Technology for 5G Cellular Communications: Theoretical Feasibility and Prototype Results," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 106–13.
[3] Y.-W. Hong, P.-C. Lan, and C.-C. Kuo, "Enhancing Physical-Layer Secrecy in Multiantenna Wireless Systems: An Overview of Signal Processing Approaches," *IEEE Signal Proc. Mag.*, vol. 30, no. 5, Sept. 2013, pp. 29–40.

[4] A. Mukherjee *et al.*, "Principles of Physical Layer Security in Multiuser Wireless Networks: A Survey," *IEEE Commun. Surveys Tutorials*, vol. 16, no. 3, 3rd Quarter 2014, pp. 1550–73.
[5] R. Bassily *et al.*, "Cooperative Security at the Physical Layer: A Summary of Recent Advances," *IEEE Signal Proc. Mag.*, vol. 30, no. 5, Sept. 2013, pp. 16–28.
[6] J. Andrews, "Seven Ways that HetNets are a Cellular Paradigm Shift," *IEEE Commun. Mag.*, vol. 51, no. 3, Mar. 2013, pp. 136–44.
[7] H. Wang, X. Zhou, and M. C. Reed, "Physical Layer Security in Cellular Networks: A Stochastic Geometry Approach," *IEEE Trans. Wireless Commun.*, vol. 12, no. 6, June 2013, pp. 2776–87.
[8] G. Geraci *et al.*, "Physical Layer Security in Downlink Multi-Antenna Cellular Networks," *IEEE Trans. Commun.*, vol. 62, no. 6, June 2014, pp. 2006–21.
[9] M. Tehrani, M. Uysal, and H. Yanikomeroglu, "Device-to-Device Communication in 5G Cellular Networks: Challenges, Solutions, and Future Directions," *IEEE Commun. Mag.*, vol. 52, no. 5, May 2014, pp. 86–92.
[10] N. Bhushan *et al.*, "Network Densification: The Dominant Theme for Wireless Evolution into 5G," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 82–89.
[11] E. Larsson *et al.*, "Massive MIMO for Next Generation Wireless Systems," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 186–95.
[12] H. Q. Ngo, E. G. Larsson, and T. L. Marzetta, "Energy and Spectral Efficiency of Very Large Multiuser MIMO Systems," *IEEE Trans. Commun.*, vol. 61, no. 4, Apr. 2013, pp. 1436–49.
[13] E. Björnson *et al.*, "Massive MIMO Systems with Non-Ideal Hardware: Energy Efficiency, Estimation, and Capacity Limits," http://arxiv.org/pdf/1307.2584v1.pdf.
[14] J. Zhu, R. Schober, and V. Bhargava, "Secure Transmission in Multicell Massive MIMO Systems," *IEEE Trans. Wireless Commun.*, vol. 13, no. 9, Sept. 2014, pp. 4766–81.
[15] N. Yang *et al.*, "Physical Layer Security of TAS/MRC with Antenna Correlation," *IEEE Trans. Info. Forensics Security*, vol. 8, no. 1, Jan 2013, pp. 254–59.
[16] T. Rappaport et al., "Millimeter Wave Mobile Communications for 5G Cellular: It Will Work!," *IEEE Access*, vol. 1, May 2013, pp. 335–49.
[17] F. Rusek, D. Persson, B. K. Lau, E. G. Larsson, T. L. Marzetta, O. Edfors, and F. Tufvesson, "Scaling Up MIMO: Opportunities and Challenges with Very Large Arrays," *IEEE Signal Process. Mag.*, vol. 30, no. 1, Jan. 2013, pp. 40–60.
[18] Z. Pi and F. Khan, "An Introduction to Millimeter-Wave Mobile Broadband Systems," *IEEE Commun. Mag.*, vol. 49, no. 6, June 2011, pp. 101–07.

Biographies

Nan Yang (nan.yang@anu.edu.au) received the Ph.D. degree in electronic engineering from the Beijing Institute of Technology, Beijing, China, in 2011. He is currently a future engineering research leadership fellow and lecturer at the Australian National University. He received the IEEE ComSoc Asia-Pacific Outstanding Young Researcher Award in 2014, the Exemplary Reviewer Certificate of *IEEE Communications Letters* in 2012 and 2013, and the Best Paper Award at the IEEE VTC in 2013. His research interests include collaborative networks, network security, massive MIMO, millimeter wave, and molecular communications.

Lifeng Wang (lifeng.wang@qmul.ac.uk) received the M.S. degree in electronic engineering from the University of Electronic Science and Technology of China, Sichuan, China, in 2012. He is currently working toward the Ph.D. degree in electronic engineering at Queen Mary University of London, London, U.K.. His research interests include MIMO, cooperative communications, cognitive radio, and physical-layer security.

Giovanni Geraci (giovanni_geraci@sutd.edu.sg) received the Ph.D. in electrical engineering from the University of New South Wales, Sydney, Australia, in 2014. He is currently a postdoctoral research fellow at the Singapore University of Technology and Design, Singapore. His research interests include wireless communications, signal processing, applied mathematics, and information technology.

Maged Elkashlan (maged.elkashlan@qmul.ac.uk) received the Ph.D. degree in electrical engineering from the University of British Columbia, Vancouver, BC, Canada, in 2006. Since 2011 he has been with the School of Electronic Engineering and Computer Science, Queen Mary University of London, London, U.K., as an assistant professor. He currently serves as an editor for *IEEE Transactions on Wireless Communications*, *IEEE Transactions on Vehicular Technology*, and *IEEE Communications Letters*. His research interests fall into the broad areas of communication theory, wireless communications, and statistical signal processing for distributed data processing, millimeter-wave communications, cognitive radio, and wireless security.

Jinhong Yuan (j.yuan@unsw.edu.au) received the Ph.D. degree in electronics engineering from the Beijing Institute of Technology, Beijing, China, in 1997. In 2000 he joined the School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, Australia, where he is currently a telecommunications professor. He has co-authored three Best Paper Awards and one Best Poster Award. He currently serves as an associate editor for *IEEE Transactions on Communications*. His current research interests include error control coding and information theory, communication theory, and wireless communications.

Marco Di Renzo (marco.direnzo@lss.supelec.fr) received the Ph.D. degree in electrical and information engineering from the University of L'Aquila, Italy, in 2007. Since January 2010 he has been a tenured researcher with the French National Center for Scientific Research and a faculty member at the Laboratory of Signals and Systems. He currently serves as an editor for *IEEE Transactions on Communications* and *IEEE Communications Letters*. His main research interests are in the area of wireless communications theory.

With the introduction of small cell deployments and D2D connections, the use of a very large number of antennas, and the exploration of the underutilized mmWave frequency spectrum, we believe that the 5G network is well positioned to meet the ever-increasing demand on data-centric applications over the next decade.

Authentication Handover and Privacy Protection in 5G HetNets Using Software-Defined Networking

Xiaoyu Duan and Xianbin Wang

ABSTRACT

Recently, densified small cell deployment with overlay coverage through coexisting heterogeneous networks has emerged as a viable solution for 5G mobile networks. However, this multi-tier architecture along with stringent latency requirements in 5G brings new challenges in security provisioning due to the potential frequent handovers and authentications in 5G small cells and HetNets. In this article, we review related studies and introduce SDN into 5G as a platform to enable efficient authentication handover and privacy protection. Our objective is to simplify authentication handover by global management of 5G HetNets through sharing of user-dependent security context information among related access points. We demonstrate that SDN-enabled security solutions are highly efficient through its centralized control capability, which is essential for delay-constrained 5G communications.

INTRODUCTION

Over the past few years, anywhere, anytime wireless connectivity has gradually become a reality and has resulted in remarkably increased mobile traffic. Mobile data traffic from prevailing smart terminals, multimedia-intensive social applications, video streaming, and cloud services is predicted to grow at a compound annual growth rate of 61 percent before 2018, and is expected to outgrow the capabilities of the current fourth generation (4G) and Long Term Evolution (LTE) infrastructure by 2020 [1]. This explosive growth of data traffic and shortage of spectrum have necessitated intensive research and development efforts on 5G mobile networks. However, the relatively narrow usable frequency bands between several hundred megahertz and a few gigahertz have been almost fully occupied by a variety of licensed or unlicensed networks, including 2G, 3G, LTE, LTE-Advanced (LTE-A), and Wi-Fi. Although dynamic spectrum allocation could provide some improvement, the only way to find enough new bandwidth for 5G is to explore idle spectrum in the millimeter-wave range of 30~300 GHz [2].

The authors are with Western University.

NETWORK ARCHITECTURE OF 5G

Due to the poor signal propagation characteristics at extremely high frequencies, future 5G networks will be heterogeneous with small cell deployment and overlay coverage, as shown in Fig. 1. Cellular networks operating at low frequencies (e.g., 2G, 3G, LTE, LTE-A) could provide wide area coverage, mobility support, and control, while small cells operating at higher frequencies guarantee high data rates in the area of spectral and energy efficiency.

This heterogeneous paradigm with multi-tier coverage in 5G not only follows the natural evolution from existing cellular technologies, but also satisfies the requirements of increased data traffic, with small cells providing very high throughput and underlying macrocells providing extensive coverage. Therefore, network densification using low-power small cells is widely considered to be a critical element toward low-cost high-capacity 5G communications.

SECURITY CHALLENGES IN 5G

Along with the advantages of 5G architecture in Fig. 1, there also come several major technical challenges. The massive deployment of small cells poses potential challenges in network management, including interference alignment, extensive backhauling, and inconsistent security mechanisms over heterogeneous networks (HetNets). Network management and service provisioning are challenging in this multi-tier model due to the increased number of base stations and complexity of network architecture. Therefore, new technologies are needed to provide intelligent control over HetNets for consistent and effective resource allocation as well as security management.

Moreover, 5G users may leave one cell and join another more frequently with reduced cell size, which could introduce excessive handover-induced latency in 5G. Future 5G applications like interactive gaming and tele-operations require 5G latency to be an order of magnitude smaller than 4G, with 1 ms target round-trip time [2]. However, due to smaller cell deployment, users and different access points (APs) in 5G need to perform more frequent mutual authentications than in 4G to prevent imperson-

0163-6804/15/$25.00 © 2015 IEEE

ation and man-in-the-middle (MitM) attacks. On the other hand, the power and resource constraints of small cell APs require low complexity and highly efficient handover authentication procedures. Therefore, faster, efficient, and robust handover authentication and privacy protection schemes need to be developed for complex 5G HetNets.

The Scope of This Article

In this article, we first introduce the 5G background and identify the challenges in 5G HetNets, especially in security management. Existing related studies are overviewed, providing a summary of the previous security solutions and state-of-the-art related technologies. Based on our survey and analysis, we believe that new solutions meeting the latency and complexity requirements of 5G HetNet communications are yet to be developed.

Based on this observation, we introduce a new 5G network structure enabled by software-defined networking (SDN) to bring intelligence and programmability into 5G networks for efficient security management. With SDN, the control logic is removed from the underlying infrastructures to a controller in the control layer [3] so that software can be implemented on the central SDN controller to provide consistent and efficient management over the whole 5G HetNet. With this paradigm, we propose an SDN-enabled user-specific secure context information transfer for efficient authentication handover and privacy protection in 5G to achieve seamless authentication during frequent handovers, while at the same time meeting the privacy and latency requirements effectively.

State of the Art in Handover Authentication and Challenges in 5G

Related Work on Handover Authentication and 5G Challenges

To support increased data traffic, 5G networks need to have high capacity and efficient security provisioning mechanisms. Densification of heterogeneous networks and massive deployment of small base stations become the natural choice for 5G. On the other hand, many applications supported by 5G, such as mobile banking and cloud-based social applications, require higher data confidentiality and reliable authentication against malicious attacks.

The common practice for secure communications in 3G and later wireless networks is based on admission control and cryptographic exchange. Figure 2 gives an overview of the handover authentication procedures between different networks and within one network [9]. The involved network components here are the user equipment (UE), access points (APs) or base stations (BSs), and an authentication server. It can be seen from Fig. 2 that mutual authentication during handover between the user and a new network (i.e., procedure 1) is realized by the pairing of specific hashing output. Each time the involved vector includes RAND, a random number known by the server, AUTH, an authentication token sent by the server, a pairwise key, and so on. For mobility within the same network (i.e., procedure 2), the current serving AP will inform the target AP of the possible handover so that the latter can retrieve the user authentication and key context from the server. In the following, we analyze existing handover authentication procedures and identify the challenges in 5G HetNets based on Fig. 2.

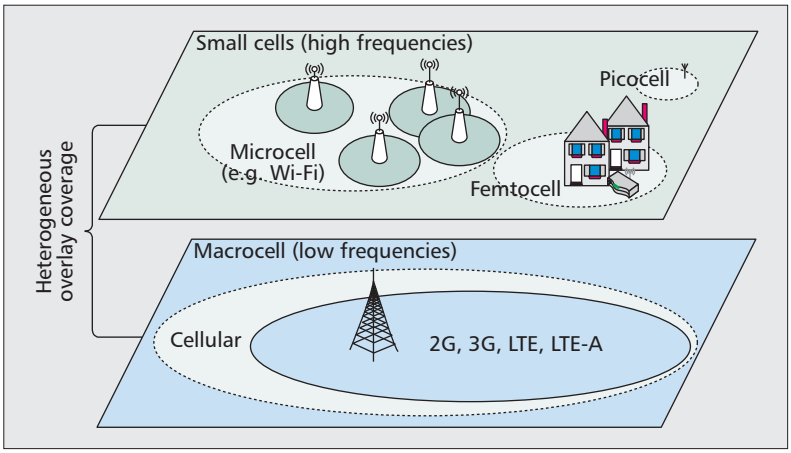

Figure 1. 5G heterogeneous network structure with densified small cells and overlay coverage.

To enable handover between different wireless networks (i.e., procedure 1 in Fig. 2), various authentication servers and protocols are involved due to the closed nature and structure of each network in a HetNet, rendering frequent establishments of trust relationships and authentications during mobility, especially in a 5G small cell scenario [2]. The Third Generation Partnership Project (3GPP) has provided specific key hierarchy and handover message flows for various mobility scenarios [10]. However, the specific key designed for handover and different handover procedures for various scenarios will increase handover complexity when applied to 5G HetNets. As the authentication server is often located remotely, the delay due to frequent enquiries between small cell APs and the authentication server for user verification may be up to hundreds of milliseconds [5], which is unacceptable for 5G communications. The authors of [6, 7] have proposed simplified handover authentication schemes involving direct authentication between UE and APs based on public cryptography. These schemes realize mutual authentication and key agreements with new networks through a three-way handshake without contacting any third party, like an authentication, authorization, and accounting (AAA) server. Although the handover authentication procedure is simplified, computation cost and delay are increased due to the overhead for exchanging more cryptographic messages through a wireless interface [5]. For the same reason, carrying a digital signature is secure but not efficient for dynamic 5G wireless communications.

For handover within the same network (i.e., procedure 2 in Fig. 2), existing security mechanisms utilize complex context transfer, and it has

When introducing SDN into 5G networks, the SDN controller will have global control over the network, while SDN switches simply follow data forwarding instructions from the controller. Applications are implemented on top of the controller to define the behavior of the switches and APs, thus creating a reconfigurable 5G HetNet.

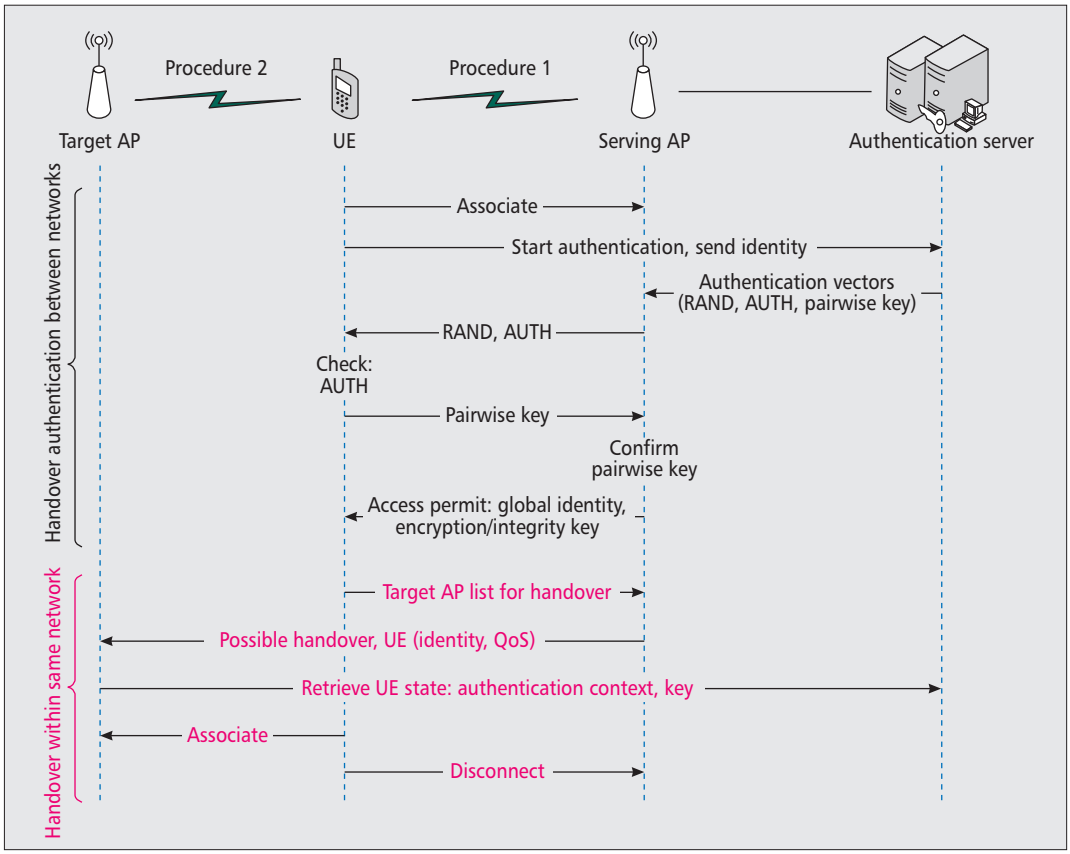

Figure 2. Authentication processes of handover procedure 1, between different networks, and handover procedure 2, within the same network.

been found that most of the handover latency is due to the scanning time for identifying the target AP and round-trip time to the authentication server. Related work in [8] proposed a user-assisted authentication context transfer scheme, by which the current AP transfers a signed authentication certificate as a security context to the user, and then to the target AP through the user. The UE is actively involved in handover authentication with its existing connections with the current and next target APs to reduce latency. However, mutual trust between APs is assumed in these solutions, which could be infeasible for 5G HetNets due to the lack of direct interfaces between different networks. In addition, the transferred security context, which is just a combination of identity and signature, may not be secure enough to prevent 5G wireless communication from potential attacks.

In light of these challenges, robust and efficient handover authentication and secure context information transfer is crucial in securing 5G networks. The unique link characteristics experienced by each UE can be explored as a security context to accelerate authentication handover. Such user-specific attributes include physical layer attributes (clock skew, signal strength, channel state information), location, and even moving speed and direction [11], some of which have already been reported to APs for the purpose of resource allocation and seamless handover. It is believed that by taking advantage of these unique attribute combinations as non-cryptographic solutions, authentication can be faster, more robust, and less complex compared to widely used cryptographic exchange mechanisms [12].

Software-Defined-Networking-Enabled 5G Networks

Software-defined networking [3] is considered as a radical new network structure to centralize network management, and enable innovation through network programmability in meeting the needs of emerging applications. One main feature of SDN is decoupling the control plane and data plane by taking control logic from the underlying switches and routers to the centralized SDN controller in the control plane.

When introducing SDN into 5G networks, the SDN controller will have global control over the network, while SDN switches will simply follow data forwarding instructions from the controller. Applications are implemented on top of the controller to define the behavior of the switches and APs, thus creating a reconfigurable 5G HetNet, as shown in Fig. 3. The separation of data forwarding switches and the control plane enables easier implementation of new protocol and functions, consistent network policy, as well as straightforward network management.

In supporting SDN-enabled 5G, appropriate SDN protocols, such as Openflow and Simple Network Management Protocol (SNMP), will be added to base stations, access points, and wireless switches through an external standardized application programming interface (API) [4].

Importantly, OpenFlow is in charge of data path control, and SNMP can be used for device control. As the SDN controller is just a program running on a server, it can be placed anywhere in the 5G network — even in a remote data center.

An SDN-based 5G network structure enables flexible ubiquitous connection, fast rerouting, and real-time network management with the software controller. Users are able to access network services anywhere and anytime regardless of the network type [4] (e.g., Wi-Fi, 3G, LTE, LTE-A) as long as these networks belong to the same operator or there are agreements between operators. Furthermore, consistent authentication and privacy protection are also manageable.

In this article, we explore SDN as a promising platform to introduce intelligence into 5G and address the security challenges. Specifically, we discuss SDN-enabled authentication handover, which provides control over HetNet infrastructures and helps the network to reduce redundant authentications across HetNets. Handover authentication thus becomes a more controlled and prepared process instead of multiple independent procedures. By sharing secure context information along moving direction of the user and choosing multiple network paths to transmit data concurrently, the SDN structure is capable of facilitating 5G security provisioning more efficiently. In doing so, user-specific attributes are utilized as the shared security context to reduce handover complexity. To further achieve privacy protection, SDN-enabled data transmission over different network paths in 5G HetNets is also investigated in order to guarantee privacy.

SDN-Enabled 5G Authentication Handover

In this section, we introduce SDN into 5G to enable the proposed authentication handover scheme in coping with the frequent handover authentication in small cells and HetNets, as shown in Fig. 4. We implement an authentication handover module (AHM) in the SDN controller to monitor and predict the location of users, and then prepare the relevant cells before the user arrives to guarantee seamless handover authentication. Using a traffic flow template (TFT) filter [13] (source/destination IP addresses and port numbers) and related quality of service (QoS) description, secure context information (SCI) is collected by the AHM to share along a projected user moving path (i.e., from cell A to cell B, C in Fig. 4). The relevant cell APs thus prepare resource in advance and ensure seamless user experience during mobility.

Specifically, user specific attributes including identity, location, direction, round-trip time (RTT), and physical layer characteristics have been considered as reliable SCI to assist secure handover in 5G networks, instead of using complex cryptographic exchange mechanisms. As a non-cryptographic method, user-specific attributes are able to simplify the authentication procedure by providing the unique fingerprint of the specific device without additional hardware and computation cost [12]. In this article, we focus on using user-specific attributes as SCI (location, direction, etc.) to realize SDN-enabled authentication handover. Based on the proposed authentication context handover, security in SDN-enabled 5G networks becomes a monitored seamless procedure instead of multiple independent verifications, which could significantly reduce the possibility of impersonation and MitM attacks.

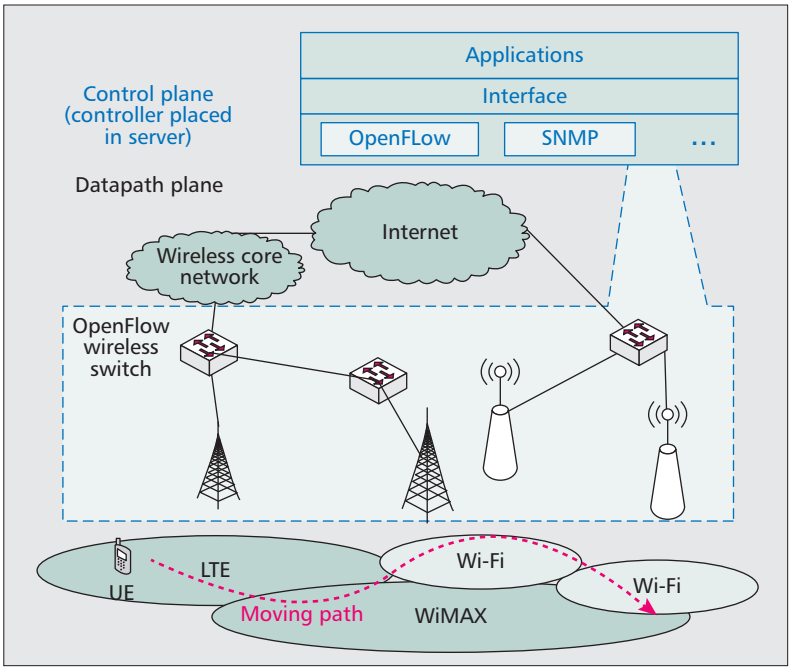

Figure 3. SDN-enabled 5G wireless HetNet structure with control plane design.

More precisely, the way in which the SDN controller shares the user's SCI to next cell APs along the predicted path is just like a trustworthy introduction from a previous AP before handover. The future cell APs thus finish authentication with the user quickly and begin to monitor the user to prepare service according to the SCI. As the trace of the user is monitored, the risk of impersonation is significantly, if not entirely, reduced. More importantly, there would be risk of service disruption in previous networks if the connection between APs and the authentication server is broken. Under similar network conditions, however, our mechanism will not lose global network connectivity because a new AP is monitoring the user, which can help the controller retrieve the necessary information according to the pre-shared SCI. Thus, the SDN-enabled security handover possesses high levels of tolerance to network failures. In the following, a description of the authentication handover mechanism in terms of assumptions and designs is presented in detail.

Assumptions and Design Goals

We assume that the SDN controller is a program running in a mobile operator's data center with an AHM for user authorization. The AHM is in charge of both authentication and handover, which maintains user information specifying what the user can access. The AHM also pos-

Due to the reduced cell size in 5G HetNets, users might move through multiple small cells before completing one communication session. Thus, the privacy protection is more challenging in 5G due to the possible involvement of untrusted or compromised APs during handover.

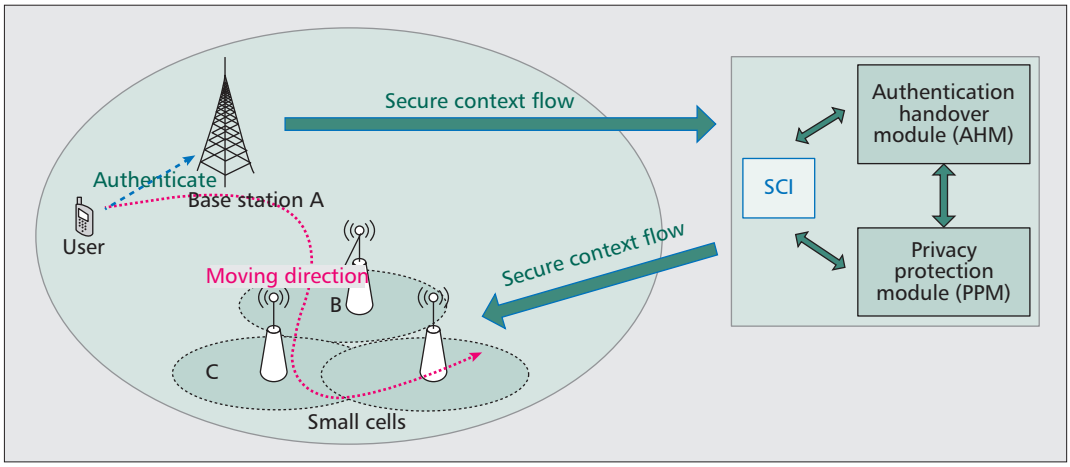

Figure 4. SDN enabled secure context information transfer between 5G UE, APs and AHM in SDN controller.

sesses a master public-private key pair (K, K^{-1}), with a public key K that is known to users and APs. Both APs and UEs need to be verified before gaining access to network services to reduce security risks.

Our design goal for the authentication handover mechanism is to accelerate authentication in 5G HetNets by enabling SCI transfer using SDN. In further reducing the overall authentication delay, the AHM in the controller could periodically authenticate the APs in off-peak times using its master key to avoid leakage of privacy caused by compromised APs. If certified, a key pair (K_N, K_N^{-1}) with a signature $[K_N, T]_{K^{-1}}$ is distributed to the AP, where T is the timeout of the signature; if the AP is detected as compromised, it will be blacked out from further operation. This way, some of the authentication procedures are moved to off-peak times and relieves the SDN controller burden.

SDN-Enabled Authentication Handover Mechanism Design

With the assumptions and design goals described above, we can design the SDN-enabled authentication handover mechanism. User-specific SCI, such as ID, physical layer attributes, location, speed, and direction, can be collected and shared easily with SDN flow-based forwarding [3]. According to the UE location information from SCI, the SDN controller uses an ascending index to indicate the sequential order of next cells in the moving direction. Once authenticated by one cell AP, an appropriate combination of user attributes is then shared as SCI by the SDN controller along this user's future path. This way, the UE is able to enjoy seamless service without complex operation during authentication handover, thus saving time for data communications.

For example, we assume that user U is in cell A, and the future cells are B and C, as shown in Fig. 4. The authentication procedure between user U and cell A follows the commonly used authentication protocol [10], and the proposed SDN-enabled authentication handover procedure is described in Algorithm 1.

The SCI attributes in the proposed SDN-enabled authentication handover could include identity, physical layer attributes, location, moving speed, and direction. The number of attributes to be used is based on the security level of the information requested. For example, if the user is requesting banking or email services, a higher security level can be achieved by transferring more SCI attributes; if it is just Internet browsing or video gaming, the security level can be lower, and few SCI attributes are needed.

The aforementioned authentication handover method requires no changes to the existing UE and AP hardware, and significantly simplifies the authentication procedure and reduces handover latency through a non-cryptographic technique. By predicting the user moving path and shifting the authentication of APs to off-peak times, the SDN-enabled 5G networks can always be well prepared for other service requests. Moreover, operators can choose to switch off/on lightly loaded cells if the users approaching these cells are not going to exceed a certain threshold according to the SCI information to save more energy.

SDN-Enabled 5G Privacy Protection

Data privacy means the right of network users to seclude themselves from prying and eavesdropping. Due to the reduced cell size in 5G HetNets, users might move through multiple small cells before completing one communication session. Thus, the privacy protection is more challenging in 5G due to the possible involvement of untrusted or compromised APs during handover. Existing privacy protection schemes use complex key agreements and interactions or additional watermarking to protect data privacy. Such cryptographic methods bring computation burden and complexity to both the AP and client sides [9], which is undesirable for 5G low-power small cell infrastructures. On the other hand, privacy protection requires that no link can be established

State(A, U): Authenticated.
State(B, U): Not Authenticated.
State(C, U): Not Authenticated.
AHM → **B**: (*index* = 1, *ID*, *SCI*)
AHM → **C**: (*index* = 2, *ID*, *SCI*)
Ascending index number shows the direction of user movement. ID is the identity of *U* and *SCI* is the secure context information of *U*.
B → **A**: Handoff REQ(*ID*, *SCI*).
When *B* discovers *U* in its coverage, *B* sends handoff request to *A* until receives reply from *A*.
A → **B**: Handoff ACK(*ID*, *SCI'*).
A replies with handoff acknowledgement. *SCI'* is the secure context information which is more recent than previous shared *SCI*.
B → **U**: Update REQ().
After matching *SCI'* from *A* with *U*, *B* authenticates *U* and starts to associate with *U*.
U → **B**: Update ACK(*SCI''*).
Here *U* is connected with *B*. *SCI''* is the latest secure context information.
State(B,U): Authenticated.
B → **AHM**: Update(*SCI''*).
B updates the UE secure context information to AHM. AHM then shares secure information to next cell APs according to the location and direction information in new *SCI''*.
C → **B**: *C* keeps on monitoring *U* and follows similar procedure.

Algorithm 1. User-SCI-based authentication handover.

1: **procedure** PDO(*n*)
2: T_s: delay threshold
3: $V_{sn} = b_n \min(t_r; T_s)$: size in bytes to be transferred in nearby Wi-Fi, Femtocell or cellular within T_s
4: **for** $d_1 < V_{s1}$, $d_2 < V_{s2}$, … $d_n < V_{sn}$ and $d = d_1 + d_2 + \ldots + d_n$ **do**
5: Encrypt d_1, d_2, … d_n separately, send them on n networks concurrently and update d
6: **end for**
7: Receiver decrypt $d_1 \sim d_n$ using private key and re-organize data
8: **end procedure**

Algorithm 2. Partial data offloading over different SDN-controlled network paths.

between information and the owner, while authentication requires an identity provided for the purpose of authentication. Previously, these contradictory requirements were met through a trusted third party. However, multiple enquiries to the remote third party cause a network bottleneck, which is not suitable for 5G low-latency communications.

We introduce an SDN-enabled privacy protection scheme, which employs partial transmission over different SDN-controlled network paths to guarantee privacy and offload traffic in 5G cellular networks at the same time. With the proposed privacy protection scheme, SDN controller is able to choose multiple network paths to transmit different parts of the data stream (i.e., partial transmission) according to the HetNet coverage. The number of network paths is decided by the sensitivity level of the data stream. As long as the UE has been authenticated and is covered by the HetNets (e.g., Wi-Fi, femtocell, or cellular), the induced data stream can be routed through these network backhauls under the control of an SDN controller. Only the receiver can decrypt the data using its private key and then re-organize the data stream coming from multiple network paths, which avoids privacy leakage via compromised APs. Moreover, the proposed scheme is able to realize traffic offloading through the other network paths, which is desirable given the fact that a 5G cellular network will be flooded by a huge volume of mobile traffic [1]. Simply by choosing nearby Wi-Fi or femtocells as different paths for data offloading, the traffic load of a 5G cellular network is relieved through either the unlicensed band of Wi-Fi or reusing the femtocell's band. The proposed SDN-enabled privacy protection mechanism is described in Algorithm 2.

In Algorithm 2, n is the number of network paths that an SDN controller chooses for data transmission, and d_n is the different part of data that will be transmitted in the nth network concurrently. t_r is the data transfer time within the involved networks. T_s is the delay threshold of 5G applications, which means to achieve concurrent privacy protection, this kind of service needs to be finished before T_s to guarantee user experience. For example, email transfer can tolerate long latency, while real-time video and two-way gaming have a very low delay threshold. b_n is the bandwidth allocated by the SDN controller according to the traffic situation of different networks, and V_{sn} is the volume of data that can be transferred in the multiple paths (i.e., offloading networks) within the application delay threshold.

More importantly, the number of paths n here is decided by a trade-off between privacy level, offloading revenue, and system complexity, which is reconfigurable and can easily be set up through an SDN controller application by 5G operators. User privacy protection thus becomes programmable and under the control of SDN, which is especially desirable for future highly diverse communication requirements and application needs.

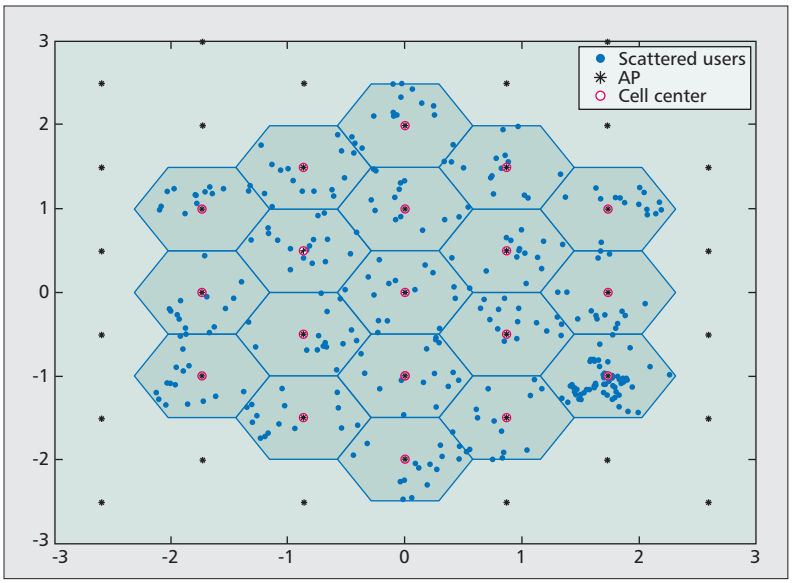

Figure 5. Simulation layout of 5G small cells with proportional axis (1 = 300 m).

Cell layout	Hexagonal grid, 19 cell sites, with wrap-around technique
Cell radius	150m
User mobility speed	3 km/h
User mobility direction	Random
Total number of users	570

Table 1. Simulation parameters of 5G networks.

Performance Analysis

MATLAB simulations of a 5G network with commonly used hexagonal cells are adopted to evaluate the performance of the aforementioned mechanisms in terms of the secure level and latency. A total of 19 small cells in Fig. 5 with an inter-site distance (i.e., distance between two APs) of 300 m is considered in the simulation. Users are randomly distributed around APs, while each UE takes a random walk and changes direction every 5 s. The wrap-around technique (i.e., users moving out of the predefined service area are assumed to enter the area from the other side of the network) is used to avoid boundary effects. The specific simulation parameters are listed in Table 1.

In simulating the proposed SDN-enabled authentication handover, we consider the separation distance between UE and APs, and the moving direction of the UE as the transferred SCI to verify the reliability of the proposed SCI-based authentication handover scheme. From the simulation results, we find that during the monitored user handover process, the probability that any two users have the same distance (with accuracy to the first decimal) to the closest AP is 44 percent. When it comes to the same AP, the probability of two users having the same distance to this AP decreases to 11 percent. Combined with moving direction, signal strength, channel state information, and other user-specific attributes, the probability of UEs with the same SCI could be reduced to virtually 0. Therefore, we believe that the SDN-enabled authentication handover mechanism using SCI transfer is robust to guarantee security with enough SCI attributes. Moreover, it is flexible in setting a security level by different combinations of user-specific attributes.

Authentication handover delays from SDN-enabled handover and the traditional methods are simulated and compared in evaluating the latency performance of the proposed schemes. Without loss of generality, we assume that the data of each user following Poisson arrivals and new users initiate the authentication process when the UE is on the move. In simulating the proposed authentication handover, user-specific SCI is collected and transferred to relevant cells on the projected moving path of the UE under the coordination of the SDN controller. On the other hand, traditional authentication handover protocol requires separate authentication in each network involved in the handover. Here we use two publicly available OpenFlow controllers as representatives to show the performance [14], NOX-MT and Beacon. NOX-MT is a multi-threaded successor of NOX, while Beacon is a Java controller built by David Erickson at Stanford [3].

Figure 6 shows the comparison of authentication delay vs. 5G network utilization rates. Here network utilization is defined as the ratio of total data arrival rate and controller processing rate. Network utilization rate is used as it reflects the different load situations of the network. We can see from Fig. 6 that when the network load is fairly low, authentication delay is not a problem for all different methods. With more arrivals and increased network load, SDN-enabled authentication handover still keeps the latency under 1 ms most of the time, which meets the 5G latency requirement. NOX-MT- and Beacon-enabled solutions perform 30 and 14.29 percent better than traditional handover authentication protocol in latency reduction with the commonly used deployment of an eight-core machine, 2 GHz CPUs, and 32 switches in [14]. It is obvious that the SDN-enabled authentication handover and privacy protection scheme meet the critical latency requirement in 5G, while maintaining the SDN flexibility, programmability, and data offloading capability in further improving the energy efficiency and network management of 5G networks.

Conclusion

With the upcoming multi-tier architecture and small cell deployment, challenges emerge in security provisioning and privacy protection in 5G heterogeneous networks. 5G network security handover needs to be fast, with low complexity due to the reduced cell size and stringent latency constraint. In this article, we review the existing studies and identify current challenges on authentication handover and privacy protection in 5G. In addressing these challenges, we

propose SDN-enabled authentication handover and privacy protection through sharing of user-specific security context information among related access points. The proposed SDN-enabled solution not only provides a reconfigurable network management platform, but also simplifies authentication handover in achieving reduced latency. The performance of the proposed schemes have been demonstrated through numerical simulations and examples. We expect that more progress could be made by using emerging SDN-enabled 5G architecture and non-cryptographic techniques to address the 5G challenges of reduced cell size and coexistence of heterogeneous networks. Many interesting related topics, including network complexity, security performance under different attacks, and effective use of security context information, could be explored for SDN-enabled 5G security mechanisms.

References

[1] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013–2018," http://tinyurl.com/b9berc, 2014.

[2] J. Andrews *et al.*, "What Will 5g Be?," *IEEE JSAC*, vol. 32, no. 6, 2014, pp. 1065–82.

[3] B. A. A. Nunes *et al.*, "A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks," *IEEE Commun. Surveys and Tutorials*, vol. 99, Feb. 2014, pp. 1–18.

[4] K.-K. Yap *et al.*, "Blueprint for Introducing Innovation into Wireless Mobile Networks," *Pro. ACM SIGCOMM Wksp.*, 2010, pp. 25–32.

[5] D. He *et al.*, "Secure and Efficient Handover Authentication based on Bilinear Pairing Functions," *IEEE Trans. Wireless Commun.*, vol. 11, no. 1, Jan. 2012, pp. 48–53.

[6] J. Choi and S. Jung. "A Handover Authentication Using Credentials Based on Chameleon Hashing," *IEEE Commun. Letters*, vol. 14, no. 1, 2010, pp. 54–56.

[7] J. Cao *et al.*, "A Simple and Robust Handover Authentication between HeNB and eNB in LTE Networks," *Computer Networks*, vol. 56, no. 8, 2012, pp. 2119–31.

[8] L. Cai, S. Machiraju, and H. Chen, "CapAuth: A Capability-Based Handover Scheme," *Proc. INFOCOM*, 2010, pp. 1–5.

[9] L. Chen, J. Ji, and Z. Zhang, *Wireless Security: Models, Threats, and Solutions*, Higher Education Press, 2013.

[10] 3GPP TS 33.401 V11.5.0. 3rd Generation Partnership Project; Technical Specification Group Service and System Aspects; 3GPP System Architecture Evolution (SAE); Security Architecture (Rel 11), 2012.

[11] D. He *et al.*, "Security and Efficiency in Roaming Services for Wireless Networks: Challenges, Approaches, and Prospects," *IEEE Commun. Mag.*, vol. 51, no. 2, Feb. 2013, pp. 142–50.

[12] K. Zeng, K. Govindan, and P. Mohapatra, "Non-Cryptographic Authentication and Identification in Wireless Networks," *IEEE Wireless Commun.*, vol. 17, no. 5, Oct. 2010, pp. 56–62.

[13] 3GPP, Technical Specification Group Services and System Aspects; Policy and Charging Control Architecture, tech. spec. 3G TS 23.203 ver. 9.5.0 (2010-06).

[14] A. Tootoonchian *et al.*, "On Controller Performance in Software-Defined Networks," *Proc. HotICE*, 2012.

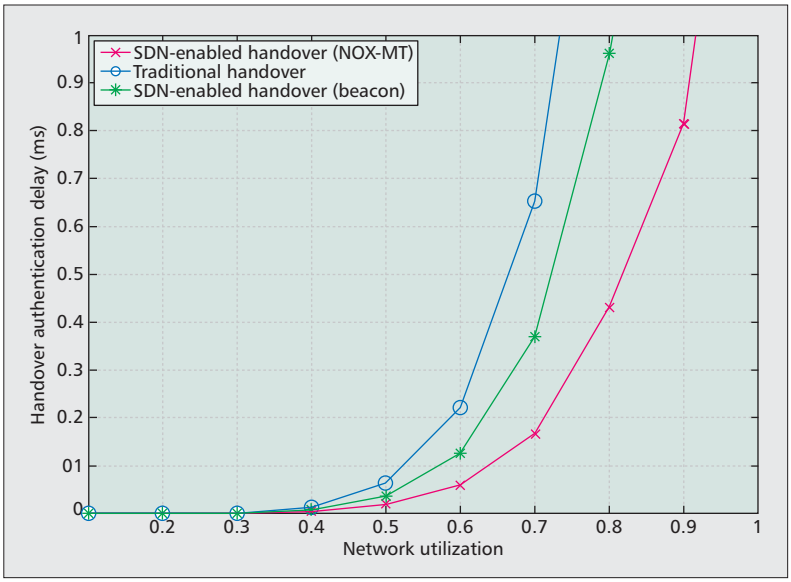

Figure 6. Comparison of authentication delays vs. network utilization rates.

Biographies

Xiaoyu Duan (xduan8@uwo.ca) is a Ph.D. candidate at the Department of Electrical and Computer Engineering, Western University, Canada. She received a B.Sc. in communication engineering from Tianjin University in 2010 and an M.Sc. in signal and information processing from Beijing University of Posts and Telecommunications, China, in 2013. Her research interests include software-defined networking, traffic offloading, self-organizing networks, and communication security in 5G heterogeneous networks.

Xianbin Wang [S'98, M'99, SM'06] (xianbin.wang@uwo.ca) is a professor at Western University and Canada Research Chair in Wireless Communications. He received his Ph.D. degree in electrical and computer engineering from National University of Singapore in 2001. Prior to joining Western, he was with Communications Research Centre Canada as a research scientist/senior research scientist between July 2002 and December 2007. From January 2001 to July 2002, he was a system designer at STMicroelectronics, where he was responsible for system design for DSL and Gigabit Ethernet chipsets. His current research interests include adaptive wireless systems, 5G networks, communications security, and distributed computing systems. He has 200 peer-reviewed journal and conference papers on various communication system design issues, in addition to 24 granted and pending patents and several standard contributions. He is an IEEE Distinguished Lecturer. He has received three IEEE Best Paper Awards. He currently serves as an Associate Editor for *IEEE Wireless Communications Letters*, *IEEE Transactions on Vehicular Technology*, and *IEEE Transactions on Broadcasting*. He was also an Editor for *IEEE Transactions on Wireless Communications* between 2007 and 2011. He has been involved in a number of IEEE conferences including GLOBECOM, ICC, WCNC, VTC, ICME, and CWIT, in different roles such as Symposium Chair, Tutorial Instructor, Track Chair, TPC Chair, and Session Chair.

Securing Software Defined Networks: Taxonomy, Requirements, and Open Issues

Adnan Akhunzada, Ejaz Ahmed, Abdullah Gani, Muhammad Khurram Khan, Muhammad Imran, and Sghaier Guizani

ABSTRACT

The emergence of SDNs promises to dramatically simplify network management and enable innovation through network programmability. Despite all the hype surrounding SDNs, exploiting its full potential is demanding. Security is still the key concern and is an equally striking challenge that reduces the growth of SDNs. Moreover, the deployment of novel entities and the introduction of several architectural components of SDNs pose new security threats and vulnerabilities. Besides, the landscape of digital threats and cyber-attacks is evolving tremendously, considering SDNs as a potential target to have even more devastating effects than using simple networks. Security is not considered as part of the initial SDN design; therefore, it must be raised on the agenda. This article discusses the state-of-the-art security solutions proposed to secure SDNs. We classify the security solutions in the literature by presenting a thematic taxonomy based on SDN layers/interfaces, security measures, simulation environments, and security objectives. Moreover, the article points out the possible attacks and threat vectors targeting different layers/interfaces of SDNs. The potential requirements and their key enablers for securing SDNs are also identified and presented. Also, the article gives great guidance for secure and dependable SDNs. Finally, we discuss open issues and challenges of SDN security that may be deemed appropriate to be tackled by researchers and professionals in the future.

INTRODUCTION

The emergence of the software defined networking (SDN) paradigm has created great potential and hope to overcome the need for flexible, secure, reliable, and well managed next-generation networks. The revolutionary concept of SDN has brought radical change to the traditional vertical integration of the network by decoupling the forwarding hardware (data plane) from the control logic of the network (control plane) [1]. Subsequently, just the switches and routers are held responsible for forwarding the traffic; however, the control functionality is simply shifted to a centralized logical controller. Moving the control logic to an external entity known as an SDN controller provides an abstract view of the underlying network resources to achieve smooth facilitation of the programming of forwarding hardware. Moreover, the abstraction of flow broadly unifies the behavior of different SDN agents. Obviously, these remarkable features of SDN provide a more flexible, programmable, vendor-agnostic, cost-effective, and innovative network architecture. In spite of all these exciting features of SDN, industry observers are apprehensive about the security of SDNs. The security of SDNs is still considered the topmost priority, and a key concern and an equally arresting challenge have recently begun to receive the attention they deserve. Industry experts strongly believe that security issues surrounded by SDNs must be thoroughly addressed.

Besides, the architecture of SDNs poses new external and internal threats and vulnerabilities [1]. Predominantly, the integrity and security of SDNss remain unproven when it comes to the placement of management functionality in a single centralized virtual server. Subsequently, compromising the whole network through a single point of failure is much easier. Moreover, it becomes the primary potential attack target. The programmability aspect of SDNs also makes them more vulnerable to a number of malicious code exploits and attacks. Furthermore, the abstraction of different available flows and underlying hardware resources at the SDN controller significantly supports harvesting intelligence from the existing resources. Afterward, it can be effortlessly used for further attacks, exploitations, and particularly reprogramming the entire network. Likewise, the southbound interface of an SDN can also easily be targeted with diverse denial of service and side channel attacks. Equally important, configuration errors of SDNs can have more serious consequences than in traditional networks. Besides, SDN agents can also potentially be targeted for injecting false flows. Keeping in view the SDN features and architecture, cyber-

Adnan Akhunzada, Ejaz Ahmed, and Abdullah Gani are with the University of Malay.

Muhammad Khurram Khan and Muhammad Imran are with King Saud University.

Sghaier Guizani is with Alfaisal University.

The authors extend their sincere appreciations to the Deanship of Scientific Research at King Saud University for its funding this Prolific Research Group (PRG-1436-16).

0163-6804/15/$25.00 © 2015 IEEE

attacks launched through SDNs can have even more devastating and larger effects than using simple networks.

Since security is not considered initially as part of SDN design, each layer of an SDN has its own security implications and requirements. Moreover, establishing trust throughout an SDN is even more critical. Likewise, the network essentially needs a dynamic forensic remediation and robust policy frameworks ensuring the right direction of the controller. Although security should be built in as part of SDN architecture, it must also be delivered as a service to ensure the privacy and integrity of all the connected resources. Some researchers claim that we are still far away from secure and dependable SDN architecture. On the contrary, it is also complementary to say that SDN can be better used to enhance and implement security; meanwhile, security of the SDN itself becomes a priority. SDN certainly necessitates a simple, cost-effective, scalable, and efficient secure environment.

The contributions of this survey are manifold:

- A critical discussion on the state-of-the-art SDN security solutions is given. These solutions depict the current state of SDNs in terms of security.
- The distinguishing aspect of our work is the classification of surveyed security solutions by devising a thematic taxonomy based on SDN layers/ interfaces, security measures, simulation or testbed environment, and security objectives.
- The possible attacks and threat vectors targeting various layers/ interfaces of the SDNs are identified and highlighted.
- The potential security implications and requirements with their key enablers for secure and dependable SDNs are also identified and presented.
- Finally, we present open security issues raised in SDNs for security researchers and practitioners around the globe.

The remainder of this article is organized as follows. We introduce a simplified overview of SDN architecture to provide the fundamental background to the reader. We then present the state-of-the-art security solutions. Next, thematic taxonomy of SDN security solutions is devised. We discuss the security threats and possible attacks on SDNs. We also discuss the requirements and key enablers for SDN security. We highlight open issues in securing SDNs, and the article is concluded.

A Simplified View of SDN Architecture

This section provides a brief fundamental discussion of SDN architecture so that readers can better comprehend the security concerns with respect to SDN architecture. SDN is an emerging networking paradigm that separates the control plane from the data plane and provides programming ability on the control plane [1] . The most simplified view of the SDN architecture mainly comprises three planes with their corresponding connected interfaces, as shown in Fig. 1.

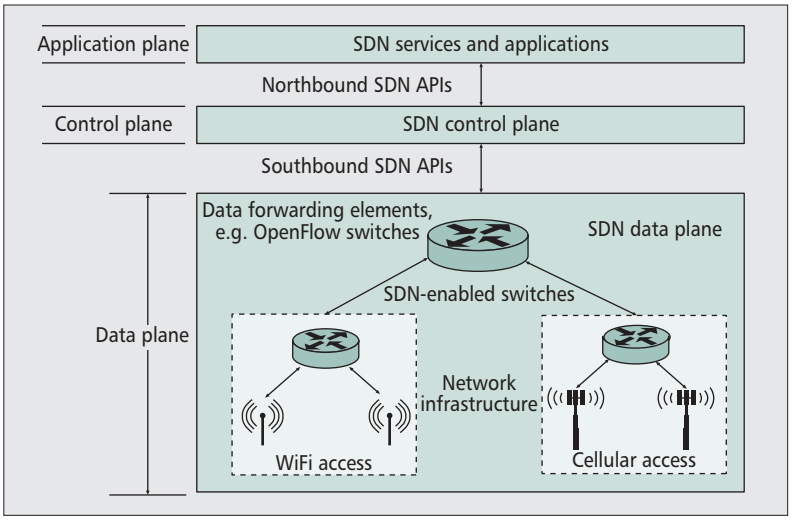

Figure 1. A simplified view of SDN architecture.

The application plane is also known as the application layer, which provides a set of services and applications such as an intrusion detection system (IDS), intrusion prevention system (IPS), deep packet inspection (DPI), load balancers, security monitoring, and access controls. The second most important component of the network is the control plane, which is also known as the control layer of SDN. The control layer is the central layer and comprises the controller. The controller is a software platform and is considered the brain of an SDN. This is the central decision point and is responsible for establishing and terminating flows and paths in SDNs. The management functionality of SDN is simply placed in the central logical controller, which also facilitates the network's programmability. This layer also provides an abstraction of the underlying resources.

Moreover, the data plane comprises the underlying network infrastructures and is known as the infrastructure layer of SDN. This layer comprises the forwarding hardware such as switches and routers. Since the control functionality is placed in the controller, the underlying hardware is only held responsible for forwarding. The infrastructure layer implements the management functionality of the controller through SDN-enabled switches to forward the data, collect the network information, and send it to the control layer. The southbound interface is an application programming interface (API) that provides a link between the control layer and the infrastructure layer. However, the northbound interface enables communication between the control layer and application layer.

State-of-the-Art SDN Security Solutions

In this section, we present state-of-the-art security solutions for securing SDNs. Although SDNs are not mature enough and security is not considered as part of the initial design, the literature clearly shows that two opinions exist in the research community. One trend in research is

SDN is an emerging networking paradigm that separates the control plane from the data plane and provides programming ability on the control plane. The most simplified view of the SDN architecture mainly comprises of three planes with their corresponding connected interfaces.

Security solution classification	Security solutions	SDN layer/interface				
		Application layer	Northbound interface	Control layer	Southbound interface	Data layer
Secure design	FRESCO	✓	✓	✓	✓	
	FortNox		✓	✓	✓	
Security audit	Verificare		✓	✓	✓	✓
	SDN Debugger	✓			✓	
Security enforcement policy	FLOVER	✓	✓	✓	✓	
	PermOF	✓	✓	✓	✓	
	VeriFlow	✓	✓	✓	✓	
Security enhancement	FleXam	✓		✓	✓	✓
	CloudWatcher	✓	✓	✓	✓	
	L-IDS	✓		✓	✓	✓
Security analysis	OpenWatch	✓		✓	✓	✓
	AVANT-GUARD			✓	✓	✓
	Header Space Analysis	✓		✓	✓	✓

Table 1. Comparison of state-of-the-art SDN security solutions.

more curious about securing SDNs. On the contrary, the other school of thought believes the use of SDNs improve and enhance security. Table 1 presents a comparative summary of the state-of-the-art SDN security solutions.

Some of the state-of-the-art solutions for securing SDNs are discussed below.

Secure Design of SDN

The efforts put forward for a secure design of SDN are extremely limited. Shin *et al.* propose FRESCO, a security-specific application development framework for OpenFlow networks [2]. FRESCO facilitates exporting the application programming interface (API) scripts, which enables security experts to develop threat detection logic and security monitoring as programming libraries. However, the framework uses FortNox, a security enforcement kernel [3]. It is an enforcement engine responsible for avoiding rule conflicts arising from different security authorizations. The research work in this category is entirely based on the above two proposals (FRESCO, FortNox) for secure design of SDNs. The first proposal is a major contribution toward secure programming, and has a direct impact on the application layer, control layer, and the interfaces between the two layers except the data layer. The second proposal is more toward rule conflicts and authorization having more of an effect on the control layer, and south and northbound interfaces. It does not, however, improve the security of the application and infrastructure layers.

Implementation of Satisfactory Audit

R. Skowyra *et al.* developed a model that satisfies all requirements of a system design [4]. They discuss an infrastructure tool to specify and analyze a real environment without relying on previous knowledge of formal languages or logic. The authors further give an example of using an OpenFlow-based network of learning switches to enable communication between mobile nodes. This proposal considers the verification of network correctness and specification modeling while considering the scalability issues of OpenFlow networks. Another contribution is presented in [5], which allows the software developers of the SDN to trace the root cause of bugs by reconstructing the series of events causing that particular bug. The packet back-trace assists SDN programmers in resolving logical errors, helps switch implementers to resolve the protocol compatibility errors, and facilitates network operators in submitting complete bug reports to vendors.

Enforcement of Security Policy

Enforcement of security policy is a serious issue in the dynamic environment of SDNs, and the research community has given considerable attention to this area. Son *et al.* [6] propose FLOVER, a model checking system that verifies the flow policies against the network's security policies. D. Kreutzer *et al.* [1] discuss other major contributions in this particular area. The VeriFlow scheme is used for verification of real-

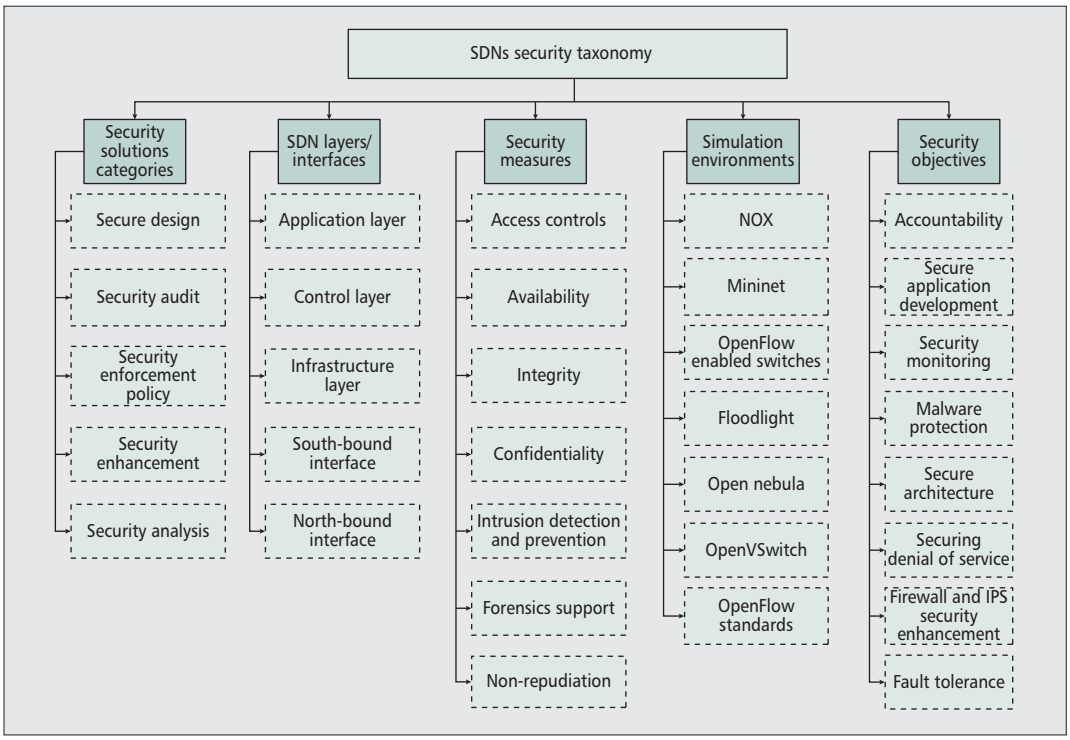

Figure 2. Taxonomy of SDN security solutions.

The major security objectives are auditing and accountability. The other security objectives include rapid designing and development of secure applications, monitoring for security purposes, and malware protection to avoid stealthy scanning and propagation.

time invariants. The paper also presents flow-based policy enforcement using language-based security. Moreover, the authors also discuss the verification of the isolation of program traffic. They also present use of binary decision diagrams for handling the misconfiguration of intra-switch for a single flow.

Another major security contribution is of PermOF [7], a fine-grained permission system that comprises a set of OF-specific permissions and a runtime isolation mechanism for applying the permissions. The set of OF-specific permissions are designed considering four different aspects:

- Threat model
- Controller implementation API set
- Application functional requirements
- Control messages in OpenFlow

The proposed isolation mechanism isolates the controller and applications in a thread container. The applications cannot call controller procedures or directly refer to the memory of the kernel. The application and operating system are also isolated by introducing a shim layer called an access control layer between them. The shim layer is controlled by the kernel of the controller.

Security Enhancement

Sajad *et al.* propose FleXam [8], a sampling extension for OpenFlow that provides access to the controller of OpenFlow to get packet-level information. FleXam enables the controller to sample the packets stochastically or deterministically considering the application requirements. Consequently, such applications can directly run on a small network's controller. Moreover, FleXam eliminates flow setup time and reduces the control plane load. Other prominent solutions for security enhancement are CloudWatcher [9] and L-IDS [1]. CloudWatcher is a framework for monitoring clouds, whereas L-IDS, a learning intrusion detection system, is a security service embedded to protect mobile devices in a particular location.

Security Analysis

J. Wang *et al.* [10] propose a systematic approach to detect and resolve conflicts in an SDN firewall by checking firewall authorization space and flow space. The approach searches the flow paths in the whole network and checks the paths against all firewall deny rules to determine the conflicts with the firewall deny rules. The conflict resolution strategies vary with the operations involved in the flow entries and flow rules. The effectiveness and efficiency of the proposed approach is investigated using header space analysis.

S. Shin *et al.* [11] proposed two significant changes in SDN. One extension is to the data plane, called connection migration, which significantly minimizes data-to-control-plane interactions, which increase during denial-of-service (DoS) attacks on the southbound interface. Another extension, called an actuating trigger, is to expedite the responsiveness to the changing flow dynamics within the SDN data plane. Actuating triggers are introduced over the statistics collection services of the data plane. Another credible solutions is OpenWatch [12], an adaptive method for flow counting to detect anomalies in SDN.

Taxonomy of SDN Security

The taxonomy is based on the literature of SDN security as shown in Fig. 2. The existing solutions can be categorized based on the following

Security threats	Protection techniques	Security requirements	Affected functionalities	Targeted layers/interfaces
Operating system alteration	Trusted computing	System integrity protection	Application management	All layers
Software framework alteration	Trusted computing	System integrity protection	Application management	All layers
Software failure	High assurance	Robustness, system integrity protection	All functionalities	All layers
Hardware failure	High assurance	Robustness, system integrity protection	All functionalities	Control layer, data layer
Configuration data alteration	Data integrity functionality in SDN middleware	Data integrity protection	Resource management, application management	Control layer, control-data interface, and data layer
Configuration data extraction	Data integrity functionality in SDN	Confidentiality protection	Data management	Control layer, control-data interface, and data layer
Unauthorized access to SDN services	Deploying secure administration module	Identities verification, ensuring system integrity	All functionalities	All layers and interfaces
User data alteration	Data integrity functionality in SDN	Ensuring data integrity	Data management	Data layer
Masquerading as authorized SDN controller	Use of digital signatures for SDN software modules	Ensuring system integrity, Identities verification, accountability	Application management	Control layer, control-data interface, and data layer

Table 2. Security threats in SDNs.

parameters: solution categories, SDN layers/interfaces, security measures, simulation environment, and security objectives. Secure design is the primary issue of SDN, although it is not considered as part of the initial design. There are very few proposals on secure design of the SDN. The contributions in this area are still limited, and it deserves more comprehensive research attention. Researchers have also contributed toward auditing of SDN environment for security and accountability purposes. Moreover, major contributions of the researcher are based on security enforcement policy, and this particular area really incurs attention in the dynamic SDN environment. Besides, much of the work is done on security enhancement using SDN, and very few contributions are on security analysis. The security solutions also address issues of different layers/interfaces of SDN. Different layers/interfaces are targeted for defense against various possible attacks and exploitations. However, these security research contributions defending each layer/interface are mainly based on the broad mechanisms of security. The security mechanisms are given as access control, authentication, authorization, encryption, intrusion detection, intrusion prevention, and recovery. The taxonomy also presents classification based on the simulation/emulation environment. The majority of the networks are designed following the OpenFlow (OF) standards to conduct their corresponding experiments. The major security objectives are auditing and accountability. The other security objectives include rapid design and development of secure applications, monitoring for security purposes, and malware protection to avoid stealthy scanning and propagation. Besides, securing the architecture of OF networks and defense against different DoS attacks are considered objectives.

Security Threats and Possible Attacks in SDN

In this section, we present some of the possible threats and attacks in SDN that are presented in Tables 2 and 3, respectively. Operating system alteration represents the destruction or alteration of components or the complete operating system of SDN elements such as a controller or forwarder nodes. The operating system can be made secure by ensuring system integrity, which can be achieved by implementing trusted computing. The threat can target all layers of SDN and can affect the management of running services and applications in SDNs. Software framework alteration identifies the destruction or alteration of middleware and components of the software framework. Similar to the operating system alteration threat, the software framework alteration can also be made secure by protecting system integrity through trusted computing. Software framework alteration also affects all layers of SDN.

The software failure threat represents a general software failure in any of the software components comprising the software framework, applications, and operating system. The threat can be mitigated by employing high assurance techniques and ensuring the robustness of a sys-

SDN layers/interfaces		Application layer	Northbound interface	Control layer	Southbound interface	Data layer
Policy enforcement related attacks				✓	✓	✓
Availability related attacks		✓	✓	✓		
Authorization related attacks				✓	✓	✓
Authentication related attacks		✓	✓	✓		
Data alteration related attacks				✓	✓	✓
Nasty applications	Fake rule insertion	✓	✓	✓		
	Hijacking the controller			✓	✓	✓
Side-channel attacks						✓

Table 3. Possible security attacks in SDNs.

The threat can be mitigated by deploying a secure administration module and ensuring system integrity. The security requirement for mitigating the threat is identification verification. The threat can affect all functionalities, and can target all layers and interfaces of SDN.

tem. The entire set of functionalities is affected by the software failure threat, and the threat can target all layers of SDN architecture. The hardware failure threat represents the generic failure of hardware in any of the components. Similar to the software failure threat, the hardware failure threat can be reduced by employing high assurance techniques and ensuring the robustness of a system. The entire set of functionalities is affected by the hardware failure threat, which can target control and data layers.

The configuration data alteration threat represents the destruction or alteration of configuration data that is required by SDN to perform different functions. Configuration data can be removed or modified from the SDN platform. The threat can be mitigated by ensuring data integrity in SDN middleware. The threat can target the control layer, control-data interface, and data layer, and can affect resource and application management. The configuration data extraction threat is an eavesdropping threat, where an attacker gathers configuration data that can be used in subsequent attacks. Configuration data extraction requires confidentiality protection and ensuring data integrity functionality in SDN. The threat targets the control layer, control-data interface, and data layer. Unauthorized access to SDN services identifies a security breach, where an authorized SDN entity can access services of SDN for which the entity does not have the proper access level. The threat can be mitigated by deploying a secure administration module and ensuring system integrity. The security requirement for mitigating the threat is identification verification. The threat can affect all functionalities and can target all layers and interfaces of SDN.

User data alteration is a threat that represents the destruction or alteration of user data such as customized profiles of user traffic. The user data alteration threat can be mitigated by ensuring data integrity. The threat can affect the data management and target the data layer. Masquerading as an authorized SDN controller identifies the activation of malicious software on an SDN platform such as the controller platform. The threat can be mitigated by use of digital signatures of SDN software modules. The threat mitigation requires the assurance of system integrity, identities verification, and accountability. Application management is affected by the threat activation, which can target the control layer, control-data interface, and data layer of SDN.

Apart from security threats, there are a number of possible attacks in SDN. The checks in Table 3 show these security attacks affecting the corresponding layers/interfaces in the following example scenarios. Even so, these security issues may potentially affect each layer/interface of the SDN. Availability-related attacks refer to various DoS attacks. For example, the communication flooding attack is possible between the switch and the controller, and will affect all the corresponding three layers. However, this flooding attack can also be generated through switch flow tables, which will ultimately affect the data layer only. Moreover, due to the dynamic SDN environment, security policy enforcement related attacks will affect the upper three layers, if the underlying two layers are not protected by Transport Layer Security (TLS) security or other authentication techniques. However; the existing higher version of OpenFlow protocol uses TLS between the data and control layer. Furthermore, authorization related attacks can lead to illegal access to the controller, which will certainly affect the lower three layers. Unauthentic applications can cause damage to the corresponding three higher layers. Data alteration can be done through the modification of flow rules and will cause damage to the data layer only. However, a malicious application can target all the corresponding layers, and this will remain a challenge for SDNs. Malicious applications can be used to insert fake rules affecting the upper three layers/interfaces and can also be used to hijack the controller, which will affect the layers down the stream. There is also a possibility of side channel attacks. For instance, an input buffer can be used for the discovery of flow rules, and analyzing the time of packet processing may lead to discover the forwarding policy.

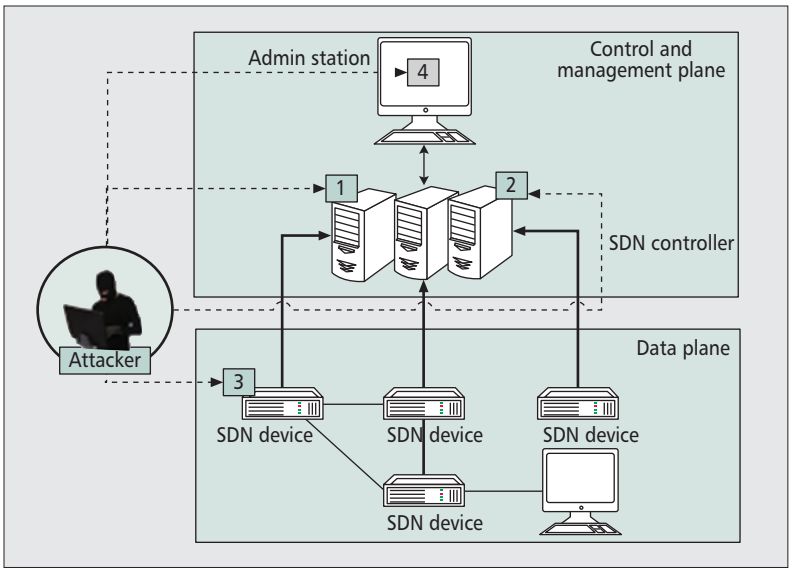

Figure 3. An attacker searching for potential targets.

Requirements and Key Enablers for SDN Security

To build a secure SDN environment, it is essential to ensure the security of each and every component of SDN. Following are some of the essential security requirements for securing the SDN key components. Figure 3 depicts the attacker searching for the potential components of SDNs to be compromised.

Securing the SDN Controller

Securing the SDN controller is the top priority. The SDN controller is responsible for the overall management of the network. Moreover, the controller is a central decision point. The compromise of a centralized SDN controller can simply lead to the disaster of the entire network. Besides, being a single point of failure, it serves as a potential target for attackers. The availability of the SDN controller is of serious concern for the whole network. The SDN controller, as a software platform, essentially supports a hacker in reconfiguring the complete network. By spoofing the address of an SDN controller, an attacker can simply take over the entire network by means of a fake controller. This key component deeply needs defense in depth, which may include the protection of the system containing the SDN controller from physical and external threats. While ensuring the availability of the controller, it must be protected from different DoS and distributed DoS (DDoS) attacks. The operating system must be secure, having no patches and back-door accounts and open doors at the same time, such as vulnerable open ports, services, and protocols.

Protecting the Flow Paradigm of the SDN

SDN is grounded on flow-based forwarding and can certainly ensure end-to-end communication security. The flow paradigm acts as the soul of an SDN, and it must be protected. A successful injection of bogus flow may lead to a total network disaster. The flow abstraction shown by the controller may easily lead to harvesting the intelligence of the connected resources, which can be used in further attacks and exploitations. An up-to-date access control mechanism should be deployed in the network. Moreover, flows should be encrypted to avoid injecting malicious flows. Proper authentication and authorizations should be implemented to avoid side channel attacks.

Fortifying SDN Agents

An SDN agent's security is essential as it constitutes the environment. To compromise a strong entity such as an SDN controller, an attacker, as part of an attack strategy, may start to reach the target by compromising any vulnerable agent of SDN. Moreover, the consequences of the attack must be serious, particularly when the attacker is a "man at the end" [13]. For instance, Link Layer Discovery Protocol packets with forged source addresses can lead the SDN controller to install flow rules grounded on bogus information. Moreover, many existing switches, as part of the SDN infrastructure layer, are by default in listener mode, which may easily lead to launching malicious connections. Besides, injecting a false flow at any SDN agent can lead to its distribution to numerous agents that ultimately cause serious disturbance. The security of SDN agents requires deploying the latest identity management, threat isolation, and mitigation techniques. Moreover, an SDN agent also requires physical security. Also, IPS, IDS, and firewalls should be actively deployed.

Hardening Application Programming Interfaces and Communication Channels

Application programming interfaces (APIs) can also be a potential target for attackers. Most importantly, southbound APIs can easily be targeted for different DoS attacks to make the whole network unavailable. Building malicious APIs by skilled programmers is more critical, and the trend is already in the security research community. The communication channel between each layer must be well protected; for example, in the OpenFlow protocol, it is protected by TLS. However, OpenFlow must not be the only protocol considered for SDN; there can be other options. The security measures include secure coding, deployment of integrity checks, and, most important, digital signing of the code. Moreover, the communication channels can be hardened using TLS security.

Open Issues for Securing SDNs

Security plays a vital role in tremendously deploying SDNs across different networks. The use of SDNs in network virtualization has brought many security issues, and as soon as it proliferates, it becomes increasingly important and more vibrant [14, 15]. This is in spite of the fact that the level of security of virtual networks compared to traditional networks is not shown by the research community; also, there are not adequate measures to secure them. This section discusses the unaddressed open issues surrounding SDNs when deployed for network virtualiza-

tion. However, to prove this relevant open issue, we consider discussing these issues by illustrating a real experience of OpenFlow already used for network virtualization.

SDN Controller-Specific Security Issues in Virtual Environments

A controller will always remain a potential target for the attacker and most probably the first line of attack. Following are some of the unaddressed scenarios in targeting a logical controller in a virtual environment as the most significant part of an SDN, as demonstrated in Fig. 4. Consider the particular scenario of using OpenVirtex, a special controller used for creating virtual networks. In the case of using OpenVirtex, some controllers in Fig. 4, such as floodlight and others, are placed on the end user side. The virtualization layer of OpenVirtex demonsterates the concept of programmable virtual networks. Moreover, the physical layers simply show the physical hardware that can be specified with the user's topology and addressed using OpenVirtex. Although it has many advantages, at the same it is exploitable. The diagram is used to clearly depict the following issues related to SDN while creating a virtual network.

Denial of Service Attack — Take an example of a POX controller, a special controller placed on the end user side for creating a particular virtual network using OpenVirteX. A POX controller possesses critical knowledge of the network and is prone to many attacks, more specifically DoS attacks. An adversary can simply generate a huge number of flows, rendering network breakdown or improper functioning.

Spoofing Attack — We keep in view the same scenario. Since the floodlight is already aware of the IP address of the OpenVirteX, the floodlight controller can simply forge the IP address of the main virtual controller (OpenVirteX) to launch a simple and easy spoofing attack.

Malicious Injection — Malicious injections can be made by simply following the existing Field Rewrite problem, which allows the end user to change their VLAN ID tag subject in particular circumstances. This situation simply creates an opportunity for a nasty controller to inject packets into another slice. Moreover, using OpenVirtex instead of FlowVisor does a good job of addressing space isolation, but at the same time, it does not implement action isolation. Subsequently, a controller can set any type of action in the flow entry without any control of the controller in this particular case.

Conclusions

The emergence of SDN is imposing new requirements for network security due to newly deployed infrastructural entities and architectural components. In order to meet the newly imposed network security requirements, several solutions for securing SDNs have been proposed. A discussion of state-of-the-art security solutions is presented to help the reader understand the recent efforts in securing the SDNs. We provide a tutorial on research efforts made in this direction. In this article, we also classify the state-of-the-art security solutions by devising a thematic taxonomy considering SDN layers/interfaces, security mechanisms, simulation environments, and security objectives. The article discusses the possible attacks and threats targeting different layers/interfaces of SDNs. Moreover, the requirements for securing the SDNs are also identified and presented. Finally, we discuss the open research issues to give researchers directions for future research.

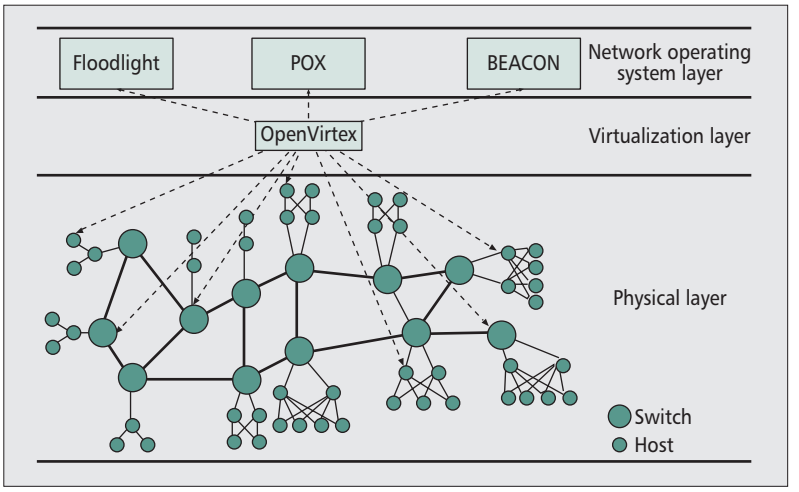

Figure 4. Attack scenario of creating virtual environment using OpenVirtex.

Acknowledgments

This work is fully funded by Bright Spark Unit, the University of Malaya, Malaysia, and partially funded by the Malaysian Ministry of Higher Education under the University of Malaya High Impact Research Grant UM.C/625/1/HIR/MOE/FCSIT/03.

References

[1] D. Kreutz *et al.*, "Software-Defined Networking: A Comprehensive Survey," *Proc. IEEE*, vol. 103, vol. 1, pp. 14–76.
[2] S. Shin *et al.*, "Fresco: Modular Composable Security Services for Software-Defined Networks," *Internet Society NDSS*, 2013.
[3] P. Porras *et al.*, "A Security Enforcement Kernel for OpenFlow Networks," *Proc. 1st Wksp. Hot Topics in Software Defined Networks*, ACM, 2012, pp. 121–26.
[4] R.W. Skowyra *et al.*, "Verifiably-Safe Software-Defined Networks for CPS," *Proc. 2nd ACM Int'l. Conf. High Confidence Networked Systems*, ACM, 2013, pp. 101–10.
[5] N. Handigol *et al.*, "Where is the Debugger for My Software-Defined Network?," *Proc. 1st Wksp, Hot Topics in Software Defined Networks*, ACM, 2012, pp. 55–60.
[6] S. Son *et al.*, "Model Checking Invariant Security Properties in OpenFlow," *Proc. IEEE ICC, 2013*, 2013, pp. 1974–79.
[7] X. Wen *et al.*, "Towards a Secure Controller Platform for Openflow Applications," *Proc. 2nd ACM SIGCOMM Wksp. Hot topics in Software Defined Networking*, 2013, pp. 171–72.
[8] S. Shirali-Shahreza and Y. Ganjali, "FleXam: Flexible Sampling Extension for Monitoring and Security Applications in Openflow," *Proc. 2nd ACM SIGCOMM Wksp. Hot Topics in Software Defined Networking*, 2013, pp. 167–68.
[9] S. Shin and G. Gu, "CloudWatcher: Network Security Monitoring Using OpenFlow in Dynamic Cloud Networks (or: How to Provide Security Monitoring as a Service in Clouds?)," *2012 20th IEEE Int'l. Conf. Proc. Network Protocols*, 2012, pp. 1–6.
[10] J. Wang *et al.*, "Towards a Security-Enhanced Firewall Application for OpenFlow Networks," *Cyberspace Safety and Security*, Springer, 2013, pp. 92–103.

OpenFlow must not be the only protocol considered for SDN; there can be other options. The security measures include secure coding, deployment of integrity checks, and, most importantly, digital signing of the code. Moreover, the communication channels can be hardened using TLS security.

[11] S. Shin *et al.*, "AVANT-GUARD: Scalable and Vigilant Switch Flow Management in Software-Defined Networks," *Proc.2013 ACM SIGSAC Conf. Computer & Communications Security*, 2013, pp. 413–24.

[12] Y. Zhang, "An adaptive Flow Counting Method for Anomaly Detection in SDN," *Proc. 9th ACM Conf. Emerging Networking Experiments and Technologies* 2013, pp. 25–30.

[13] A. Akhunzada *et al.*, "Man-At-The-End Attacks: Analysis, Taxonomy, Human Aspects, Motivation and Future Directions," *J. Network and Computer Applications*, 2014.

[14] L. Wei *et al.*, "Security and Privacy for Storage and Computation in Cloud Computing," *Info. Sciences*, vol. 258, 2014, pp. 371–86.

[15] Q. Duan, Y. Yan, and A. V. Vasilakos, "A Survey on Service-Oriented Network Virtualization Toward Convergence of Networking and cloud Computing," *IEEE Trans. Network and Service Management*, vol. 9, no. 4, 2012, pp. 373–92.

Biographies

Adnan Akhunzada (a.adnan@siswa.um.edu.my) is currently a Ph.D. fellow and an active senior researcher in the Centre for Mobile Cloud Computing, University of Malaya, Malaysia. He had a great experience teaching international modules of the University of Bradford, United Kingdom. He is a senior lecturer at CIIT, Islamabad. He has published several high impact research journal papers. His current research interests include secure design and modeling of software defined networks, man-at-the-end attacks, lightweight cryptography, human attacker attribution and profiling, and remote data auditing.

Ejaz Ahmed (ejazahmed@ieee.org) is a Ph.D. candidate at the University of Malaya. Before that, he worked as a research associate at CogNet (Cognitive Radio Research Lab) SEECS, NUST Pakistan, and CoReNet (Center of Research in Networks and Telecom), Maju, Islamabad, Pakistan. Currently, he is an active researcher at the Centre for Mobile Cloud Computing Research (C4MCCR), University of Malaya, Kuala Lumpur, Malaysia. His areas of interest include software-defined networks, cognitive radio setworks, and mobile cloud computing.

Abdullah Gani (abdullahgani@ieee.org) is an associate professor at the Department of Computer System and Technology, University of Malaya. His academic qualifications were obtained from UK's universities: Bachelor's and Master's degrees from the University of Hull, and a Ph.D. from the University of Sheffield. He has published more than 100 academic papers in conferences and respected journals. His research interests include self-organized systems, reinforcement learning, and wireless-related networks.

Muhammad Khurram Khan (mkhurram@ksu.edu.sa) is working at the Center of Excellence in Information Assurance, King Saud University, Saudi Arabia. He has edited seven books and proceedings published by Springer-Verlag and IEEE. He has published more than 200 papers in international journals and conferences. He is an inventor of 10 U.S./PCT patents. He is Editor-in-Chief of a well-reputed journal, *Telecommunication Systems* (Springer), and a member of several editorial boards. His research interests include cybersecurity, biometrics, multimedia security, and digital authentication.

Muhammad Imran (cimran@ksu.edu.sa) has worked as an assistant professor in the Deanship of e-Transactions and Communication at King Saud University since 2011. His research interests include mobile ad hoc and sensor networks, cognitive radio ad hoc networks, WBANs, multihop wireless networks, and fault-tolerant computing. He has published more than 45 research papers in reputable international journals, conferences, and workshops. He currently serves as an Editor for the *International Journal of Information Technology and Electrical Engineering*.

Sghaier Guizani (sguizani@alfaisal.edu) is an assistant professor at the Electrical Engineering Department, Alfaisal University, Riyadh, Saudi Arabia. He received his B.S. from the State University of New York at Binghamton in 1990, his M.S. from North Carolina State University in 1992, and his Ph.D. from the University of Quebec, Trois Rivieres, Canada, in 2006, all in electrical and computer engineering. His research interests are in the areas of wireless communication, computer networks, computer security, RFID, and optical fiber communication systems.

Evolving Defense Mechanism for Future Network Security

Haifeng Zhou, Chunming Wu, Ming Jiang, Boyang Zhou, Wen Gao, Tingting Pan, and Min Huang

ABSTRACT

The past few years have witnessed revolutionary advances in network technology. Along with new techniques such as SDN come lots of new network security challenges. Conventional network security mechanisms are incompetent to overcome these challenges, since they are built on a static network configuration that facilitates attackers in finding the weaknesses of a network. In this article, we conceive a novel conceptual network security mechanism, the evolving defense mechanism (EDM), to resolve current and future security problems. EDM is based on a bio-inspired idea of network configuration variations. According to the security requirements of the system, the user, and the network security state, EDM selects an efficient network configuration variation strategy to prevent corresponding security threats. Combined with SDN implementation, EDM resolves security problems from a new angle and is capable of evolving with new network security technology. We sketch a way to implement EDM and present its reference framework, which serves as an ecosystem and coexisting environment for various kinds of network configuration variations. The proposed mechanism avoids the deficiency of conventional mechanisms and has potential to cope with emerging security threats.

INTRODUCTION

The Internet has reached almost every corner of the world, and provides great benefits in our daily lives. However, it also exposes us to security and privacy problems. Currently, network technology is experiencing a revolutionary development, generating new networks such as software-defined networks (SDNs). The SDN is based on a framework that is completely different from the current Internet framework. The separation of control and data planes in SDNs brings numerous advantages, such as programmability and better evolvability [1]. However, it also contributes to more complex security problems due to the need to protect controllers, switches, and the communications between them [2]. Even worse, the pervasive network infrastructure sharing in data centers augments existing security threats, since a successful intrusion brings about more serious financial damage. Unfortunately, a conventional network security mechanism is incapable of overcoming the increasingly complex and severe security problems, since it is built on a static network configuration, such as the static assignment of IP addresses. This inherent deficiency facilitates attackers in finding weaknesses of a network and thus gives rise to a situation in which attacks are much easier than defenses. Consequently, we need to consider the network security from a new perspective.

The biological system has evolved over billions of years, which makes it the most complex and large-scale system. The similarities between the biological and network systems shed light on the development of new network security mechanisms. Researchers have made significant progress in bio-inspired solutions for networking [3]. Currently, biologically inspired ideas have been used in the design of the next generation network [4, 5]. However, these research works focus on the overall principal design of the future network without discussing network security mechanisms. The existing research on network security mechanisms from a biological perspective reveals the importance of cooperation between different security devices after observation of the cooperation of different components of biological immune systems [3]. Nevertheless, these researchers ignore the deficiency of the static network configuration. To overcome this problem, the moving target defense (MTD) [6] is proposed to safeguard the network from attacks by using dynamic and random network configurations. However, three critical problems in MTD inhibit its development in practice. First, since the dynamic and random network configuration will introduce extra operational cost, MTD fails to make precise estimations of the cost and profit. Second, MTD lacks the evolutional capability to deal with future security problems. Evolutional capability is extremely important because future attack techniques may even find a way to cope with original dynamic network configurations. Third, MTD also fails to consider the different efficacy of different combinations of dynamic network configurations to eliminate the same security threat. This kind of consideration is beneficial to minimization of the cost. In this article we systematically take all these aspects into consideration. In addition, the

Haifeng Zhou, Chunming Wu, Boyang Zhou, Wen Gao, Tingting Pan, and Min Huang are with Zhejiang University.

Ming Jiang is with Hangzhou Dianzi University.

0163-6804/15/$25.00

From the perspective of network security, we can see the biological system as a network system and see the change of a certain environmental condition as a certain kind of atomic network attack. Similarly, the change extent of the environmental condition corresponds to the threat extent of the atomic attack.

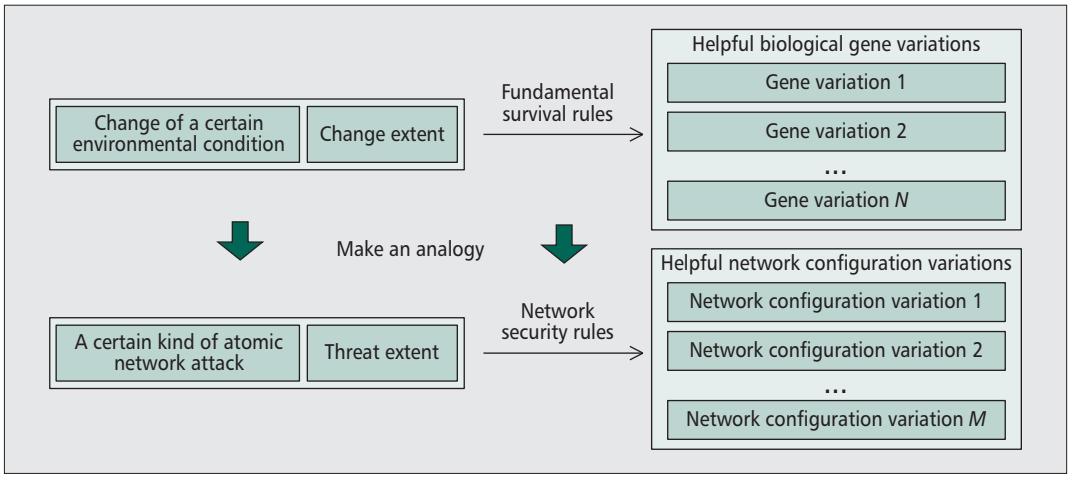

Figure 1. Inspiration from the biological system.

preliminary trials of MTD in an SDN environment is made in [7, 8], whereas only the dynamic and random configurations of IP addresses and routes are implemented.

In this article, we conceive a novel conceptual network security mechanism, the evolving defense mechanism (EDM), based on up-to-date security techniques. EDM adopts the idea of network configuration variations, inspired by biological gene variation and its critical function for the continuation of the biology system in an ever changing environment. EDM is equipped with processes to detect ongoing attacks and select an efficient network configuration variation strategy to eliminate corresponding security threats according to the security requirements of the system, the user, and the detected network security state. The reference framework based on SDN is then designed to offer an ecosystem and coexisting environment for various network configuration variations. Benefiting from the logically centralized control plane and current development of security techniques in the SDN environment such as monitor technique [9] and intrusion detection system (IDS) [10], EDM is promising for implementation. Even more, due to the programmability of SDN, EDM can be progressively augmented by new kinds of network configuration variations and variation strategies, which makes it evolvable as the future development of network security technology.

Inspiration from the Biological System

The biological system always has effective solutions to cope with ever changing environmental conditions. In this process biological gene variation plays a dominant role. We note that some helpful gene variations of a creature make it possible to survive from the change of a certain environmental condition, such as temperature, humidity, and water PH level, over a period of time. In addition, the fundamental survival rules in the biological system play a role in mapping from the changed environmental condition and its change extent to the creature's helpful gene variations.

From the perspective of network security, we can see the biological system as a network system and see the change of a certain environmental condition as a certain kind of atomic network attack, as shown in Fig.1. Similarly, the change extent of the environmental condition corresponds to the threat extent of the atomic attack. In addition, the network security rules play a role in mapping from the atomic attack and its threat extent to helpful network configuration variations, such as variations of IP addresses and routes, used for eliminating the threat.

Based on the concept of the network configuration variation, we draw inspiration from genetic engineering. Implanted with helpful genes in some crops' chromosomes, these crops' resistibility against drought, frost, diseases, and pests can be improved. The implanting process, based on prior knowledge of the functions of the helpful genes, helps those crops achieve beneficial gene variations. Based on the same idea, if we have prior knowledge of the set of helpful network configuration variations against a certain kind of atomic attack, we can forbid the atomic attack effectively and efficiently by performing the variations.

Evolving Defense Mechanism

Principles

The general processes in EDM are shown in Fig. 2. In the first process, EDM detects ongoing network attacks. Current common network attacks are always composed of seven attack steps: scanning, fingerprinting, correlation, coordination, attack action, reporting, and propagating [6]. Each attack step represents a certain category of the above mentioned atomic attacks, which has a similar attack pattern and thus can be eliminated by similar kinds of network configuration variations. EDM is designed to perceive ongoing network attacks by detecting abnormal events caused by the atomic attacks of each attack step (or category). Detection can be performed by multiple approaches, such as monitoring techniques [9], an intrusion detection system (IDS) [10], software sensors embedded in network infrastructure, and firewalls. By this process, EDM can obtain information, including the

ongoing abnormal events from each of the seven attack steps and the involved infrastructure.

Figure 3 illustrates the second process. First, the detection results are delivered to the pre-established table of threat types to determine the threat types of the atomic attacks. The threat types play a role in further categorizing the atomic attacks of a certain attack step. If several kinds of atomic attacks of the attack step are the same threat type, it means that the same network configuration variations can be used to eliminate them. Second, according to the information including the threat types and the involved infrastructure, EDM then finds the corresponding item in the pre-established table of network configuration variation strategies, as shown in Fig. 3. It thus obtains the network configuration variation strategy, including the types of helpful network configuration variations and the variation attributes. The establishment of the tables of network configuration variation strategies and threat types is extremely critical, since it determines the efficiency and efficacy of EDM. A potential establishment approach is introduced in the next section.

In addition, the selection of network configuration variation strategies also needs to take the security requirements from the system and the user into consideration. Even though there is no security threat detected, an essential network configuration variation strategy should be performed all the time according to the demands of the system and the user. When security threats are detected, they tell us whether the essential network configuration variation strategy needs to contain more kinds of network configuration variations or needs higher requirements for them, aiming to eliminate ongoing attacks. The selection of the essential variation strategy required by the system and the user is similar to the selection approach introduced here.

The third process is to produce network operations according to the selected network configuration variation strategy, the involved infrastructure, and the network situation (situation of network resources, topology, etc.).

The last process is to perform the operations in these involved network elements in turn.

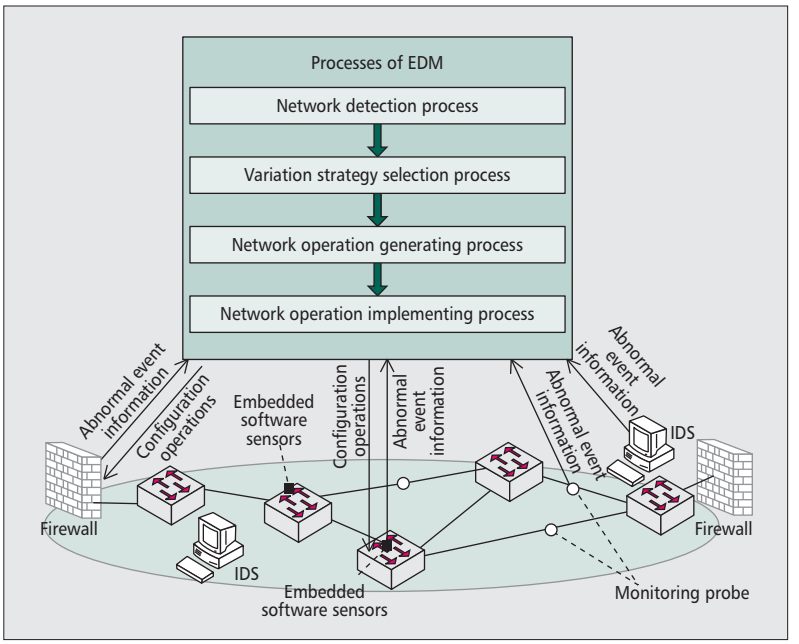

Figure 2. Four processes in EDM.

Network Configuration Variation

We now introduce five common kinds of network configuration variations.

Variation of IP addresses: This variation type brings a dynamic and random configuration of IP addresses. Since current attackers scan a network and target vulnerable hosts by their IP addresses, the dynamic and random configuration of IP addresses with enough change frequency is competent to invalidate the scan. It is also effective to prohibit the denial of service (DoS) and the propagation of viruses and worms.

Variation of routes: This can protect critical communications from monitoring and masquerading. It is easy to understand when we recall the frequency hopping communications used by the army to keep communications safe.

Variation of host responses: This variation type is to intercept and even modify host responses to confuse attackers. Without accurate host information such as operating system (OS) types and versions, running services, and opening ports, attackers are unable to find vulnerable hosts.

Variation of encryption methods: This is used to protect critical communications, such as communications between administrators, controllers, and switches in SDNs, or communications between customers, administrators, and infrastructure in a modern data center. It largely reduces the risk of being cracked from having only one encryption approach.

Variation of authentication approaches: Changing authentication approaches further strengthens the security of access to crucial infrastructure such as controllers in SDNs and vital data resources such as configuration management data in data centers.

More kinds of network configuration variations can be developed on SDN controllers and progressively added to EDM. Furthermore, if EDM is equipped with reconfigurable network devices, it can even reconfigure the architecture of these devices for the sake of security.

Evolvability

EDM is capable of evolving by integrating new network configuration variations and variation strategies. Specifically, new network configuration variations can be developed as special security applications on SDN controllers. By estimating the cost and the profit, new network configuration variation strategies with better performance are permitted to replace original ones in the table of network configuration variation strategies, which makes EDM more efficient in dealing with the original security threats. Furthermore, newly developed network configuration variations can be used to prevent new emerging security threats. Consequently, EDM has the potential to cope with current and future security threats by continually improving itself.

More kinds of network configuration variations can be developed on SDN controllers and progressively added to EDM. Furthermore, if EDM is equipped with reconfigurable network devices, it can even reconfigure the architecture of these devices for the sake of security.

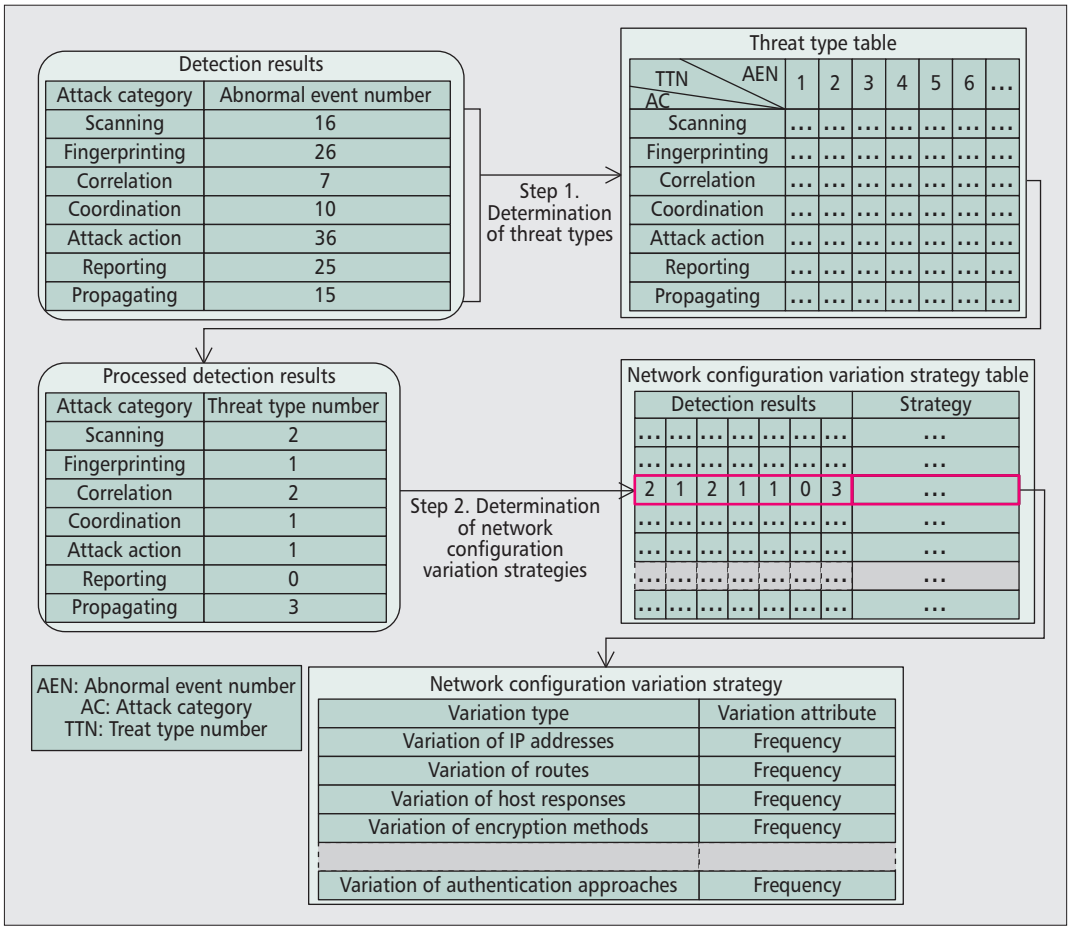

Figure 3. Selection process of the network configuration variation strategy.

Reference Framework

Figure 4 illustrates the reference framework of EDM based on the current SDN controller framework. This framework serves as a coexisting environment for various kinds of network configuration variations. It is mainly composed of four modules: *security information detection, variation strategy selection, variation strategy libraries*, and *configuration generator*. These four modules are able to evolve independently. Other modules, such as *topology manager* and *routing computation*, are common modules of SDN controllers [1]. *Southbound protocol libraries* include protocols like OpenFlow, I2RS, and NETCONF [1].

We clarify the interior processing by the following introduction of the four highlighted modules.

1) Security Information Detection — This module is used to detect abnormal events in a network caused by ongoing atomic attacks. After detection, abnormal network event information is uploaded from such network elements as SDN switches to this module on SDN controllers by the southbound protocol. EDM thus can obtain the abnormal event numbers and involved infrastructure.

Currently, we can achieve this detection via four approaches. First is the network monitor in the SDN environment [9]; second, the IDS in SDN [10]; third, embedding software sensors in network infrastructure (e.g., switches, hosts) to detect abnormal behavior; and finally, the firewall developed as a security application on SDN controllers.

2) Variation Strategy Selection — The primary function of this module is to make the selection of network configuration variation strategies according to the security requirements from the system, the user, and the network security state. First, an abnormal event information from *security information detection* is delivered to *variation strategy libraries* for the selection of a network configuration variation strategy. This selected network configuration variation strategy is used to eliminate ongoing attacks. Meanwhile, the system and the user also have their demands of network configuration variation strategies for the sake of security.

3) Variation Strategy Libraries — The first function of this module is to maintain and update the tables of network configuration variation strategies and threat types. The second function is to transform abnormal event numbers delivered from *security information detection* into corresponding threat type numbers. The third is to search the corresponding item in the table of network configuration variation strategies according to the threat type numbers.

In addition, EDM needs five steps to estab-

lish the two tables, as shown in Fig. 5. First, it collects existing representative network attacks. Second, it decomposes each of these attack samples into atomic attacks. Third, it analyzes and records the features of the atomic attacks of each attack sample. Fourth, it categorizes all the atomic attacks according to similar features, and marks them with different threat type numbers for identification. The table of threat types is thus established. Fifth, it finds an effective and efficient network configuration variation strategy for each threat type of the atomic attack, and thus obtains the table of network configuration variation strategies.

4) Configuration Generator — The first function of this module is to offer a sharing mechanism of mutated network configurations among various kinds of network configuration variations, aiming to avoid any potential collisions between them. For instance, if the variations of IP addresses and routes work together, the variation of IP addresses should share the mutated IP addresses with the other one. The second function is to generate network configuration operations according to the selected network configuration variation strategy, the involved infrastructure, and the actual network situation. This module needs to reorder these network operations for avoiding collisions between them and examines whether network resources are enough to implement all the operations. Through southbound protocols, the network operations are finally delivered from the controller to the involved network elements.

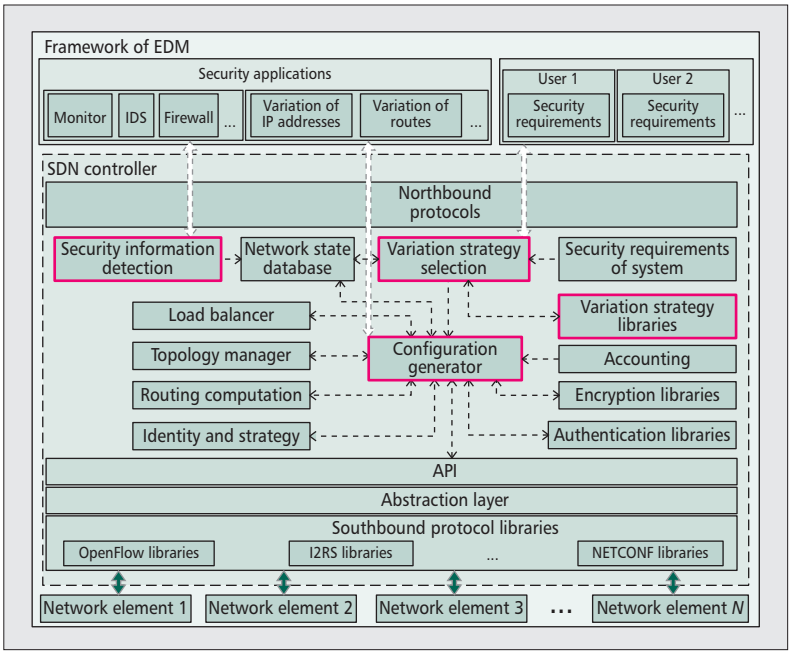

Figure 4. EDM's reference framework.

Effectiveness Analysis

If various kinds of network configuration variations are well coordinated and potential collisions between them are avoided, the effectiveness of EDM ultimately depends on the effectiveness of each kind of network configuration variation. Currently, dynamic and random configurations of IP addresses and routes (i.e., the variation of IP addresses and of routes) have validated their effectiveness in [7, 8]. The former invalidated 99 percent of external scanners and protected 90 percent of network hosts from even zero-day unknown worms [7]. The latter protected up to 90 percent of flow packets from persistent attackers, compared to original least-cost single-path routing schemes [8]. The other three variation types listed in this article are intuitively helpful for preventing corresponding attacks, while the actual efficacy needs further research. We have made a preliminary implementation of EDM and show its effectiveness by a use case. The variation of IP addresses and of routes are implemented in this use case and performed at the same time without any collision. In the following subsections, we introduce what EDM can protect and the use case, respectively.

What EDM Can Protect

Equipped with the listed five variation types, EDM is capable of taking effect in the following five aspects:

1) Interrupting internal and external reconnaissance. Due to variation of IP addresses bringing dynamically changing IP addresses and variation of host responses disturbing attackers, EDM is capable of making the current reconnaissance invalid.

2) Protecting critical communications from monitoring and masquerading. Variations of IP addresses, routes, and communication encryption methods are able to interrupt these kinds of threats.

3) Protecting critical data resources, infrastructure, and administration from security threats. Variation of IP addresses is capable of hiding these critical network nodes from attackers. Furthermore, variations of communication encryption methods and authentication approaches are able to further strengthen their security by preventing monitoring and masquerading.

4) Protecting infrastructure from denial of service (DoS). The initiation of DoS needs specific target IP addresses marked by current attackers. Consequently, variation of IP addresses is able to invalidate it.

5) Restricting the expansion of damage. Even when some network nodes are conquered or compromised, expansion of the attacker's territory is much harder due to the fact that variation of IP addresses is capable of disrupting the communications between the conquered or compromised network nodes. Furthermore, the frequent trials of the communications will raise the precaution of IDS.

A Use Case Study

We now show a use case based on our preliminary implementation of EDM. We implement variation of IP addresses [7] and variation of routes as special security applications on an OpenDaylight controller [1], and use Mininet [1] to generate a network of OpenFlow switches and hosts, as shown in Fig. 6. First, according to the security requirements from the users of hosts 1

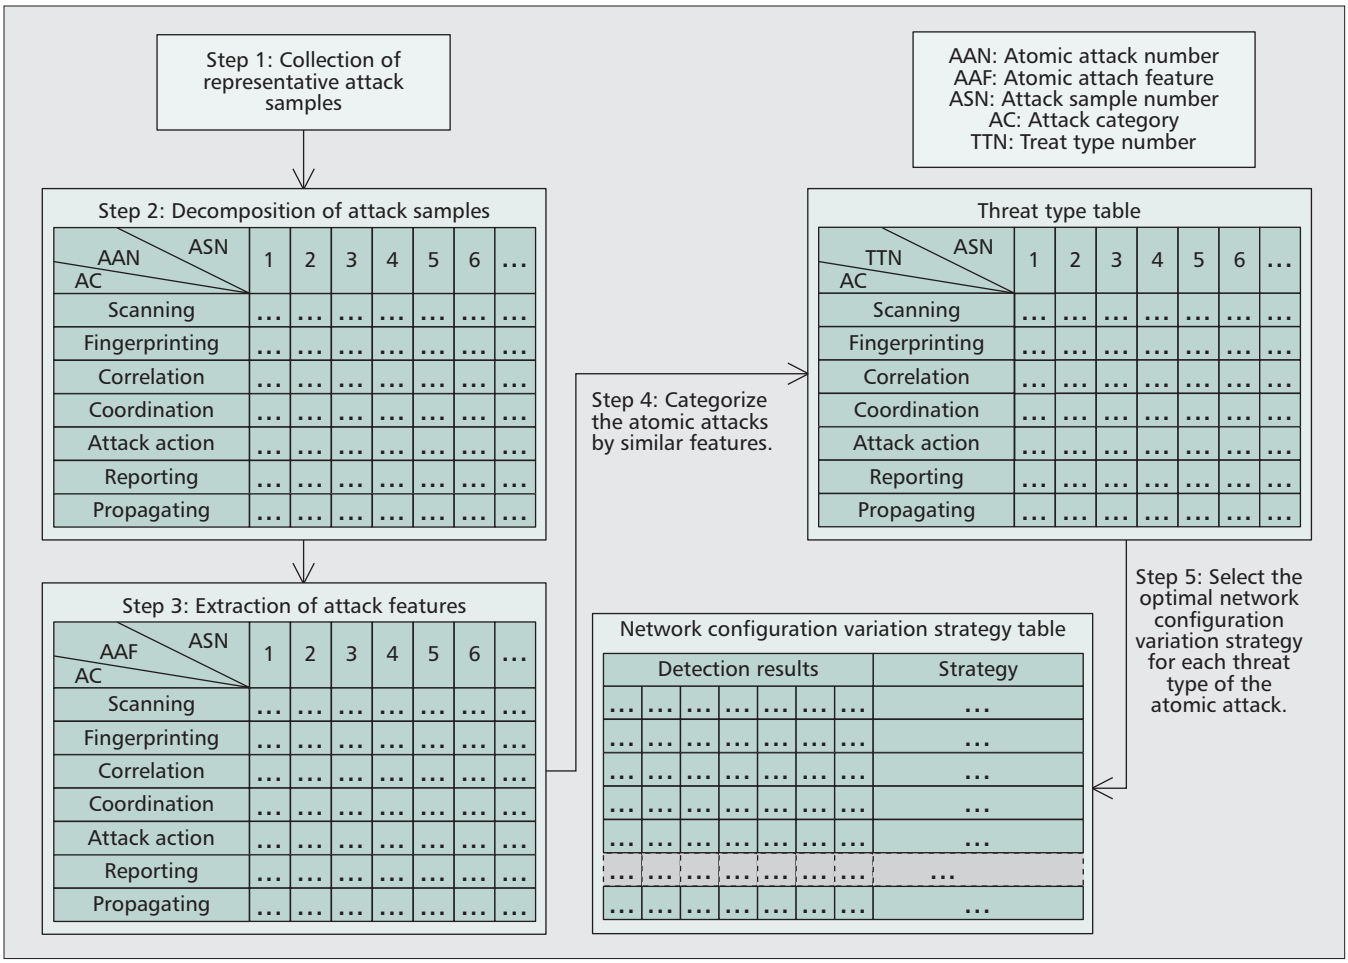

Figure 5. Establishment processes of the tables of network configuration variation strategies and threat types.

and 2, the variation of routes is performed on the TCP communication between these two hosts with a change frequency of 20 s, aiming to protect the communication from monitoring and masquerading. The first variation of routes happens in the 20th second. Second, from the 30th second, the compromised hosts 2, 3, and 4 start to send numerous UDP packets to links 2, 3, and 4 toward host 5, which is one kind of DoS. The communication between hosts 1 and 5 is thus disrupted, and its bandwidth dramatically decreases to zero, as shown in Fig. 6. Third, from the 50th second, the variation of IP addresses is carried out according to the security requirement from the network security state detected by IDS. According to the variation method of IP addresses in [7], the communications based on real IP addresses without being authorized are disconnected. As a result, hosts 2, 3, and 4 are unable to congest the network, because their connections to host 5 based on real IP addresses are not authorized, and the involved flow table entries in OpenFlow switches 2, 3, and 4 are removed. Consequently, from the 50th second, the bandwidth between hosts 1 and 5 gradually recovers. To avoid incorrect disruptions when the variation of IP addresses starts occurring, the table of authorized connections is maintained in the controller. In addition, no collision happens between these two variations of IP addresses and routes.

Challenges and Research Directions

The major challenge of EDM is the efficient establishment of tables of network configuration variation strategies and threat types, which involves the decomposition of representative network attacks, the categorization of the atomic attacks, and research on effective network configuration variation strategies for each kind of atomic attack.

To overcome this challenge, three research directions are identified. First, research on effective network configuration variations for different kinds of atomic attacks is expected. Second, a reasonable evaluation model of network configuration variation strategies is needed. This evaluation model is used to reveal the profit and cost for a network configuration variation strategy to prohibit a corresponding kind of atomic attack. The last is research on improving current security techniques, such as monitors, IDS, and firewalls, for identifying different kinds of atomic attacks.

Conclusion

In this article we present a novel conceptual network security mechanism for future network security. This mechanism is based on the idea of

network configuration variations, originating from a biological inspiration. We introduce the four processes of this mechanism and present its reference framework based on the framework of SDN controllers. We also analyze its effectiveness with a use case. Due to current developments in network security technology, it is promising to be implemented.

Acknowledgment

The authors acknowledge the following funding support: the National Basic Research Program of China (no. 2012CB315903), the National High Technology Research Program of China (no. 2015AA016103), the Program for Zhejiang Leading Team of Science and Technology Innovation (no. 2013TD20), the National Natural Science Foundation of China (no. 61379118), the Program of Science and Technology Commission of Shanghai (no. 13DZ1108800), and the Research Fund of ZTE Corporation.

References

[1] T. D. Nadeau and K. Gray, *SDN: Software Defined Networks*, O'Reilly, 2013.
[2] D. Kreutz, F. Ramos, and P. Verissimo, "Towards Secure and Dependable Software-Defined Networks," *Proc. ACM SIGCOMM HotSDN '13 Wksp.*, 2013, pp. 55–60.
[3] M. Meisel, V. Pappas, and L. Zhang, "A Taxonomy of Biologically Inspired Research in Computer Networking," *Computer Network*, vol. 54, no. 6, 2010, pp. 901–16.
[4] S. Balasubramaniam *et al.*, "Biological Principles for Future Internet Architecture Design," *IEEE Commun. Mag.*, vol. 49, no. 7, 2011, pp. 44–52.
[5] P. Lio and D. Verma, "Biologically Inspired Networking," *IEEE Network*, vol. 24, no. 3, 2010, p. 4.
[6] S. Jajodia *et al.*, *Moving Target Defense*, Springer, 2011.
[7] J. H. Jafarian, E. Al-Shaer, and Q. Duan, "Openflow Random Host Mutation: Transparent Moving Target Defense Using Software Defined Networking," *Proc. ACM SIGCOMM HotSDN '12 Wksp.*, 2012, pp. 127–32.
[8] J. H. Jafarian, E. Al-Shaer, and Q. Duan, "Formal Approach for Route Agility Against Persistent Attackers," *Proc. ESORICS*, 2013, pp. 237–54.
[9] S. Shirali-Shahreza and Y. Ganjali, "FleXam: Flexible Sampling Extension for Monitoring and Security Applications in Openflow," *Proc. ACM SIGCOMM HotSDN '13 Wksp.*, 2013, pp. 167–68.
[10] K. Giotis *et al.*, "Combining OpenFlow and sFlow for an Effective and Scalable Anomaly Detection and Mitigation Mechanism on SDN Environments," *Computer Networks*, vol. 62, 2014, pp. 122–36.

Biographies

Haifeng Zhou (zhouhaifeng@zju.edu.cn) is a Ph.D. candidate at the College of Computer Science, Zhejiang University, Hangzhou, China. His research interests include network security, software-defined networks, bio-inspired future networks, reconfigurable networks, and network traffic engineering.

Chunming Wu (wuchunming@zju.edu.cn) received his Ph.D. degree in computer science from Zhejiang University in 1995. He is currently a professor at the College of Computer Science, Zhejiang University, and director of the NGNT laboratory. His research fields include Internet QoS provision, reconfigurable network technology, network virtualization, the architecture of next generation Internet, and network security.

Ming Jiang (jmzju@163.com) received his Ph.D. degree in computer science from Zhejiang University in 2004. He is currently a professor at the College of Computer Science, Hangzhou Dianzi University, China. His research interests include network virtualization, Internet QoS provisioning, the differentiated services network, and network congestion control.

Boyang Zhou (zby_zju@163.com) received his Ph.D. from the College of Computer Science, Zhejiang University. His research fields include software-defined networks, the architecture of future Internet, and reconfigurable networks.

Wen Gao (gavingao@zju.edu.cn) is a Ph.D. candidate in the College of Computer Science, Zhejiang University. His current research interests include network security, network management, software-defined networks, and reconfigurable networks.

Tingting Pan (pantingting@zju.edu.cn) is an M.S. candidate in the College of Computer Science, Zhejiang University. Her research interests include software-defined networks, network security, and network virtualization.

Min Huang (huangmin@zju.edu.cn) is a B.S. candidate at the College of Computer Science, Zhejiang University. His current research interests include software-defined networks and network security.

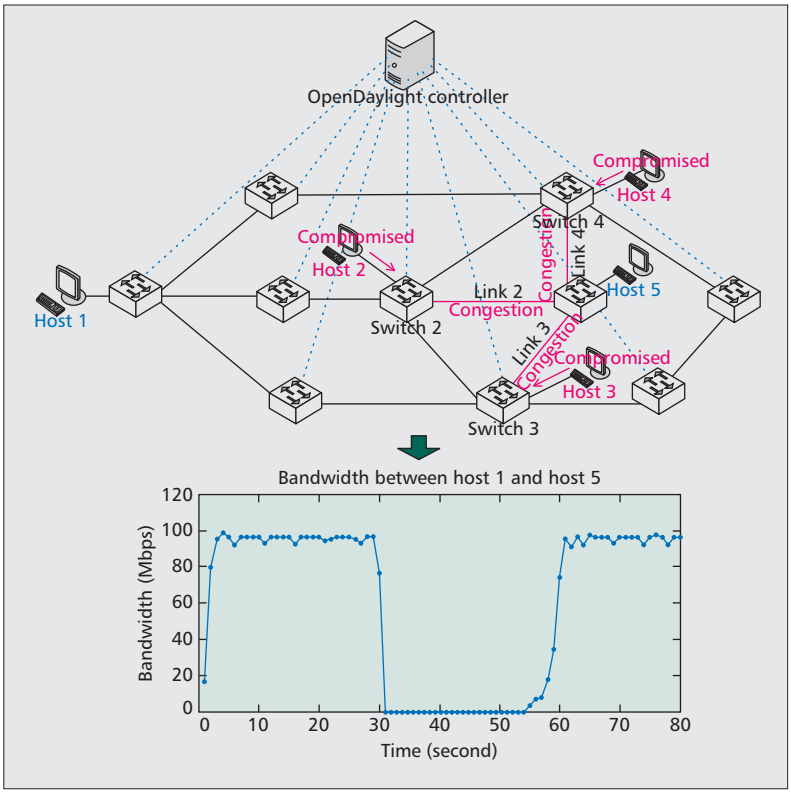

Figure 6. A use case of EDM.

Distributed Denial of Service Attacks in Software-Defined Networking with Cloud Computing

Qiao Yan and F. Richard Yu

ABSTRACT

Although software-defined networking (SDN) brings numerous benefits by decoupling the control plane from the data plane, there is a contradictory relationship between SDN and distributed denial-of-service (DDoS) attacks. On one hand, the capabilities of SDN make it easy to detect and to react to DDoS attacks. On the other hand, the separation of the control plane from the data plane of SDN introduces new attacks. Consequently, SDN itself may be a target of DDoS attacks. In this paper, we first discuss the new trends and characteristics of DDoS attacks in cloud computing environments. We show that SDN brings us a new chance to defeat DDoS attacks in cloud computing environments, and we summarize good features of SDN in defeating DDoS attacks. Then we review the studies about launching DDoS attacks on SDN and the methods against DDoS attacks in SDN. In addition, we discuss a number of challenges that need to be addressed to mitigate DDoS attached in SDN with cloud computing. This work can help understand how to make full use of SDN's advantages to defeat DDoS attacks in cloud computing environments and how to prevent SDN itself from becoming a victim of DDoS attacks.

INTRODUCTION

Cloud computing has emerged as a hotspot in both academia and industry due to its essential characteristics, such as on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. Currently, security issues have been regarded as the dominant barrier in the development of cloud computing [1]. Security requirements of cloud computing include confidentiality, integrity, availability, accountability, and privacy-preservability. Among these security requirements, availability is crucial since the core function of cloud computing is to provide on-demand service of different levels. If a certain service is no longer available or the quality of service cannot meet the service level agreement (SLA), customers may lose faith in the cloud system [1]. Denial of service (DoS) attacks and distributed denial of service (DDoS) attacks are the main methods to destroy the availability of cloud computing. In DoS or DDoS attacks, an attacker attempts to make a machine or network resource unavailable to its intended users [2]. DoS attacks are sent by one person or system, while DDoS attacks are sent by two or more persons or systems. An attacker may be a real person or a group of zombies that are controlled by an attacker. An attacker has the capability to send large volume packets to the target with spoofed source IP addresses.

Although some excellent work has been done to defeat DDoS attacks in traditional computing environments, DDoS attacks are becoming more prevalent in cloud computing environments. Moreover, we have started to see new forms of attack based on the new characteristics of cloud computing, such as the emergence of new economic denial of sustainability (EDoS) attacks [1].

Recently, software defined networking (SDN) has attracted much interest as a new paradigm in networking [3]. Although SDN brings numerous benefits by decoupling the control plane from the data plane, there is a *contradictory relationship* between SDN and DDoS attacks. On one hand, the capabilities of SDN (e.g. software-based traffic analysis, logical centralized control, global view of the network, and dynamic updating of forwarding rules) make it easy to detect and to react to DDoS attacks rapidly. On the other hand, the separation of the control plane from the data plane introduces new attacks. Consequently, SDN itself may be a target of DDoS attacks. Indeed, potential DDoS vulnerabilities exist across the SDN platform [4]. For example, an attacker can take advantage of the characteristics of SDN to launch DDoS attacks against the control layer, infrastructure layer, and application layer of SDN.

In this article we first discuss the new trends and characteristics of DDoS attacks in cloud computing environments. We show that SDN brings us a new chance to defeat DDoS attacks in cloud computing environments, and we summarize good features of SDN in defeating DDoS attacks. Then we review the studies about

Qiao Yan is with Shenzhen University.

F. Richard Yu is with Carleton University.

0163-6804/15/$25.00 © 2015 IEEE

launching DDoS attacks on SDN and the methods against DDoS attacks in SDN. In addition, we discuss a number of challenges that need to be addressed to mitigate DDoS attached in SDN with cloud computing.

To the best of our knowledge, the contradictory relationship between SDN and DDoS attacks has not been well addressed in previous work. Essentially, it is the unique dynamics associated with SDN and DDoS attacks that present unique challenges beyond the existing works. We believe that the initial steps we have taken here help understand how to make full use of SDN's advantages to defeat DDoS attacks in cloud computing environments and how to prevent SDN itself from becoming a victim of DDoS attacks.

The rest of the article is organized as follows. We present the new trends of DDoS in cloud computing environments. Some good features of SDN in defeating DDoS attacks are discussed. We review the work about launching DDoS attacks on SDN. Some open research issues are presented. Finally, we conclude this study.

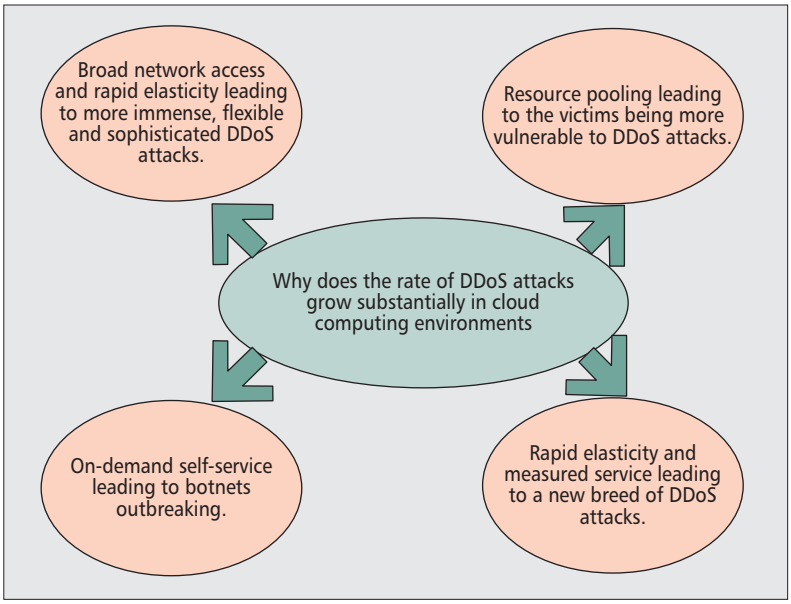

Figure 1. The reasons that the rate Of DDoS attacks in cloud computing environments grows substantially.

DDoS Attacks in Cloud Computing Environments are Growing

In this section we explain the reasons why the rate of DDoS attacks in cloud computing environments has grown substantially based on the essential characteristics of cloud computing, including on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service, as shown in Fig. 1.

On-Demand Self-Service Leading to Botnets Outbreak

One major reason why the rate of DDoS attacks in cloud computing has grown substantially is the emergence and development of botnets. Botnets are networks that are formed by bots or machines compromised by malware. Large-scale botnets (e.g. Srizbi, Kraken/Bobax, and Rustock) have gained notoriety for performing DDoS attacks.

It remains fairly complex to infect a sufficient number of machines in a short time frame in traditional networks. But the on-demand self-service capabilities of the cloud could be used by hackers to instantly create a powerful botnet. With cloud computing, malware-as-a-service is being used for launching DDoS attacks. Because of competition among suppliers, the price of malware-as-a-service has been falling rapidly. Cheap prices make it easier to use botnets to launch large-scale DDoS attacks in cloud computing environments than in traditional networks.

Broad Network Access and Rapid Elasticity Leading to More Immense, Flexible, and Sophisticated DDoS Attacks

With cloud computing's capabilities of broad network access and rapid elasticity, attackers can not only launch immense DDoS attacks, but also produce more flexible and more sophisticated DDoS attacks by using heterogeneous thin or thick client platforms, which are discussed in the following.

More Immense DDoS Attacks in Cloud Computing: The size and frequency of DDoS attacks have grown dramatically as attackers take advantage of botnets and other high-speed Internet access technologies to overwhelm their victim's network infrastructure. In March 2013, Spamhaus, an organization that maintained lists of spammers, came under a massive DNS reflection DDoS attack. The greatest attack traffic was reportedly as high as 300Gbps.

More Flexible DDoS Attacks in Cloud Computing: Because of cloud computing's capabilities of broad network access, mobile devices such as smartphones and tablets are expected to become a significant launching platform for DDoS attacks. The rising bandwidth and processing power and the lack of security of mobile devices make them an ideal platform for hackers to compromise for DDoS attack campaigns.

More Sophisticated DDoS Attacks in Cloud Computing: DDoS attacks are becoming larger and more frequent, and they are becoming more sophisticated as they pinpoint specific applications (e.g. DNS, HTTP or VoIP) or a weak point in the victim's system design. Although sophisticated DDoS attacks require more understanding of the attacked system, they usually use less traffic and are harder to detect.

Resource Pooling Leading to the Victims More Vulnerable to DDoS Attacks

In cloud computing, virtualization technology and multi-tenant infrastructure on one hand make attackers launch DDoS attacks more easily, and on the other hand cause the victims to be more vulnerable to DDoS attacks.

• Virtualization technology makes attackers launch DDoS attacks more easily: Virtualization technology can be used by attackers to preset for DDoS attacks before launching attacks. Virtual

With rapid elasticity and measured service, adopters of the cloud service model are charged based on a pay-per-use basis of the cloud's server and network resources. With this model, a conventional DDoS attack on server and network resources is transformed in a cloud environment to a new breed of attack.

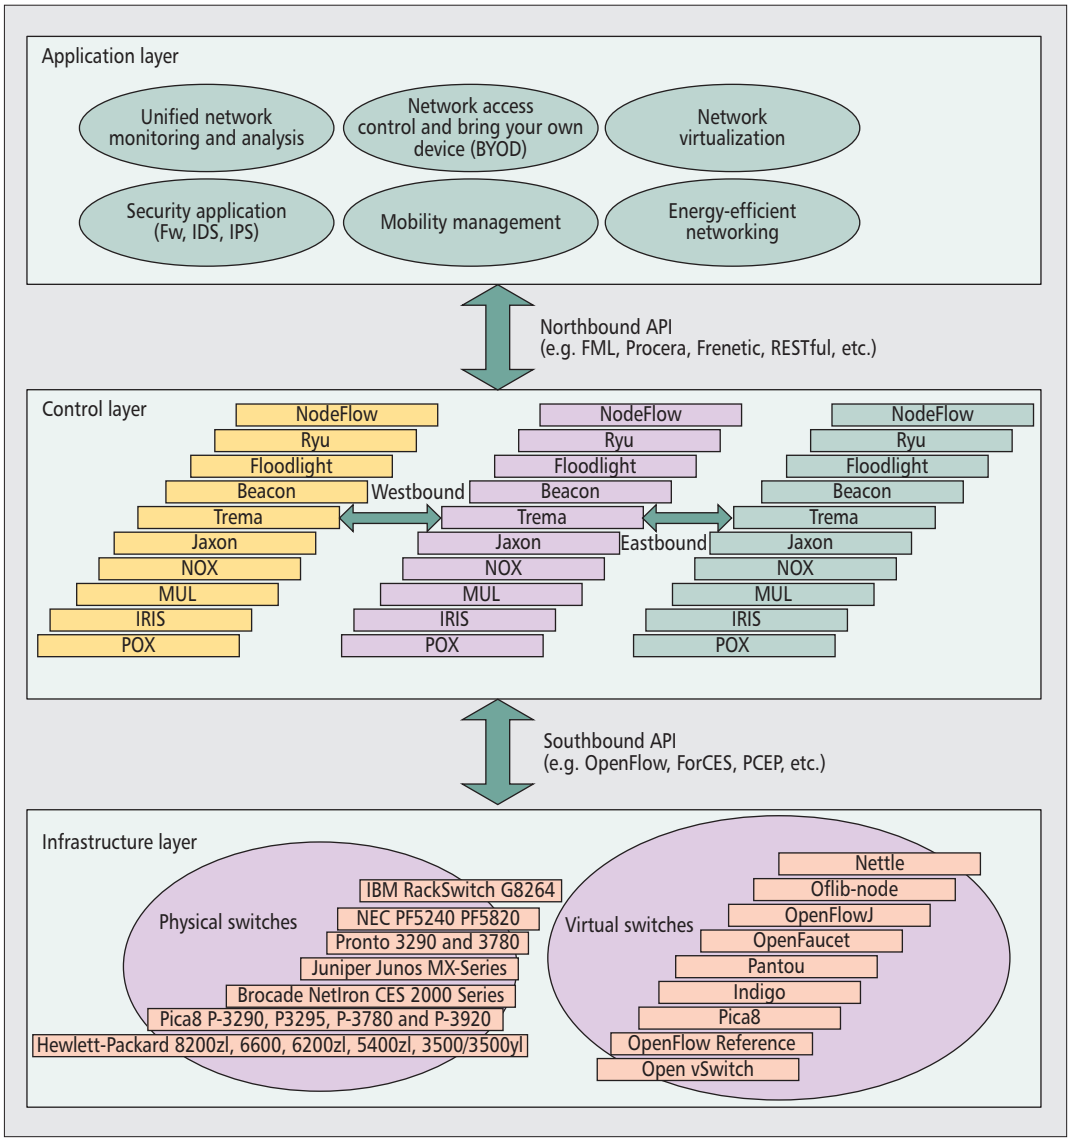

Figure 2. High-level overview of the SDN architecture.

machines can be built using little memory or disk space to launch more attacks with less costs.

• Virtualization technology and multi-tenant infrastructure cause the victims to be more vulnerable to DDoS attacks: Researchers have shown that on a DoS attack, the performance of a web server hosted in a virtual machines can degrade by up to 23 percent, while that of a non-virtualized server hosted on the same hardware degrades by only eight percent [5]. Since the cloud computing environment is inherently a multi-tenant infrastructure, an attack against a single customer is actually an attack against all customers in that given cloud.

Rapid Elasticity and Measured Service Leading to a New Breed of DDoS Attacks

With rapid elasticity and measured service, adopters of the cloud service model are charged on a pay-per-use basis of the cloud's server and network resources. With this model, a conventional DDoS attack on server and network resources is transformed in a cloud environment into a new breed of attack that targets the cloud adopter's economic resources, e.g. economic denial of sustainability (EDoS) attacks [1].

The goal of an EDoS attack is to deprive the victims (i.e. regular cloud customers) of their long-term economic viability. An EDoS attack succeeds when it causes financial burden on the victim. For example, attackers who act as legal cloud service clients continuously send requests to a website hosting in cloud servers to consume bandwidth, which bills to the cloud customer owning the website. It seems to the web server that this traffic does not reach the level of service denial, and it is difficult to distinguish EDoS attack traffic from other legitimate traffic [1].

Is SDN a Silver Bullet for Defeating DDoS Attacks?

Enterprises have enthusiastically embraced cloud computing, which offers an effective way to reduce capital expenditure (CapEx) and operational expenditure (OpEx) [1]. However, security

and privacy issues become a critical concern. As mentioned before, DDoS attacks are becoming the biggest threat to the availability of cloud computing. Traditional DDoS attacks mitigating mechanisms are meeting with various difficulties. SDN, as a new paradigm for enabling innovation in networking research and development, provides us with a new way of thinking to solve the problem. In this section we first introduce SDN and OpenFlow. Then we discuss the good features of SDN in defeating DDoS attacks.

What is Software-Defined Networking

SDN is currently attracting significant attention from both academia and industry. The Open Networking Foundation (ONF) is a nonprofit consortium dedicated to the development, standardization, and commercialization of SDN. ONF has provided the most explicit and well received definition of SDN as follows: "In the SDN architecture, the control and data planes are decoupled, network intelligence and state are logically centralized, and the underlying network infrastructure is abstracted from the applications" [6].

ONF presents a high-level architecture for SDN that is vertically split into three main functional layers: the infrastructure layer, the control layer, and the data layer (see Fig. 2).

- Infrastructure layer: Also known as the data plane, it consists mainly of forwarding elements (FEs), including physical switches and virtual switches. These switches are accessible via an open interface to switch and forward packets.
- Control layer: Also known as the control plane, it consists of a set of software-based SDN controllers providing a consolidated control functionality through open APIs to supervise the network forwarding behavior through an open interface.
- Application layer: It mainly consists of the end-user business applications. Examples of such business applications include network virtualization, mobility management, security applications, and so on.

SDN is often linked to the OpenFlow protocol. OpenFlow is an open protocol, which is proposed to standardize the communication between the switches and the controller in an SDN architecture.

SDN is closely related network function virtualization (NFV). Although both SDN and NFV aim at increasing the agility and flexibility of networks and decreasing complexity and cost, they use different methods. In SDN, control planes are separated from data planes, while in NFV, network devices are replaced by software. SDN and NFV are not dependent on one another, but one can be benefit from the other.

Good Features of SDN in Defeating DDoS Attacks

SDN has many good features, and these good features offer many benefits for defeating DDoS attacks, as shown in Fig. 3.

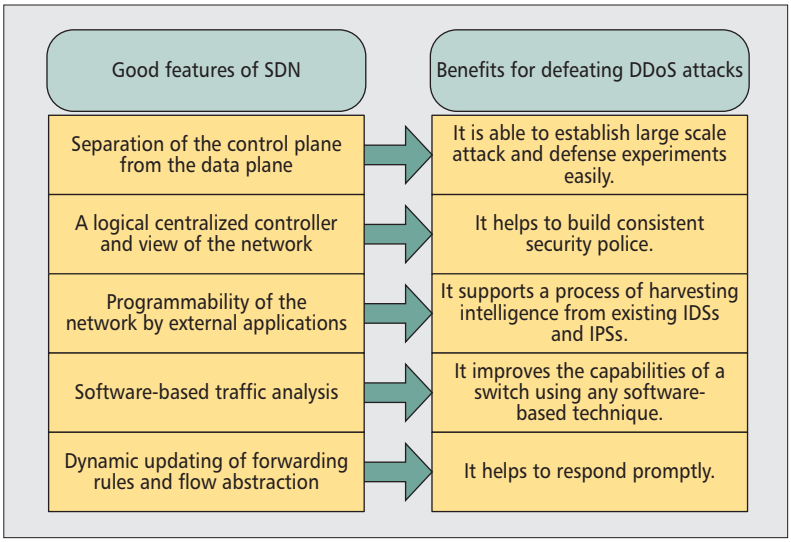

Figure 3. Good features of SDN in defeating DDoS attacks.

• Separation of the control plane from the data plane: DDoS attacks are not a new problem. Since Yahoo, Amazon, and other well-known web sites were subjected to DDoS attacks in 2000, researchers have presented many methods to mitigate DDoS attacks. But in traditional networks, researchers cannot experiment with their ideas on a large scale in a real network setting, hence the performance of the presented algorithms cannot be well tested and verified. SDN decouples the data plane from the control plane, and thus makes it possible to easily establish large scale attack and defense experiments. The high configurability of SDN offers clear separation among virtual networks, permitting experimentation in a real environment [3]. Progressive deployment of new ideas can be performed through a seamless transition from an experimental phase to an operational phase [3]. This feature of SDN offers great convenience in putting forward new thoughts and methods for DDoS attack mitigation.

• A centralized controller and view of the network: The controller has network-wide knowledge of the system and global views to build consistent security policies and to monitor or analyze traffic patterns for potential security threats. Centralized control of SDN makes it possible to dynamically quarantine compromised hosts and authenticate legitimate hosts based on the information obtained through requesting end hosts and remote authentication dial in user service (RADIUS) servers for users' authentication information and system scanning during registration [3]. In a multi-tenant model such as cloud computing, distinguishing tenants' activities and provisioned resources plays an important role in anomaly detection. TaheriMonfared *et al.* [7] proposed a method to build the per-tenant view by use of an OpenFlow controller. The controller provides a unified view of the network, and is aware of the tenant logic. The monitoring node communicates with the controller to build a per-tenant view of the network and generates monitoring information for each tenant.

• Programmability of the network by external applications: The programmability of SDN supports a process of harvesting intelligence from existing intrusion detection systems (IDSs) and

SDN holds great promise in terms of mitigating DDoS attacks in cloud computing environments. However, the security of SDN itself remains to be addressed. Many security issues may happen in SDN, such as unauthorized access, data leakage, malicious applications, configuration issues.

intrusion prevention systems (IPSs) [4]. More intelligent algorithms can be flexibly used based on different DDoS attacks. Within the infrastructure-as-a-service (IaaS) clouds, to prevent vulnerable virtual machines from being compromised in the cloud, Chun-Jen Chung *et al.* [8] proposed a multi-phase distributed vulnerability detection, measurement, and countermeasure selection mechanism called NICE. The proposed framework leverages OpenFlow network programming APIs to build a monitor and control plane over distributed programmable virtual switches in order to significantly improve attack detection and mitigate attack consequences.

- Software-based traffic analysis: Software-based traffic analysis greatly enables innovation, as it can be performed using all kinds of intelligent algorithms, databases, and any other software tools. Motivated by the flexibility of the SDN architecture and the observation that most mobile malware requires Internet connections, Jin and Wang designed a system that detects mobile malware through real-time traffic analysis using the SDN architecture [9].
- Dynamic updating of forwarding rules and flow abstraction: Dynamic updating of forwarding rules assist in the prompt response to DDoS attacks. Based on the traffic analysis, new or updated security policy can be propagated across the network in the form of flow rules to block the attack traffic without delay. Yu *et al.* [10] proposed a memory-efficient system for distributed and collaborative per-flow monitoring, called DCM. DCM uses Bloom filters to represent monitoring rules using a small size of memory. It utilizes SDN's ability to dynamically update forwarding rules to install a customized and dynamic monitoring tool into the switch data plane [10].

Are DDoS Attacks a Nightmare for SDN?

SDN holds great promise in terms of mitigating DDoS attacks in cloud computing environments. However, the security of SDN itself remains to be addressed. Many security issues may happen in SDN, such as unauthorized access, data leakage, malicious applications, configuration issues, etc. [4]. This article focuses on DDoS attacks. In this section, we first discuss how SDN itself may be a target of DDoS attacks. Then we provide an overview of available solutions to this problem.

Possible DDoS Attacks on SDN

SDN itself may be a target of DDoS attacks. Since SDN is vertically split into three main functional layers — infrastructure layer, control layer, and application layer — potential malicious DDoS attacks can be launched on these three layers of SDN's architecture. Based on the possible targets, we can classify the DDoS attacks on SDN into three categories: application layer DDoS attacks, control layer DDoS attacks, and infrastructure layer DDoS attacks, as shown in Fig. 4.

Application Layer DDoS Attacks: There are two methods to launch application layer DDoS attacks: attack applications, or attack the northbound API. Since isolation of applications or resources in SDN is not well solved, DDoS attacks on one application can affect other applications.

Control Layer DDoS Attacks: The controllers could potentially be seen as a risk of single point of failure for the network, so they are a particularly attractive target for DDoS attacks in the SDN architecture. The following methods can launch control layer DDoS attacks: attacking the controller, the northbound API, the southbound API, the westbound API, or the eastbound API. For example, many conflicting flow rules from different applications may cause DDoS attacks on the control plane. Within the operation of SDN, the data plane will typically ask the control plane to obtain flow rules when the data plane sees new network packets that it does not know how to handle [6]. There are two options for the handling of a new flow when no flow match exists in the flow table: either the complete packet or a portion of the packet header is transmitted to the controller to resolve the query. With a large volume of network traffic, sending the complete packet to the controller would occupy high bandwidth.

Infrastructure Layer DDoS Attacks: There are two methods to launch infrastructure layer DDoS attacks: attack switches or attack the southbound API. For example, if only header information is transmitted to the controller, the packet itself must be stored in node memory until the flow table entry is returned. In this case, it would be easy for an attacker to execute a DoS attack on the node by setting up a number of new and unknown flows. As the memory element of the node can be a bottleneck due to high cost, an attacker could potentially overload the switch memory (e.g. targeting to exhaust TCAMs). The generated fake flow requests can produce many useless flow rules that need to be held by the data plane, thus making it difficult for the data plane to store flow rules for normal network flows [6].

To demonstrate the feasibility of DDoS attacks, a new SDN network scanning prototype tool (named SDN scanner) is proposed in [11] to remotely fingerprint networks that deploy SDN. This method can be easily operated by modifying existing network scanning tools (e.g. ICMP scanning and TCP SYN scanning). The attack can be conducted to an SDN network by a remote attacker, and it can significantly degrade the performance of an SDN network without requiring high performance or high capacity devices.

Porras *et al.* [12] show that OpenFlow applications may contradict or override one another, incorporate vulnerabilities, or possibly be written by adversaries. In the worst case, an adversary can use the deterministic OpenFlow application to control the state of all OpenFlow switches in the network [12]. A rule conflict is said to arise when the candidate OpenFlow rule enables or disables a network flow that is otherwise inversely prohibited (or allowed) by existing rules [12]. Hackers may use rule conflict to launch DDoS attacks.

Because DDoS attacks use forged source IP addresses or faked traffic, simple authentication mechanisms could mitigate forged or faked traffic flows. But if an attacker assumes the control

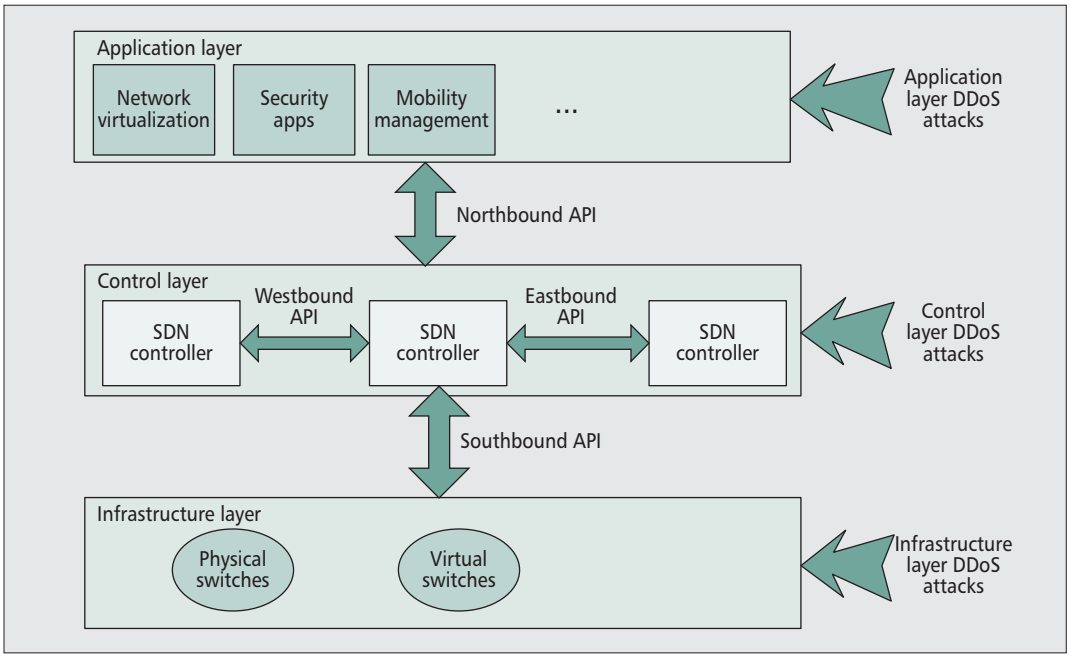

SDN itself may be a target of DDoS attacks. Since SDN is vertically split into three main functional layers, including infrastructure layer, control layer and application layer, potential malicious DDoS attacks can be launched on these three layers of SDN's architecture.

Figure 4. Potential DDoS attacks can be launched on these three layers of SDN's architecture.

of an application server that stores the details of many users, it can easily use the same authenticated ports and source MAC addresses to inject authorized, but forged, flows into the network [13].

Although OpenFlow provides optional support for encrypted transport layer security (TLS) communication and a certificate exchange between the switches and the controller(s), using TLS/SSL does not per se guarantee secure communications. The security of those communications is as strong as its weakest link, which could be a self-signed certificate, a compromised certificate authority, or vulnerable applications and libraries [13]. Moreover, the TLS/SSL model is not enough to establish and assure trust between controllers and switches. After an attacker gains access to the control plane, it may be capable of aggregating enough power force (in terms of the number of switches under its control) to launch DDoS attacks.

Available Solutions

We summarize possible DDoS attacks on SDN and available solutions in Table 1.

FortNox is a new security policy enforcement kernel as an extension to the open source NOX OpenFlow controller, which mediates all OpenFlow rule insertion requests [12]. FortNOX implements role-based authentication to determine the security authorization of each OpenFlow application (rule producer), and enforces the principle of least privilege to ensure the integrity of the mediation process.

For security, OpenFlow provides optional support for encrypted transport layer security (TLS) communication and a certificate exchange between the switches and the controller(s) [14], and the use of oligarchic trust models with multiple trust-anchor certification authorities (e.g. one per sub-domain or per controller instance) is a possibility [13]. Moreover, securing communications with threshold cryptography across controller replicas (where the switch will need at least n shares to get a valid controller message) may be helpful. Additionally, the use of dynamic, automated, and assured device association mechanisms may be considered, in order to guarantee trust between the control plane and data plane devices [13].

The use of IDSs with support for runtime root-cause analysis could help identify abnormal flows [13]. This could be coupled with mechanisms for dynamic control of switch behavior (e.g. rate bounds for control plane requests).

AVANT-GUARD is a new framework to advance the security and resilience of OpenFlow networks with greater involvement from the data plane [15]. It addresses two security challenges for SDN-enabled networks. The first goal is to secure the interface between the control plane and the data plane, and shield it from saturation attacks by a connection migration technique on the data plane. The second goal is to improve responsiveness so that security applications can efficiently access network statistics to respond to threats by creating actuating triggers when a predefined trigger condition is detected.

Open Problems

There are many open research problems that are still not well investigated and need to be addressed by future research efforts. In this section we discuss some of the most important open research issues to mitigate DDoS attacks in cloud computing environments by use of SDN.

How to Defeat Application Layer DDoS Attacks Using SDN

According to new research by Gartner, there will be noticeable growth in the incidence of application layer DDoS attacks. Access to payload information is crucial for application DDoS

Possible DDoS attacks	Attack implementation methods	Available solutions
Application layer DDoS attacks	By attacking application	FortNOX [12]
	By attacking northbound API	
Control layer DDoS attacks	By attacking controller	Transport Layer Security (TLS) [14] FortNOX [12] AVANT-GUARD [15]
	By attacking northbound API	
	By attacking southbound API	
	By attacking westbound API	
	By attacking eastbound API	
Infrastructure layer DDoS attacks	By attacking switch	Transport Layer Security (TLS) [14] AVANT-GUARD [15]
	By attacking southbound API	

Table 1. Possible DDoS attacks on SDN and available solutions.

attack mitigation. Moreover, this information needs to be obtained at considerably reduced latencies and with reasonable cost.

Current SDN architectures only provide the visibility and control on L2-L4. Thus, defeating application layer DDoS attacks may not benefit from the current OpenFlow implementation. Major efforts need to be spent in this area in order to to extend traffic intelligence to Layer 4 to Layer 7 with good trade-offs between performance and security.

How to Defeat Mobile DDoS Attacks Using SDN

With the number of smart devices increasing, popular apps will be installed and millions of their instances can be running at the same time. Both the mobile devices and the apps can be used to initiate DDoS attacks. Based on the current trend of usage of mobile devices and cloud computing, we believe the battlefield of DDoS attacks and defense will shift from the traditional network to the mobile cloud computing environment. Because mobile networks use super proxies, the simple filter method based source IP addresses may not be used since it will also block legitimate traffic. Although some efforts have been made to extend SDN capability to mobile devices for many network problems (e.g. QoS, virtualization, and fault diagnosis), more research needs be done to defeat mobile DDoS attacks using SDN.

How to Implement Multiple Locations Defensive Methods

Many multiple locations defensive methods have been presented in traditional networks. Multiple locations defense is comprised of multiple defense nodes deployed at various locations such as the source, the destination, or the networks [2]. For instance, detection can be done at the victim side and the response can be initiated and distributed to other nodes by the victim. So we believe with widely deployment of SDN in carrier networks, there are many research opportunities to implement multiple locations defensive methods using SDN to defeat DDoS attacks.

How to Cooperate Among the Key Defensive Points

Since attackers cooperate to perform successful attacks, defenders must also form alliances and collaborate with each other to defeat DDoS attacks [2]. A cooperation defense mechanism is the best way to combat DDoS attacks, and many methods have been proposed in traditional networks. Cooperation among the key defensive points can be greatly beneficial to attack prevention, detection, and response. The feature of global view and dynamic updating of forwarding rules of SDN will greatly reduce the cost of cooperation. However, this topic has not been well researched in SDN.

How to Build a DDoS Attacks Tolerant System Using SDN

Since it is very difficult to accurately detect DDoS attacks and prevent them in a timely manner, a DDoS attacks tolerant system may be more realistic. A DDoS attacks tolerant system is a system designed with a fault-tolerant design approach, and it can operate correctly despite attacks. For instance, the system may provide service that meets the requirements of a service-level agreement (SLA) even under an attack by triggering automatic mechanisms to regain and recover the compromised services and resources. A DDoS attacks tolerant system often has some essential properties such as redundancy, diversity, and independence. These properties are easier to implement in SDN networks than in traditional networks. Although some efforts on building a DDoS attacks tolerant system have been made, how to use SDN characteristics to realize attack tolerant systems is a new direction that needs to be addressed by future research efforts.

Conclusions

In this article we first discussed the reasons why DDoS attacks are becoming more prevalent in cloud computing environments. Since SDN could be a good tool to defeat DDoS attacks in cloud computing environments, we presented some good features of SDN in defeating DDoS attacks. After that we discussed how SDN may be a victim of DDoS attacks. We reviewed the studies about how to launch DDoS attacks on SDN and how to deal with this problem. We also discussed some significant open problems.

In summary, SDN creates a very fascinating dilemma: a promising tool to defeat DDoS attacks, versus a vulnerable target of DDoS attacks. How to make full use of SDN's advantages to defeat DDoS attacks in cloud computing environments and how to prevent SDN itself from becoming a victim of DDoS attacks are an urgent problem that needs to be addressed. This article attempted to briefly explore the current technologies related to SDN and DDoS attacks,

and we discussed future research that may be beneficial in these issues.

References

[1] Z. Xiao and Y. Xiao, "Security and Privacy in Cloud Computing," *IEEE Commun. Surveys & Tutorials*, vol. 15, no. 2, 2013, pp. 843–59.

[2] S. T. Zargar, J. Joshi, and D. Tipper, "A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks," *IEEE Commun. Surveys & Tutorials*, vol. 15, no. 4, 2013, pp. 2046–69.

[3] W. Xia *et al.*, "A Survey on Software-Defined Networking," *IEEE Commun. Surveys & Tutorials*, 2014, to be published.

[4] S. Scott-Hayward, G. O'Callaghan, and S. Sezer, "SDN Security: A Survey," *Proc. IEEE SDN for Future Networks and Services (SDN4FNS)*, 2013, pp. 1–7.

[5] R. Shea and J. Liu, "Performance of Virtual Machines under Networked Denial of Service Attacks: Experiments and Analysis," *IEEE Systems J.*, vol. 7, no. 2, 2013, pp. 335–45.

[6] S. Sezer *et al.*, "Are We Ready for SDN? Implementation Challenges for Software-Defined Networks," *IEEE Commun. Mag.*, vol. 51, no. 7, 2013.

[7] A. TaheriMonfared and C. Rong, "Multi-Tenant Network Monitoring Based on Software Defined Networking," *Proc. OTM Conf. Move to Meaningful Internet Systems*, 2013.

[8] C.-J. Chung *et al.*, "Nice: Network Intrusion Detection and Countermeasure Selection in Virtual Network Systems," *IEEE Trans. Dependable and Secure Computing*, vol. 10, no. 4, July 2013, pp. 198–211.

[9] R. Jin and B. Wang, "Malware Detection for Mobile Devices Using Software-Defined Networking," *Proc. IEEE 2nd GENI on Research and Educational Experiment Wksp. (GREE)*, 2013, pp. 81–88.

[10] Y. Yu, Q. Chen, and X. Li, "Distributed Collaborative Monitoring in Software Defined Networks," arXiv preprint arXiv:1403.8008, 2014.17

[11] S. Shin and G. Gu, "Attacking Software-Defined Networks: A First Feasibility Study," *Proc. 2nd ACM SIGCOMM Wksp. Hot Topics Software Defined Networking*, 2013, pp. 165–66.

[12] P. Porras *et al.*, "A Security Enforcement Kernel for OpenFlow Networks," *Proc. 1st Wksp. Hot Topics in Software Defined Networks*, 2012, pp. 121–126.

[13] D. Kreutz, F. Ramos, and P. Verissimo, "Towards Secure and Dependable Software-Defined Networks," *Proc. 2nd ACM SIGCOMM Wksp. Hot Topics in Software Defined Networking*, 2013, pp. 55–60.

[14] B. Nunes, M. Mendonca, X.-N. Nguyen, K. Obraczka, and T. Turletti, "A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks," *IEEE Commun. Surveys & Tutorials*, vol. 16, no. 3, Third Quarter 2014, pp. 1617–34.

[15] S. Shin, V. Yegneswaran, P. Porras, and G. Gu, "AvantGuard: Scalable and Vigilant Switch Flow Management in Software-Defined Networks," *Proc. ACM SIGSAC Conf. Computer & Commun. Security*, 2013, pp. 413–24.

Biographies

QIAO YAN (yanq@szu.edu.cn) is a professor at the College of Computer Science and Software Engineering at Shenzhen University, Shenzhen, China. She received her Ph.D. degree in information and communication engineering from Xidian University, Xi'an, China, in 2003. From 2013 to 2014 she worked at Carleton University, Ottawa, Canada, as a visiting scholar. Her research interests are in network security, cloud computing, and software-defined networking. Her current focus is research and development of security of software defined networking.

F. RICHARD YU (richard.yu@carleton.ca) is an associate professor at Carleton University, Canada. He received the IEEE Outstanding Leadership Award in 2013, Carleton Research Achievement Award in 2012, the Ontario Early Researcher Award (formerly Premier's Research Excellence Award) in 2011, the Excellent Contribution Award at IEEE/IFIP TrustCom 2010, the Leadership Opportunity Fund Award from Canada Foundation of Innovation in 2009, and the Best Paper Awards at IEEE ICC 2014, Globecom 2012, IEEE/IFIP TrustCom 2009, and the Int'l Conference on Networking 2005. His research interests include cross-layer design, security, green IT, and QoS provisioning in wireless networks. He serves on the editorial boards of several journals, including *IEEE Transactions on Vehicular Technology* and *IEEE Communications Surveys and Tutorials*. He has served on the Technical Program Committee (TPC) of numerous conferences, such as the TPC co-chair of IEEE INFOCOM-MCV'15, Globecom'14, WiVEC'14, INFOCOM-MCC'14, Globecom'13, GreenCom'13, CCNC'13, INFOCOM-CCSES'12, ICC-GCN'12, VTC'12S, Globecom'11, INFOCOM-GCN'11, INFOCOM-CWCN'10, IEEE IWCMC'09, VTC'08F, and WiN-ITS'07, and as the publication chair of ICST QShine'10, and the co-chair of ICUMT-CWCN'09.

SDN brings a very fascinating dilemma: a promising tool to defeat DDoS attacks, versus a vulnerable target of DDoS attacks. How to make full use of SDN's advantages to defeat DDoS attacks in cloud computing environments and how to prevent SDN itself becoming a victim of DDoS attacks are an urgent problem that needs to be addressed.

De-Anonymizing and Countermeasures in Anonymous Communication Networks

Ming Yang, Junzhou Luo, Zhen Ling, Xinwen Fu, and Wei Yu

ABSTRACT

To fulfill global requirements for network security and privacy, anonymous communication systems have been extensively investigated and deployed over the world to provide anonymous communication services for users. Nonetheless, diverse de-anonymizing techniques have been proposed to compromise anonymity and impose a severe threat to anonymous communication systems. In this article, we classify the existing de-anonymizing techniques and provide an overview of these techniques. In addition, corresponding countermeasures are studied to mitigate the risks posed by these de-anonymizing techniques.

INTRODUCTION

With rising concerns about privacy and security, Internet users employ numerous ways to encrypt their network traffic. According to the recent statistics of data networks in [1], encrypted traffic in North America has doubled since a year ago while more than quadrupling in Europe and Latin America. Currently, a number of people not only employ traditional encryption methods (e.g., SSL/TLS — Secure Sockets Layer/Transport Layer Security) to preserve the privacy of content in traffic, but also leverage anonymous communication networks to further protect their communication privacy and security.

To provide comprehensive anonymous communication service, researchers developed diverse anonymous communication systems (e.g., the onion routing based system Tor). In light of anonymous applications, anonymous communication systems can generally be categorized into two groups: message based (i.e., high-latency) and flow based (i.e., low-latency) systems.

Email and e-voting are the classic message-based anonymity applications, and have been well studied over the past decade. Because of the increasing need for anonymity over prevalent applications (e.g., web browsing and instant messaging), flow-based anonymity systems have been extensively studied and deployed around the world.

Various low-latency anonymous communication systems have been developed and deployed. By using a searching service, a list of free HTTP or SOCKS proxies can easily be found around the Internet. Additionally, one of the most popular anonymous communication systems (Tor) provides anonymous communication services for hundreds of thousands of Internet users and carries terabytes of traffic each day. As of September 2014, there were more than 6000 Tor routers voluntarily deployed around the world to contribute their bandwidth to the entire Tor network.

A number of de-anonymizing techniques have been investigated to compromise users' communication anonymity. Particularly, traffic-analysis-based de-anonymizing techniques are the primary threats to anonymous communication systems. In this article, we provide an overview of existing de-anonymizing techniques and countermeasures in flow-based anonymous communication systems. To be specific, we first model two categories of anonymous communication systems and introduce their basic anonymizing techniques. We then categorize the de-anonymizing techniques into four groups from different perspectives and elaborate on these techniques from each groups. Existing and possible countermeasures for these de-anonymizing techniques are also studied to improve the anonymity service provided by anonymous communication systems.

The rest of this article is organized as follows. We first introduce two types of classic anonymous communication models. Then we categorize the de-anonymizing techniques and introduce them in detail. Next the corresponding countermeasures are investigated. Finally, we conclude this article.

ANONYMOUS COMMUNICATION MODEL

In this section, we briefly introduce two categories of flow-based anonymous communication models in terms of the length of

Ming Yang, Junzhou Luo, and Zhen Ling are with Southeast University.

Xinwen Fu is with the University of Massachusetts Lowell.

Wei Yu is with Towson University.

0163-6804/15/$25.00 © 2015 IEEE

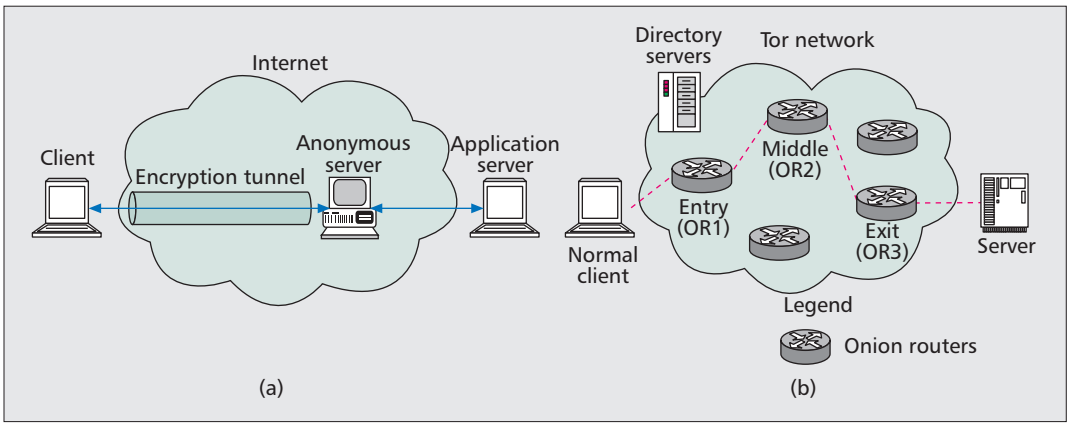

Figure 1. Anonymous communication model: a) single-hop anonymous communication model; b) multihop anonymous communication model (Tor)to compromise users' communication anonymity.

Currently, VPNs and encrypted tunnels to single-hop proxies are typical single-hop anonymous communication systems. These types of systems provide better performance. Nonetheless, a compromised single-hop proxy will result in the exposure of all clients' traffic through the proxy.

anonymous communication links, including single-hop and multihop anonymous communication models.

Single-Hop Anonymous Communication Model

The single-hop anonymous communication model has been widely applied over the Internet due to its effectiveness and efficiency. Broadly speaking, there are three basic components in this model, including a client, an anonymous server, and an application server. Figure 1a illustrates a traditional single-hop anonymous communication model. A client first installs anonymous service software. Then, by using the installed software, the client can establish an encryption tunnel to a specific anonymous server and delivers application data associated with the user to the anonymous server over this tunnel. The anonymous server decrypts the data and forward it to the destination application server. The application server supports Transmission Control Protocol (TCP) applications (e.g., FTP servers and web servers) and receives the data from the anonymous server. After that, the corresponding data is transmitted to the server. Because the application server cannot identify the real IP address of the client, the client can anonymously communicate with the remote server.

Currently, virtual private networks (VPNs) (e.g., OpenVPN and CiscoVPN) and encrypted tunnels to single-hop proxies (e.g., OpenSSH) are typical single-hop anonymous communication systems. These types of systems provide better performance. Nonetheless, a compromised single-hop proxy will result in the exposure of all clients' traffic through the proxy. To overcome this issue, the multihop anonymous communication model has been developed to mitigate this risk.

Multihop Anonymous Communication Model

Again, to address the security issue of the single-hop anonymous communication model, multihop anonymous communication systems such as Tor and JonDonym have been proposed. Because Tor is the most popular multihop anonymous communication system, we take it as an example to depict the organization of a multihop anonymous communication system. Broadly speaking, a Tor network consists of four components: a client, onion routers, directory servers, and server. Figure 1b illustrates the basic architecture of a Tor network. A client needs to install Tor software (i.e., *Onion Proxy*, OP) to pack the data from the client into a Tor *cell*, which is a basic transmission unit in the Tor network. *Onion routers* (ORs) are used to relay data between clients and servers. The directory servers collect all of the information associated with the ORs, including IP address, port, public key, and so on. Servers provide TCP application services such as web and FTP services.

To anonymously communicate with the remote server through Tor, the client builds a multihop path using a source routing mechanism and communicates with the remote server through the established path in the Tor network. First, the client downloads all of the OR information from the directory servers and selects several high-performance ORs. The number of selected ORs in a path is denoted as the path length. The default path length is three, which is hard-coded in the Tor client software. The path is also referred to as a *circuit* in the Tor network. Upon choosing the appropriate ORs, an OP will establish a one-hop circuit to the first OR (i.e., entry OR) and negotiate a symmetric key. After that, it will extend this tunnel to the following two ORs, referred to as the middle and exit ORs, respectively, and negotiate symmetric keys with them. Once the circuit is completely built, multi-TCP streams from the client are multiplexed into this circuit.

It is worth noting that the multihop anonymous communication model is more secure than the single-hop anonymous communication model. In the case of Tor, because any respective OR in the circuit cannot link the client and server, the system can provide better anonymous service to the users. Nonetheless, the data from users is relayed through more nodes in a multihop anonymous system, which definitely increases the end-to-end latency.

End-to-end attacks explored in existing work mainly focus on correlation attacks to confirm the communication relationship between senders and receivers. In contrast, single-end attacks focus on fingerprinting-based attacks to identify the victim's accessed web page.

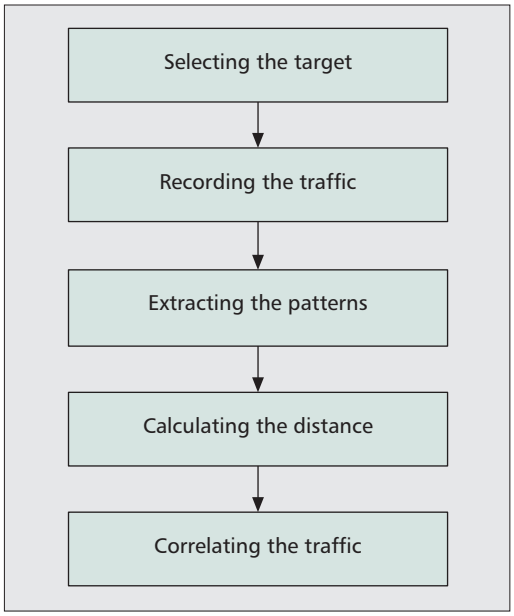

Figure 2. Workflow of end-to-end passive attacks.

De-Anonymizing Techniques to Compromise Users' Communication Anonymity

In this section, we first categorize de-anonymizing techniques into two groups from two different perspectives. We then elaborate on these techniques individually.

Category of De-Anonymizing Techniques

According to existing network-traffic-analysis-based de-anonymizing techniques, we introduce two dimensions of attacks.

- Passive and active attacks: The adversary can passively monitor the victim's traffic or actively manipulate the traffic.
- Single-end and end-to-end attacks: The adversary conducts attacks by monitoring or controlling related devices at either the sender or receiver side, or at both the sender and receiver sides.

End-to-end attacks explored in existing work mainly focus on correlation attacks to confirm the communication relationship between senders and receivers. In contrast, single-end attacks focus on fingerprinting-based attacks (e.g., website fingerprinting attacks) to identify the victim's accessed web page. In the following, we discuss end-to-end attacks and single-end attacks in detail.

End-to-End Attacks

End-to-end attacks are designed to correlate the communication relationship between clients and servers using either passive or active attack methods. To carry out attacks, the adversary should control or monitor the devices (e.g., routers or Tor entry and exit nodes) at both the sender and receiver sides. In the following, we discuss passive and active end-to-end attacks, respectively, and then summarize the advantages and disadvantages of these attacks.

End-to-End Passive Attack — The object of the end-to-end passive attack is to record traffic passively and evaluate the similarity between the sender's outbound traffic and the receiver's inbound traffic based on statistical measures. Figure 2 illustrates the basic workflow of end-to-end passive attacks. This type of technique can exploit traffic features (e.g., packet counter, traffic pattern correlation, timing correlation). For example, the adversary can simply count the number of outgoing packets in several time intervals at the output link of the sender and then count the number of arrival packets in the same time interval at the input link of the receiver. Then a distance function can be applied to compute the distance between these two links in terms of traffic features.

The primary advantage of end-to-end passive attacks is stealth because the traffic will only be monitored. Nonetheless, the true positive rate is low, while the false positive rate is high. Accordingly, the adversary needs a sufficient amount of time to observe the traffic and discover traffic pattern similarities between senders and receivers in order to reduce the number of errors associated with the attack. In addition, to improve the true positive rate and reduce the false positive rate, end-to-end active attacks have been proposed to manipulate traffic in order to generate a desired signal.

End-to-End Active Attack — The basic idea of this attack is that the adversary can manipulate traffic at the sender or receiver side by embedding a special signal in the victim's traffic. Then the traffic at the receiver or sender side is monitored in order to recognize the signal and confirm the communication relationship between the sender and the receiver. This class of attacks is also referred to as *watermarking-based attacks* [2].

Because the adversary can exploit various features of distinct layers to embed watermarking into network traffic, we present these attacks from three different layers: *network layer*, *protocol layer*, and *application layer*.

At the network layer, the adversary can exploit such features as the traffic rate [2], packet delay interval [4], and packet size to embed a signal into target traffic. For example, Yu *et al.* [2] proposed that an adversary could interfere with traffic from a sender and shape its traffic rate pattern. In this way, an invisible direct sequence spread spectrum (DSSS) signal can be embedded in the traffic. Then the embedded signal along with the traffic transmitted through the anonymous communication network arrives at the receiver. After that, the adversary can recognize the signal and compromise the anonymity between the sender and the receiver.

Wang *et al.* [4] investigated packet delay interval centroid-based watermarking techniques. Assume that the arrival distribution of packets in a time interval $[0, T)$ is uniform. By having the adversary intentionally delay each packet within this interval T, packets can uniformly exhibit values in the range $[a, T)$ so that the mean of packet arriving times in this time interval is $T + a/2$. In order to embed the signal in the traffic, the adversary first chooses two

groups of time intervals. Denote the two original groups as A and B and the delayed groups as A' and B'. To encode a 1 bit, the packets in the time interval of group A are carefully delayed, and the mean of these two groups can be derived as

$$E(A') - E(B) = \frac{T+a}{2} - \frac{T}{2} = \frac{a}{2}. \quad (1)$$

To encode a 0 bit, the packets in the time interval of group B are altered, and the adversary can adjust the mean of these two groups as

$$E(A) - E(B') = \frac{T}{2} - \frac{T+a}{2} = -\frac{a}{2}. \quad (2)$$

By adjusting the time interval centroid, a series of binary signal bits can be embedded into the traffic.

In addition, an adversary can vary the packet size to embed a signal into the victim's traffic. For example, the adversary controls a web server and manipulates the size of the response HTTP packets. A specific packet length can be mapped into a single hex bit. By altering the length of several packets, the adversary can encode a message into the traffic. Although the packet length is partially padded at the single-hop proxy, the adversary can still infer the packet length in order to recover the original signal at the client side and confirm the communication relationship between the client and the server. Additionally, to keep this attack invisible, the adversary needs to keep both the distribution and the self-similarity of the original packet size. To this end, the adversary needs to deliberately select appropriate packets and alter their sizes.

At the protocol layer, the watermarking attack can employ different protocol features of anonymous communication systems. For example, Ling *et al.* [5] deeply explored the communication protocol of Tor and discovered that Tor employs the counter mode of the Advanced Encryption Standard (AES-CTR) to encrypt and decrypt Tor cells. Consequently, each Tor node, including the Tor client in a circuit, maintains a local counter to synchronize the counter values with each other in order to correctly encrypt or decrypt the cells. Nonetheless, this attack exploits the feature of the counter synchronization mechanism in a multihop path, and disturbing the counter value at some node along this path incurs encryption/decryption failure of the Tor cell. To achieve this goal, the adversary first needs to control both the Tor exit and entry nodes, and then have the ability to operate the Tor cell at the entry node. Specifically, the adversary can replay, delete, or insert a cell to the target circuit at the entry node. Replaying a cell or inserting a faked cell will result in increasing the counter at both the middle and Tor exit nodes, whereas deleting a cell will decrease the counter. These operations can make the counter value at the middle and exit Tor nodes unsynchronized and cause cell decryption failure at the exit node. Because this type of decryption failure is fairly rare in a normal circuit, the adversary can use this unique feature to detect whether a manipulated cell passes through its controlled Tor exit node. Once this decryption failure is recognized at the exit node, the adversary at the entry node knows the source of the circuit, while the conspirator at the exit node learns the destination of the circuit. In this way, the adversary can trivially confirm the communication relationship between the source and destination of the circuit.

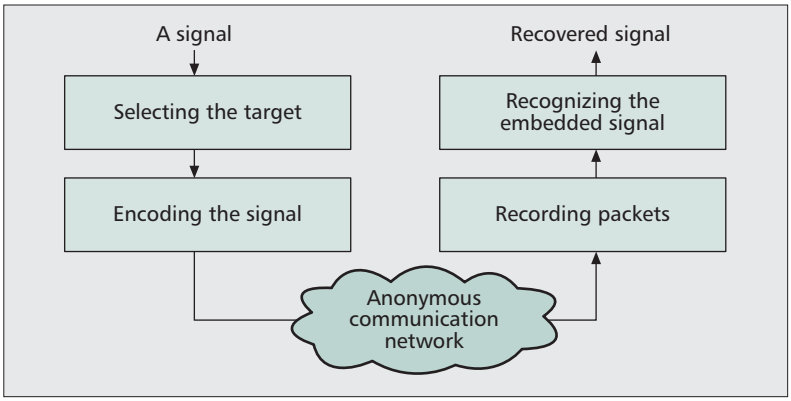

Figure 3. Workflow of end-to-end active attacks [3].

Additionally, the adversary can exploit the defects in data integrity verification in multihop anonymous communication systems [5]. For example, the Tor cell is encrypted in an onion-like fashion during the transmission in the circuit. Consequently, unlike the Tor client and exit node, which can obtain the plaintext of the cell to check the integrity of the data, entry and middle nodes cannot verify the integrity of the cell. If the adversary tampers with the content of the ciphertext of the cell at an entry node, it results in cell decryption failure at the exit node due to a lack of data integrity verification at each node along the entire multihop path. Likewise, the adversary at the exit node can leverage decryption failure as a signal to correlate the communication relationship between the sender and the receiver.

The adversary can also use the Tor protocol characteristic (i.e., the size of each Tor cell is equal [3]). The adversary can control the number of transmitted cells during each time slot at a Tor node in order to encode signal bits. Particularly, to encode a 1 bit, the exit node will collect three Tor cells in the circuit queue and flush them out into the circuit. In addition, a single cell will be sent into the circuit to encode the 0 bit. To avoid the problem of adjacent signals being merged together, a delay interval can be introduced between signals. When the cells that carry the signals arrive at the entry node, a sophisticated signal recovery algorithm can be applied to decode and recover the signals in terms of the number of the received cells in the circuit queue. Once the signal is detected, the communication relationship between the sender and the receiver can be correlated.

At the application layer, the adversary at the server side can inject special content into the victim's web response traffic in order to force the client to generate special traffic patterns as a signal. Then the adversary at the client side can observe this signal and confirm the communication relationship between the sender and the receiver. Specifically, once an

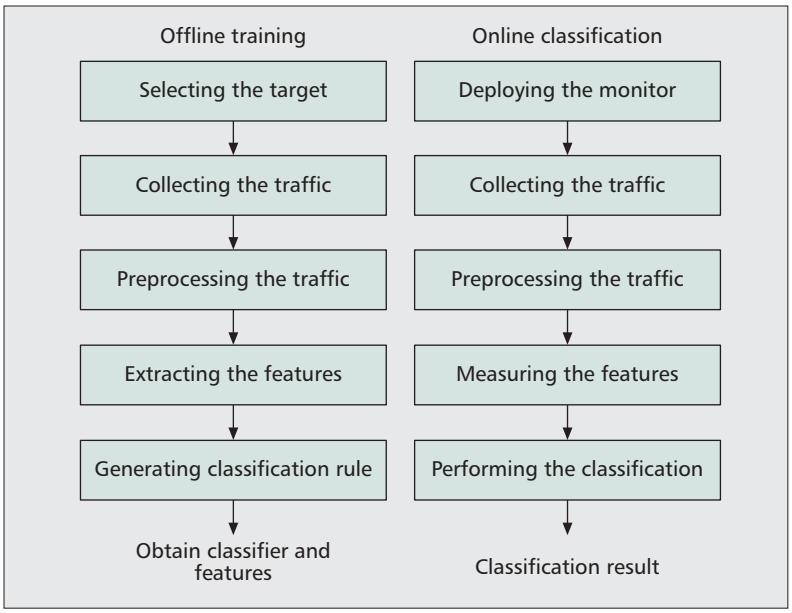

Figure 4. Workflow of single-end passive attacks.

adversary discovers the target web response traffic that passes through its exit node, malicious web links of empty images can be injected so that the browser at the client side will be forced to generate a specific traffic pattern to download these links [6]. Then the conspirator at the entry node will inspect the traffic to detect the desired pattern. Upon discovering the expected traffic pattern, the adversary can compromise the anonymity between the sender and the receiver. In this case, the adversary takes advantage of the HTTP application features to generate a signal. Moreover, similar techniques can be used to conduct such kinds of application layer attacks. For example, a piece of Javascript code can also be injected into the victim's web response traffic as a signal generator, which is executed at the victim's browser to generate signal traffic.

Single-End Attacks

The adversary who performs single-end attacks needs to control or monitor traffic passing through devices at the sending or receiving side so as to compromise the users' communication privacy and security. Single-end passive attacks extract the pattern of traffic, referred to as a fingerprint, and infer the content of traffic at the application layer (e.g., users' accessed websites). In addition, single-end active attacks can actively inject content into traffic at the application layer in order to force the client to directly send a signal to the adversary and expose the real IP address of the client. In the following, we elaborate on these attacks.

Single-End Passive Attack — The idea of single-end passive attacks is to monitor traffic between the victim and the anonymous proxy and identify the real accessed web pages by comparing a prospective traffic pattern with pre-collected web page fingerprints. This type of attack is also referred to as a Website Fingerprinting (WF) attack.

There are two phases of this attack:
- Offline training
- Online classification

Figure 4 illustrates the basic workflow of a single-end passive attack. In the offline training phase, the adversary first needs to select several websites of interest and set up the victim's environment to emulate the procedure of the victim's browsing activities. Then, the adversary will browse the websites one by one and collect the website traffic. Furthermore, the collected data should be preprocessed in order to remove noise. For example, the accuracy of the website fingerprinting attack will be affected because advertisement links on web pages are dynamic. Hence, by using some preprocessing strategies, this type of noise will be filtered. Additionally, the adversary should extract appropriate features from the preprocessed traffic. These features should be carefully selected to exhibit the most effective patterns, which are usually hidden in the traffic. According to existing attacks, various features have been explored to effectively conduct attacks, including packet length distribution, traffic volume, total time, traffic direction, packet length order, up/downstream bytes, bytes in traffic bursts [7], etc. Finally, the adversary chooses a proper classifier to generate a classification rule by using the collected sample data. In existing attacks, various classifiers have been investigated, including Bayes classifiers, multinomial naive-Bayes classifiers, Support Vector Machines (SVMs), decision trees, etc.

In the second phase, the adversary can record real traffic and launch an attack to identify the victim's accessed web pages. First, the adversary needs to deploy a monitoring tool and then silently collect the victim's traffic between the client and the anonymous proxy. After obtaining real traffic, the adversary will preprocess the traffic in order to remove noise. Moreover, the adversary measures the features from the processed traffic and performs the attack by using the classification rule to identify accessed web pages.

Single-End Active Attack — This type of attack actively inserts malicious code into non-encrypted traffic at the server side so that the code arrives and executes at the victim's host. This is done in order to bypass the installed client of an anonymous communication system and directly establish a connection to a malicious server. To this end, the adversary should control the non-encrypted link between the proxy and the remote server. For example, in the case of the Tor network, an adversary who controls a Tor exit node can arbitrarily inject or modify content of non-encrypted traffic. After assuming control of a non-encrypted link, the adversary can inject diverse software instances into the link, including Flash, Javascript, ActiveX Controls, and Java. Once these software types are executed in the browser, they will bypass the local proxy settings in the browser and directly create a connection to a specific remote server in order to expose the real IP address of the client. In addition, the adversary can take advantage of browser exploits to conduct this type of attack to compromise a victim's anonymity.

To prevent themselves from experiencing active single-end attacks, users should disable active content systems such as Flash, ActiveX Controls, Java, and Javascript to avoid malicious code executing in the browser. Alternatively, a transparent proxy can be deployed at the client side in order to ensure that all of the traffic is directed into the anonymous communication system.

Countermeasure

To mitigate the threat posed by de-anonymizing techniques, there have been a number of research efforts on developing various countermeasures to defend against these attacks. Broadly speaking, the countermeasures can be deployed from three perspectives: network layer, protocol layer, and application layer. In the following, we discuss these diverse countermeasures.

Because network traffic characteristics can be exploited to de-anonymize the communication between users, a basic idea of defense that can be used to thwart attacks at the network layer is to remove the features of traffic associated with users, including packet size distribution, packet order, traffic volume, traffic time, and so on. To be specific, *packet padding techniques* can be used to pad packet sizes in order to remove the packet length feature from features such as packet size and packet order. Intuitively, the size of each packet can be padded into the same size (e.g., maximum transmission unit, MTU). Additionally, various sophisticated strategies have been studied to effectively and efficiently pad the packet size [8]. To obfuscate the traffic time, delay can be intentionally added between each packet to increase the traffic time. Furthermore, *dummy traffic techniques* can be applied to inject dummy packets into users' original traffic in order to obfuscate the traffic volume. In addition, *traffic morphing techniques* can be used to vary current traffic patterns to look like other traffic patterns. For example, to thwart a website fingerprinting attack, the web server can first select a target page and then mimic the packet size distribution of that target web page. Generally speaking, defense techniques at the network layer are more general and can be used in various anonymous communication systems, although they can incur high transmission overhead.

At the protocol layer, protocol-level padding and dummy techniques can be used to hide traffic features associated with users. As a matter of fact, secure shell (SSH), TLS, and IPsec apply such protocol-level padding techniques to align plaintext with block cipher boundaries, thereby obfuscating the packet size to some degree. To further improve security, a random amount of padding can be selected [7]. Additionally, protocol-level dummy techniques can be used. For example, Tor does not commonly employ the functionality of padding cells for circuit-level padding purposes because it can significantly decrease the performance of the circuit. Protocol-level padding and dummy techniques could be designed to reduce the overhead incurred to some degree. Nonetheless, it should be carefully designed; otherwise, it could be used to conduct an MTU end-to-end active attack.

At the application layer, HTTP features and background traffic (i.e., decoy web pages) can be exploited to remove traffic features from user flows. For example, HTTP pipelining and HTTP ranges can be used to adjust both incoming and outgoing packet sizes [9]. Moreover, changing the order of the HTTP requests at the client side can vary the traffic pattern to some extent. To apply background traffic techniques at the application layer, a decoy web page can be silently loaded in the background while a user is browsing a target web page. This type of defense technique can only be used for some specific applications (e.g., HTTP) and cannot be widely applied for diverse applications [3].

Hybrid techniques can be deployed at the different layers to provide a comprehensive solution to obstruct various attacks. Moreover, the trade-off between security and performance should be carefully studied [8] to ensure efficient and secure defenses. To further understand the interaction between attacks and countermeasures, we can adopt communication theory by modeling attacks as transmissions of binary messages through a noisy communication channel. By leveraging such a channel model, we can conduct a holistic investigation of the impact of attacks and the effectiveness of countermeasures on various anonymous channels.

To mitigate the threat posed by de-anonymizing techniques, there have been a number of research efforts on developing various countermeasures to defend against these attacks. Broadly speaking, the countermeasures can be deployed from three perspectives: network layer, protocol layer, and application layer.

Conclusion and Future Work

In this article, we have studied two classes of anonymous communication systems (i.e., single-hop and multihop systems). We have described the organization and basic mechanisms of these anonymous communication systems. We have then elaborated on various attacks, including end-to-end confirmation attacks and single-end analysis attacks. To mitigate these attacks, we have also studied possible countermeasures at different layers.

Further research on attacks and corresponding countermeasures should be investigated. For example, Juarez *et al.* [10] argue that current website fingerprinting attacks may not be effective in practice because previous work makes unrealistic assumptions on the models of the adversary, client setting, and website. Therefore, more practical website fingerprinting attacks should be studied. In addition, active website fingerprinting attacks have not been well investigated in existing work. This could be a new research direction in this field. The potential approaches to performing active website fingerprinting attacks could be actively used to modulate the victim's web traffic pattern. In this way, the adversary could trivially infer the various features (e.g., web objects) in the web response traffic. In addition, it makes the attacks more accurate and practical. Finally, the arms race of developing new attacks to compromise users' anonymity and corresponding countermeasures to fight against these attacks will continue. It is critical to establish a theoretical foundation (e.g., channel modeling) capable of studying interactions between various attack and countermeasure mechanisms.

The arms race of developing new attacks to compromise users' anonymity and corresponding countermeasures to fight against these attacks will continue. It is critical to establish a theoretical foundation capable of studying interactions between various attack and countermeasure mechanisms.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under grants 61272054, 61402104, and 61320106007, China National High Technology Research and Development Program under Grant No. 2013AA013503, by U.S. NSF grants 1116644, 1350145, and CNS 1117175, Jiangsu Provincial Key Laboratory of Network and Information Security under grant BM2003201, and Key Laboratory of Computer Network and Information Integration of Ministry of Education of China under grant 93K-9. Any opinions, findings, conclusions, and recommendations in this article be are those of the authors and do not necessarily reflect the views of the funding agencies.

References

[1] Sandvine, "Global internet phenomena report (1h 2014)," https://www.sandvine.com/downloads/general/global-internetphenomena/2014/1h-2014-global-internet-phenomenareport.pdf, 2014.

[2] W. Yu *et al.*, "DSSS-Based Flow Marking Technique for Invisible Traceback," *Proc. 2007 IEEE Symp. Security and Privacy*, Berkeley, CA, May 2007, pp. 18–32.

[3] Z. Ling *et al.*, "A New Cell-Counting-Based Attack Against Tor," *IEEE/ACM Trans. Net.*, vol. 20, no. 4, 2012, pp. 1245–61.

[4] X. Wang, S. Chen, and S. Jajodia, "Network Flow Watermarking Attack on Low-Latency Anonymous Communication Systems," *Proc. IEEE Symp. Security & Privacy*, Berkeley, CA, May 2007, pp. 116–30.

[5] Z. Ling *et al.*, "Protocol-Level Attacks Against Tor," *Computer Networks*, vol. 57, no. 4, Mar. 2013, pp. 869–86.

[6] X. Wang *et al.*, "A Potential HTTP-Based Application-Level Attack Against Tor," *Future Generation Computer System*, vol. 27, no. 1, Jan. 2011, pp. 67–77.

[7] K. P. Dyer *et al.*, "Peek-a-Boo, I Still See You: Why Efficient Traffic Analysis Countermeasures Fail," *Proc. 2012 IEEE Symp. Security and Privacy*, San Francisco, CA, May 2012, pp. 332–46.

[8] X. Cai *et al.*, "A Systematic Approach to Developing and Evaluating Website Fingerprinting Defenses," *Proc. 21th ACM Conf. Computer and Commun. Security*, Scottsdale, AZ, Nov. 2014, pp. 227–38.

[9] X. Luo *et al.*, "HTTPOS: Sealing Information Leaks with Browser-Side Obfuscation of Encrypted Flows," *Proc. 18th Network and Distributed System Security Symp.*, San Diego, CA, Feb. 2011, pp. 1–20.

[10] M. Juarez *et al.*, "A Critical Evaluation of Website Fingerprinting Attacks," *Proc. 21th ACM Conf. Computer and Commun. Security*, Scottsdale, AZ, Nov. 2014.

Biographies

Ming Yang (yangming2002@seu.edu.cn) received his M.Sc. and Ph.D. degrees in computer science and engineering in 2002 and 2007, respectively, from Southeast University, Nanjing, P. R. China. He is currently an associate professor in the School of Computer Science and Engineering, Southeast University. His research interests include network security and privacy.

Junzhou Luo [M] (jluo@seu.edu.cn) is a full professor in the School of Computer Science and Engineering, Southeast University. He received his B.S. degree in applied mathematics from Southeast University in 1982, and then got his M.S. and Ph.D. degrees in computer networks, both from Southeast University in 1992 and 2000, respectively. His research interests are next generation networks, protocol engineering, network security and management, grid and cloud computing, and wireless LANs. He is a member of the IEEE Computer Society and co-chair of the IEEE SMC Technical Committee on Computer Supported Cooperative Work in Design.

Zhen Ling (zhenling@seu.edu.cn) is a lecturer at the School of Computer Science and Engineering of Southeast University. He received his B.S. degree (2005) and Ph.D. degree (2014) in computer science from Nanjing Institute of Technology, China and Southeast University, respectively. He joined the Department of Computer Science at the City University of Hong Kong from 2008 to 2009 as a research associate, and then joined the Department of Computer Science at the University of Victoria from 2011 to 2013 as a visiting scholar. His research interests include network security, privacy, and forensics.

Xinwen Fu (xinwenfu@cs.uml.edu) received his B.S. and M.S. degrees in electrical engineering from Xian Jiaotong University, China, and the University of Science and Technology of China in 1995 and 1998, respectively. He obtained his Ph.D. degree in computer engineering from Texas A&M University, College Station, in 2005. He is an associate professor in the Department of Computer Science, University of Massachusetts Lowell. His current research interests include network security and privacy, digital forensics, wireless networks, and network QoS.

Wei Yu (wyu@towson.edu) is currently an associate professor with the Department of Computer and Information Sciences, Towson University. He received his B.S. degree in electrical engineering from Nanjing University of Technology in 1992, his M.S. degree in electrical engineering from Tongji University, Shanghai, China in 1995, and his Ph.D. degree in computer engineering from Texas A&M University in 2008. He received the U.S. National Science Foundation (NSF) Early CAREER Award in 2014 and University System of Maryland (USM) Regents' Faculty Award for Excellence in Scholarship, Research, or Creative Activity in 2015. His research interests include cyberspace security, computer networks, and cyber-physical systems.

1st IEEE International 5G Summit
Tuesday, May 26, 2015

www.5Gsummit.org
#PRINCETON5G

Friend Center at Princeton University
35 Olden Street, Princeton, NJ




To help drive innovation, standards development and rapid deployment of emerging technology in areas such as SDN/NFV, 5G, IoT, Big Data and Cybersecurity, IEEE Communications Society will hold a series of high impact one day summits.

The 1st IEEE International 5G Summit will be held at Princeton University on Tuesday, May 26, 2015. This one day event will provide a platform for industry leaders, innovators, and researchers from the university and academic community to collaborate and exchange ideas on 5G.

Speakers include key industry leaders and eminent personalities from AT&T, Cisco, China Mobile, Columbia University, Google, Huawei, Intel, Keysight Technologies, NIST, National Instruments, Princeton University, Qualcomm, Rutgers and Verizon Wireless.

Register online at www.5Gsummit.org.

Keynote Speakers

Marian Croak
Google

Joe Cozzolino
Cisco

Edward G. Amoroso
AT&T

Princeton / Central Jersey 65th Anniversary IEEE 1947 - 2012 Section

CTiF Center for TeleInFrastruktur

IEEE COMMUNICATIONS SOCIETY

IEEE Region 1

ENERGY HARVESTING COMMUNICATIONS: PART 1






Chau Yuen *Maged Elkashlan* *Yi Qian* *Trung Q. Duong* *Lei Shu* *Frank Schmidt*

Over the last decade, energy harvesting has emerged as a promising approach to enable self-sufficient and self-sustaining operation for low-cost devices in energy-constrained networks by scavenging energy from the ambient environment to power up devices.

In wireless sensor networks, small, wireless, autonomous sensors usually operate at ultra-low power. If these wireless sensors, which spread throughout homes or factories, in buildings or even outdoors to monitor all kinds of environmental conditions, are powered by energy harvesting, there are no batteries to replace and no laborious cost associated with replacing them. As such, wireless sensor networks can be deployed in hard-to-reach areas to provide ubiquitous coverage.

In cellular networks, energy harvesting can be used to provide power in many elements of a telecom network, saving considerable cost in electricity supply and providing low maintenance monitoring. Powering base stations with wind or solar power allows telecom networks to expand beyond the limit of the power grid. The possibility of redistribution of the renewable energy in smart grid allows further efficient utilization, although it leads to many challenges as well.

Another important focus of this research area is RF energy harvesting. RF energy is currently broadcasted from billions of radio transmitters around the world, including mobile telephones, handheld radios, mobile base stations, and television/ radio broadcast stations. The ability to harvest RF energy, from ambient or dedicated sources, enables wireless charging of low-power devices and provides significant benefits to product design, usability, and reliability.

This Feature Topic includes 10 accepted papers, which address a number of critical and relevant issues studied in the emerging area of energy harvesting communications. We hope this Feature Topic is able to help readers obtain better understanding of some key issues in energy harvesting and drive more research interest.

This Feature Topic starts with an article, "Smart RF Energy Harvesting Communications: Challenges and Opportunities" by Deepak Mishra *et al.*, which explores various communication strategies that can complement the RF harvesting hardware advances toward the realization of energy harvesting communication networks.

The article "A General Utility Optimization Framework for Energy-Harvesting-Based Wireless Communications" by Hang Li *et al.* introduces a general utility optimization framework for energy-harvesting-based wireless communication systems subject to a novel type of energy usage constraint.

The article "Application of Smart Antenna Technologies in Simultaneous Wireless Information and Power Transfer," written by Zhiguo Ding *et al.*, proposes an application of smart antenna technologies, MIMO and relaying, to simultaneous wireless information and power transfer systems, and provides some future research challenges for the design of those systems.

The article "RF-Powered Cognitive Radio Networks: Technical Challenges and Limitations" by Lina Mohjazi *et al.* presents an overview of the architecture of cognitive radio networks that operate based on intended and unintended RF energy harvesting.

The article "Provisioning Quality of Service to Energy Harvesting Wireless Communications" by Xiaojing Chen *et al.* develops a dynamic string tautening method to produce the most energy-efficient schedule that adapts to the bursty characteristics of wireless traffic and energy harvesting.

The article "Increasing Sustainability and Resiliency of Cellular Networks Infrastructure by Harvesting Renewable Energy" by Andres Kwasinski *et al.* discusses the use of harvested renewable energy to power cellular base stations to reduce the carbon footprint of cellular infrastructure and to enable the deployment of cellular service in areas that lack electrification infrastructure.

The article "Wireless Powered Communication: Opportunities and Challenges," written by Suzhi Bi *et al.*, presents an overview of state-of-the-art RF enabled wireless energy transfer technology and it applications to wireless communications, with highlights on the key design challenges, solutions, and opportunities ahead.

The article "Fundamental Limits of Energy Harvesting

Communications" by Omur Ozel *et al.* surveys recent results in the literature and point to open problems in the fields of communication theory, information theory, signal processing, and networking.

The article "Enhancing Wireless Information and Power Transfer by Exploiting Multi-Antenna Techniques," written by Xiaoming Chen *et al.*, provides a tutorial on various aspects of multi-antenna-technique-based wireless information and power transfer, with a focus on tackling the challenges through parameter optimization and protocol design.

The article "Green Delivery: Proactive Content Caching and Push with Energy-Harvesting-Based Small Cells" by Sheng Zhou *et al.* proposes a new access network framework that enables efficient content delivery via energy-harvesting-based small cells to provide more multicast opportunities.

BIOGRAPHIES

CHAU YUEN (yuenchau@sutd.edu.sg) received his B. Eng and Ph.D. degrees from Nanyang Technological University, Singapore, in 2000 and 2004, respectively. He was a postdoctoral fellow at Lucent Technologies Bell Labs, Murray Hill, New Jersey, during 2005. He was a visiting assistant professor at Hong Kong Polytechnic University in 2008. During the period of 2006–2010, he worked at the Institute for Infocomm Research, Singapore, as a senior research engineer. He joined Singapore University of Technology and Design as an assistant professor in June 2010. He serves as an Associate Editor for *IEEE Transactions on Vehicular Technology* and was awarded Top Associate Editor for three consecutive years. In 2012, he received the IEEE Asia-Pacific Outstanding Young Researcher Award. He has held positions on several conference organizing committees, and is on Technical Program Committees of various international conferences.

MAGED ELKASHLAN received his Ph.D. degree in electrical engineering from the University of British Columbia, Canada, in 2006. From 2006 to 2007, he was with the Laboratory for Advanced Networking at the University of British Columbia. From 2007 to 2011, he was with the Wireless and Networking Technologies Laboratory at the Commonwealth Scientific and Industrial Research Organization (CSIRO), Australia. He also held an adjunct appointment at the University of Technology Sydney, Australia, between 2008 and 2011. In 2011, he joined the School of Electronic Engineering and Computer Science at Queen Mary, University of London, United Kingdom, as an assistant professor. His research interests include millimeter wave communications, energy harvesting, cognitive radio, and wireless security. He currently serves as an Editor for *IEEE Transactions on Wireless Communications*, *IEEE Transactions on Vehicular Technology*, and *IEEE Communications Letters*. He received Best Paper awards at IEEE ICC '14, International Conference on Communications and Networking in China in 2014, and IEEE VTC-Spring 2013. He received the Exemplary Reviewer Certificate of *IEEE Communications Letters* in 2012.

YI QIAN [M'95, SM'07] is an associate professor in the Department of Electrical and Computer Engineering, University of Nebraska-Lincoln (UNL). Prior to joining UNL, he worked in the telecommunications industry, academia, and the government. Some of his previous professional positions include serving as a senior member of scientific staff and technical advisor at Nortel Networks, a senior systems engineer and technical advisor at several startup companies, an assistant professor at the University of Puerto Rico at Mayaguez, and a senior researcher at the National Institute of Standards and Technology. His research interests include information assurance and network security, network design, network modeling, simulation and performance analysis for next generation wireless networks, wireless ad hoc and sensor networks, vehicular networks, smart grid communication networks, broadband satellite networks, optical networks, high-speed networks, and the Internet. He has a successful track record in leading research teams and publishing research results in leading scientific journals and conferences. Several of his recent journal articles on wireless network design and wireless network security are among the most accessed papers in the IEEE Digital Library. He is the current Chair of the Communications and Information Security Technical Committee in the IEEE Communications Society. He is an IEEE Distinguished Lecturer.

TRUNG Q. DUONG received his Ph.D. degree in telecommunications systems from Blekinge Institute of Technology (BTH), Sweden, in 2012, and then continued working at BTH as a project manager. Since 2013, he has joined Queen's University Belfast, United Kingdom, as a lecturer (assistant professor). He held visiting positions at Polytechnic Institute of New York University and Singapore University of Technology and Design in 2009 and 2011, respectively. His current research interests include cooperative communications, cognitive radio networks, green communications, physical layer security, massive MIMO, cross-layer design, mmWave communications, and localization for radios and networks. He has been a TPC chair for several IEEE international conferences and workshops, including most recently the IEEE GLOBECOM '13 Workshop on Trusted Communications with Physical Layer Security. He currently serves as an Editor for *IEEE Communications Letters* and Wiley *Transactions on Emerging Telecommunications Technologies*. He served as Lead Guest Editor of the Special Issue on Location Awareness for Radios and Networks of the *IEEE Journal on Selected Areas in Communications*, Lead Guest Editor of the Special Issue on Secure Physical Layer Communications of *IET Communications*, Guest Editor of the Special Issue on Green Media: Toward Bringing the Gap between Wireless and Visual Networks of *IEEE Wireless Communications*, Guest Editor of the Special Issue on Millimeter Wave Communications for 5G of *IEEE Communications Magazine*, Guest Editor of the Special Issue on Cooperative Cognitive Networks of the *EURASIP Journal on Wireless Communications and Networking*, and Guest Editor of the Special Issue on Security Challenges and Issues in Cognitive Radio Networks of the *EURASIP Journal on Advances Signal Processing*. He was awarded the Best Paper Award at IEEE VTC-Spring '13 and the Exemplary Reviewer Certificate of *IEEE Communications Letters* in 2012.

LEI SHU [M] received his B.Sc. degree in computer science from South Central University for Nationalities, China, in 2002, his M.Sc. degree in computer engineering from Kyung Hee University, Korea, in 2005, and his Ph.D. degree from the Digital Enterprise Research Institute, National University of Ireland, Galway, Ireland, in 2010. Until March 2012, he was a specially assigned researcher in the Department of Multimedia Engineering, Graduate School of Information Science and Technology, Osaka University, Japan. He is a member of IEEE IES, IEEE ComSoc, EAI, and ACM. In October 2012, he joined Guangdong University of Petrochemical Technology, China, as a full professor. In 2013, he started to serve Dalian University of Technology as a Ph.D. supervisor in the College of Software, Beijing University of Posts and Telecommunications as a Master's supervisor in information and communication engineering, Wuhan University as a Master's supervisor in the College of Computer Science, guest professor at Tianjin University of Science and Technology, and a guest researcher at Guangzhou Institute of Advanced Technology, Chinese Academy of Sciences. Meanwhile, he is also working as vice-director of the Guangdong Provincial Key Laboratory of Petrochemical Equipment Fault Diagnosis, China. He is the founder of the Industrial Security and Wireless Sensor Networks Lab. HIs research interests include wireless sensor networks, multimedia communication, middleware, fault diagnosis, and security. He has published over 230 papers in related conferences, journals, and books in the area of sensor networks. Currently, his H-index is 21 in Google Citation. Total citations of his papers by other people are more than 1600. He developed an open source wireless sensor networks simulator, NetTopo, to evaluate and demonstrate algorithms. NetTopo has been downloaded more than 3420 times over the past three years, and is widely used by international researchers and students. He was awarded the MASS 2009 IEEE TCs Travel Grant and the Outstanding Leadership Award of EUC 2009 as Publicity Chair, GLOBECOM 2010, and ICC 2013, the ComManTel 2014 Best Paper Award, and the Outstanding Service Award of IUCC 2012 and ComcomAP 2014. He also received a few more awards from the Chinese government: Top Level Talents in "Sailing Plan" of Guangdong Province, China, and Outstanding Young Professor of Guangdong Province, China. He has been serving as Editor-in-Chief for *EAI Endorsed Transactions on Industrial Networks and Intelligent Systems*, and Associate Editor for *IEEE Access*, ACM/Springer *Wireless Networks*, *Journal of Network and Computer Applications*, *Transactions on Emerging Telecommunications Technology*, and several other publications. He has served as Co-Chair for many international conferences. He has obtained more than 4 million RMB in research grants since October 2012.

FRANK SCHMIDT is a pioneer in energy harvesting and the visionary in the management team of EnOcean. As chief technology officer he is responsible for the overall technical orientation, patent related activities, as well as the relationship management with educational, research and scientific organizations. Before joining EnOcean he was at the Central Research Department of Siemens AG, where he created self-powered wireless sensor technology as early as 1995. He has been granted more than 40 patents for his energy harvesting inventions and is the author of numerous technical publications in this field. He is a physicist and studied at the Technical University of Chemnitz, Germany.

Smart RF Energy Harvesting Communications: Challenges and Opportunities

Deepak Mishra, Swades De, Soumya Jana, Stefano Basagni, Kaushik Chowdhury, and Wendi Heinzelman

ABSTRACT

RF energy harvesting (RFH) is emerging as a potential method for the proactive energy replenishment of next generation wireless networks. Unlike other harvesting techniques that depend on the environment, RFH can be predictable or on demand, and as such it is better suited for supporting quality-of-service-based applications. However, RFH efficiency is scarce due to low RF-to-DC conversion efficiency and receiver sensitivity. In this article, we identify the novel communication techniques that enable and enhance the usefulness of RFH. Backed by some experimental observations on RFH and the current state of the art, we discuss the challenges in the actual feasibility of RFH communications, new research directions, and the obstacles to their practical implementation.

INTRODUCTION

In recent times, RF energy harvesting (RFH) has emerged as a promising technology for alleviating the node energy and network lifetime bottlenecks of wireless sensor networks (WSNs). The RF radiation pattern is generally wide-angled; radio waves can simultaneously carry information and energy, and the radiation directivity can be electronically steerable. These features have been exploited in multihop energy transfer (MHET) as well as combining it with data transfer over the same RF signal (called simultaneous wireless information and power transfer, or SWIPT), without requiring critical alignment of the nodes [1, 2]. Besides, RFH has found applications in cognitive radio networks, wireless body area networks, and other wireless charging systems [3].

In RFH, RF wave radiations in the frequency range 3 kHz to 300 GHz are used as energy carriers. The amount of electrical energy that can be harvested is dependent on the power being emitted from the RF source, the antenna gains of the RF source and the receiving device, the distance of the receiving antenna from the RF source antenna, path loss exponent, and the RF-to-DC rectification efficiency η_{RF-DC}. The received electrical power is

$$P_R^{DC} = (\eta_{RF-DC})P_R,$$

where P_R is the received RF power that can be calculated using the Friis transmission equation.

RF energy sources can be classified into two categories:

- *Ambient RF source*: Ambient RF energy sources are not actually dedicated to RF energy transfer (RFET), and this RF energy is freely available. The frequency range of ambient RF transmission is 0.2–2.4 GHz, and this includes most of the radiations from domestic appliances (e.g., television, Bluetooth, WiFi).
- *Dedicated RF source*: This on-demand supply generally has a relatively higher power density due to directional transmission, and it is used to recharge nodes that require predictable and high amounts of energy. The energy transfer is done in the license-free industrial, scientific, and medical (ISM) frequency bands.

As RFH from dedicated RF sources, also known as RFET, is fully controllable, it is better suited for supporting applications with quality of service (QoS) constraints.

While RFET shows several promising directions and has an advantage over non-radiative wireless energy transfer in terms of relaxed coupling/alignment requirements [4], RFH suffers from various losses, including path loss, energy dissipation, shadowing, and fading. The problem is compounded by low energy reception sensitivity, restriction of maximum RF energy radiation due to human health hazards, and sharply decreasing RF-to-DC conversion efficiency at low receive powers. These RFH constraints place additional challenges compared to wireless data transfer because the information reception sensitivity is higher by a few orders of magnitude (typically –60 dBm in data reception vs. –10 dBm in RFH). This implies that with the current state of devices and RF circuits technologies, some applications may have limited practical utility. For example, the wireless energy plus

Deepak Mishra and Swades De are with IIT Delhi.

Soumya Jana is with IIT Hyderabad.

Stefano Basagni and Kaushik Chowdhury are with Northeastern University.

Wendi Heinzelman is with the University of Rochester.

0163-6804/15/$25.00 © 2015 IEEE

data transfer paradigm in two-hop decode-and-forward relay mode may not work with conventional inter-node distances (a few tens of meters) because the currently realizable energy transfer range is only on the order of 1 m.

The first goal of this article is to provide an overview of the recent developments toward improving RFH efficiency. Some novel approaches we discuss include increasing RFH circuit efficiency, multi-path energy routing (MPER), multi-antenna energy transmission, distributed beamforming techniques, and protocol-based optimization for cooperative energy transmission. Second, we provide experimental insight on the practical implementations of some of these strategies for enhancing RFH efficiency and their corresponding implications, followed by the current theoretical practices in this regard. Next, driven by experimental insights and practical system parameter considerations, the article highlights the challenges and the allied systems research opportunities for various aspects of RFH communications.

Improved RFH Circuits

The general architecture of a RFH unit is shown in Fig. 1.

The harvested energy is used to run a low-power micro-controller that processes the data from the application unit and controls the node's overall operation including information transmission and reception. The effectiveness of the RFH circuit is mainly determined by the RF-to-DC conversion efficiency and the DC output voltage. The conversion efficiency depends on the effectiveness of the antenna in collecting RF power, the precision of the matching circuit in energy conversion in the chosen frequency range, and the choice of the number of stages and diodes in the multiplier circuit.

Improving RF-to-DC Conversion Efficiency

Maximum power transfer from the antenna to the voltage multiplier can be realized when the antenna output impedance and load impedance are conjugates of each other (impedance matching). For the RFH circuit to work efficiently at low input power, diodes with low turn on voltage are used in the voltage multiplier circuit. The number of multiplier stages also has a significant impact, as a higher number of stages provides higher load voltage but reduces the load current in the process, whereas a lower number of stages provide a faster charging, but the load voltage is significantly lowered. As the received input RF power is very low, Dickson topology comprising multiple stages of parallel capacitors is used for high RF-to-DC conversion efficiency [5].

Because of the nonlinearity of the diode characteristics, the energy conversion efficiency sharply reduces at low input RF power [5, 6]. A dual-stage design — one with seven stages that works well for low input RF power and the other with 10 stages for higher input RF power — was proposed in [5], and an optimization framework was used to decide on the switchover point between the two stages, resulting in about 20 percent efficiency improvement over the commercially available Powercast harvester. Recently, in [7] it was shown that the efficiency can be further improved by using a resonating L type matching circuit along with a low pass filter (LPF) at the last stage. The resonator circuit exhibits resonant behavior at a specific frequency that can strengthen the weak RF power signals significantly. The LPF at the last stage of the harvesting circuit reduces the output harmonics and ripples in the output voltage to increase the output DC voltage.

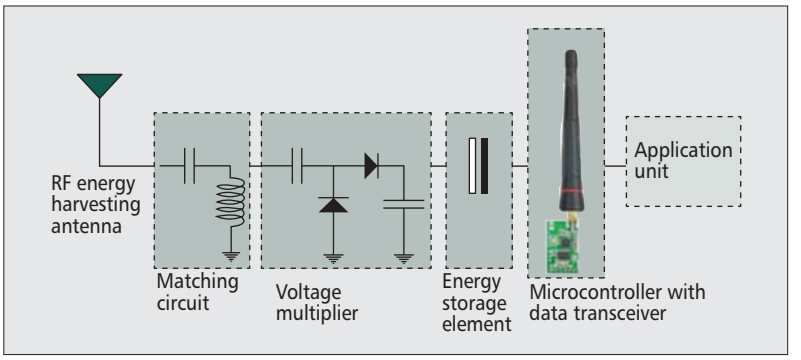

Figure 1. RF energy harvesting node.

Scalable Rectenna Array

As the incident ambient RF waves vary in frequency, power density, polarization, and incidence angles, scalable rectenna arrays with optimized power management circuits have been discussed in [8] for increasing the RFH efficiency. For maximum DC power generation, inserting a transitional DC-DC converter with peak power tracking that can reconfigure the equivalent DC load of the rectenna array with varying input RF power has been suggested. Also, it has been noted that the amount of harvesting power can be maximized by optimizing the antenna cover area on printed circuit boards by placing a greater number of antenna patches.

While these are some of the RFH circuit and hardware related developments, the main focus in this article is the advances and opportunities involving communication systems.

Experimental Insight on Smart RFH Communications

As noted earlier, RFH performance is limited by low energy reception sensitivity, low conversion efficiency at low input power, and the maximum allowable RF radiation power. In this section, we present some experimental observations on RFET with a special focus on multi-path energy routing (MPER), which provides efficient RFH communication by overcoming these hardware-based shortcomings. MPER helps improve RFH efficiency by first collecting the dispersed or dissipated RF energy transmitted by the RF source with the help of energy routers, and then directing it to the desired sensor node via paths other than the direct single hop path (Fig. 2a). These "energy routers" can be part of the network or may be introduced as optimally positioned

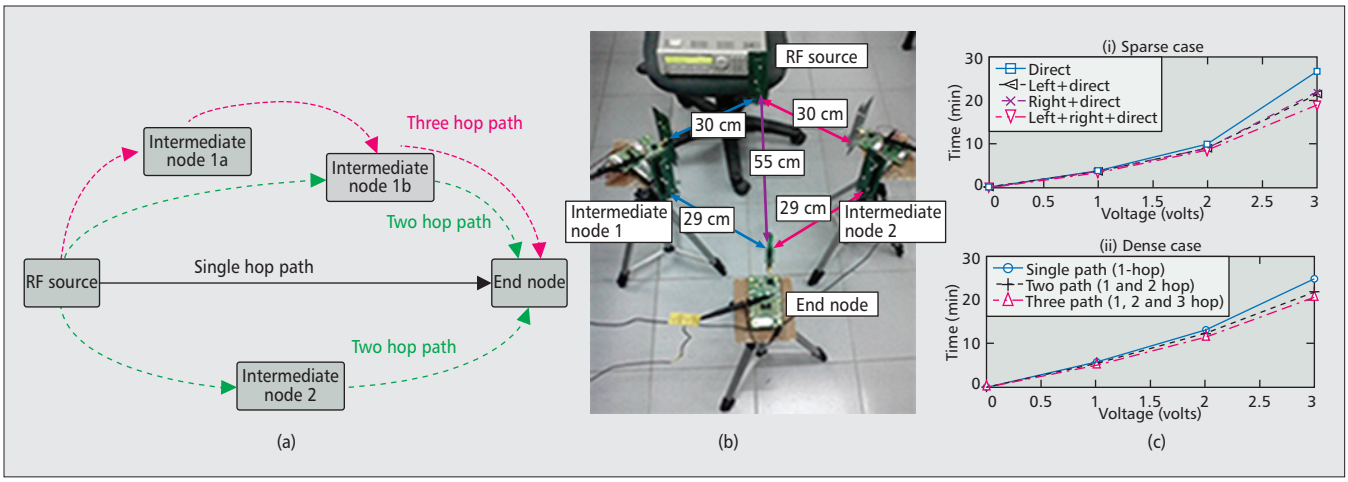

Figure 2. Multipath energy routing: a) block diagram of MPER; b) experiment setup for MPER in a sparse network; c) charging time comparison.

dummy nodes. MPER is based on the principle that multihop energy transfer (MHET) is beneficial to the energy transfer process.

MHET can improve energy harvesting efficiency by deploying relay nodes close to the target sensor node. The gain in MHET is achieved because of reduced path loss from the relay node to the end node, and improved RF-to-DC efficiency due to higher received power. Under the same principle, MHET can provide RFH range extension, which helps in implementing smarter RFH communications as it reduces the gap between the energy transfer and data transfer ranges. Feasibility studies in [2] showed that under certain optimum distance conditions, MHET can provide energy and time gains over direct energy transfer (DET). The improved performance of MHET over DET was experimentally demonstrated in [9]. The MHET experimental studies have been further extended to more generalized cases of MPER, which are discussed next.

MPER in a Sparse Deployment Scenario

In a sparse deployment, the intermediate nodes' presence does not cause any blocking or shadowing to the direct line of sight (LOS) path to the end node. However, the intermediate node is in a disadvantageous position because it can receive a lesser amount of the signal. The MPER experimental setup is shown in Fig. 2b, where, to improve gain, two intermediate nodes are symmetrically placed on either side of the LOS path, constituting a three-path energy transfer.

The system specifications for the experimental setup are: HAMEG RF synthesizer transmitting +13 dBm at 915 MHz via a 6 dBi Powercast patch antenna; intermediate nodes composed of a Powercast P1110 EVB, a Mica2 mote, two Powercast 6 dBi patch antennas; and an end node comprising a Powercast P1110 EVB and a Powercast 1 dBi dipole antenna. Further details on the experimental setup can be found in [9]. The intermediate nodes store the energy harvested via a 6 dBi antenna in a 50 mF capacitor for running the Mica2 mote. They forward the energy in the form of data packets from the modified Mica2 mote to the end node via another 6 dBi antenna every time the capacitor is fully charged. For efficient RFET, the Mica2 mote has been reprogrammed to transmit packets continuously one after the other during the energy transmission state [9].

The MPER performances in the two-path scenarios (left+direct, right+direct) as well as in the three-path scenario (direct+left+right) are shown in Fig. 2c.i. Compared to DET, time saving to charge the end node's capacitor up to 3 V is about 18 and 28 percent, respectively, in the two-path and three-path cases. The energy gain is the same as the time gain, as energy and time are proportional for a constant power source.

MPER in a Dense Deployment Scenario

In a dense deployment, charging one sensor node directly using LOS RF energy transmission may not be very efficient because of blocking/shadowing caused by the neighboring nodes. Hence, these intermediate nodes can be made to act like energy routers for the end node by adding transmission capabilities to them. Here, recharging multiple nodes simultaneously can also improve the overall system efficiency, as the sensor nodes near to the target sensor node can collect the otherwise dissipated energy.

The system specifications for the experimental setup are similar to the sparse scenario, except that a Hittite RF Synthesizer transmitting +23 dBm was used as the RF source, and the end node receives energy via a 6 dBi PCB patch antenna to overcome the blocking loss due to lower inter-node distances. The performance comparison in this case is shown with respect to the number of hops, which also demonstrates the feasibility of three-hop energy transfer. Referring to Fig. 2a, the intermediate nodes present are node 1a and node 1b. In the two-hop path, node 1a does not participate in RFET.

The representative results as plotted in Fig. 2c.ii show that both two-path (1-hop and 2-hop) and 3-path (1-hop, 2-hop, and 3-hop) MPER provide time gains of around 12 and 18 percent, respectively, over DET for charging the end node up to 3 V.

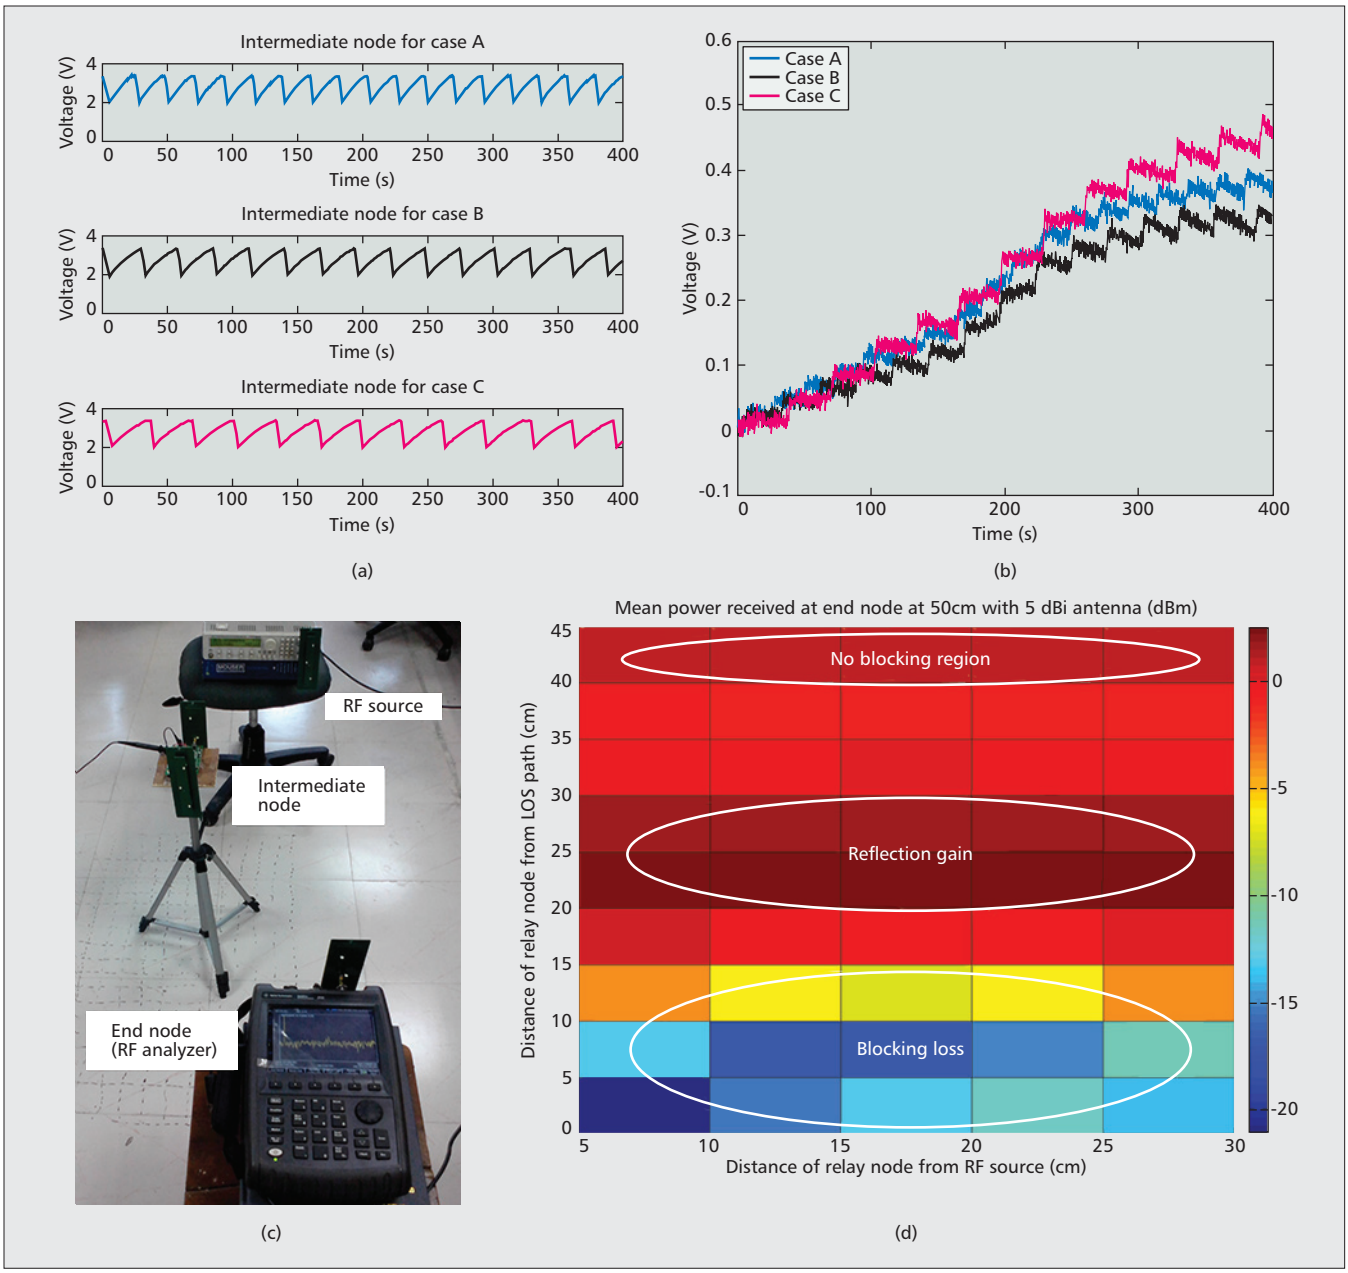

Figure 3. Effect of intermediate node placement: a) number of on-off cycles comparison; b) contribution of relay ($V_{on} - V_{off}$); c) experiment setup; d) blocking characterization.

Optimal Relay Placement

As demonstrated earlier, MHET can improve RFET efficiency. However, this improvement is strongly influenced by the placement of relay nodes. To show the effect of relay node position, consider the following three cases in a 2-hop RFET (Fig. 3c) with RF source transmitting at +13 dBm over Powercast directional antennas at 915 MHz from a distance of 30 cm to the end node.

Case A: The relay node is closer to the RF source, so it harvests energy at a faster rate and forwards energy more frequently (with a higher number of on/off cycles; Fig. 3a). However, the energy received per cycle at the end node is very low due to a higher path loss.

Case B: The relay node is at the midway point to the end node. It harvests less energy compared to case A over a given time. Also, the energy received per forwarding cycle by the end node is lesser than case C due to higher path loss.

Case C: The relay node is closer to the end node. In this case the node harvests the least amount of energy over a given time, but the energy received per cycle by the end node is the highest due to minimum path loss. The overall performance of case C was noted to be the best, as shown in Fig. 3b. The average energy gain was noted to be about 10 percent over the worst case [9].

The above experimental study, however, calls for optimization formulations to find the optimum intermediate node positions to maximize RFET efficiency under different deployment

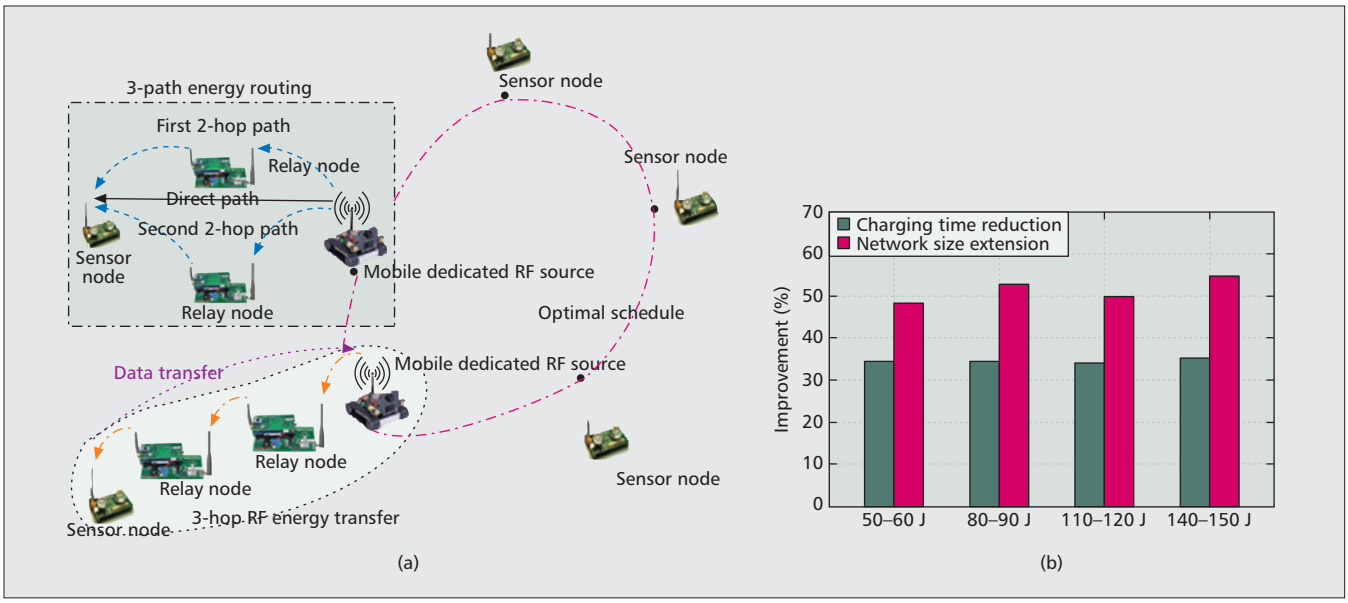

Figure 4. Networking consequence of improved RFH efficiency: a) quicker charging using dedicated RF source with MPER; b) network size expansion.

scenarios. This requires characterization of the blocking losses. For the setup in Fig. 3c, the blocking region characteristics are shown in Fig. 3d. It is clear that the intermediate node can cause significant blocking loss. Interestingly, there lies an intermediate region between the blocking and nonblocking region, which can provide energy gain due to reflection.

RF Charging Time Characterization

To characterize the RFET and MPER, it is necessary to characterize the capacitor charging from an RF source. A recent study [10] has shown that RF charging is different from the conventional constant voltage charging. It is a special case of constant power charging, because the RF power received for recharging the capacitor is fixed for an RF source transmitting constant power from a fixed distance. It provides the analytical expression for the time required to charge a capacitor having an initial voltage up to some threshold value depending on the energy requirement of the sensor node. Experimental validation of the RF charging equations using a Powercast P1110 EVB has also been provided. The derived charging time equation in [10, Eq. 13] as a function of received DC power P_R^{DC}, circuit parameters (capacitance and resistance), and the lower and upper voltage limits (corresponding to the initial and final charge stored in the capacitor) can be used for analyzing the RFH efficiency.

Network Size Expansion: Consequence of MPER

We now discuss how the energy efficiency improvements provided by improved RFH techniques via MPER can provide network size expansion. Consider the network model shown in Fig. 4a, where a mobile node (called the integrated data and energy mule, or IDEM) with dedicated RF source and MPER capability has the objective of providing uninterrupted network operation by regularly visiting the nodes to recharge them and collect field data [2]. Network size can be defined as *the number of nodes that can be served by a single mobile dedicated RF source in such a way that none of the nodes ever runs out of energy*. This network size extension is achieved by quicker charging of the nodes via advanced RFH circuits (20 percent gain [5]) and RFH communication (total 30 percent gain [9]) techniques. In total, we consider an energy harvesting efficiency improvement of around 50 percent (Table 1), leading to 50 percent more received DC power. The corresponding charging time gain is obtained using [10, Eq. 13].

In [11], it was shown that the average energy consumption per sensing cycle in a node in a pollution monitoring application increases from 50–60 J to 140–150 J as the number of sensors per node is increased from one to four, which is an increment of nearly 30 J per additional sensor. The node energy consumption is a random variable because it depends on the pollution level. Thus, the network size depends on the average charging time, which in turn depends on the average energy consumption per node. In order to serve the maximum number of nodes, the IDEM should spend the minimum time traveling so that the length of the overall tour is minimized.

We have considered sensor nodes that are uniformly randomly deployed over a square field of 5 km × 5 km. The IDEM speed is assumed to be 5 m/s. The charging time parameters, based on the experimental observations, are: charging distance d = 0.328 m; RF transmit power = 3 W; operating frequency = 915 MHz; transmitter and receiver antenna gains of 6 dBi; capacitor value C = 20 F with ESR R = 0.16 Ω. For simplicity, we have considered that each node will be visited only once in a cycle. Hence, the IDEM should follow the shortest

S. No.	Strategy		Gain (%)	Challenges	Opportunities
1	**Improved RFH circuit design**		20–30% [5, 7]	Low efficiency for very low RF inputs; Hardware constraints; Supporting wide-band and multi-band operation	It plays the central and most significant role; can provide efficient ambient RFH; bridges the gap between data and energy sensitivity of receiver; scalable rectenna array
2	**Multipath energy routing (MPER)**	(a) MPER in sparse networks	10–20%	Lower efficiency due to lower transmit power of the relay; Cost of deploying dummy nodes	Can provide RFET range extension, if single-hop energy cannot be received due to path loss and lower receiver sensitivity
		(b) MPER in dense networks	10–30%	Lower inter-node distances; Node deployment not suitable to higher order MHET	Improves RFH efficiency by simultaneous charging of multiple nodes, MHET by using the nodes causing blocking of DET as energy routers
3	**Relay node optimizations**	(a) Optimal relay placement	5–10% [9]	Non-convex Optimization problem	Effectiveness of MPER and MHET is strongly affected by relay placement
		(b) Cooperative relaying [15]	Not quantified	Relay selection has to solve non-trivial reliable data-efficient energy transfer trade-off due to huge discrepancy in data and energy reception sensitivity	Can boost harvesting efficiency, meet QoS requirements by using relay nodes by exploiting the beamforming and diversity gains
4	**Beamforming**	(a) Distributed beamforming [13]	Not quantified	Overhead cost involved in the phase and frequency synchronization of the carrier signals generated by the local oscillators of differently located RF energy transmitters	Significant RFH efficiency improvement by cooperative transmission of distributed and independent transmitters; sophisticated digital implementation of optimal frequency and phase estimator instead of analog phase locked loop can further increase efficiency
		(b) Energy beamforming [1]	Not quantified	Nontrivial tradeoffs in allocating communication resources for optimizing interference levels and RFH efficiency; Form factor constraints	Energy allocation based on estimated CSI can provide a higher harvesting efficiency via energy beamforming
5	**Protocol-based optimizations**	(a) MAC	>100% [14]	Optimal energy transfer v/s data communication trade-off	Integration among efficient energy harvesting, multi-antenna transmission, data communication, resource management, and signal processing
		(b) Routing	Yet to study	Joint optimization of RFH and networking parameters; Factors like low receiver sensitivity, propagation losses, judicious utilization of the nearby sensor nodes (energy routers), varying residual energy at different nodes	Optimal joint routing and recharging scheme for mobile dedicated RF source(s) can lead toward uninterrupted network operation [2]

Table 1. Strategies for improving RFH efficiency.

Hamiltonian cycle, which has been found by solving the Traveling Salesman Problem using a genetic algorithm.

The simulation results are based on an average of 30 runs. The percentage improvement in charging time and network size achieved due to the implementation of RFH communication techniques is shown in Fig. 4b for all four cases. The results show that on average, there is about 35 percent reduction in charging time of the nodes and about 50 percent increase in the number of nodes than can be served by a single IDEM. Thus, RFH efficiency improvement can not only prolong the network lifetime but also provide network expansion.

Theoretical Advances on RFH Communications

We now discuss the recent developments on RFH communication that are primarily theoretically driven. These include novel methods at the physical layer as well as the upper communication layers.

Multiple Antenna Transmission

Single-antenna transmitters with omnidirectional radiation cause significant path loss with increasing transmission distance due to beam spreading. Multi-antenna transmission can achieve spatial multiplexing as in multiple-input multiple-output (MIMO) systems, by employing beamforming techniques (Fig. 5) to improve the RFH efficiency in long-distance energy transfer by exploiting large antenna array gain. This enables faster charging without any increase in transmit power. This RFET method of concentrating the RF waves in the direction of the intended receiver is called energy beamforming, which was first considered in [1] for SWIPT in multiuser downlink.

An issue associated with beamforming gains is channel state information (CSI) feedback. In [12], it was shown that energy beamforming based on accurate CSI feedback can provide higher energy transfer efficiency. However, this is at the cost of significant time overhead incurred at the receiver. A longer channel estimation duration can provide a

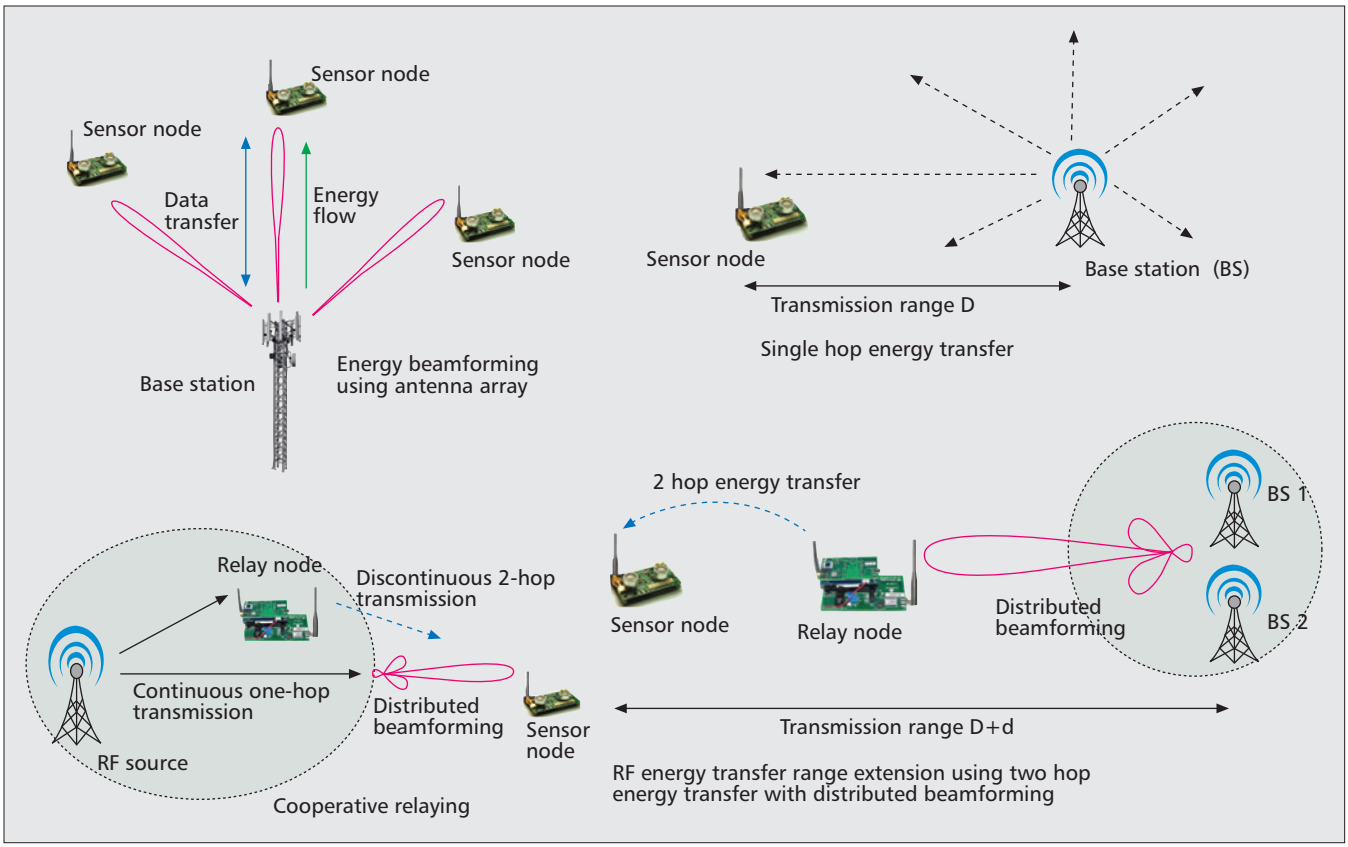

Figure 5. Beamforming techniques for the enhancement of RF harvesting efficiency.

more accurate CSI, but it also shortens the energy transfer duration, which leads to less harvested energy.

Distributed Beamforming of Multiple Transmitters

A fully wireless distributed beamforming prototype based on a software defined radio platform was proposed in [13], where several nodes fine-tune their data transmissions in a coordinated fashion so as to form a large virtual antenna array that directs the beam toward the receiver, thereby increasing the data rate and transmission range. Frequency and phase synchronizations were made using receiver feedback packet waveform and payload using extended Kalman filtering and a 1-bit feedback algorithm, respectively, which is an overhead cost. Along similar lines, collaborative beamforming of distributed RF energy transmitters can provide improved energy efficiency due to increased received power [14]. Cooperative beamforming has the potential to enhance energy efficiency by adjusting the carrier phase of each energy transmitter in such a manner that it can compensate for the path difference between the energy waves arriving at the target node, causing constructive interference. As a result, there can be a maximum of N^2 times power reception of RF power for N cooperative RF energy transmitters due to the increased directivity. Figure 5 shows that cooperative distributed beamforming can provide range extension, which can be further aided by 2-hop RFET. However, the major underlying challenge is the overhead cost required for frequency, phase, and time synchronization for high-frequency carrier signals.

Cooperative Relaying

Selection of a relay node among various relays strongly affects the performance of cooperative relaying that can provide improved energy transfer efficiency and better data transfer reliability. In [15] a stochastic-scale geometry approach has been adopted to study the impact of cooperative density and relay selection to analyze the fundmental trade-off between information transfer efficiency in terms of outage probability performance and RFH efficiency in SWIPT applications. But the author's objective of SWIPT to the same node is a very challenging task because of very different sensitivities to data and energy reception process. Cooperative energy relaying is useful in RFET, when the inter-nodal distances are small. Here, the end node can simultaneously receive from both the RF source and the relay node(s), which is conceptually similar to MPER. In fact, the energy gains can be further increased by distributed beamforming of continuous transmission of the RF source and the discontinuous transmission of relay nodes, as shown in Fig. 5.

Protocol Based Optimization

In [14] RFH is posed as a Medium Access Control (MAC) objective to maximize the RF energy transfer rate while minimizing interference to data communication. The proposed RF-MAC protocol tackled several challenges, namely time of energy transfer, priority between data and

energy transfer, multiple transmitter charging and choice of frequency for transmission. Categorization of different energy transmitters into two groups with varying transmission frequencies based on their phase differences, helped improve the RFH efficiency. It was shown that the RF-MAC protocol maintains a balance between the efficient RFH and data transfer by outperforming the classical modified carrier sense multiple access (CSMA) protocol in terms of both average harvested RF energy and the average network throughput.

The Way Forward: Research Challenges

Under the prism of the strategies discussed so far for implementing efficient RFH communication, we now discuss the challenges that lie ahead in the practical implementation of these techniques and their further extensions. These strategies are summarized in Table 1, indicating the corresponding gains (wherever available), the challenges in their implementation, and the opportunities associated with them.

Circuits and hardware constraints: RFH circuits suffering from quiescent losses at input RF power levels below –20 dBm is a limiting factor to ambient RFH and the practical implementation of SWIPT. Thus, there is an urgent need to narrow the gap between the receiver sensitivities for data and energy. With the advanced ultra-low-power electronics and custom circuits for ultra-low-power RF scavenging, it is possible to overcome this limitation in the future by integrating improved hardware with scalable approaches as in [8].

Optimizations on MPER: Energy gains provided by MPER can be improved significantly by relay node optimizations, like selecting the relay node's position, transmit power, capacitor size, optimal store and forward energy duration as decided by the charging and discharging levels, and so on. Furthermore, the optimal relay placement on a Euclidean 2D plane in a sparse network (no blocking of DET) is a nonconvex and highly nonlinear optimization problem. The problem in the dense deployment case is even more challenging, as it includes the blocking characterization of the relay node. Thus, it has to tackle the trade-off among blocking loss, reflection gain, and path loss. These formulations are some of our ongoing research.

Constraints on joint energy and data transfer: Although it has been assumed that the receiver is able to harvest energy and decode information simultaneously from the same RF signal, it is not feasible over practical data communication ranges. Hence, two practical approaches, time switching and power splitting, have been proposed in [1] for implementation of SWIPT. However, in spite of several virtues of cooperative relaying (cooperative diversity, efficient energy, and reliable data transfer), due to the huge discrepancy in the receiver's data and energy sensitivities, utilizing these assets for SWIPT is still an open issue. Also, as accurate CSI estimation can significantly affect both information and energy transfer efficiency, a key challenge is to balance time resources for channel estimation and SWIPT in multi-user MIMO systems. Distributed beamforming can overcome the form factor constraints of energy beamforming or conventional MIMO by forming a virtual MIMO or antenna array system, and provide benefits like increased directivity and spectral efficiency, and enhanced spatial diversity. However, there are underlying synchronization bottlenecks, which is an open research area.

Protocol-level challenges: As far as protocol-based optimization is concerned, there is a need to consider the practical limitations discussed above. To this end, the most important challenge is to have a protocol architecture (MAC+Routing) that jointly optimizes efficient energy transfer and reliable data transfer while taking into account various parameters. Some of the critical parameters of interest are relay node placement for efficient MPER, RF charging time characterization, cooperation among the participating nodes for data and energy transfer, interference minimization, and collaborative transmission of multiple transmitters.

Some of the critical parameters of interest are relay node placement for efficient MPER, RF charging time characterization, cooperation among the participating nodes for data and energy transfer, interference minimization, and collaborative transmission of multiple transmitters.

Concluding Remarks

This article has explored various communication strategies that can complement RFH hardware advances toward the realization of energy harvesting communication networks. The outlined experience on hardware implementation of MPER has revealed that while the energy routing concept is practically realizable and efficient, the energy transfer range is still low. While some concepts, such as ambient RFH-driven communication and joint energy and data transfer, may have to wait due to the significant asymmetry in energy and data transfer ranges with present-day technologies, future strategies, such as multi-antenna transmission, distributed beamforming, cooperative relaying, RF-MAC, and routing optimizations, are a few promising beacons to extend the benefits of RFH. Physical challenges include time, phase, and frequency synchronization of the independent transmitters for achieving beamforming gain. Likewise, there are challenges to the upper layer strategies, which have been summarized in Table 1. By overcoming these challenges, the combined effect of these strategies can make RFH-assisted network communication a popular technology.

Acnowledgment

This work was supported jointly by the Department of Electronics and Information Technology (DeitY) and the National Science Foundation awards "GENIUS: Green sEnsor Networks for aIr qUality Support" (grant numbers DeitY 13(2)/2012-CC&BT and NSF CNS 1143681).

References

[1] R. Zhang and C. K. Ho, "MIMO Broadcasting for Simultaneous Wireless Information and Power Transfer," *IEEE Trans. Wireless Commun.*, vol. 12, no. 5, May 2013, pp. 1989–2001.
[2] S. De and R. Singhal, "Toward Uninterrupted Operation of Wireless Sensor Networks," *IEEE Comp. Mag.*, vol. 45, no. 9, Sept. 2012, pp. 24–30.
[3] N. Shinohara, *Wireless Power Transfer via Radiowaves*, Wiley, 2014.

[4] A. Kurs *et al.*, "Wireless Power Transfer via Strongly Coupled Magnetic Resonances," *Science*, vol. 317, no. 5834, July 2007, pp. 83–86.
[5] P. Nintanavongsa *et al.*, "Design Optimization and Implementation for RF Energy Harvesting Circuits," *IEEE J. Emerging Sel. Topics Circuits and Sys.*, vol. 2, no. 1, Mar. 2012, pp. 24–33.
[6] Powercast: http://www.powercastco.com.
[7] S. Agrawal *et al.*, "Realization of Efficient RF Energy Harvesting Circuits Employing Different Matching Technique," *Proc. Int'l. Symp. Quality Electronic Design*, Santa Clara, CA, Mar. 2014, pp. 754–61.
[8] Z. Popovic *et al.*, "Scalable RF Energy Harvesting," *IEEE Trans. Microwave Theory Tech.*, vol. 62, no. 4, Apr. 2014, pp. 1046–56.
[9] K. Kaushik *et al.*, "Experimental Demonstration of Multi-Hop RF Energy Transfer," *IEEE PIMRC*, London, U.K., Sept. 2013, pp. 538–42.
[10] D. Mishra, S. De, and K. R. Chowdhury, "Charging Time Characterization for Wireless RF Energy Transfer," *IEEE Trans. Circuits and Sys. II, Exp. Briefs*, vol. PP, no. 99, 2015.
[11] P. Gupta *et al.*, " Feasibility Analysis on Integrated Recharging and Data Collection in Pollution Sensor Networks," *Proc. Nat. Conf. Commun.*, New Delhi, India, Feb. 2013.
[12] G. Yang, C. K. Ho, and Y. L. Guan, " Dynamic Resource Allocation for Multiple-Antenna Wireless Power Transfer," *IEEE Trans. Signal Processing*, vol. 62, no. 14, July 2014, pp. 3565–77.
[13] F. Quitin *et al.*, "A Scalable Architecture for Distributed Transmit Beamforming with Commodity Radios: Design and Proof of Concept," *IEEE Trans. Wireless Commun.*, vol. 12, no. 3, Mar. 2013, pp. 1418–28.
[14] M. Naderi, P. Nintanavongsa, and K. Chowdhury, "RF-MAC: A Medium Access Control Protocol for Re-Chargeable Sensor Networks Powered by Wireless Energy Harvesting," *IEEE Trans. Wireless Commun.*, vol. 13, no. 7, July 2014, pp. 3926–37.
[15] I. Krikidis, "Simultaneous Information and Energy Transfer in Large-Scale Networks with/without Relaying," *IEEE Trans. Commun.*, vol. 62, no. 3, Mar. 2014, pp. 900–12.

Biographies

Deepak Mishra [S'13] received his B.Tech degree in electronics and communication engineering from Guru Gobind Singh Indraprastha University, Delhi, India, in 2012. He is currently pursuing his Ph.D. degree from the Department of Electrical Engineering, Indian Institute of Technology (IIT) Delhi. His research interests include energy harvesting, wireless rechargeable sensor networks, and mobility management in ad hoc sensor networks.

Swades De [S'02, M'04, SM'14] received a Ph.D. in electrical engineering from the State University of New York Buffalo in 2004. He is currently an associate professor in the Department of Electrical Engineering, IIT Delhi. His research interests include performance study, resource efficiency in wireless networks, broadband wireless access, and communication and systems issues in optical networks. He is currently an Associate Editor for *IEEE Communications Letters* and Springer's *Photonic Network Communications*.

Soumya Jana received his B.Tech. (Honors) degree in electronics and electrical communication engineering from IIT Kharagpur in 1995, his M.E. degree in electrical communication engineering from the Indian Institute of Science, Bangalore, in 1997, and his Ph.D. degree in electrical engineering from the University of Illinois at Urbana-Champaign in 2005. He is currently an assistant professor in the Department of Electrical Engineering, IIT Hyderabad. His research interests include statistical signal processing, information theory, and immersive multimedia.

Stefano Basagni [SM] holds a Ph.D. in electrical engineering from the University of Texas at Dallas (December 2001) and a Ph.D. in computer science from the University of Milano, Italy (May 1998). Since 2002 he has been an associate professor at the Department of Electrical and Computer Engineering at Northeastern University, Boston, Massachusetts. His current research interests concern research and implementation aspects of mobile networks and wireless communications systems, radio and acoustic sensor networking, definition and performance evaluation of network protocols, and theoretical and practical aspects of distributed algorithms. He has published over six dozen refereed technical papers and book chapters that are highly cited (his h-index is currently 29, with over 7500 citations to his works). He is also co-editor of three books. He serves as a member of the Editorial Boards and Technical Program Committees of ACM and IEEE journals and international conferences. He is a Senior Member of the ACM (including the ACM SIGMOBILE), and a member of American Society for Engineering Education (ASEE) and Council on Undergraduate Research (CUR).

Kaushik Chowdhury has been an assistant professor in the Electrical and Computer Engineering Department at Northeastern University since 2009. He graduated with B.E. in Electronics Engineering with distinction from VJTI, Mumbai University, India, in 2003. He received his M.S. in Computer Science from the University of Cincinnati, OH, in 2006, and Ph.D. from the Georgia Institute of Technology, Atlanta, GA, in 2009. His M.S. thesis was given the outstanding thesis award jointly by the ECE and CS departments at the University of Cincinnati. He received the Best Paper Award at the IEEE ICC Conference in 2009, 2012, 2013, as well as the Best Paper award in the ICNC Conference in 2013. His expertise and research interests lie in wireless cognitive radio ad hoc networks, energy harvesting, and intra-body communication. He is currently an area editor for the Elsevier Ad Hoc and Computer Communications journals, and the Chair for the IEEE Technical Committee on Simulation.

Wendi Heinzelman is a Professor in the Department of Electrical and Computer Engineering at the University of Rochester, with a secondary appointment in the Computer Science Department at Rochester. She also currently serves as the Dean of Graduate Studies for Arts, Sciences & Engineering. She received a B.S. degree in Electrical Engineering from Cornell University in 1995 and M.S. and Ph.D. degrees in Electrical Engineering and Computer Science from MIT in 1997 and 2000, respectively. Her current research interests lie in the area of wireless communications and networking, mobile computing, and multimedia communication. She is currently an Associate Editor for Elsevier Ad Hoc Networks and Information Director for ACM Transactions on Sensor Networks. She is a member of Networking Networking Women (N^2 Women) and the Society of Women Engineers (SWE), a Distinguished Scientist of ACM Sigmobile, and a Senior Member of the IEEE Communications Society and the IEEE Signal Processing Society.

A General Utility Optimization Framework for Energy-Harvesting-Based Wireless Communications

Hang Li, Jie Xu, Rui Zhang, and Shuguang Cui

ABSTRACT

In the near future, wireless communication systems are expected to achieve more cost-efficient and sustainable operations by replacing conventional fixed power supplies such as batteries with energy harvesting devices, which could provide electric energy from renewable energy sources (e.g., solar and wind). Such EH power supplies, however, are random and unstable in nature, and as a result impose new challenges on reliable communication design and have triggered substantial research interest in EH-based wireless communications. Building on existing works, in this article, we develop a general optimization framework to maximize the utility of EH wireless communication systems. Our framework encapsulates a variety of design problems, such as throughput maximization and outage probability minimization in single-user and multiuser setups, and provides useful guidelines to the practical design of general EH-based communication systems with different assumptions regarding the knowledge of time-varying wireless channels and EH rates at the transmitters.

INTRODUCTION

Energy harvesting (EH) is expected to have abundant applications in future wireless communication networks to power transceivers by utilizing environmental energy such as solar, thermal, wind, and kinetic energy. Since renewable energy is generally clean and cheap, EH offers various benefits compared to conventional energy supplies such as fossil-fuel-based generators. For example, in cellular networks, solar panels and wind farms have been deployed to power base stations, thus lowering the expenses of energy bills as well as reducing the level of carbon dioxide emissions. Besides, in wireless sensor networks, EH has been considered as a good substitute for the traditional battery, prolonging the operation time to almost infinity, at least theoretically.

Despite the above advantages, the introduction of EH also imposes new challenges on the communication system design. Specifically, the random and intermittent characteristics of renewable energy impose a new type of EH constraint: the available energy at an EH communication node up to any time is bounded by its accumulatively harvested energy at that time. This is in contrast to conventional communication systems with fixed energy sources, in which the available energy at any time is either unbounded or only limited by the remaining energy in the storage device (e.g., battery). In addition, wireless communication channels often fluctuate more substantially and dynamically than practical EH rates (e.g., the channel changes on the order of milliseconds while the EH rate changes on the order of seconds or minutes), while channel fading is the main challenge faced in the design of reliable wireless communications. Due to both the new EH constraints and the multi-timescale channel/EH-rate variations, it is a challenging problem to jointly optimize the communication scheduling and energy management in EH-based wireless communications.

Given both the advantages and challenges mentioned above, EH-based wireless communications have drawn significant research attention in recent years. Various EH-oriented transmission policies have been proposed for different applications [1, references therein]. Building on prior works, we aim to develop a general utility optimization framework in this article to further reveal the fundamental limits and design principles of EH-based wireless communications, and apply this framework to solve some selected problems in single-user and multiuser communication setups. Finally, we discuss some future research directions that we deem worthy of further investigation.

SINGLE-USER EH-BASED COMMUNICATIONS: SYSTEM MODEL AND OPTIMIZATION FRAMEWORK

First, we consider in this section a point-to-point wireless communication system as depicted in Fig. 1, which consists of one transmitter powered by an energy harvester and one receiver with a reliable power supply. In practice, the coherence

Hang Li and Shuguang Cui are with Texas A&M University. Shuguang Cui is also a Distinguished Adjunct Professor at King Abdulaziz University in Saudi Arabia.

Jie Xu is with the National University of Singapore.

Rui Zhang is with the National University of Singapore and the Institute for Infocomm Research, ASTAR.

This work was supported in part by DoD with grant HDTRA1-13-1-0029, by NSF with grants CNS-1343155, ECCS-1305979, and CNS-1265227, by grant NSFC-61328102, and by NUS with grant R-263-000-679-133.

0163-6804/15/25.00 © 2015 IEEE

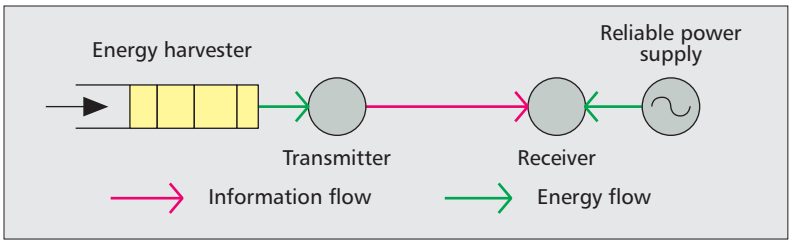

Figure 1. A point-to-point communication link with an EH transmitter and a receiver with reliable power supply.

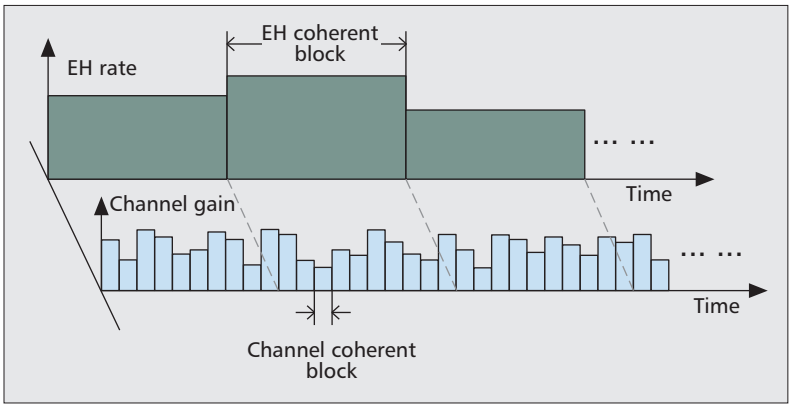

Figure 2. Time variation in the EH process vs. a wireless channel.

time of EH processes is often much larger than that of wireless channels, as previously mentioned. Therefore, a block-based quasi-static EH model is practically valid, where the EH rate remains constant within each EH coherent block and may change from one block to another, and at the same time each EH block spans over many communication channel coherent blocks, as shown in Fig. 2.

For the purpose of exposition, we consider wireless data transmissions over a finite horizon of $M \geq 1$ EH blocks. Each EH block is further divided into $N \geq 1$ communication blocks each of one unit time and a constant channel gain. Let $E(m) \geq 0$ denote the EH rate in the mth EH block, and $h(n, m) \geq 0$ the channel power gain of the (n, m)th communication block[1] with $1 \leq n \leq N$ and $1 \leq m \leq M$. Furthermore, we use $P(n, m) \geq 0$ to denote the power consumption at the transmitter in the (n, m)th communication block, which in general includes both the transmission power and the circuit power overhead. Assuming an ideal energy storage device (i.e., with infinite capacity and no energy leakage) employed at the transmitter, we have the *EH constraints* on the scheduled power consumptions $\{P(n, m)\}$; that is, the energy accumulatively consumed up to any communication block (n, m): $\sum_{j=1}^{m-1}(\sum_{i=1}^{N} P(i, j)) + \sum_{i=1}^{n} P(i, m)$, should be no larger than the energy accumulatively harvested by then, that is, $N\sum_{j=1}^{m-1} E(j) + nE(m)$ [2, 3].

To characterize the communication quality of service (QoS) measured at the receiver, we define a general utility function $U_{n,m}(P(n, m))$, which is dependent on the allocated power $P(n, m)$ at the transmitter for the (n,m)th block. Note that in general the utility function $U_{n,m}(P(n, m))$ also depends on the channel power gain $h(n, m)$ of the (n, m)th block. For example, the utility function can be defined more explicitly as throughout [2–4], non-outage probability [6, 7], or other performance metrics such as end-to-end distortion in an EH-based estimation system [8], which could be either deterministic or statistical average based on the availabilities of the channel state information (CSI), that is, $\{h(n, m)\}$, and the energy state information (ESI), $\{E(m)\}$, at the transmitter, or CSIT and ESIT, respectively.

Thus, the overall utility maximization problem over the M EH blocks could be formulated as (P1) depicted in Fig. 3. Among all different assumptions about the CSIT and ESIT, there are four cases of primary interest in this article, listed below.

Case 1: non-causal CSIT and ESIT. At the beginning of the transmission, the transmitter perfectly knows the past, current, and future CSI and ESI. This case approximates the practical scenario when the transmitter can accurately predict the future CSI (e.g., slowly varying channels in low-mobility applications) and the future ESI (e.g., based on historical data in a periodically varying energy environment). The optimal solution in this case provides a performance upper bound for all other CSIT/ESIT availability cases.

Case 2: causal CSIT and ESIT. At the beginning of each EH/communication block, the transmitter knows the past and current CSI/ESI, as well as the statistical information (e.g., distributions) of future CSI/ESI. In general, the solution of this case achieves the lowest utility among the first three cases considered herein.

Case 3: causal CSIT and non-causal ESIT. This is a hybrid model based on cases 1 and 2, in which all ESI is perfectly known at the beginning of the transmission, while only the past and current CSI is known.

Case 4: no CSIT and non-causal/causal ESIT. During the transmission, the transmitter does not have any CSI, and only has statistical information on the CSI. The non-causal or causal ESIT is defined as that in case 1 or 2 above.

Note that in all the above cases, we assume that at each communication block, the receiver perfectly knows the CSI in that block.

Optimal Solution for Single-User EH-Based Communications

In this section, we consider two example utility functions of the general utility maximization problem defined in problem (P1) in Fig. 3: the throughput maximization and non-outage probability maximization (or equivalently, outage probability minimization). In the literature, these two utilities have been commonly adopted to characterize the performance limits of communications over fading channels for delay-sensitive and delay-tolerant applications, respectively, assuming that the transmitter always has backlogged data to send. Based on these two utilities, we reveal the structure of the optimal power allocation solution at the EH transmitter to problem (P1).

[1] For notational convenience, we use (n, m) to denote the nth communication block of the mth EH block.

Throughput Maximization

For throughput maximization, the utility function $U_{n,m}(P(n, m))$ is generally modeled as the achievable rate at the (n, m)th communication block, which is a positive non-decreasing and concave function of the transmit power $P(n, m)$ after ignoring the circuit power consumption at the transmitter (for the more general case with additional circuit power consumption considered, please refer to [4, 5]). For example, the Shannon capacity formula is widely used to model $U_{n,m}(P(n, m))$ [2–5], that is, $U_{n,m}(P(n, m)) = \log_2(1 + h(n, m)P(n, m))$, in bits per second per Hertz. We present the optimal power allocation solution to problem (P1) as follows for different CSIT/ESIT cases introduced in the previous section.

First, consider the extreme case, case 1, for obtaining the throughput upper bound. Given the non-causal CSIT and ESIT in this case, problem (P1) is a convex optimization problem since the objective function is concave and the constraints are all affine. By using the Karush-Kuhn-Tucker (KKT) conditions, it is shown in [2, 3] that the optimal power solution exhibits a *staircase water-filling* structure, where the water level is a non-decreasing and staircase function over EH blocks. Particularly, if the additive white Gaussian noise (AWGN) channel model is assumed (i.e., $h(n, m)$'s are constant over all n's and m's), the transmitter should follow a non-decreasing staircase strategy for power allocation as depicted in Fig. 4b.

As for case 2, since both the CSI and ESI are only known causally at the transmitter, we are interested in maximizing the expected throughput over the randomness of the EH and channel variations, instead of its exact value as in case 1. Note that due to the possible statistical correlation of EH rates and channel gains over time, the EH communication system may be considered as a dynamic system, for which some analytic models such as Markov decision process (MDP) and queueing system are applicable. The optimal power allocation could be derived based on the dynamic programming (DP) technique,[2] which balances the trade-off between the exact throughput achieved in the current block vs. the expected throughput over future blocks based on the knowledge of current CSIT and ESIT and their distributions. In general, DP may incur exponentially growing computation complexity as the number of EH/communication blocks increases. Therefore, suboptimal solutions with lower complexity have attracted a great deal of attention [2, 3], where certain throughput performance may be compromised to reduce the complexity.

For cases 3 and 4, to the best of our knowledge, the throughput maximization problems have not been studied in the literature yet. Here, we provide brief discussions on these two cases to motivate future investigations. With causal CSIT and non-causal ESIT in case 3, the transmitter is aimed to maximize the expected throughput over the randomness of channel realizations subject to a set of deterministic EH constraints. In general, such a problem can be optimally solved by the DP technique similar as in case 2, based on which the transmitter needs to decide its power allocations with the updated CSI block by block. By exploiting DP and the deterministic EH constraints, it may be feasible to further obtain insightful and well structured solutions, at least under some specific channel distributions, such as independent and identically distributed (i.i.d.) channel gains, which is an interesting problem worth pursuing.

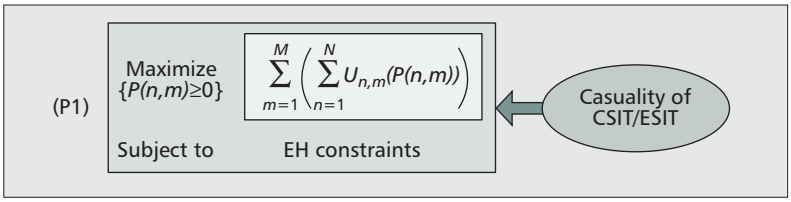

Figure 3. The general utility maximization problem.

In case 4, the throughput optimization is generally a very challenging problem that remains open. The reason is that due to the lack of CSIT, the transmitter cannot adapt its transmit power and hence rate based on the instantaneous CSI; therefore, the receiver may fail to decode the information if the channel is too weak at that block. Thus, the total throughput may not be achievable in general. Despite this difficulty, one tractable case of this problem is when the number of communication blocks is sufficiently large within each EH block (say, $N \to \infty$). In this case, the ergodic capacity is achievable by letting the transmit power remain constant over each EH block and using a sufficiently long capacity-achieving code. With such a transmission scheme, the throughput utility achieved by each EH block can be expressed explicitly (although not in closed form for general fading distributions); hence, problem (P1) can be formulated as an ergodic throughput maximization problem over M EH blocks with non-causal or causal ESIT, which could be solved in a similar way as cases 1 and 2.

Outage Probability Minimization

Outage occurs when there is a failure in decoding the data packet at the receiver, which is mainly due to the fact that the received signal undergoes a "deep" channel fading such that its power is not sufficient to combat the receiver noise at the decoder. In EH-based wireless communications, besides the channel fading, the uncertainty in the amount of harvested energy could be another source for transmission outage due to insufficient transmit power (or even insufficient circuit power; see [7]) at the EH transmitter.

To better address the outage issue, let $Q_{n,m}(P(n, m))$ denote the outage probability function of the (n, m)th block, which depends on both the transmit power $P(n, m)$ and the channel power gain $h(n, m)$. For example, by assuming the Shannon capacity as the achievable rate, and that the required transmission rate at all NM blocks is constant, the outage probability $Q_{n,m}(P(n, m))$ at the (n, m)th block can be expressed as the probability that the achievable rate is less than the required transmission rate [6]. Following this definition, the utility function

[2] Please refer to [2, 3] for more detailed discussions about the DP technique.

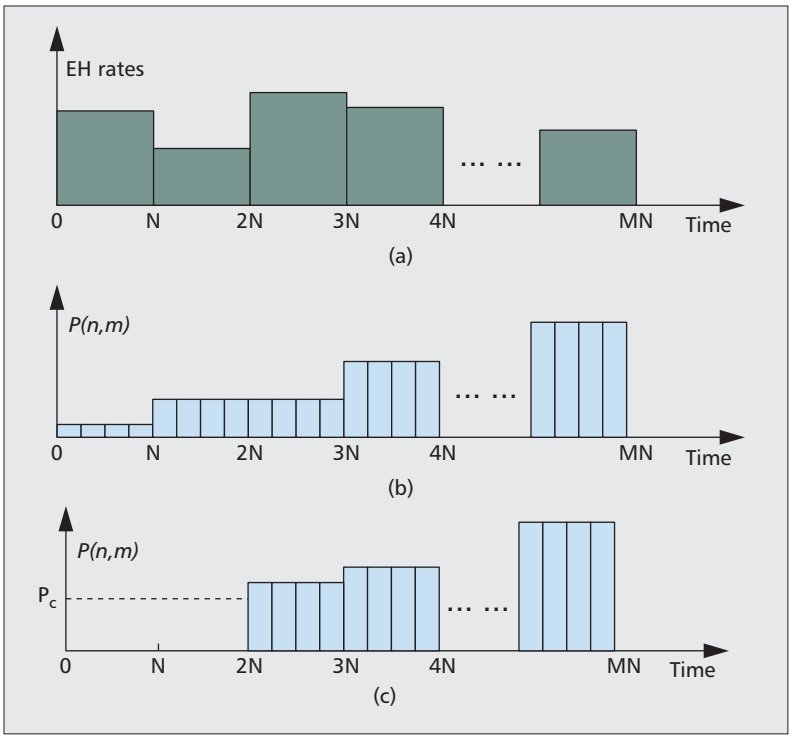

Figure 4. Optimal transmit power allocations in a point-to-point channel subject to EH constraints: a) EH rates at the transmitter; b) the optimal power allocation for throughput maximization in case 1 with an AWGN channel; c) the optimal power allocation for outage probability minimization in case 3 with Weibull fading channel.

$U_{n,m}(P(n, m))$ could be modeled by the non-outage probability for the (n, m)th block, that is, $U_{n,m}(P(n, m)) = 1 - Q_{n,m}(P(n, m))$. It is thus straightforward to see that problem (P1) is equivalent to minimizing the average outage probability over the transmissions.

As the optimal solution structure for the previous throughput maximization crucially relies on the concavity of the rate function, the optimal power allocation for the outage probability minimization problem also critically depends on the properties of the outage probability function $Q_{n,m}(P(n, m))$. Here, $Q_{n,m}(P(n, m))$ is determined based on the availability of CSIT or the probability distribution of the channel gain $h(n, m)$, as is specified for different cases in the following. For simplicity, we assume in this subsection that the channel gains $\{h(n, m)\}$ are i.i.d. over different communication blocks.

First, we consider case 4 without CSIT. The outage probability function $Q_{n,m}(P(n, m))$ is generally non-increasing over transmit power and satisfies one of the following two properties [6]: 1) it is convex over transmit power on $[0, +\infty)$; or 2) there exists a critical point $P_c > 0$ such that it is concave over $[0, P_c]$ and convex over $[P_c, +\infty)$. Such results hold for a large class of fading channels, including Weibull, Rician, Nakagami, and double Rayleigh fading. Under the above properties of $Q_{n,m}(P(n, m))$, we study the optimal power allocation in case 4 with non-causal ESIT. If the outage probability function satisfies Property 1), we can directly solve problem (P1) based on the similar techniques used for throughput maximization [2, 3] as discussed in the previous subsection. However, if the outage probability function satisfies Property 2), the problem becomes non-convex, but is solvable by considering two regions: the concave region $[0, P_c]$ and convex region $[P_c, +\infty)$. It is shown in [6] that the optimal solution follows a "save-then-transmit" structure with non-decreasing power allocation as depicted in Fig. 4c. Specifically, at the beginning of the transmission, the transmitter keeps harvesting energy until the available power becomes larger than P_c, say, from time 0 to $2N$ in Fig. 4c. After that, the transmit power should be allocated non-decreasingly over time similar to staircase water-filling.

For case 4 with causal ESIT, problem (P1) becomes minimizing the expected outage probabilities over the randomness of EH rates, which can be solved by standard DP techniques [6]. It is interesting to note that since EH rates vary per EH block spanning N communication blocks, the DP-based optimal solution should obtain the optimal power allocation for all the N constituting communication blocks at the beginning of the current EH block by balancing the trade-off between minimizing the outage probabilities in the current EH block vs. those in the future EH blocks.

For the other three cases (i.e., cases 1, 2, and 3), to our best knowledge, how to find the optimal power allocation for the outage probability minimization has not yet been investigated. In the following, we provide some intuitions that may be helpful to solve this problem under different cases.

Considering case 1 with non-causal CSIT and ESIT, the outage probability function $Q_{n,m}(P(n, m))$ becomes an indicator function of the transmit power $P(n, m)$ for any given (n, m)th block; that is, $Q_{n,m}(P(n, m)) = 1$ if the achievable rate with power $P(n, m)$ is strictly less than the required transmission rate; otherwise, $Q_{n,m}(P(n, m)) = 0$. With this simplification, the average outage probability minimization is equivalent to minimizing the sum-value of the indicator functions over $1 \le n \le N$ and $1 \le m \le M$. Although such a problem is non-convex in general, it has the following structures. First, it is evident that the allocated transmit power at any communication block should be either zero for an outage event or the minimum required power for a non-outage transmission (with the resulting achievable rate equal to the required transmission rate), since the energy would be wasted otherwise. Another fact is that the objective value of each communication block is either zero or one. With these two observations, we conjecture that the optimal power allocation in this case could be found by first ordering the channel gains from the best to the worst, and then allocating the non-zero transmit power iteratively based on the order of channel gains subject to the EH constraints.

Finally, for cases 2 and 3 with causal CSIT, the outage probability minimization corresponds to minimizing the expected outage probability over either the randomness of EH rates (case 2) or deterministic EH constraints (case 3). In general, the optimal power allocations for these two cases can be derived via standard DP techniques, and in each communication block, the transmit-

ter should decide its current power allocation based on the binary objective value for outage. Interestingly, since the current CSIT is always known, the transmit power at each block (n, m) should be either zero for an outage event or the minimum for a non-outage transmission, similar to case 1 above. This decision is made by balancing the trade-off between avoiding the outage in the current block vs. minimizing the future outage probabilities.

Multiuser EH-Based Communications

So far, we have studied the optimal transmission policy for utility maximization in a single-user EH communication system. However, in practical wireless systems, multiple users with independent or shared energy sources may communicate over the same spectrum. In such systems, besides fading channels and time-varying EH rates of different users, their mutual interference among themselves is a new challenge to be tackled in multiuser utility optimization. In this section, we present promising communication and/or energy cooperation approaches among EH users to maximize the system-level utility. First, we consider the classic three-node relay channel with EH source and relay, and then we discuss other multiuser setups with EH transmitters.

Energy and Communication Cooperation in a Relay Channel

We consider a three-node relay channel with half-duplex and orthogonal transmissions, where the relay node transmits (to the destination node) and receives (from the source node) over two different frequency bands. As shown in Fig. 5, the source and relay nodes transmit with the power drawn from their own EH devices, while the destination is powered by a constant energy source. Similar to the single-user case, in the (n, m)th communication block, denote by $E_S(m)$ and $E_R(m)$ the EH rates, and $P_S(n, m)$ and $P_R(n, m)$ the power allocations at the source and relay nodes, respectively. Then the source and relay nodes are each subject to individual EH constraints similar to those in the single-user scenario. Note that when the EH rates $\{E_S(m)\}$ and $\{E_R(m)\}$ are independent (e.g., one uses a solar panel and the other uses a wind turbine), the source and relay nodes may have very different energy availabilities at any given time. It is thus important to jointly optimize the power allocations at the source and relay nodes based on their available CSI and ESI to maximize system utility.

For simplicity, in the rest of this subsection, we only consider the end-to-end throughput (from source to destination) of the relay channel shown in Fig. 5 as the utility function, and furthermore focus on the case with non-causal CSIT and ESIT at both the source and relay nodes (i.e., case 1 for the single-user setup) with time-invariant (AWGN) channels. We divide our discussions into two scenarios: in the first scenario, the source and relay transmit with their individually harvested energy, while in the second scenario, the source and relay are allowed to share their harvested energy via a new technique called wireless energy transfer (WET) [9, references therein].

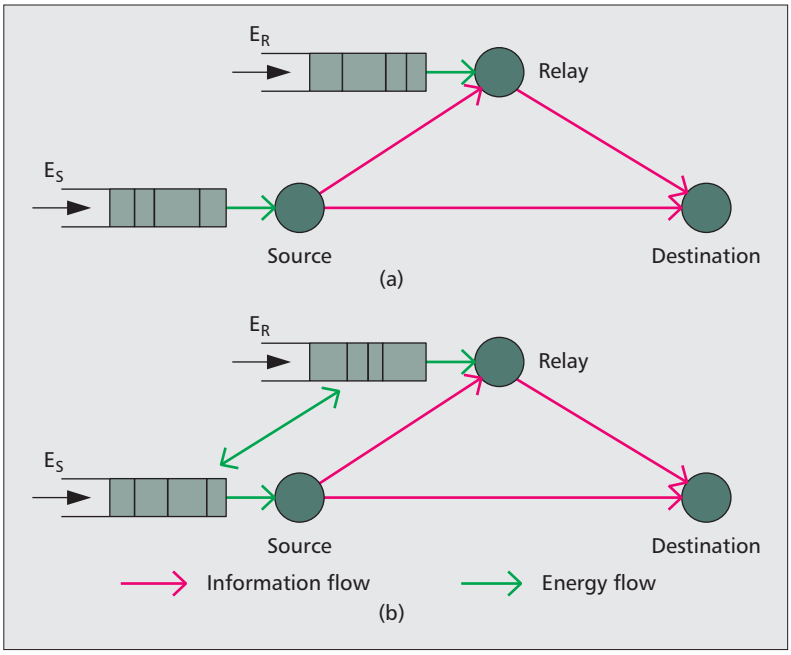

Figure 5. A three-node relay channel with EH source and relay: a) joint power allocation without energy sharing; b) joint power allocation with energy sharing

Joint Power Allocation without Energy Sharing — This case is shown in Fig. 5a, where the source and relay nodes adapt their power allocations in a cooperative manner based on their individual EH rates over transmissions to maximize the end-to-end throughput. This problem is solved in [10] by assuming the decode-and-forward operation of the relay. The optimal joint power allocation is shown to depend on whether a new *decoding constraint* is imposed by the destination or not, as discussed next.

For delay-constrained traffic, a decoding constraint may be applied at the destination such that the source message needs to be decoded immediately after it is transmitted. In other words, the relay needs to decode and forward its received signal from the source without delay. In this case, the achievable rate for each transmission and relaying is limited by both the available energy at the source and that at the relay at a given time. Therefore, the optimal source and relay power allocations are coupled, which could be obtained by a two-dimensional (for both the source and relay) search algorithm [10].

On the other hand, for delay-tolerant traffic, the destination could tolerate arbitrary decoding delays provided that all source messages are decoded at the end of the transmission (i.e., the decoding constraint is much more relaxed). Consequently, the relay is allowed to store its decoded source messages and forward them to the destination at a later time based on its energy availability. Due to the relaxed decoding constraint, it is shown in [10] that the optimal power allocations at the source and relay could be

In practical wireless systems, multiple users with independent or shared energy sources may communicate over the same spectrum. In such systems, besides fading channels and time-varying EH rates of different users, their mutual interference among themselves is a new challenge to be tackled in the multiuser utility optimization.

decoupled and separately optimized. For both the delay-constrained and delay-tolerant cases, the optimal source and relay power allocations should follow a non-decreasing staircase[10], which can be considered as an extension of the one-dimensional (for source only) case for the point-to-point EH system discussed earlier.

Joint Power Allocation with Energy Sharing — Energy sharing via WET is a new approach that allows different EH nodes to exchange energy between each other with certain losses, such that the nodes with excessive energy can share part of their energy to other nodes with insufficient energy in order to balance the energy availabilities among the nodes for further improving the system utility.

As shown in Fig. 5b, the energy sharing capability offers the source and relay nodes more flexibility to manage their energy usage. For example, when the source node is excessive in energy but the relay node has insufficient energy, the source could either increase its own transmit power to boost the achievable rates of both the source-relay and source-destination links, or transfer a portion of its energy to the relay (with certain energy transfer losses) to increase the rate of the relay-destination link and vice versa. In general, the joint power allocation with additional energy sharing between the source and relay could improve the throughput. In [11], by assuming a two-hop relay channel (i.e., ignoring the source-destination link in Fig. 5b) with one-way energy sharing from the source to the relay, the KKT optimality conditions are used to find the optimal control of the transmit power levels at the source and relay, as well as the energy transferred from the source to the relay, to maximize the end-to-end throughput.

Other Setups

In fading channels, the source and relay should use adaptive power allocations to explore the channel variations as well as different energy availabilities among the nodes. In this case, it is anticipated that energy sharing could achieve higher gains than in the time-invariant channel case discussed above. On the other hand, when different CSI and/or ESI availability (e.g., causal vs. non-causal) cases are considered at the source and relay, the optimal power allocation (e.g., for throughput maximization or outage probability minimization) becomes far more complicated than the point-to-point case discussed earlier, with many open questions unsolved yet.

In addition to the simple three-node relay channel, other multiuser systems, such as the multiple access channel [11], two-way channel [11], and coordinated multipoint (CoMP) systems [12] with distributed EH-based transmitters, have been investigated with or without energy sharing. In EH-based CoMP systems [12], a new joint communication and energy cooperation approach is proposed in which distributed EH-powered base stations cooperatively design their transmit signals and the energy amounts exchanged between them to maximize the downlink sum throughput. Instead of using WET, the authors in [12] proposed to implement energy sharing between base stations by leveraging the smart power grid infrastructure with bidirectional power transfer, which in general has much higher efficiency than WET.

Extensions and Future Research Directions

Besides the above studies on EH-based wireless communications, there have been numerous other valuable works on this new and high-potential subject, which are not discussed in this article due to the page limitation. In this section, we briefly present some recent extensions in this area as well as some promising directions for future work.

The Case with an EH Receiver — Most existing studies on EH-based communications have considered EH at the transmitter side only, by assuming that the receivers are powered by a stable energy supply. When the receiver is powered by EH, a similar EH constraint as in the EH transmitter case needs to be applied, but with a key difference in that the energy used at the receiver is for decoding the signal instead of sending it at the transmitter. As a preliminary work in this new direction, [13] showed that the detection and decoding operations dominate the energy cost for EH receivers, and the energy cost is non-decreasing over both the sampling rate and the decoding complexity. Thus, the communication rate should be designed by taking into account the energy availability at both the EH transmitter and EH receiver.

Cross-Layer Design — So far, we have been focusing on the physical (PHY) layer design issues in EH-based communications by assuming the presence of backlogged data for transmission at all times. However, in practical systems, the data arrives at the transmitter with random timing and amounts in general. In such cases, the transmitter needs to deal with the uncertainties in both energy and data arrivals, and it is thus beneficial to jointly schedule the energy usage and data packet transmission based on the channel conditions [1]. As another example, consider a wireless network with multiple EH transmitters sharing the same limited channel resources for communications, for which there is a necessity for the design of energy-aware medium access control (MAC) to optimize the system throughput. A preliminary work on this problem is presented in [14], where a distributed opportunistic scheduling scheme is proposed to jointly design the access control and power management for EH transmitters. In general, a cross-layer design approach should be further investigated to achieve more efficient operation of EH-based communication systems.

Hybrid Energy Sources and Imperfect Energy Storage Devices — Due to the random and intermittent characteristics of practical EH sources, using renewable energy alone may not be sufficient to provide reliable operation of wireless systems with large power demands (e.g., in base stations). To maintain their reliable operation, it is wise to use hybrid energy sources

by efficiently integrating renewable energy with conventional energy (e.g., fuel generators). On the other hand, energy storage devices (ESDs) with imperfect charging-discharging efficiency and a finite capacity may be employed in the system. In general, how to optimally design the energy management policies with hybrid energy sources and/or imperfect ESDs to achieve the maximum utility in EH-based communications still remains largely open, while some initial results have been obtained [15].

RF EH with Dedicated WET — In addition to the conventional EH sources such as wind and solar power as well as ambient RF transmissions, deploying dedicated power transmission nodes in the network for delivering controllable energy over the air to distributed communication devices (e.g., sensors) has drawn growing interest recently. The devices can either harvest RF energy from the signal transmitted by the power transmission nodes, decode the information in it, or even use part of the energy harvested to decode the information and the remaining energy to transmit or relay other information [16]. The RF-signal-enabled WET is a very promising technique for powering low-power wireless communication devices such as those in sensor networks and personal/body area networks, even with its practically limited energy transfer efficiency, which actually could be alleviated by some new techniques such as highly directional massive multiple-input multiple-output (MIMO) [16]. Clearly, such WET powered communication brings a new avenue for future research on EH-based systems.

Conclusions

We have introduced a general utility optimization framework for energy-harvesting-based wireless communication systems subject to a new type of energy usage constraint. Under this framework, the solutions for a variety of design problems of high practical interest have been discussed, including throughput maximization and outage probability minimization under different practical assumptions on the channel and energy state information. A new design paradigm for multiuser EH communication systems with joint energy and communication cooperation is also discussed. Promising research directions for extensions are also highlighted. We hope that this article provides a timely overview of the state-of-the-art results on the fundamental design principles for EH-based wireless communications, and serves as an inspiring key that leads to more fruitful results in future works.

References

[1] D. Gunduz *et al.*, "Designing Intelligent Energy Harvesting Communication Systems," *IEEE Commun. Mag.*, vol. 52, no. 1, Jan. 2014, pp. 210–16.

[2] C. K. Ho and R. Zhang, "Optimal Energy Allocation for Wireless Communications with Energy Harvesting Constraints," *IEEE Trans. Signal Proc.*, vol. 60, no. 9, Sept. 2012, pp. 4808–18.

[3] O. Ozel *et al.*, "Transmission with Energy Harvesting Nodes in Fading Wireless Channels: Optimal Policies," *IEEE JSAC*, vol. 29, no. 8, Sept. 2011, pp. 1732–43.

[4] J. Xu and R. Zhang, "Throughput Optimal Policies for Energy Harvesting Wireless Transmitters with Non-Ideal Circuit Power," *IEEE JSAC*, vol. 32, no. 2, Feb. 2014, pp. 322–32.

[5] O. Ozel, K. Shahzad, and S. Ulukus, "Optimal Energy Allocation for Energy Harvesting Transmitters with Hybrid Energy Storage and Processing Cost," *IEEE Trans. Signal Proc.*, vol. 62, no. 12, June 2014, pp. 3232–45.

[6] C. Huang, R. Zhang, and S. Cui, "Optimal Power Allocation for Outage Probability Minimization in Fading Channels with Energy Harvesting Constraints," *IEEE Trans. Wireless Commun.*, vol. 13, no. 2, Feb. 2014, pp. 1074–87.

[7] S. Luo, R. Zhang, and T. J. Lim, "Optimal Save-Then-Transmit Protocol for Energy Harvesting Wireless Transmitters," *IEEE Trans. Wireless Commun.*, vol. 12, no. 3, Mar. 2013, pp. 1196–1207.

[8] Y. Zhao, B. Chen, and R. Zhang, "Optimal power Allocation for an Energy Harvesting Estimation System," *Proc. IEEE ICASSP*, Vancouver, Canada, May 26–31, 2013, pp. 4549–53.

[9] X. Lu, P. Wang *et al.*, "Wireless Network with RF Energy Harvesting: A Contemporary Survey," to appear, *IEEE Commun. Surveys and Tutorials*.

[10] C. Huang, R. Zhang, and S. Cui, "Throughput Maximization for the Gaussian Relay Channel with Energy Harvesting Constraints," *IEEE JSAC*, vol. 31, no. 8, Aug. 2013, pp. 1469–79.

[11] B. Gurakan *et al.*, "Energy Cooperation in Energy Harvesting Communications," *IEEE Trans. Commun.*, vol. 61, no. 12, Dec. 2013, pp. 4884–98.

[12] J. Xu and R. Zhang, "CoMP Meets Smart Grid: A New Communication and Energy Cooperation Paradigm," to appear, *IEEE Trans. Vehic. Tech.*

[13] H. M.-Doost and R. D. Yates, "Fading Channels in Energy-Harvesting Receivers," *Proc. 48th Annual Conf. Info. Sci. and Sys.*, Princeton, NJ, Mar. 19–21, 2014, pp. 1–6.

[14] H. Li *et al.*, "Distributed Opportunistic Scheduling for Wireless Networks Powered by Renewable Energy Sources," *Proc. IEEE INFOCOM*, Toronto, Canada, Apr. 27–May 2, 2014, pp. 898–906.

[15] Y. K. Chia, S. Sun, and R. Zhang, "Energy Cooperation in Cellular Networks with Renewable Powered Base Stations," *IEEE Trans. Wireless Commun.*, vol. 13, no. 12, Dec. 2014, pp. 6996–7010.

[16] R. Zhang and C. K. Ho, "MIMO Broadcasting for Simultaneous Wireless Information and Power Transfer," *IEEE Trans. Wireless Commun.*, vol. 12, no. 5, May 2013, pp. 1989–2001.

Biographies

Hang Li [S'13] (david_lihang@tamu.edu) received his B.E. and M.S. degrees from Beihang University, Beijing, China, in 2008 and 2011, respectively. He is currently a Ph.D. student at Texas A&M University. His current research interests include energy-harvesting-based communications, energy-aware scheduling, and applications of optimal stopping theory and dynamic programming.

Jie Xu [S'12, M'13] (elexjie@nus.edu.sg) received his B.E. and Ph.D. degrees in electronic engineering and information science from the University of Science and Technology of China, Hefei, in 2007 and 2012, respectively. He is currently a research fellow with the Department of Electrical and Computer Engineering, National University of Singapore. His current research interests include smart grid, multiuser MIMO, energy efficiency, and energy harvesting in wireless communication.

Rui Zhang [S'00, M'07, SM'15] (elezhang@nus.edu.sg) received his Ph.D. degree from Stanford University in 2007. He is now an assistant professor with the ECE Department of the National University of Singapore. His current research interests include multiuser MIMO, cognitive radio, energy-efficient and energy-harvesting-based wireless communications, and wireless information and power transfer. He was the recipient of the 6th IEEE ComSoc Asia-Pacific Best Young Researcher Award in 2011. He is now an Editor for *IEEE Transactions on Wireless Communications* and *IEEE Transactions on Signal Processing*.

Shuguang Cui [S'99, M'05, SM'12, F'14] (cui@ece.tamu.edu) received his Ph.D from Stanford in 2005. He is now an associate professor at TAMU. His current research interests are data oriented large-scale information analysis and system design. He was selected as the Thomson Reuters Highly Cited Researcher and listed in the World's Most Influential Scientific Minds by Sciencewatch in 2014. He was the recipient of the IEEE Signal Processing Society 2012 Best Paper Award. He is a ComSoc Distinguished Lecturer.

The RF signal enabled WET is a very promising technique for powering low-power wireless communication devices such as those in sensor networks and personal/body area networks, even with its practically limited energy transfer efficiency, which actually could be alleviated by some new techniques such as highly directional massive MIMO.

Application of Smart Antenna Technologies in Simultaneous Wireless Information and Power Transfer

Zhiguo Ding, Caijun Zhong, Derrick Wing Kwan Ng, Mugen Peng, Himal A. Suraweera, Robert Schober, and H. Vincent Poor

ABSTRACT

Simultaneous wireless information and power transfer (SWIPT) is a promising solution to increase the lifetime of wireless nodes and hence alleviate the energy bottleneck of energy constrained wireless networks. As an alternative to conventional energy harvesting techniques, SWIPT relies on the use of radio frequency signals, and is expected to bring some fundamental changes to the design of wireless communication networks. This article focuses on the application of advanced smart antenna technologies to SWIPT, including multiple-input multiple-output and relaying techniques. These smart antenna technologies have the potential to significantly improve the energy efficiency and also the spectral efficiency of SWIPT. Different network topologies with single and multiple users are investigated, along with some promising solutions to achieve a favorable trade-off between system performance and complexity. A detailed discussion of future research challenges for the design of SWIPT systems is also provided.

INTRODUCTION

In wireless power transfer, a concept originally conceived by Nikola Tesla in the 1890s, energy is transmitted from a power source to a destination over the wireless medium. The use of wireless power transfer can avoid the costly process of planning and installing power cables in buildings and infrastructure. One of the challenges for implementing wireless power transfer is its low energy transfer efficiency, as only a small fraction of the emitted energy can be harvested at the receiver due to severe path loss and the low efficiency of radio frequency (RF)–direct current (DC) conversion. In addition, early electronic devices, such as first generation mobile phones, were bulky and suffered from high power consumption. For the aforementioned reasons, wireless power transfer had not received much attention until recently, although Tesla had already provided a successful demonstration to light electric lamps wirelessly in 1891.

In recent years, a significant amount of research effort has been dedicated to reviving the old ambition of wireless power transfer, which is motivated by the following two reasons [1, 2]. The first reason is the tremendous success of wireless sensor networks (WSNs) which have been widely applied for intelligent transportation, environmental monitoring, etc. However, WSNs are energy constrained, as each sensor has to be equipped with a battery that has a limited lifetime in most practical cases. It is often costly to replace these batteries and the application of conventional energy harvesting (EH) technologies relying on natural energy sources is problematic due to their intermittent nature. Wireless power transfer can be used as a promising alternative to increase the lifetime of WSNs. The second reason is the now widespread use of low-power devices that can be charged wirelessly. For example, Intel has demonstrated the wireless charging of a temperature and humidity meter as well as a liquid-crystal display using the signals of a TV station 4 km away [4].

This article considers the combination of wireless power transfer and information transmission, a recently developed technique termed simultaneous wireless information and power transfer (SWIPT), in which information carrying signals are also used for energy extraction. Efficient SWIPT requires some fundamental changes in the design of wireless communication networks. For example, the conventional criteria for evaluating the performance of a wireless system are the information transfer rates and the reception reliability. However, if some users in the system perform EH by using RF signals, the trade-off between the achievable information rates and the amount of harvested energy becomes an important figure of merit [1]. In this context, an ideal receiver, which has the capability to perform information decoding (ID) and EH simultaneously, was considered in [1]. In [2] a more practical receiver architecture was proposed, in which the receiver has two circuits to perform ID and EH separately.

This article focuses on the application of smart antenna technologies in SWIPT systems, namely multiple-input multiple-output (MIMO)

Zhiguo Ding is with Princeton University and Lancaster University.

Mugen Peng is with Princeton University and Beijing University of Posts and Telecommunications.

H. Vincent Poor is with Princeton University.

Caijun Zhong is with Zhejiang University.

Derrick Wing Kwan Ng and Robert Schober are with the University of Erlangen-Nurnberg.

Himal A. Suraweera is with the University of Peradeniya.

0163-6804/15/$25.00 © 2015 IEEE

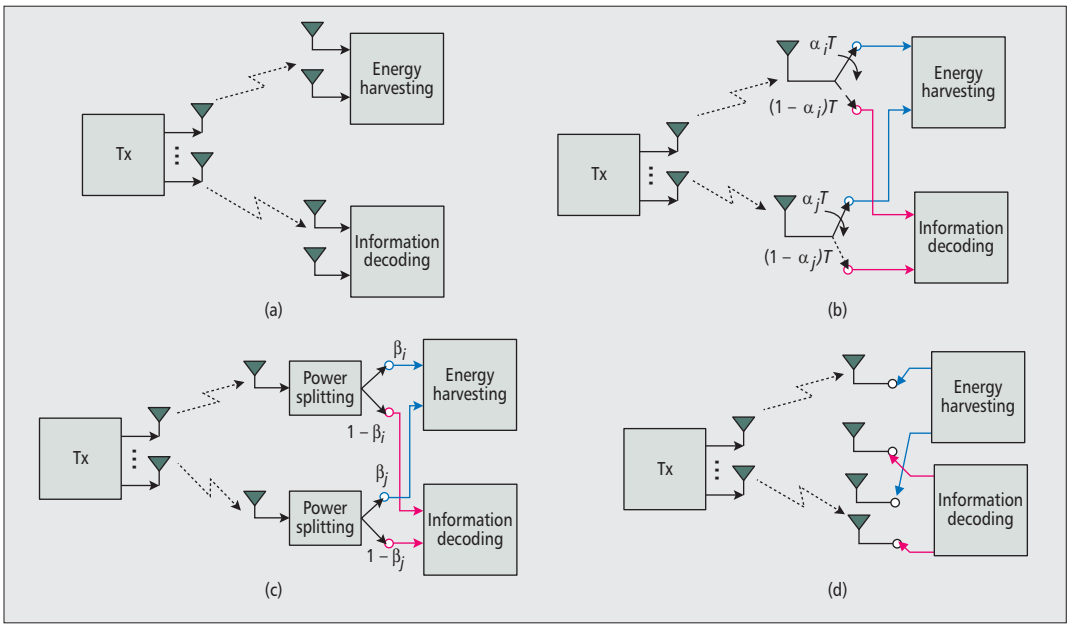

Figure 1. Illustration of the described SWIPT receiver structures. α_i denotes the time switching factor, β_i denotes the power splitting factor, i denotes the antenna index, and T denotes the transmission block duration.

By taking into account the channel statistics and the quality of service requirements regarding the energy transfer, the time switching sequence and the transmit signal can be jointly optimized for different system design objectives.

and relaying. The use of these smart antenna technologies is motivated by the fact that they have the potential to improve the energy efficiency of wireless power transfer significantly. For example, MIMO can be used to increase the lifetime of energy constrained sensor networks, in which a data fusion center is equipped with multiple antennas with which it can focus its RF energy on sensors that need to be charged wirelessly, leading to a more energy efficient solution compared to a single-antenna transmitter. Furthermore, a relay can harvest energy from RF signals from a source and then use the harvested energy to forward information to the destination, which not only facilitates the efficient use of RF signals but also provides motivation for information and energy cooperation among wireless nodes [3]. The application of smart antenna technologies to SWIPT opens up many new exciting possibilities, but also brings some challenges for improving spectral and energy efficiency in wireless systems.

The organization of this article is as follows. Some basic concepts of SWIPT are introduced first. Then the separate and joint application of MIMO and relaying in SWIPT is discussed in detail. Finally, some future research challenges for the design of multi-antenna and multi-node SWIPT systems are provided.

SWIPT: Basic Receiver Structures

In SWIPT systems, ID and EH cannot be performed on the same received signal in general. Furthermore, a receiver with a single antenna typically may not be able to collect enough energy to ensure a reliable power supply. Hence, centralized/distributed antenna array deployments, such as MIMO and relaying, are required to generate sufficient power for reliable device operation. In the following, we provide an overview of MIMO SWIPT receiver structures, namely the power splitting, separated, time-switching, and antenna-switching receivers, as shown in Fig. 1.

Separated Receiver

In a separated receiver architecture, an EH circuit and an ID circuit are implemented as two separate receivers with separated antennas, which are served by a common multiple antenna transmitter [2]. The separated receiver structure can be easily implemented using off-the-shelf components for the two individual receivers. Moreover, the trade-off between the achievable information rate and the harvested energy can be optimized based on the channel state information (CSI) and feedback from the two individual receivers to the transmitter. For instance, the covariance matrix of the transmit signal can be optimized for capacity maximization of the ID receiver subject to a minimum required amount of energy transferred to the EH receiver.

Time Switching Receiver

This receiver consists of an information decoder, an RF energy harvester, and a switch at each antenna [2]. In particular, each receive antenna can switch between the EH circuit and the ID circuit periodically based on a time switching sequence for EH and ID, respectively. By taking into account the channel statistics and the quality of service requirements regarding the energy transfer, the time switching sequence and the transmit signal can be jointly optimized for different system design objectives.

Power Splitting Receiver

Employing a passive power splitting unit, this receiver splits the received power at each antenna into two power streams with a certain power

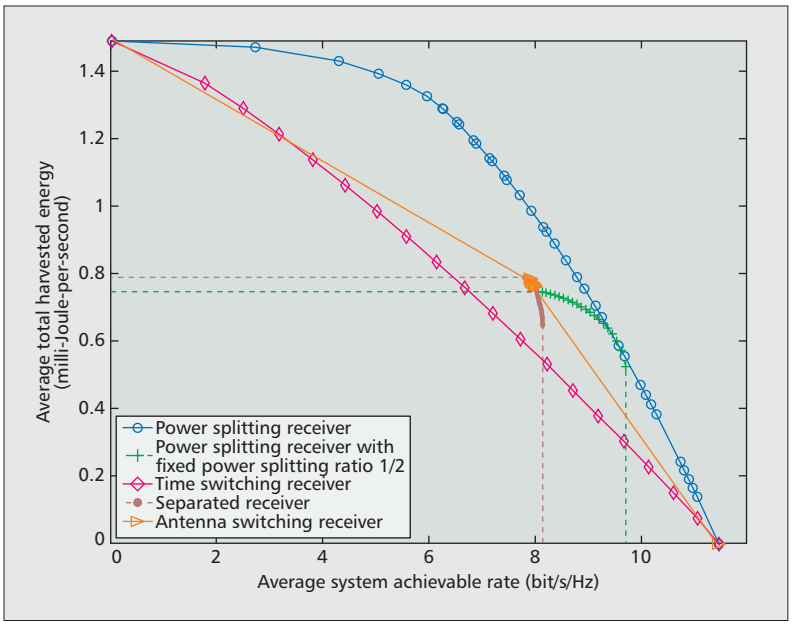

Figure 2. The trade-off region of the average total harvested energy (mJ/s) and the average system achievable rate (bit/s/Hz) for the different receivers. The carrier frequency is 915 MHz and the receiver is located 10 meters from the transmitter. The total transmit power, noise power, transceiver antenna gain, and RF-to-electrical energy conversion loss are set to 10 Watt, –23 dBm, 10 dBi, and 3 dB, respectively. The multipath fading coefficients are modelled as independent and identically distributed Rician random variables with a Rician *K*-factor of 6 dB.

splitting ratio before any active analog/digital signal processing is performed. Then the two streams are sent to an energy harvester and an information decoder, respectively, to facilitate simultaneous EH and ID [2, 5, 6]. The power splitting ratio can be optimized for each receive antenna. In particular, a balance can be struck between the system achievable information rate and the harvested energy by varying the value of the power splitting ratios. Further performance improvement can be achieved by jointly optimizing the signal and the power splitting ratios.

Antenna Switching Receiver

With multiple antennas, low-complexity antenna switching between decoding/rectifying can be used to enable SWIPT [7]. For instance, given N_R antennas, a subset of L antennas can be selected for ID, while the remaining $(N_R - L)$ antennas are used for EH. Unlike the time switching protocol which requires stringent time synchronization and the power splitting protocol where performance may degrade in case of hardware imperfections, the antenna switching protocol is easy to implement, and attractive for practical SWIPT designs. From a theoretical point of view, antenna switching may be interpreted as a special case of power splitting with binary power splitting ratios at each receive antenna.

Figure 2 illustrates the performance trade-offs of the considered SWIPT receiver structures. In particular, we show the average total harvested energy versus the average system achievable information rate in a point-to-point scenario with one transmitter and one receiver. A transmitter equipped with $N_T = 2$ antennas is serving a receiver equipped with $N_R = 2$ receive antennas. Resource allocation is performed to achieve the respective optimal system performance in each case [15]. For a fair comparison, for the separated receiver, the EH receiver, and the ID receiver are equipped with a single antenna, respectively, which results in $N_R = 2$. Besides, we also illustrate the trade-off region for a suboptimal power splitting receiver with a fixed power splitting ratio of 1/2 at each antenna. It can be observed that the optimized power splitting receiver achieves the largest trade-off region among the considered receivers at the expense of incurring the highest hardware complexity and the highest computational burden for resource allocation.

MIMO SWIPT Networks

MIMO can be exploited to bring two distinct benefits to SWIPT networks. On the one hand, due to the broadcast nature of wireless transmission, the use of additional antennas at the receiver can yield more harvested energy. On the other hand, the extra transmit antennas can be exploited for beamforming, which could significantly improve the efficiency of information and energy transfer. The impact of MIMO on point-to-point SWIPT scenarios with one source, one EH receiver, and one ID receiver was studied in [2], where the trade-off between the MIMO information rate and power transfer was characterized. The benefits of MIMO are even more obvious for the multiuser MIMO scenario illustrated in Fig. 3a. Specifically, a source equipped with multiple antennas serves multiple information receivers, where the RF signals intended for the ID receivers can also be used to charge EH receivers wirelessly. Since there are multiple users in the system, co-channel interference (CCI) needs to be taken into account, and various interference mitigation strategies can be incorporated into SWIPT implementations, e.g. block diagonalization precoding as in [8], where information is sent to receivers that are interference free, and energy is transmitted to the remaining receivers. Furthermore, it is beneficial to employ user scheduling, which allows receivers to switch their roles between an EH receiver and an ID receiver based on the channel quality in order to further enlarge the trade-off region between the information rate and the harvested energy.

The multi-source multiuser MIMO scenario illustrated in Fig. 3b is another important SWIPT application, where multiple source–destination pairs share the same spectrum and the associated interference control is challenging. Since in interference channels, interference signals and information bearing signals co-exist, issues such as interference collaboration and coordination bring both new challenges and new opportunities for the realization of SWIPT, which are very different from those in the single source-destination pair scenario. For example, with antenna selection and interference alignment as illustrated in [9], the received signal space can be partitioned into two subspaces, where the subspace containing the desired signals is used for infor-

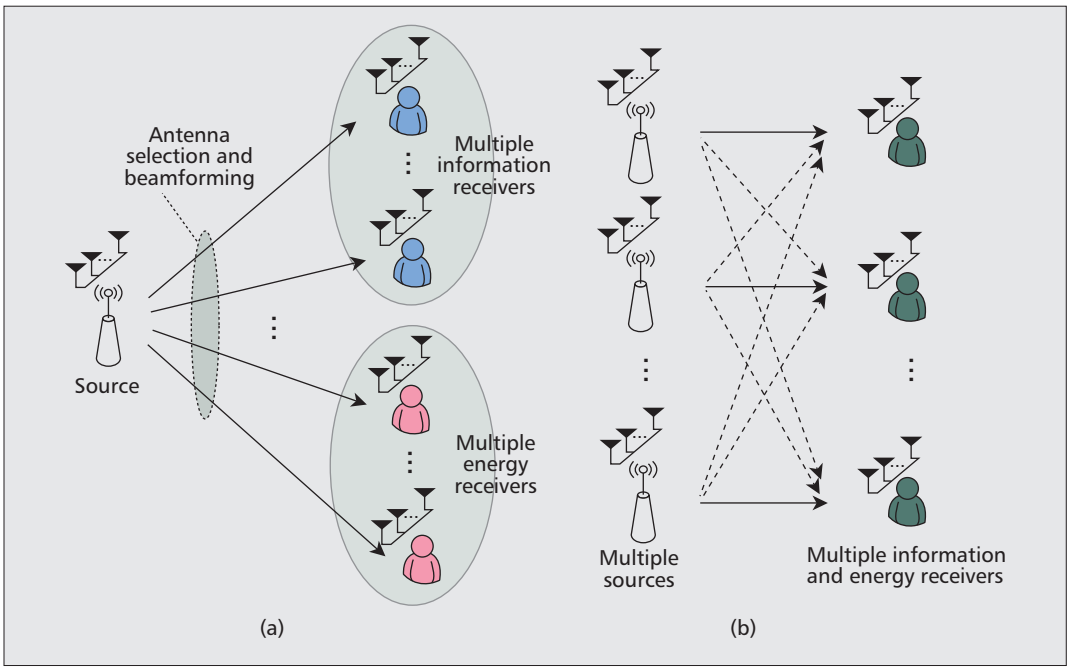

Figure 3. Two typical multiuser MIMO scenarios.

It is beneficial to employ user scheduling, which allows receivers to switch their roles between an EH receiver and an ID receiver based on the channel quality in order to further enlarge the trade-off region between the information rate and the harvested energy.

mation transfer, and the other subspace containing the aligned interference is used for power transfer. This design is a win-win strategy since the information transfer is protected from interference, and the formerly discarded interference can be utilized as an energy source. More importantly, this approach offers a new look at interference control, since the formerly undesired and useless interference can be used to enhance the performance of SWIPT systems. On the other hand, the use of RF EH introduces additional constraints to the design of transmit beamforming. Hence, the solutions well known from conventional wireless networks, such as zero forcing and maximum ratio transmission, need to be suitably modified to be applicable in SWIPT systems, as shown in [10].

Relay Assisted SWIPT Systems

Centralized MIMO as described above may be difficult to implement due to practical constraints, such as the size and cost of mobile devices. This motivates the use of relaying in SWIPT networks. In addition, the use of wireless power transfer will encourage mobile nodes to participate in cooperation, since relay transmissions can be powered by the energy harvested by the relay from the received RF signals and hence the battery lifetime of the relays can be increased. The benefits of using EH relays can be illustrated based on the following example. Consider a relaying network with one source–destination pair and a single decode-and-forward (DF) relay. SWIPT is performed at the relay by using the power splitting receiver structure shown in Fig. 1. In Fig. 4, the performance of the scheme using this EH relay is compared to that of direct transmission, i.e. when the relay is not used. As can be observed from the figure, the use of an EH relay can decrease the outage probability from 7×10^{-1} to 5×10^{-2}, a more than tenfold improvement in reception reliability, compared to direct transmission.

The performance of time sharing and power splitting SWIPT systems employing amplify-and-forward (AF) and DF relays was analyzed in [11], and the impact of power allocation was investigated in [12]. These existing results demonstrate that the behavior of the outage probability in relay assisted SWIPT systems is different from that in conventional systems with self-powered relays. For example, in the absence of a direct source–destination link, the outage probability with an EH relay decays with increasing signal-to-noise ratio (SNR) at a rate of log SNR/SNR, i.e. slower than the rate of 1/SNR in conventional systems. The reason for this performance loss is that the relay transmission power fluctuates with the source-relay channel conditions. This performance loss can be mitigated by exploiting user cooperation. For example, in a network with multiple user pairs and an EH relay, advanced power allocation strategies, such as water filling based and auction based approaches, can be used to ensure that the outage probability decays at the faster rate of 1/SNR [12]. This performance gain is obtained because allowing user pairs to share power can avoid the situation in which some users are lacking transmission power whereas the others have more power than needed.

Relay selection is an important means to exploit multiple relays with low system complexity, and the use of EH also brings fundamental changes to the design of relay selection strategies. In conventional relay networks, it is well known that the source-relay and relay-destination channels are equally important for relay selection, which means that the optimal location of the relay is the middle of the line connecting the source and the destination, i.e. (2.5 m,0) for

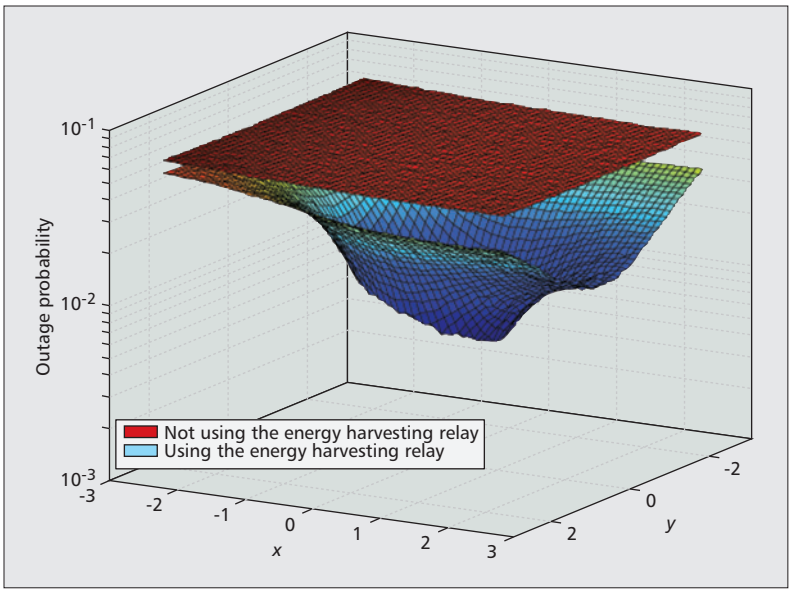

Figure 4. Outage performance of a relaying network with one source, one relay, and one destination. The source is located at (0, 0), the destination is located at (5 m,0), and the x-y plane shows the location of the relay. The carrier frequency is 915 MHz. The total transmit power, noise power, transceiver antenna gain, and RF-to-electrical energy conversion loss are set to 10 Watt, −17 dBm, 0 dBi, and 3 dB, respectively. We assume that the multipath fading coefficients are modelled as independent and identically distributed Rayleigh random variables. The targeted data rate is 0.1 bit/Hz/s. The path loss exponent is 3.

the scenario considered in Fig. 4. Nevertheless, Fig. 4 shows that an EH relay exhibits different behavior than a conventional relay, i.e. moving the relay from the source toward the middle point (2.5 m,0) has a detrimental effect on the outage probability. We note that this observation is also valid for SWIPT systems with AF relays. This phenomenon is due to the fact that in EH networks, the quality of the source-relay channels is crucial since it determines not only the transmission reliability from the source to the relays, but also the harvested energy at the relays. In [13] it was shown that the max-min selection criterion, a strategy optimal for conventional DF relaying networks, can only achieve a small fraction of the full diversity gain in relaying SWIPT systems.

The Combination of MIMO and Cooperative Relaying in SWIPT

MIMO and cooperative relaying represent two distinct ways of exploiting spatial diversity, and both techniques can significantly enhance the system's energy efficiency, which is of paramount importance for SWIPT systems. Hence, the combination of these two smart antenna technologies is a natural choice for SWIPT systems. The benefits of this combination can be illustrated using the following example.

Consider a lecture hall packed with students, in which there are many laptops/smart phones equipped with multiple antennas as well as some low-cost single-antenna sensors deployed for infrastructure monitoring. This hall can be viewed as a heterogeneous network consisting of mobile devices with different capabilities. Inactive devices with MIMO capabilities can be exploited as relays to help the active users in the network, particularly the low-cost sensors. Since the relays have multiple antennas, more advanced receiver architectures, such as antenna switching receivers, can be used. In addition, the use of these MIMO relays opens the possibility of serving multiple source-destination pairs simultaneously. In this context, it is important to note that the use of SWIPT will encourage the inactive MIMO users to serve as relays since helping other users will not reduce the lifetime of the relay batteries. Therefore, the MIMO relays can be exploited as an extra dimension for performance improvement, and can achieve an improved trade-off between the information rate and the harvested energy [7].

As discussed earlier, one unique feature of SWIPT systems is the energy efficient use of CCI, which is viewed as a detrimental factor that limits performance in conventional wireless systems. In particular, CCI can be exploited as a potential source of energy in MIMO relay SWIPT systems. To illustrate this point, let us consider the following example. An AF relay with N antennas is employed to help a single-antenna source that communicates with a single-antenna destination. The relay first harvests energy from the received RF signals with the power splitting architecture, and then uses this energy to forward the source signals. Two separate cases are considered, i.e. without CCI and with CCI. To exploit the benefits of multiple antennas, linear processing of the information stream is performed to facilitate ID. Since the optimal linear processing matrix $\mathbf{W}$ is difficult to characterize analytically, a heuristic rank-1 processing matrix $\mathbf{W}$ is adopted. As such, in the case without CCI, the processing matrix is designed based on the principle of maximum ratio transmission, i.e. $\mathbf{W} = a\mathbf{h}\mathbf{g}^{\dagger}$, where the vectors $\mathbf{h}$ of size $N \times 1$ and $\mathbf{g}$ of size $1 \times N$ are chosen to match the first and second hop channels, respectively, and a is a scaling factor to ensure the relay transmit power constraint. On the other hand, in the presence of CCI, the relay first applies the minimum mean square error criterion to suppress the CCI, and then forwards the transformed signal to the destination using maximum ratio transmission. Figure 5 illustrates the achievable ergodic rate as a function of the average strength of the CCI ρ_I, with the optimized power splitting ratio. We observe that increasing the number of relay antennas significantly improves the achievable rate. For instance, increasing the number of antennas from three to six nearly triples the rate. Moreover, we see that when the CCI is weak ($\rho_I \leq -10$dB), the rate difference is negligible compared to the case without CCI. However, when the CCI is strong, a substantial rate improvement is realized. In fact, the stronger the CCI, the higher the rate gain. For example, in some applications the relays will operate at the cell boundaries and the benefit of exploiting CCI will be significant in such situations.

Research Challenges

In the following, we discuss some research challenges for future MIMO and relay assisted SWIPT.

Energy Efficient MIMO SWIPT: Because of severe path loss attenuation, the energy efficiency of MIMO SWIPT systems may not be satisfactory for long distance power transfer unless advanced green technologies, such as EH technologies relying on natural energy sources, and MIMO resource allocation are combined. We now discuss two possible approaches to address this problem.

- EH transmitter: In this case the transmitter can harvest energy from natural renewable energy sources such as solar, wind, and geothermal heat. Then the energy harvested at the transmitter can be transferred to the desired receiver over the wireless channel, thereby reducing substantially the operating costs of the service providers and improving the energy efficiency of the system, since renewable energy sources can be exploited virtually for free. However, the time varying availability of the energy generated from renewable energy sources may introduce energy outages in SWIPT systems, and efficient new techniques have to be developed to overcome them.
- MIMO energy efficiency optimization: Energy efficient MIMO resource allocation can be formulated as an optimization problem in which the degrees of freedom in the system such as space, power, frequency, and time are optimized for maximization of the energy efficiency. By taking into account the circuit power consumption of all nodes, the finite energy storage at the receivers, the excess spatial degrees of freedom in MIMO systems, and the utilization of the recycled transmit power and the interference power, the energy efficiency optimization reveals the operating regimes for energy efficient SWIPT systems. Yet, the non-convexity of the energy efficiency objective function [6] is an obstacle in designing algorithms for achieving the optimal system performance and low-complexity but efficient algorithms are yet to be developed.

Energy Efficient SWIPT Relaying: The concepts of SWIPT and relaying are synergistic since the use of SWIPT can stimulate node cooperation and relaying is helpful to improve the energy efficiency of SWIPT. In the following, several research challenges for relay assisted SWIPT are discussed:

- Practical relaying systems suffer from spectral efficiency reduction due to half-duplex operation. One possible approach to overcome this limitation is to use the idea of successive relaying, where two relays listen and transmit in succession. When implemented in a SWIPT system, the inter-relay interference, which is usually regarded as detrimental, can now be exploited as a source of energy. Another promising solution is to adopt full-duplex transmission. In the ideal case, full-duplex relaying can double the spectral efficiency, but the loopback interference corrupts the information signal in practice. Advanced MIMO solutions can be designed to exploit such loopback interference as an additional source of energy.

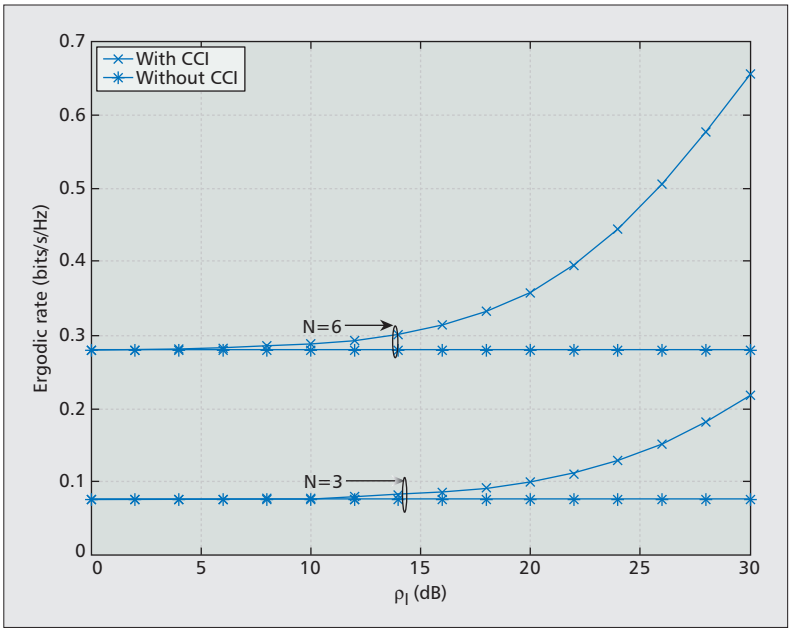

Figure 5. Achievable ergodic rate of a SWIPT relay system with a single-antenna source, a single-antenna destination, and a relay with N antennas. The distances from source to relay, relay and destination, and interferer to relay are set to 2 m, 3 m, and 5 m, respectively. The path loss exponent is 3. The total transmit power, noise power, transceiver antenna gain, and RF-to-electrical energy conversion efficiency are set to 10 Watt, 3 dBm, 0 dBi, and 80%, respectively.

- Relay assisted SWIPT is not limited to the case of EH relays, and can be extended to scenarios in which RF EH is performed at the source and/or the destination based on the signals sent by the relay. For example, in WSNs, two sensors may communicate with each other with the help of a self-powered data fusion center. For this type of SWIPT relaying, the relaying protocol needs to be carefully redesigned, since an extra phase for transmitting energy to the source and the destination is needed.
- Most existing works on SWIPT relaying have assumed that all the energy harvested at the relays can be used as relay transmission power. In practice, this assumption is difficult to realize due to non-negligible circuit power consumption, power amplifier inefficiency, energy storage losses, and the energy consumed for relay network coordination, which need to be considered when new SWIPT relaying protocols are designed. In addition, the superior performance of MIMO/relay SWIPT is often due to the key assumption that perfect CSI knowledge is available at the transceivers; however, a large amount of signalling overhead will be consumed to realize such CSI assumptions. Therefore, for fair performance evaluation, future works should take into account the extra energy cost associated with CSI acquisition [14].

Communication Security Management: Energy transfer from the transmitter to the receivers can be facilitated by increasing the transmit power of the information carrying signal. However, a higher transmit power leads to a larger susceptibility for information leakage due to the broadcast nature of wireless channels. Therefore, communication security is a critical issue in systems with SWIPT.

The application of smart antenna technologies, such as MIMO and relaying, in SWIPT systems has been investigated for different network topologies. In addition, future research challenges for the design of energy efficient MIMO and relay assisted SWIPT systems have been outlined.

• Energy signal: Transmitting an energy signal along with the information signal can be exploited for expediting EH at the receivers. In general, the energy signal can utilize arbitrary waveforms such as a deterministic constant tone signal. If the energy signal is a Gaussian pseudo-random sequence, it can also be used to provide secure communication since it serves as interference to potential eavesdroppers [5]. On the other hand, if the sequence is known to all legitimate receivers, the energy signal can be cancellated at the legitimate receivers before ID. However, to make such cancellation possible, a secure mechanism is needed to share the seed information for generating the energy signal sequence, to which MIMO precoding/beamforming can be applied.

• Jamming is an important means to prevent eavesdroppers from intercepting confidential messages; however, performing jamming also drains the battery of mobile devices. The use of SWIPT can encourage nodes in a network to act as jammers, since they can be wirelessly charged by the RF signals sent by the legitimate users. However, the efficiency of this harvest-and-jam strategy depends on the network topology, where a harvest-and-jam node needs to be located close to legitimate transmitters to harvest a sufficient amount of energy. Advanced multiple-antenna technologies are needed to overcome this problem.

Conclusions

In this article the basic concepts of SWIPT and corresponding receiver architectures have been discussed along with some performance trade-offs in SWIPT systems. In particular, the application of smart antenna technologies, such as MIMO and relaying, in SWIPT systems has been investigated for different network topologies. In addition, future research challenges for the design of energy efficient MIMO and relay assisted SWIPT systems have been outlined.

Acknowledgment

The work of Z. Ding was supported by the UK EPSRC under Grant EP/I037423/1. The work of C. Zhong was supported in part by the NSFC (61201229), the Zhejiang Science and Technology Department Public Project (2014C31051), and the Fundamental Research Funds for Central Universities (2014QNA5019). This work of M. Peng was supported in part by the NSFC (61222103), the National High Technology Research and Development Program of China (2014AA01A701), the State Major Science and Technology Special Projects (2013ZX03001001), and the Beijing Natural Science Foundation (Grant No. 4131003). R. Schober's work was supported by the Alexander von Humboldt (AvH) Professorship Program. The work of H. V. Poor was supported in part by the National Science Foundation under Grants CMMI-1435778, DMS-1118605 and ECCS-1343210.

References

[1] L. R. Varshney, "Transporting Information and Energy Simultaneously," *Proc. IEEE Int'l. Symp. Inf. Theory (ISIT)*, Toronto, Canada, July 2008, pp. 1612–16.
[2] R. Zhang and C. K. Ho, "MIMO Broadcasting for Simultaneous Wireless Information and Power Transfer," *IEEE Trans. Wireless Commun.*, vol. 12, May 2013, pp. 1989–2001.
[3] G. Zheng *et al.*, "Information and Energy Cooperation in Cognitive Radio Networks," *IEEE Trans. Wireless Commun.*, vol. 62, no. 9, May, 2014, pp. 2290–303.
[4] A. Sample and J. R. Smith, "Experimental Results with Two Wireless Power Transfer Systems," *Proc. IEEE Radio and Wireless Symp. (RWS)*, San Diego, CA, Jan. 2009, pp. 16–18.
[5] D. W. K. Ng, E. S. Lo, and R. Schober, "Robust Beamforming for Secure Communication in Systems with Wireless Information and Power Transfer," *IEEE Trans. Wireless Commun.*, vol. 13, Aug. 2014, pp. 4599–615.
[6] D. W. K. Ng, E. S. Lo, and R. Schober, "Wireless Information and Power Transfer: Energy Efficiency Optimization in OFDMA Systems," *IEEE Trans. Wireless Commun.*, vol. 12, Dec. 2013, pp. 6352–70.
[7] I. Krikidis *et al.*, "A Low Complexity Antenna Switching for Joint Wireless Information and Energy Transfer in MIMO Relay Channels," *IEEE Trans. Commun.*, vol. 62, no. 5, May 2014, pp. 1577–87.
[8] W. Wang *et al.*, "Power Allocation in Multiuser MIMO Systems for Simultaneous Wireless Information and Power Transfer," *Proc. IEEE Vehic. Tech. Conf. (VTC)*, Las Vegas, NV, Sept. 2013, pp. 1–5.
[9] B. Koo and D. Park, "Interference Alignment and Wireless Energy Transfer via Antenna Selection," *IEEE Commun. Lett.*, vol. 18, no. 4, Apr. 2014, pp. 548–51.
[10] S. Timotheou, I. Krikidis, G. Zheng, and B. Ottersten, "Beamforming for MISO Interference Channels with QoS and RF Energy Transfer," *IEEE Trans. Wireless Commun.*, vol. 13, no. 5, May 2014, pp. 2646–58.
[11] A. A. Nasir *et al.*, "Relaying Protocols for Wireless Energy Harvesting and Information Processing," *IEEE Trans. Wireless Commun.*, vol. 12, no. 7, July 2013, pp. 3622–36.
[12] Z. Ding *et al.*, "Power Allocation Strategies in Energy Harvesting Wireless Cooperative Networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 2, Feb. 2014, pp. 846–60.
[13] Z. Ding and H. V. Poor, "'User Scheduling in Wireless Information and Power Transfer Networks," *Proc. IEEE Int. Conf. on Commun. Systems (ICCS)*, Macau, China, Nov. 2014, pp. 1–5 (a journal version available at http://arxiv.org/abs/1403.0354).
[14] S. Yatawatta, A. P. Petropulu, and C. J. Graff, "Energy Efficient Channel Estimation in MIMO systems," *Proc. IEEE Int'l. Conf. Acoustics, Speech, and Signal Processing (ICASSP)*, Las Vegas, NV, Mar. 2005, pp. iv/317–iv/320.
[15] X. Lu *et al.*, "Resource Allocation in Wireless Networks with RF Energy Harvesting and Transfer," *IEEE Network*, to appear in 2014.

Biographies

Zhiguo Ding (z.ding@lancaster.ac.uk) received his Ph.D. degree from Imperial College London in 2005, and is currently a chair professor at Lancaster University, UK. His research interests include 5G communications, MIMO and relaying networks, and energy harvesting. He serves as an editor for several journals including *IEEE Transactions on Communications*, *IEEE Communication Letters*, *IEEE Wireless Communication Letters*, and *Wireless Communications and Mobile Computing*.

Caijun Zhong (caijunzhong@zju.edu.cn) received his Ph.D. degree in electrical and electronic engineering from University College London in 2010. He worked as a research fellow at Queen's University Belfast from September 2009 to September 2011. He is currently an associate professor at the Institute of Information and Communication Engineering, Zhejiang University. His research interests include MIMO systems, cognitive radio, full-duplex relaying, and energy harvesting.

Derrick Wing Kwan Ng (wingn@ece.ubc.ca) received the bachelor degree with first class honors and the Master of Philosophy degree in electronic engineering from the Hong Kong University of Science and Technology (HKUST) in 2006 and 2008, respectively. He received his Ph.D. degree from the University of British Columbia (UBC) in 2012. He is now working as a postdoctoral fellow at the Institute for Digital Communications, Friedrich-Alexander-University Erlangen-Nürnberg (FAU), Germany.

MUGEN PENG (pmg@bupt.edu.cn) received the Ph.D. degree from Beijing University of Posts & Telecommunications (BUPT), China, in 2005, where he is a full professor. During 2014 he was also an academic visiting fellow at Princeton University, Princeton, NJ, USA. His main research areas include wireless communication, radio signal processing, and convex optimizations. He received the 2014 IEEE ComSoc AP Outstanding Young Researcher Award, and the Best Paper Award at GameNets 2014, CIT 2014, ICCTA 2011, ICBNMT 2010, and IET CCWMC 2009.

HIMAL A. SURAWEERA (himal@ee.pdn.ac.lk) received his Ph.D. degree from Monash University, Australia, in 2007. Currently he is a senior lecturer at the University of Peradeniya, Sri Lanka. His research interests currently include cooperative relay networks, energy-efficient communication systems, full-duplex radios, and massive MIMO systems. He serves as an editor for several journals, including *IEEE Transactions on Wireless Communications*.

ROBERT SCHOBER (schober@LNT.de) received the Diplom and the Ph.D. degrees in electrical engineering from the University of Erlangen-Nuermberg in 1997 and 2000, respectively. Since May 2002 he has been with the University of British Columbia (UBC), Vancouver, Canada, where he is now a full professor. Since January 2012 he has been an Alexander von Humboldt Professor and the Chair for Digital Communication at the Friedrich-Alexander-University Erlangen-Nürnberg (FAU), Erlangen, Germany. He is a Fellow of the IEEE and the editor-in-chief of *IEEE Transactions on Communications*.

H. VINCENT POOR (poor@princeton.edu) is with Princeton University, where his interests are in wireless networking and related fields. He is a member of the National Academy of Engineering and the National Academy of Sciences, and a foreign member of the Royal Society. He received the IEEE ComSoc Marconi and Armstrong Awards in 2007 and 2009, respectively, and more recently the 2014 URSI Booker Gold Medal and honorary doctorates from several universities.

RF-Powered Cognitive Radio Networks: Technical Challenges and Limitations

Lina Mohjazi, Mehrdad Dianati, George K. Karagiannidis, Sami Muhaidat, and Mahmoud Al-Qutayri

ABSTRACT

The increasing demand for spectral- and energy-efficient communication networks has spurred great interest in energy harvesting cognitive radio networks. Such a revolutionary technology represents a paradigm shift in the development of wireless networks, as it can simultaneously enable the efficient use of the available spectrum and the exploitation of RF energy in order to reduce reliance on traditional energy sources. This is mainly triggered by the recent advancements in microelectronics that puts forward RF energy harvesting as a plausible technique in the near future. On the other hand, it has been suggested that the operation of a network relying on harvested energy needs to be redesigned to allow the network to reliably function in the long term. To this end, the aim of this survey article is to provide a comprehensive overview of recent development and the challenges regarding the operation of CRNs powered by RF energy. In addition, the potential open issues that might be considered for future research are also discussed in this article.

INTRODUCTION

Harvesting energy from ambient sources and converting it to electrical energy used to power devices is of increasing importance in designing green communication networks. While this approach enables more environmentally friendly energy supplies, it helps realize the vision of long-lived, self-maintained, and autonomous communication systems. In addition to well-known alternative energy sources, such as solar, wind, geothermal, and mechanical, ambient RF signals present another promising source that can be exploited in the future. A clear advantage of this technique, in comparison with other alternative energy sources, is that ambient RF sources can be consistently available regardless of time and location in urban areas. Moreover, RF energy harvesting (EH) systems can be built cheaply in small dimensions, which could be a significant advantage in the manufacturing of small and low-cost communication devices such as sensor nodes.

RF signals can be used by a node to extract information or harvest energy. Scavenging energy from RF signals is broadly known as wireless EH or wireless power transfer (WPT), as it refers to the transmission of electrical energy from a power source to one or more electrical loads without any wires. Investigating techniques for RF-powered mobile networks has received significant attention during the past few years in a number of applications such as wireless sensor networks (WSNs) and cooperative communication systems. Most recently, wireless EH has been flagged as a potential source of energy for cognitive radio networks (CRNs) [1]. The operation of CRNs requires periodical sensing and continuous decision making on the availability of spectrum for secondary users (SUs) in the system. This process, along with subsequent signal processing and data transmissions, result in high energy consumption by CRN nodes. Thus, it is desirable to find techniques that can help prolong the lifetime of CRNs. To this end, deploying RF EH becomes a notable candidate for CRNs, aimed at improving both energy and spectral efficiency of communication networks. In this approach, in addition to the identification of spectrum holes for information transfer, an SU may exploit the ambient RF power to supply an auxiliary source of energy for the CRN nodes. Furthermore, when EH is regarded as a significant source of energy for the operation of CRN nodes, it is crucial that the operation of the system is optimized in order to improve the survival of the system, taking into account the characteristics of the considered energy source. This necessitates the need for redesigning the existing techniques in CRNs in order to simultaneously optimize the EH function and better utilize the underlying RF energy source [2].

This article aims to review the state of the art of RF-powered CRNs and to survey the enabling techniques that have been proposed in recent years. The remainder of the article is organized as follows. The classification of the existing RF EH techniques are discussed. The high-level architecture of an RF-powered CRN is presented. This is followed by surveying the technical aspects that affect the performance of RF-powered CRNs. Furthermore, some of the well-known and promising existing technical solutions in the literature are surveyed. Since this research field is still in its early stages, some of the open technical challenges for possible future investigation are addressed. Finally, concluding remarks are given.

Lina Mohjazi and Mehrdad Dianati are with the University of Surrey.

Mahmoud Al-Qutayri is with Khalifa University

George K. Karagiannidis and Sami Muhaidat are with Khalifa University and Aristotle University of Thessaloniki.

0163-6804/15/$25.00 © 2015 IEEE

Classification of RF Energy Harvesting

Several methods of WPT have been introduced in the recent literature, including near-field short-range inductive or capacitive coupling, non-radiative mid-range resonance, and far-field long-range RF energy transmission. Nonetheless, the latest class of RF energy transmission in the microwave frequency band is the most recently focused technique. In such frequencies, the wavelength of the RF signal is very small, and the WPT system does not require calibration and alignment of the coils and resonators at the transmitter and receiver sides [3]. This renders the technique a suitable solution to power a large number of small wireless mobile devices over a wide geographical area.

Due to the specific communication requirements of cognitive radio nodes and the nature of RF EH, communication techniques and protocols used in traditional CRNs may not be directly used in RF-powered CRNs [4]. In particular, it is important to first identify the sources of RF energy and their different characteristics in order to understand the technical challenges faced by RF-powered CRNs. The mechanisms by which RF energy is obtained can mainly be classified into two categories: non-intended RF EH and intended RF EH. In the following subsections, we provide an overview of these two categories.

Non-Intended RF Energy Harvesting

Non-intended RF signals are ambient RF sources not originally intended for energy transfer. This includes signals radiated due to wireless telecommunication services, such as cellular systems, mobile devices, and wireless local area networks (WLANs), or from public broadcasting systems, such as TV and radio. These ambient signals, if not received by their intended receivers, are dissipated as heat, resulting in a waste of energy. Instead, they could be used as a sustainable and low-cost source from which to harvest energy [5]. A device that harvests energy from ambient RF sources can have separate antennas or antenna arrays for an RF transceiver and an RF energy harvester. Harvesting energy by this means is subject to long-term and short-term fluctuations due to radio tower service schedules, node mobility and activity patterns, and fading. Therefore, cognitive radio terminals should employ new schemes that consider the trade-off among network throughput, energy efficiency, and RF energy supply, given the dynamic availability of the RF energy.

Intended RF Energy Harvesting

This method can be divided into two types. In the first, the receiver obtains wireless power transferred from a dedicated source that only delivers power without transmitting information to it, as in directive power beamforming.[1] The second method uses the same emitted RF signal to transport energy and information simultaneously, known as simultaneous wireless information and power transfer (SWIPT) [6].

A number of receiver designs have been proposed for SWIPT. The two most adopted designs in literature are the integrated and co-located receiver designs. The co-located receiver design can be based on either time switching or power splitting [7]. A power splitting block divides the received signal into two portions, one for EH and the other for information decoding, while time switching allocates dedicated time slots to EH and the rest to data processing. By employing this approach, controllable and efficient on-demand wireless information and energy can be simultaneously provided. This permits a low-cost alternative for sustainable wireless systems without further hardware modification on the transmitter side.

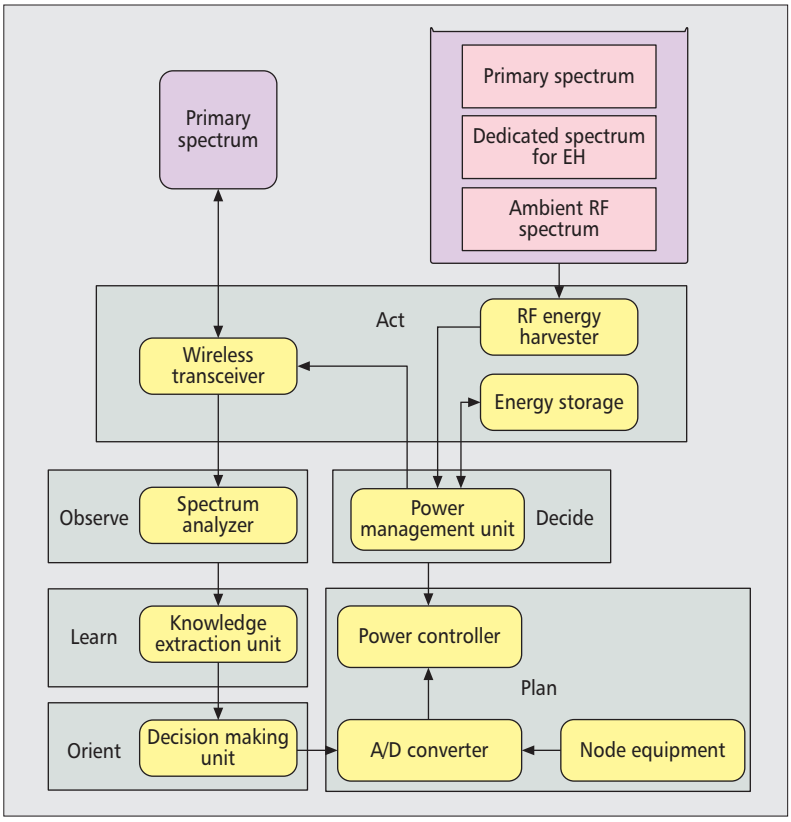

Figure 1. RF-powered CRN node operation cycle block diagram [2].

Overview of RF-Powered CRNs

There has been recent interest in exploitation of RF-based EH for CRNs. As it is the main focus of this article, in the following, we elaborate on this application in further detail. A general block diagram of the functions performed by a cognitive radio node with RF EH capability is illustrated in Fig. 1 [2]. The role of each component is described related to the major functions of a cognitive cycle, that is, observing, learning, orienting, planning, deciding, and acting, as follows:

- Wireless transceiver: a software-defined radio for data transmission and reception
- Energy storage: could be a battery or capacitor to store the harvested energy
- Power management unit: decides whether the harvested energy should be stored in energy storage or forwarded to other components

[1] The Powercast transmitter is one example that is already commercialized. Interested readers may learn more at http://www.powercastco.com/

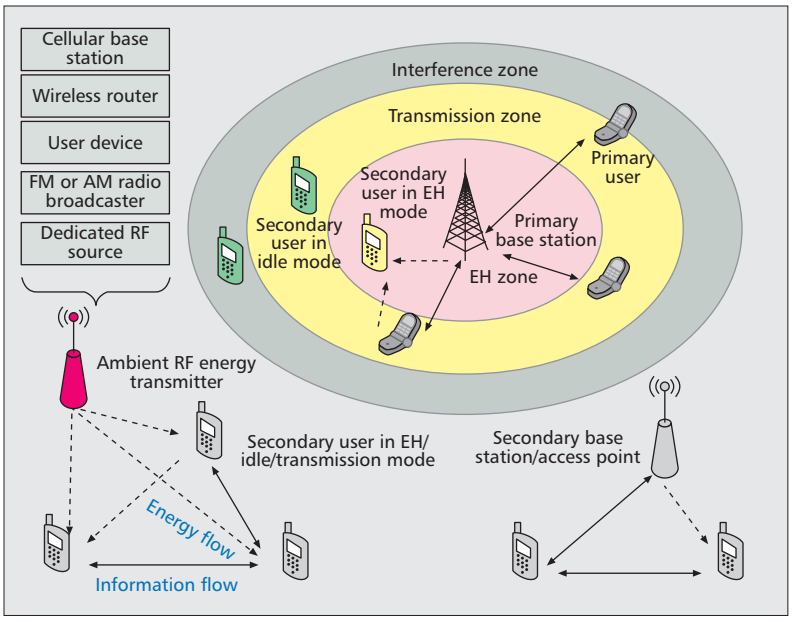

Figure 2. A general architecture of an RF-powered CRN.

- RF energy harvester: replenishes RF signals and converts them to electricity
- Spectrum analyzer: provides instantaneous analysis of the activity of spectrum usage
- Knowledge extraction unit: maintains a record about the spectrum access environment
- Decision making unit: decides on spectrum access
- Node equipment: implements device applications
- A/D converter: digitizes the analog signal produced by the node equipment
- Power controller: processes the output of the A/D converter for network applications

A general architecture of CRN powered by either ambient RF signals, energy transmitted from an intended RF source or via SWIPT, is shown in Fig. 2. When SUs harvest RF energy from the primary network, the primary base station can be associated with three zones [1] that define the SUs activity. Secondary users that are not fully charged and are located in the EH zone can harvest energy from the RF signals received from the primary base station or nearby primary users (PUs). SUs that are located inside the interference zone cannot transmit unless the spectrum is unoccupied by the PUs. Furthermore, it can be seen from Fig. 2 that the secondary network can also harvest ambient RF energy. RF-powered CRNs can adopt either an infrastructure-based or an infrastructureless communication architecture.

Technical Challenges of RF-powered CRNs

As discussed in the previous sections, CRN nodes may be powered by two different categories of RF energy sources. In this section, we provide an overview of the technical challenges that arise in both scenarios.

In the scenario where a cognitive radio node harvests energy from unintended RF energy, the energy available randomly varies over time in a random process known as the energy profile, which can be described by certain mathematical models. This inherent randomness of the energy source is a major factor that affects the performance of an EH node. On the other hand, an SU can also receive RF energy from either ambient transmissions of the primary network or a particular PU with activity known to the SU. In this case, the cognitive operation of the SU is powered solely by the RF energy from the PU. Therefore, both the occupied and idle spectra are essential for the operation of a SU. In both the aforementioned cases, the performance of a CRN is restricted by the *collision constraint*, which requires that the probability of colliding with the primary transmission is always kept below a predefined threshold. When an SU operates in a time-slotted manner, its frame structure is divided into several time slots to perform different cognitive radio tasks. The performance of each of them is directly affected by the available energy at the time it is to be executed. The total consumed energy should be equal to or less than the total harvested energy; this is called the *EH constraint* [1]. Putting those two constraints together implies the fundamental limitations on the throughput of an EH CRN.

Several studies focused on exploring the impact of EH on CRNs. A seminal work in this area is [1], which proposes a novel framework, enabling SUs to opportunistically harvest ambient RF energy as well as reuse the spectrum of PUs. Also, the transmission probability of SUs and the resulting system throughput of the CRN were derived when a stochastic-geometry model was considered. The results presented in [1] revealed key insights about the optimal network design. Moreover, the authors in [8] derived the upper bound on the achievable throughput as a function of the energy arrival rate, the temporal correlation of the primary traffic, and the detection threshold for a spectrum sensor.

We aim in this section to discuss techniques that should be revisited in order to optimize system configurations to accommodate the newly introduced requirements of RF-powered CRNs. In addition, we review the relevant solutions proposed in literature.

Mode Selection

An SU harvesting ambient RF energy usually operates in either an active or a sleep mode. In the former, it performs spectrum sensing and then data transmission if the detector decides that the PU is absent. In the latter, the SU remains silent and only harvests energy. On the other hard, when an SU needs to exploit the existence of the PU to harvest RF energy, it selects either the spectrum access mode (including sensing the idle spectrum and then transmission, or sensing the occupied spectrum and then harvesting) or the harvesting mode, which only incorporates the process of EH. There is a tradeoff for each node between utilization of the spectrum and exploitation of RF energy. The more time a node spends on sensing spectrum holes and using the opportunities for transmission, the higher the energy consumption rate

and the fewer opportunities for EH. Therefore, in order to simultaneously enhance network performance and energy utilization, an optimal mode selection policy may be investigated. Motivated by this trade-off, the work in [9] considers a cognitive radio sensor network where SUs perform either RF EH or opportunistic spectrum access at one time. Under this assumption, the authors developed an optimal mode selection policy in the framework of a partially observable Markov decision process (POMDP). Built on the concept of hybrid underlay-overlay spectrum access, the work carried out in [10] proposed a mode selection strategy where the SU can be in one of three states: transmission mode (either underlay or overlay), sleep mode, or EH mode. The objective is to find a balance between the system throughput and the harvested energy for future use.

Since the transmitted power attenuates according to the reciprocal of the distance, to ensure a certain EH efficiency, the decision to select the harvesting mode has to consider both the availability of the PU and its distance from the SU, as studied in [1].

Sensing Duration

The main question here is to determine how the duration of spectrum access is constrained by the sensing process, which is crucial to system performance. Longer sensing duration results in higher probability of true detection of the spectrum and thus lower interference caused to PUs. However, it simultaneously decreases the chances of an SU to access the spectrum. The total energy consumption behavior varies from one frame to the other according to the variation in the sensing duration. Not only does this behavior depend on the sensing duration; it is also affected by the sensing-to-transmission power ratio. Both the opportunities of accessing the idle spectrum and the energy consumed by sensing increase as the sensing duration increases. This also elevates the energy consumed by more frequent data transmissions. Nonetheless, if the sensing duration is too long, the time left for transmission becomes short; and accordingly, the total amount of energy consumption (sensing plus transmission energies) is reduced, due to the decreased opportunity for data transmission. The aforementioned conflicting factors collectively imply coming up with an optimal sensing duration that would take into account the available energy and the effect on the performance of both CR and primary networks. In [11], for example, the authors derived a mechanism that jointly optimizes the harvesting, sensing, and transmitting durations, and the number of sensed channels based on mixed-integer nonlinear programming with maximizing the achievable throughput serving as the objective function. Recently, the study of [12] suggested a new policy for determining both the sensing duration and the detection threshold that maximizes the average throughput. The proposed technique aims to find an optimal pair of sensing duration and detection threshold that can increase the spectrum access opportunities within the permissible range of collision probability for a given average harvested energy.

Detection Threshold

The performance of detecting the existence of primary signals is linked to the chosen value of the detection threshold. The choice of this value becomes even more crucial when the SU is an EH node [4]. In general, a high detection threshold increases the probability of detecting the spectrum as idle and leads to more frequent spectrum access. Not only does this increase the probability of colliding with the PU transmissions, it also causes a large waste of energy resulting from more transmissions. On the contrary, a low detection threshold alleviates unnecessary energy waste and the probability of accessing the occupied spectrum, but may in turn restrain an SU from transmitting data, even when the spectrum is idle. In [4], the authors propose a technique by which an optimal detection threshold is derived, using the probability of accessing the idle spectrum and the probability of accessing the occupied spectrum to maximize the expected total throughput while satisfying both the EH and collision constraints. They have also demonstrated that, depending on the selected threshold, the system can be characterized as a *spectrum-limited regime* and an *energy-limited regime*. In the first, the harvested energy enables continuous spectrum access, while in the second, the amount of harvested energy restricts the number of spectrum access attempts. This work was followed by that presented in [13] where they extended the problem in [14] to a joint optimization problem of a spectrum sensing policy and a detection threshold subject to the EH and collision constraints. In the framework of a POMDP, this strategy is able to achieve efficient usage of the harvested energy by exploiting the temporal correlation of the primary traffic. In addition to deriving the upper bound on the achievable throughput in [8], the authors have also explored a new technique which is able to find the optimal detection threshold that maximizes the derived upper bound.

If an SU employs SWIPT in order to simultaneously use the received RF signal to store energy and detect the presence of the PU, it is challenging to choose the optimal detection threshold. For example, in the power splitting approach, where the received signal at the SU is split into two portions, one for EH and the other for energy detection, the value of the detection threshold used in a non-EH SU receiver will not be viable. The reason is that the minimum acceptable signal energy at the input of the energy detector is divided according to the power splitting ratio. Hence, the detection threshold should correspond to the value of the received power after being split. This raises a question about the choice of energy threshold when the power splitting ratio is varying.

Energy Management

A careful allocation of power over sensing and data transmission slots is of high importance, due to its effect on the system throughput, capacity, and outage probability. In a CRN powered by ambient RF energy, the energy available at the beginning of a time slot is divided between the spectrum sensing and data transmission phases. Therefore, the harvested energy has to

The proposed technique aims to find an optimal pair of sensing duration and detection threshold that can increase the spectrum access opportunities within the permissible range of collision probability for a given average harvested energy.

In a CRN powered by ambient RF energy, the energy available at the beginning of a time slot is divided between the spectrum sensing and data transmission phases. Therefore, the harvested energy has to be efficiently expended over a specific number of time slots in order to enhance system performance.

be efficiently expended over a specific number of time slots in order to enhance system performance. The mechanism proposed in [14], for instance, enables an EH cognitive radio node to optimize its sensing and transmit energies while accounting for the detection reliability-throughput trade-off. Another method to achieve energy management is via knowledge of the previous or current statistics of the energy arrival rate, the statistical description of a PUs's activity, or the channel state information (CSI). For example, in [15], the proposed scheme allocates more energy for transmission when the channel state is good in a particular time slot. In contrast, less or no energy is allocated to a transmission slot in which the probability that the PU occupies the spectrum is anticipated to be relatively high.

The problem of energy management in a CRN applying SWIPT differs substantially from one that harvests ambient RF energy. The reason is that in some scenarios in SWIPT, the receiver has no battery to store energy, and as a result, the processes to be executed in a certain time slot directly draw energy from that available by the received RF signal. In this situation, it is challenging to optimize the parameters of the SU receiver such that energy is distributed spontaneously and efficiently between the different tasks of the cognitive cycle.

Channel Selection

Traditional channel selection schemes, which are mainly aimed at identifying the idle channels with high quality, may not be effective anymore for RF-powered CRNs. In particular, if the energy level available at the SU is low, it might select the channel that tends to be occupied by a PU and has a strong RF signal to harvest. On the other hand, if the SU has a high energy level, and there is a need for data packet transmission, it should identify the channel that is likely to be idle with a favorable channel quality. The research work reported in [16] studied a channel selection criterion that maximizes the average spectral efficiency of an SU. The proposed method jointly exploits knowledge of the PU occupancy and channel conditions, and the dependence of the decision of the SU to sense and access the PU spectrum on the probabilistic availability of energy at the SU. Similarly, in [2], the authors developed a channel selection policy used by the SU that maps the SU's state (i.e. number of packets in the data queue and the energy level in the energy storage) to the channel to be selected. This is done prior to sensing the channel and is based on statistical information such as probabilities of the channel being idle or busy, the probability of successful packet transmission if the channel is idle, and the probability of successful EH if the channel is busy.

Table 1 shows a summary of existing configuration policies for RF-powered CRNs.

Future Research for RF-Powered CRNs

CRNs may be deployed in different scenarios such as multiple-input multiple-output (MIMO), cooperative, and relaying CRNs. Existing mechanisms for conventional CRNs need to be extended, modified, or even replaced to suit the newly emerged RF-based EH technology. We focus next on discussing some issues that can be explored in the future.

Sensing Imperfections

Protecting the primary network from unbearable interference is the key to successful operation of a CRN. Therefore, a high probability of correct decisions generated by the energy detector is vital. In practice, however, those decisions are prone to errors, leading the performance of the primary network and the CRN to dramatically deteriorate. This becomes of higher concern in the presence of EH in those networks. In particular, if the channel is sensed as idle while it is actually busy, and if an SU decides to transmit, this results in unnecessary dissipation of energy, causing interference to the PU and missing a chance to harvest energy if needed. On the other hand, if the channel is sensed as busy while it is in fact idle, the SU might preserve energy but abolishes an opportunity to provide a better rate to its intended receiver. This necessitates research studies to explore the limitations caused by imperfect sensing on the performance of RF-based EH CRNs.

CRNs with Multiple Antennas

Multiple antennas in CRNs can be utilized to provide the secondary transmitter with more degrees of freedom in space in addition to time and frequency. Multi-antenna CRNs gained attraction, especially in the underlay spectrum sharing scheme, where SU and PU transmissions can be concurrent. In line with this, it is known that higher wireless energy transfer efficiencies can be achieved when multiple antennas are employed. Furthermore, in a multi-antenna RF-powered CRN, beamforming techniques can be exploited by the SU transmitter to steer RF signals toward SU receivers having different information and/or EH requirements. The problem of maximizing the SU rate subject to both the PU rate and the secondary transmitter power constraints is critical. Therefore, beamforming techniques should be redesigned to consider those conflicting objectives. The work presented in [17] is a major development in this field, where a multi-antenna EH secondary network makes use of both the spectrum and the energy of the primary network, in return to assist the primary transmissions. The main focus of this research is to design a beamforming technique that characterizes the achievable primary-secondary rate region based on power splitting and time-switching for SWIPT.

Beamforming performance optimization is tightly dependent on the acquisition of CSI. As a result, new mechanisms have to be proposed to account for the trade-off between data transmission, EH, and channel state estimation duration.

Cooperative CRNs

The concept of cooperative spectrum sensing has been proven to combat sensing errors and channel fading, and to overcome the hidden terminal problem due to shadowing. Nevertheless, conventional cooperative schemes do not take into consideration the DC power levels produced by the RF energy conversion process, which resemble

Configuration element	Literature	EH model	Constraints	Objective	Framework
Mode selection	[9]	Opportunistic EH of RF signals from primary network	1) Residual energy at the SU 2) Spectrum occupancy state partially observable to the sensor node	Maximize expected total throughput delivered by an SU sensor node over a time slot	POMDP
	[10]	EH of RF signals from primary network and ambient RF sources	1) Residual energy at the SU 2) Required transmission energy 3) Spectrum occupancy state partially observable	Enhance throughput of the SU and obtain QoS of primary network by selecting overlay or underlay transmission mode	POMDP
Sensing duration	[11]	EH from ambient RF sources	1) EH rate of the SU 2) Collision constraint to the primary network 3) Channel sensing energy cost	Optimize saving-sensing-transmitting structure that maximizes the achievable throughput of the SU	Mixed-integer nonlinear programming
	[12]	EH from ambient RF and other energy sources	1) Channel sensing and data transmission energy cost with respect to the residual energy at the SU 2) Collision constraint to the primary network	Maximize expected average throughput of the secondary network	Several optimization problems are formulated to give an insight on the joint configuration of sensing duration and threshold
Detection threshold	[4]	EH from ambient RF and other energy sources	1) Energy arrival rate 2) Channel sensing and data transmission energy cost with respect to the residual energy at the SU 3) Collision constraint to the primary network	Maximize expected total throughput of the secondary network	Deriving the probability of accessing the idle spectrum and the probability of accessing the occupied spectrum and their bounds
	[13]	EH from ambient RF and other energy sources	1) Spectrum occupancy state partially observable 2) Energy arrival rate 3) Temporal correlation of the primary traffic 4) Collision constraint to the primary network	Maximize the upper bound of the probability of accessing the idle spectrum	Unconstrained POMDP
	[8]	EH from ambient RF and other energy sources	1) Energy arrival rate 2) Channel sensing and data transmission energy cost with respect to the residual energy at the SU 3) Temporal correlation of the primary traffic 4) Collision constraint to the primary network	Maximize the upper bound of the achievable throughput	Several optimization problems are formulated to give an insight on the joint configuration of spectrum access policy and detection threshold
Energy management	[14]	EH from ambient RF and other energy sources	1) Energy arrival rate 2) Residual energy at the SU	Maximize expected total throughput of the secondary network	Markovian decision process
	[15]	EH from ambient RF and other energy sources	1) Observed information (harvested energy, fading CSI, spectrum occupancy state) in the past and present only	Maximize expected total throughput of the secondary network	Sliding window approach
Channel selection	[16]	EH from ambient RF and other energy sources	1) Probabilistic availability of energy at the SU 2) Channel conditions 3) Primary network belief state	Maximize expected total throughput of the secondary network	POMDP
	[2]	EH from RF signals of the primary network	1) Number of packets in the data queue 2) Residual energy at the SU	Maximize the long-term average throughput of the SU	Markovian decision process

Table 1. Summary of proposed techniques for RF-powered CRNs.

the only source of energy available at the CR terminal. To be more specific, an SU might refrain from participating in the process of spectrum sensing because it does not receive sufficient RF energy due to its distance from the PU. However, the more SUs that participate in sensing, the better spectrum discovery outcome is guaranteed and the more energy will be consumed. As a consequence, centralized cooperative spectrum scheduling, in which a cognitive base station or a fusion center decides which SUs should participate in the sensing process and which channels to sense, should take into account the amounts of harvested energy at the SUs. In addition, the distances between a PU transmitter and different SUs are often different. Also, the signal propagation environment differs from a PU transmitter to different SUs, making both the signal-to-noise ratio (SNR) and the harvested energy from the same primary signal dissimilar at different SU receivers. Therefore, new cooperative mechanisms that fit this environment is thus essential.

The recent interest in simultaneously achieving spectrum and energy efficiency has led to the concept of RF-powered CRNs. Integrating the capability of EH into the functionality of cognitive radio devices infer nontrivial challenges in their designs.

CRNs with Relays

In a CRN, a single or multiple relay(s) assist the SU source to sense and/or transmit data to the SU destination. All the CRN nodes or only the relay/s might be RF-based EH. In the second scenario, relays harvest energy from the SU source, the PU, or both. Under this setting, the quality of relaying the data to the SU destination is directly affected by the power received at the relay(s) from the SU source or the PU signals. This problem seems to be even more complex if the relay(s) and the SU source deploy SWIPT. In such a case, both the SU source and the relay(s) have to precisely select their receiver parameters (power splitting or time switching ratios) in order to optimize the overall system performance, while satisfying their energy needs. As a consequence, more research focus has to be directed toward exploring new relaying protocols and relay selection schemes.

Conclusions

The recent interest in simultaneously achieving spectrum and energy efficiency has led to the concept of RF-powered CRNs. Integrating the capability of EH into the functionality of cognitive radio devices infer nontrivial challenges on their designs. This article presents an overview of the architecture of CRNs that operate based on RF energy harvesting. Mainly, two methods by which CRNs can harvest RF energy were discussed: intended and non-intended RF energy harvesting. Several factors that do not exist in non-RF-powered CRNs impose fundamental limitations on their performance. As a result, the article lists key configuration parameters that need to be redesigned to achieve a desirable balance between the energy availability constraint and the system performance. Furthermore, the article surveys promising techniques that can enable successful spectrum sensing, spectrum access, and spectrum management in RF-powered CRNs. Finally, some open technical challenges that may be studied in the future are addressed.

References

[1] S. Lee, R. Zhang, and K. Huang, "Opportunistic Wireless Energy Harvesting in Cognitive Radio Networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 9, Sept. 2013, pp. 4788–99.

[2] X. Lu *et al.*, "Dynamic Spectrum Access in Cognitive Radio Networks with RF Energy Harvesting," *IEEE Wireless Commun.*, vol. 21, no. 3, June 2014, pp. 102–10.

[3] N. Shinohara, "The Wireless Power Transmission: Inductive Coupling, Radio Wave, and Resonance Coupling," *Wiley Interdisciplinary Reviews: Energy and Environment*, vol. 1, no. 3, 2012, pp. 337–46.

[4] S. Park, H. Kim, and D. Hong, "Cognitive Radio Networks with Energy Harvesting," *IEEE Trans. Wireless Commun.*, vol. 12, no. 3, Mar. 2013, pp. 1386–97.

[5] C. Valenta and G. Durgin, "Harvesting Wireless Power: Survey of Energy-Harvester Conversion Efficiency in Far-Field, Wireless Power Transfer Systems," *IEEE Microwave Mag.*, vol. 15, no. 4, June 2014, pp. 108–20.

[6] L. R. Varshney, "Transporting Information and Energy Simultaneously," *IEEE Int'l. Symp. Info. Theory*, Toronto, Canada, July 2008, pp. 1612–16.

[7] R. Zhang and C. Ho, "MIMO Broadcasting for Simultaneous Wireless Information and Power Transfer," *IEEE Trans. Wireless Commun.*, vol. 12, May 2013, pp. 1989–2001.

[8] S. Park and D. Hong, "Achievable Throughput of Energy Harvesting Cognitive Radio Networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 2, Feb. 2014, pp. 1010–22.

[9] S. Park *et al.*, "Optimal Mode Selection for Cognitive Radio Sensor Networks with RF Energy Harvesting," *IEEE 23rd Int'l. Symp. Personal Indoor and Mobile Radio Communications*, Sept. 2012, pp. 2155–59.

[10] M. Usman and I. Koo, "Access Strategy for Hybrid Underlay-Overlay Cognitive Radios with Energy Harvesting," *IEEE Sensors J.*, vol. 14, no. 9, Sept. 2014, pp. 3164–73.

[11] S. Yin *et al.*, "Optimal Saving-Sensing-Transmitting Structure in Self-Powered Cognitive Radio Systems with Wireless Energy Harvesting," *IEEE ICC '13*, June 2013, pp. 2807–11.

[12] W. Chung *et al.*, "Spectrum Sensing Optimization for Energy-Harvesting Cognitive Radio Systems," *IEEE Trans. Wireless Commun.*, vol. 13, no. 5, May 2014, pp. 2601–13.

[13] S. Park and D. Hong, "Optimal Spectrum Access for Energy Harvesting Cognitive Radio Networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 12, Dec. 2013, pp. 6166–79.

[14] A. Sultan, "Sensing and Transmit Energy Optimization for an Energy Harvesting Cognitive Radio," *IEEE Wireless Commun. Lett.*, vol. 1, no. 5, Oct. 2012, pp. 500–03.

[15] X. Gao *et al.*, "An Online Energy Allocation Strategy for Energy Harvesting Cognitive Radio Systems," *Int'l. Conf. Wireless Commun. Signal Processing*, Oct. 2013, pp. 1–5.

[16] J. J. Pradha, S. Kalamkar, and A. Banerjee, "Energy Harvesting Cognitive Radio with Channel-Aware Sensing Strategy," *IEEE Commun. Lett.*, vol. 18, no. 7, July 2014, pp. 1171–74.

[17] G. Zheng, Z. Ho, E. Jorswieck, and B. Ottersten, "Information and Energy Cooperation in Cognitive Radio Networks," *IEEE Trans. Signal Proc.*, vol. 62, no. 9, May 2014, pp. 2290–2303.

Biographies

Lina Mohjazi (l.mohjazi@surrey.ac.uk) received her B.Eng degree in electrical and electronic/communication engineering from the UAE University, UAE, in 2008, and her M.Sc. by research degree in communications engineering from Khalifa University, UAE, in 2012. Since October 2013, she has been a Ph.D. student at the University of Surrey, United Kingdom. Her main research interests include cognitive radio networks, energy harvesting communication systems, and physical layer optimization.

Mehrdad Dianati (m.dianati@surrey.ac.uk) is a reader (associate professor) in communication and networking systems at the Institute of Communication Systems (ICS) of the University of Surrey. His research area mainly includes wireless access networks and connected/autonomous vehicles. He also has nine years of industrial experience as a software/hardware developer and Director of R&D. He is currently an Associate Editor for *IEEE Transactions on Vehicular Technology*, *IET Communications*, and Wiley's *Journal of Wireless Communications and Mobile*.

George K. Karagiannidis (geokarag@ieee.org) received a Ph.D. degree in ECE from the University of Patras, Greece, in 1999. In 2004, he joined the faculty of Aristotle University of Thessaloniki, Greece, where he is a professor in the ECE Department and director of the Digital Telecommunications Systems and Networks Laboratory. In 2014, he joined Khalifa University, UAE, where is currently a professor in the ECE Department and coordinator of the ICT Cluster. Since January 2012 he has been Editor-in Chief of *IEEE Communications Letters*.

Sami Muhaidat (muhaidat@ieee.org) received his Ph.D. degree in electrical and eomputer engineering from the University of Waterloo, Canada. He is currently an associate professor at Khalifa University and a visiting professor in the Department of Electrical and Computer Engineering, University of Western Ontario, Canada. He currently serves as an Editor for *IEEE Communications Letters* and an Associate Editor for *IEEE Transactions on Vehicular Technology*.

Mahmoud Al-Qutayri (mqutayri@kustar.ac.ae) is a professor in the Department of Electrical and Computer Engineering and Associate Dean for Graduate Studies at Khalifa University. He received his B.Eng., M.Sc., and Ph.D. degrees from Concordia University, Canada, University of Manchester, United Kingdom, and the University of Bath, United Kingdom, all in electrical and electronic engineering, in 1984, 1987, and 1992, respectively. He has published numerous technical papers in peer reviewed international journals and conferences, and coauthored a book. His fields of research include embedded systems, wireless sensor networks, cognitive radio, and mixed-signal circuits.

APRIL 2015

free Tutorials Now™

IEEE ComSoc content sponsored by:

GL Communications Inc.

LTE-A Radio Network Planning Challenges & Opportunities

Long Term Evolution (LTE) is evolving into its advanced stage (LTE-A). This evolution introduces substantial network complexity which impacts network care and implementation, and provides a new avenue for innovation. Increased network complexity means that mobile broadband vendors should provide not only telecom equipment, but also network design, planning, integration and optimization expert services. Network design of LTE and LTE-A is challenged by issues such as beam forming, multicarrier, HetNet as well as new architecture and implementation concepts.

LTE-A brings network architecture evolution and is expected to interoperate seamlessly with other new and legacy technologies. Consequently, the questions around the role of Wi-Fi offloading and mobility approaches, carriers balancing, network backhaul, network liquid applications, and core network mobility versus routing are very exciting. This forum will discuss LTE-A network planning and implementation challenges and the ability to overcome them through methods, tools, and current industry practices.

LIMITED TIME ONLY AT >> WWW.COMSOC.ORG/FREETUTORIALS

IEEE COMSOC WEBINAR

available on demand

Sponsor content provided by:

InterDigital

Small Cells Big Gains: Increasing Cellular Capacity

Small cells are likely to become increasingly important to wireless carriers' network deployment plans since they can support dramatic improvements to network capacity and coverage.

However, the deploying of small cell technologies raises a number of issues that have strategic consequence for operators, especially since it involves spectrum assets. The application of small cells could provide a critical commercial difference between carriers since they have much greater capacity than macrocells.

We will look at this topic both from a technological comparison between adding more cells versus adding more spectrum in achieving higher capacity and what is actually possible in terms of the current regulatory framework.

LIMITED TIME ONLY AT >> WWW.COMSOC.ORG/WEBINARS

IEEE COMMUNICATIONS SOCIETY

For this and other sponsor opportunities,
please contact Mindy Belfer // 732-562-3937 // m.belfer@ieee.org

Provisioning Quality-of-Service to Energy Harvesting Wireless Communications

Xiaojing Chen, Wei Ni, Xin Wang, and Yichuang Sun

ABSTRACT

Energy harvesting (EH) is an innovative way to build long-term and self-sustainable wireless networks. However, an inconstant EH rate may have an adverse effect on the quality-of-service (QoS) of wireless traffic, such as packet delay and error. In this article we discuss techniques that provide QoS to EH powered wireless communications. A new "dynamic string tautening" method is presented to produce the most energy efficient schedule with substantially lower complexity, compared to convex optimization techniques. The method adapts to the bursty arrivals of wireless traffic and harvested energy, and ensures that delay-sensitive data will be delivered by deadline. Comprehensive designs of EH powered transmitters are also discussed, where the EH rate, battery capacity, and deadline requirement can be jointly adjusted to leverage QoS and the cost.

Xiaojing Chen is with EMW Lab, Fudan University.

Wei Ni is with CSIRO.

Xin Wang is with EMW Lab, Fudan University and Florida Atlantic University.

Yichuang Sun is with the University of Hertfordshire.

Work in this article is supported by the China Recruitment Program of Global Young Experts, the Program for New Century Excellent Talents in University under Grant No. NCET-13-0144, the Innovation Program of Shanghai Municipal Education Commission under Grant No. 13SG06, the National Science and Technology Major Project of the Ministry of Science and Technology of China under Grant No. 2012ZX03001013, and a research grant from Okawa Foundation.

ENERGY HARVESTING IN WIRELESS COMMUNICATIONS

Energy harvesting (or scavenging) is a process of capturing and converting ambient energy into usable electrical energy. A large number of external energy sources have potential to be harvested. They are [1]:

- Natural (renewable) energy, e.g. wind, water flow, ocean currents, and the sun.
- Mechanical energy, e.g. vibration and mechanical stress and strain.
- Thermal energy, e.g. waste energy from furnaces, heaters, and friction.
- Light energy, e.g. natural and artificial light.
- Electromagnetic energy, e.g. inductors, coils, and transformers.
- Energy from the human body, e.g. a combination of mechanical and thermal energy naturally generated by people when walking, sitting, climbing, and running.
- Energy from other sources, such as chemical and biological sources.

It is reported in [2] that environmental and kinetic energy harvesting, based on light, thermal, and motion, are the most promising techniques. Given a typical transfer efficiency of 10 percent, the energy that can be harvested in an outdoor daylight environment is about 1 mW/cm^2 [2]. This is at the same order of magnitude that carefully designed ultra-low-power micro-controller circuits typically consume.

In wireless communication systems, environmental EH is a critical component to build self-sustainable networks, such as wireless sensor networks in remote human-unfriendly environments [3]. Reducing carbon footprint by harvesting energy from renewable sources to power wireless transmissions is the key to implement self-sustainability. On the other hand, quality-of-service (QoS), e.g. delay and packet error rate, is crucial to many wireless applications. For example, sensory data are delay critical in bushfire and flood monitoring applications. However, to provide QoS to EH powered wireless transmissions, three critical challenges arise.

The first critical challenge in providing QoS to EH powered wireless transmissions is the unreliable nature of EH. Many EH technologies significantly rely on the environment where the devices are located [2]. The energy harvested may fluctuate dramatically with the time of day. This has strong impact on the reliability/availability of wireless links powered by the energy, and hence the QoS. Existing techniques developed to address the challenge are limited to delay-tolerant (bandwidth-demanding) traffic [4–6]. They are unable to provide QoS to delay-sensitive traffic.

The second critical challenge in providing QoS to EH powered wireless transmissions is to increase the energy efficiency of wireless transmissions. In general, the current energy transfer rate of EH is low, e.g. 15 to 20 percent for solar [2]. Maximizing the energy efficiency of wireless transmissions is therefore important to make the insufficient energy harvested meet the energy requirement of wireless transmissions, thereby reducing the probability of energy outage and QoS violations. Again, bandwidth-demanding delay-tolerant traffic has been the focus in existing techniques of optimizing transmission energy efficiency, e.g. the technique proposed in [7]. Those techniques cannot apply to delay-sensitive applications, where QoS is required.

0163-6804/15/$25.00 © 2015 IEEE

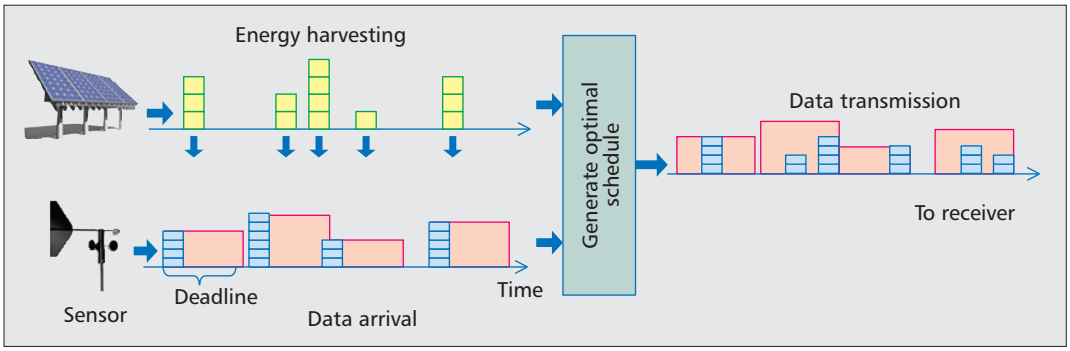

Figure 1. Illustration on EH powered wireless transmission, where solar panels are used to harvest energy. The transmitter decides the optimal transmission schedule, which adapts to the bursty characteristics of sensory data and EH, and ensures data to be transmitted by their deadlines.

Another critical challenge in providing QoS to EH powered wireless transmissions is the energy consumption of wireless circuitry. Part of the consumption is on transmission. Ultra-low-power circuitry has recently been developed to reduce active current, operating voltage, and pin leakage [2]. The current is about 16 mA when transmitting [2]. However, this is still non-negligible, given the low EH level of milliwatts. To save the transmission energy, an "on-off" mode was developed for bandwidth-demanding traffic, first in time-invariant channels [8] and then extended to time-varying channels [9, 10]. However, the mode is unable to address the QoS requirement of delay-sensitive traffic. Another part of circuit power consumption is on signal processing, especially for generating the optimal transmission schedules in delay-sensitive applications. Convex optimization techniques, such as the one proposed in [6], are in general unsuitable due to high search complexity.

Note that upper layer protocols, such as resource reservation protocol (RSVP) at the transport layer, have been widely used to provide QoS to traffic flows travelling over networks. QoS is provided by prioritizing traffic flows and adjusting bandwidths at every intermediate (routing/switching) node. However, this cannot address the aforementioned three critical challenges that reside on each individual wireless link and are caused by unreliable power sources.

In this article we present a new method that jointly addresses the three critical challenges and provides QoS to EH powered wireless transmissions. A new "dynamic string tautening" method is developed to produce the most energy efficient schedule with substantially reduced complexity, and ensure that delay-sensitive data is delivered by deadline. Comprehensive designs of EH powered transmitters are also discussed, where the EH rate, battery capacity, and deadline requirement are jointly adjusted to leverage QoS and the cost.

QoS Provisioning in EH Communications

In this article we focus on QoS provisioning on a single wireless link where the transmitter is powered by environmental EH techniques, as illustrated in Fig. 1. Delay sensitive sensory data arrive at the transmitter in bursts. The energy that the transmitter harvests is also bursty, due to the constantly changing ambient energy source. We focus on single-link QoS provisioning because it is the key and fundamental issue of EH powered wireless networks. Our results on the single link can be extrapolated to real network topologies and scenarios, as will be discussed later.

Referring to [3], the EH process can be modelled as a discrete sequence, where every element E_i ($i = 0, 1, \ldots, N$) is the energy (in joules) instantly harvested from ambient sources. We also model the bursty data arrival as a discrete sequence, where every element A_j ($j = 1, \ldots, M$) is the number of newly arrived packets.

We can similarly model a discrete sequence to represent the strict deadline requirements of the packets, where every element of the sequence D_k ($k = 0,1, \ldots, K$) indicates the number of packets that must be delivered so far. If any packet reaches its deadline but is undelivered yet, it will be discarded by the EH powered transmitter.

Every element of the sequences is tagged with a time stamp, which indicates when the event occurs (i.e. new energy is harvested, new packets arrive, or deadlines are met). E_i, A_j, and D_k are tagged τ_i, t_j, and μ_k, respectively. Of course, these time stamps may not overlap between difference sequences, as the EH process and the data arrival process are unnecessary to be synchronous in practice.

Arranging the time stamps of all three sequences in an ascending order, we can combine the three sequences into one sequence. The time interval between any consecutive two elements of the combined sequence is referred to as an "epoch." Within an epoch, the status of the three processes of EH, data arrival, and data deadlines does not change. We need to optimize the transmission schedule of each epoch for QoS provisioning, so that the overall optimality of the transmission schedule across the entire EH powered transmission process can be guaranteed.

Clearly, three constraints on generating the optimal transmission schedule arise due to causality:

- The total number of packets required to be delivered is the number of arrived packets.
- At any instant, the number of the transmitted packets must be no larger than the number of the arrived packets.

Within an epoch, the status of the three processes of EH, data arrival, and data deadlines does not change. We need to optimize the transmission schedule of each epoch for QoS provisioning, so that the overall optimality of the transmission schedule across the entire EH powered transmission process can be guaranteed.

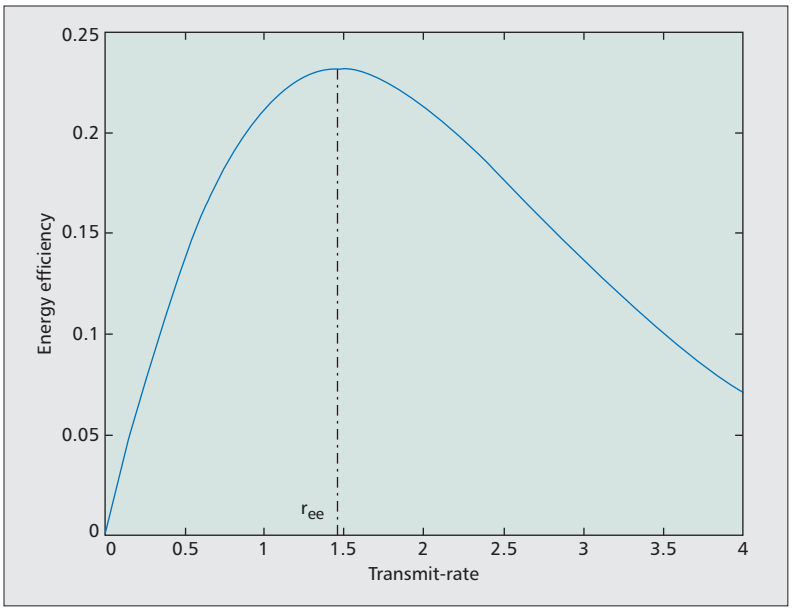

Figure 2. Bits-per-Joule energy efficiency versus transmission rate r, which is plotted based on Shannon's formulation $P(r) = 1/|h|^2\{\exp(r) - 1]$, where $P(r)$ is the power required to transmit the rate of r, and h is the channel coefficient of the link from the transmitter to the receiver.

- The total amount of energy consumed up to a time must be no greater than the energy accumulatively harvested so far.

Modelling EH powered transmission in such a way reveals that the EH process and the data arrival process have not only the same time sequence nature, but also yield causalities that need to be imposed as constraints to transmission schedules. Existing studies have been extensively conducted to address the problem of QoS provisioning to a bursty data process in the presence of unlimited energy [11]. The similarities between the processes indicate that there is a good opportunity to solve the QoS provisioning problem in EH wireless communications by extending those existing techniques.

In this article we introduce a feasible transmission schedule which, extended by our recent work [12], is able to provide QoS to bursty traffic adapting to the EH progress. We will confirm the optimality of the schedule, and show its superiority in computational efficiency.

QoS Provisioning under a Reliable Power Supply

In our recent work [12] we derived the most energy efficient transmission schedule to provide QoS to bursty data, given a constant and sufficient power supply. Here we summarize the conclusion of that work, which will provide the key insight and guidance to design the optimal EH powered transmission schedules, as will be described later.

In the case of a reliable power supply, the EH process/sequence is absent. An epoch is the interval between two consecutive time stamps in the combined sequence of data arrival and deadlines. The important conclusion we draw is that *the most energy efficient transmission schedule only adopts one of the following three policies per epoch i*:

- "On" policy, where the transmitter sends at the data rate higher than r_{ee} for the entire epoch i,
- "Off" policy, where there is no transmission in epoch *i*,
- "First-on-then-off" policy, where the transmitter sends at the data rate of r_{ee} for a time period less than the duration of the epoch *i*,

where r_{ee} is the data rate that maximizes the number of bits to be transmitted using a joule of energy. The number of bits to be transmitted using a joule is typically quasi-concave [7], as shown in Fig. 2. Therefore, r_{ee} can be efficiently obtained by a simple bi-sectional search.

In the case where a large number of packets reach their transmission deadline, the transmitter should send those packets all the time through until the deadline with the required power; see the "on" policy. In other cases, the transmitter should send deadline-reaching packets with the most energy efficient data rate. After the transmission completes, the transmitter turns off into a standby state (see the "first-on-then-off" policy). By doing these, the energy consumption of the transmitter is minimized, while the QoS of the bursty data traffic can be guaranteed.

The optimality of the policies in terms of energy efficiency can be rigidly proved through a judicious change of variables, which converts the original, non-convex problem of minimizing the transmit energy into a convex program [13]. The Karush-Kuhn-Tucker (KKT) conditions can then be employed to solve the convex problem and confirm that the policies satisfy the conditions.

The conclusion also indicates that the most energy efficient transmission schedule is event driven. The transmitter switches between the policies only on the arrival of new data bursts or on the data deadlines.

QoS Provisioning under EH

In the presence of EH, it can be shown that the most energy efficient policies described earlier still apply, because the energy efficiency of transmission is critical to EH powered systems. The key difference is that, in this case, the transmission data rate and duration depend on the availability of energy harvested from external unstable sources.

We define the time stamps of the combined time sequence of the EH, data arrival, and deadline processes, as a unified timeline $\{\sigma_0, \sigma_1, \ldots, \sigma_{N+M+K}\}$ (as described earlier). Epoch i is the interval from σ_{i-1} to σ_i with the duration of $L_i = \sigma_i - \sigma_{i-1}$. The construction of the unified timeline is due to the fact that the most energy efficient policies are event driven, as pointed out earlier. On the other hand, EH is an event that can activate switching policies. For example, when the transmitter runs out of energy, it stays "off" even if there are data to transmit. After new energy is harvested, the transmitter should turn "on" and proceed with either Policy 1 or 3.

Given the unified timeline, we can derive that the optimal transmission schedule based on harvested energy should comply with the following rule [13, 14]:

In the optimal transmission schedule, the trans-

mission data rate only changes at the instant when the data causality, deadline, or energy causality is met with equality. To be specific, the rate change happens after the epoch where:

- *There are no undelivered data.*
- *The transmitter runs out of energy.*
- *All deadline-approaching data are delivered.*

We note that at any instant, the number of delivered packets must not be larger than the number of packets that have arrived so far due to causality.

Also note that transmissions are interrupted if the transmitter runs out of energy, and will not be resumed until sufficient energy is harvested. In this case, data whose deadlines are within the transmission interruptions are dropped. Proper designs of the transmitter, including the EH rate and battery capacity, adapting to the data arrival process and the QoS (to be specific, delay) requirement are able to reduce the dropped data to lower than a required level. Details will be discussed later.

The optimality of the rule can be confirmed, because the rule was derived by first formulating the optimization problem to minimize the energy consumption under the constraints of an inconstant EH rate, bursty data arrival, and strict data deadlines. We then write the Lagrangian of the constrained optimization problem, and finally solve it with the KKT optimality conditions.

Following the rule, a computationally efficient algorithm with a temporally linear complexity can be developed to find the optimal transmission schedule that provides QoS to delay-sensitive, bursty data in EH systems.

Our new algorithm can be visualized as "string tautening," which has a distinguishing feature of changing feasible solution regions imposed by the dynamic EH process. In other words, the current direction of tautening the string depends on the past directions during a process of determining the optimal transmission schedule. The reason is because the past transmission data rates, indicated by the past tautening directions, affect the current energy level in the battery, which in turn determines the current direction to tauten the string. In contrast, existing "string tautening" methods have fixed solution regions [15], and therefore cannot address the dynamic EH problem.

Figure 3 illustrates our proposed "string tautening" process, where at any instant the data arrival curve plots the amount of data generated for transmission and the deadline (minimum data departure) curve plots the amount of data reaching their deadline. There are a number of EH curves. They are produced sequentially, one curve each time the conditions in the optimal rule are met. Each EH curve plots the maximum amount of data that can be transmitted at future instants, given the energy harvested and the data transmitted so far.

A closed feasible solution region is presented. The deadline curve provides the lower boundary of the feasible solution region. The (dotted) upper boundary of the region is provided by the lower of the data arrival curve and the dynamic EH curves, so that the optimal transmission schedule can satisfy both causalities of data and energy, as well as the deadline requirements (i.e. QoS).

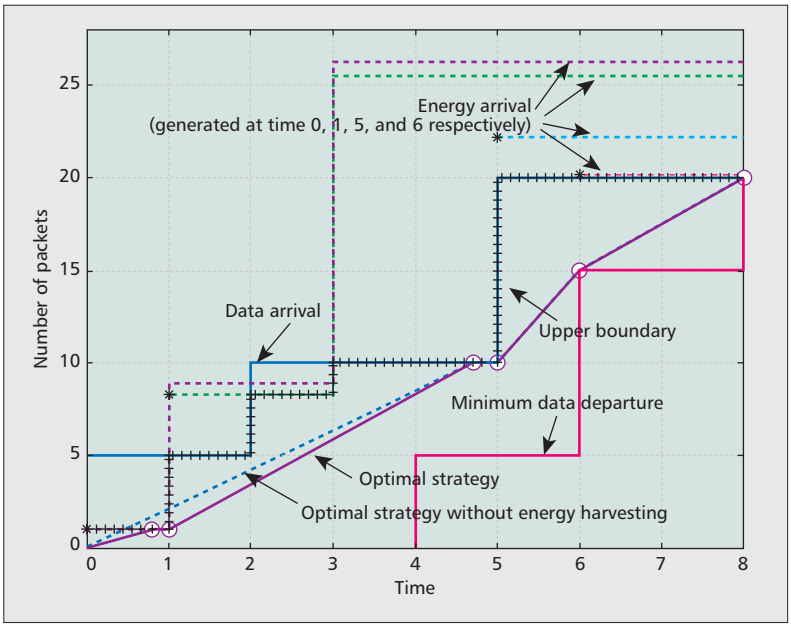

Figure 3. An illustrative example of the proposed "string tautening" method and the achieved optimal transmission schedule, which is the most energy efficient to provision QoS to EH wireless communications.

It is also possible that there are multiple closed solution regions in a single optimal transmission schedule, but not shown here. In this case, the upper boundary curve crosses and moves underneath the lower boundary during some periods of the transmission schedule. No transmissions will take place during the periods due to insufficient energy, as noted earlier. The periods become infeasible solution regions.

Seven epochs are demonstrated in the figure. Different transmission policies are adopted across the epochs, as highlighted in the figure. In the first epoch between time 0 and 1, the "first-on-then-off" policy is adopted, and the slope of the string is the most energy efficient transmission rate r_{ee}, as specified in the policy. In the next three epochs until time 5, the policy is also adopted throughout the epochs, with the most energy efficient transmission rate. In the last two epochs, the "on" policy is separately adopted, and the slope of the string is chosen such that the deadlines can be met.

It is worth pointing out that every EH curve is always underneath the previous curves, as shown in the figure. One reason is that the previous curves cannot foresee the future transmissions due to causality. Another reason is that every EH curve gives a tight energy budget based on the energy consumed so far. New transmissions will only make the budget tighter. The conclusion we can draw is that the optimal transmission schedule is able to leverage the current transmission rate, as well as the energy saved for future use.

For comparison purpose, we also plot the most energy efficient transmission schedule in the case of constant and sufficient power supply, employing [12]. We can see that the optimal transmission schedule in this case is in general above the one we developed for EH communica-

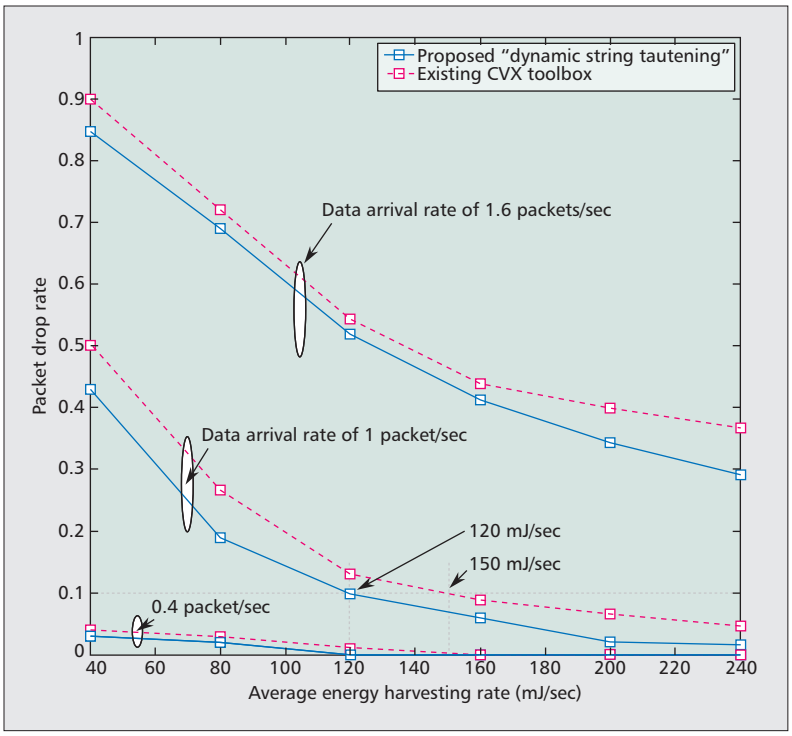

Figure 4. EH rate versus packet drop rate, where we assume the transmitter has unlimited battery capacity, the average data arrival rate is 0.4, 1.0, and 1.6 packets/sec, and the deadline is 2.5 seconds for every packet.

tions. This is because, in the presence of sufficient power supply, the upper boundary of the feasible solution region for the optimal schedule is provided solely by the data arrival curve, and therefore the region is larger than the one in the EH case. However, this optimal schedule does not apply to the EH case, because the string can cross and go beyond the EH curves, invalidating the schedule, as demonstrated in Fig. 3 when the time is 1.

Designing EH Transmitters with Guaranteed QoS

So far we have discussed the most energy efficient schedule to guarantee the deadlines of EH wireless transmissions, given any configuration on data arrival rate, EH capability, and battery capacity. However, a substantial amount of packets may still be dropped at the transmitter under a poor configuration, for example, excessively small battery or high data arrival.

A joint and holistic design of the data arrival, EH capability, and battery capacity is required to balance the deadline requirement and the tolerable packet drop rate. To demonstrate this, simulations are carried out, where we assume the circuit power consumption $\rho = 30$ mW. The transmit power is optimally achieved by using the new "string tautening" method, as described earlier. We also assume the gain of the wireless channel $|h|^2$ to be –20 dB during the optimal transmission. In practice, the channel may vary with the time. Our optimal schedule also applies to the time-varying case. Details will be discussed later.

In the simulations, the data arrival and the EH are modelled as two independent Poisson processes. The average data arrival rate is 1 packet/sec (unless otherwise specified), and the average EH rate ranges from 40 mJ/sec to 240 mJ/sec. It is worth mentioning that the optimal rule/schedule is general, and is applicable to any stochastic processes of data arrival and EH.

For comparison purposes, we also simulate the use of standard convex optimization techniques, specifically the interior point method, to generate the optimal transmission schedule. The interior point method is effective and has been extensively used to solve optimization problems with convex structures like the one discussed in this paper. The method has been implemented in the MATLAB CVX toolbox.

We note that the standard interior point method is able to produce the exact same optimal schedule in the case where reliable and sufficient power supply is available, as described earlier. However, the standard method requires a substantially higher computational complexity and sequentially higher circuit power consumption, than our "string tautening" method, because the interior point method requires matrix operations, high-order multiplications, and repeated iterations. It typically has a polynomial complexity higher than $\mathcal{O}(G^3)$ (where G is the number of epochs). In contrast, our "string tautening" method only requires linear complexity $\mathcal{O}(G)$ to adjust the direction of the "string," as shown in Fig. 3. In the case of EH, the higher circuit power consumption of the standard convex optimization methods would reduce the number of packets that can be optimally transmitted each time due to limited energy, and compromise QoS.

A careful design of the EH rate, adapting to the data arrival, is critical to leverage traffic QoS (i.e. the delay and the packet drop rate), as well as the cost of the transmitter. Figure 4 plots the average packet drop rate with the growth of the EH rate. We can see that the growth of the EH rate is critical to reducing the packet drop rate, especially when the data arrival rate is high. When the data arrival rate is 1.6 packets/sec, increasing the EH rate from 40 mJ/sec to 240 mJ/sec is able to reduce the packet drop rate from 85 percent to 29 percent.

Our "string tautening" method described earlier has substantially lower packet drop rates than the CVX programs. For a data arrival rate of 1 packet/sec and an EH rate of 80 mJ/sec, our method is able to achieve a packet drop rate of 19 percent, while the CVX incurs a packet drop rate of 28 percent. The reason is that the proposed method consumes much less energy to produce the optimal schedule than the CVX. As a result, more energy can be saved to transmit data and reduce the packet drop rate.

Our "string tautening" method is also able to relieve the requirement of EH. Consider a target packet drop rate of 10 percent for the data arrival rate of 1 packet/sec. By the figure, we can design the required EH rate to be 120 mJ/sec for our "string tautening" method, while the CVX based transmitter is required to harvest 150 mJ/sec energy on average. Typically, the EH rate is proportional to the size of the energy collector, e.g.

a solar panel. As a result, our method can be equipped with much lighter (by up to 25 percent) and therefore cheaper EH devices.

A realistic deadline requirement is important to alleviate the packet drop rate. Figure 5 shows the average packet drop rate with the growth of the deadline. We can see that given a target packet drop rate of 10 percent, a deadline requirement of 2.5 seconds can be satisfied by employing an EH rate of 120 mJ/sec. In contrast, a deadline requirement of 2 seconds requires a substantially higher EH rate of 200 mJ/sec. In other words, the required EH rate can be reduced by 40 percent by increasing the deadline by 25 percent. The cost of the EH devices decreases substantially.

A proper design of battery capacity is also crucial, adapting to the EH rate and deadline requirement. Figure 6 plots the packet drop rate with the battery capacity. It is interesting to see that the curves of different EH rates intersect with each other. The higher the EH rate, the more the packet drop rate changes with the battery capacity. The reason is that, when the battery is small, a high EH rate will lead to energy overflow. The overflow energy cannot be used for transmissions, and therefore is wasted. With the growth of the battery, the probability of energy overflow decreases. More energy can be collected for transmissions, and the packet drop rate can be significantly improved.

Consider the target packet drop rate of 10 percent. We can design by checking the figure that the required battery is 2.3 J for the EH rate of 160 mJ/sec, and 2.7 J for the EH rate of 120 mJ/sec. It is noted that in the case of the EH rate of 80 mJ/sec, the target packet drop rate cannot be achieved by enlarging the battery. The achieved packet drop rate stops decreasing when the battery is larger than 1.5 J, because the harvested energy is so small, and a battery of 1.5 J is sufficient to store the energy. However, the energy is insufficient to transmit the data, and the packet loss is high.

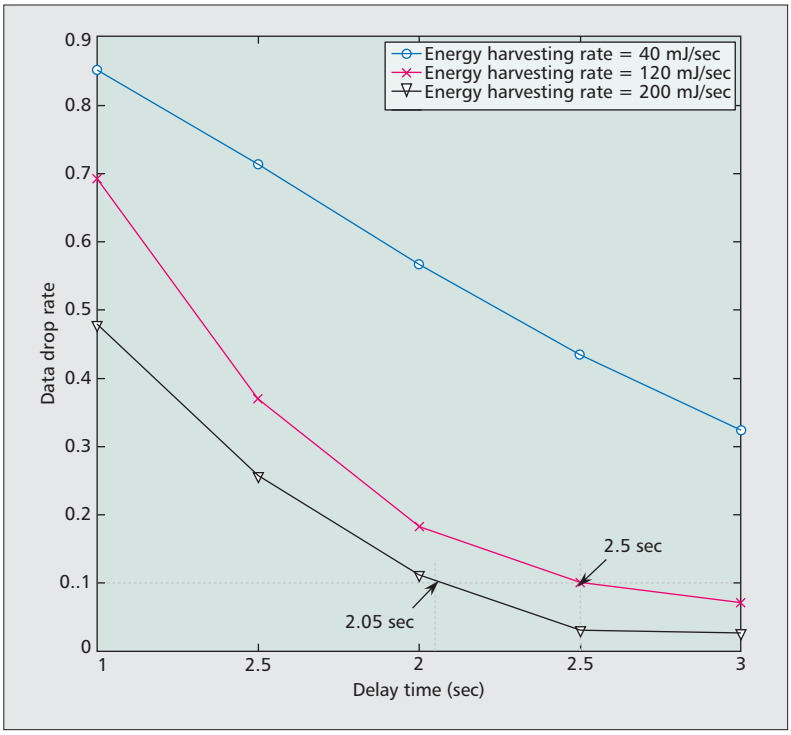

Figure 5. Delay requirement versus packet drop rate of our proposed "string tautening" scheme, where we assume the transmitter has unlimited battery capacity and the data arrival rate is 1 packet/sec.

Other Considerations on Practical Implementations

In this section we discuss the practical implementation aspects of the new "dynamic string tautening" method. The communication protocols running on each EH powered link, and connecting multiple links into a network, are described, because of their importance in real implementations.

Communication Protocol

Communication protocols are another key aspect of designing EH wireless communication systems, as mentioned earlier. It enables the receiver to report $|h|^2$, which is necessary for the transmitter to generate the energy efficient transmission rate and schedule, as discussed earlier. An adequate communication protocol is also important in many practical cases where the wireless channel fluctuates with changes in the environment and temperature.

Given the limited energy of the EH wireless systems, the protocol can be designed such that timeslots are preallocated for the receiver to update the transmitter with the channel gain $|h|^2$. The EH powered transmitter is aware of the preallocated timeslots. It switches from the transmit mode or the sleep mode to the receive mode at the timeslots, and will resume transmissions or return to sleep after that. The interval between two consecutive timeslots is less than the coherence time of the channel, i.e. the channel gain is stable during the interval. By this means, the energy that the transmitter requires to receive the channel updates can be substantially reduced, compared to maintaining a standby receive channel.

Given the protocol, the optimal transmission schedule described earlier can be generated by puncturing out the preallocated timeslots, constructing a new timeline, and carrying out "string tautening" over the new timeline. The data arrivals within a preallocated timeslot are aggregated to the timepoint where the beginning and the end of the preallocated timeslot join in the new timeline. Thus are the energy harvested and the deadlines reached within the timeslot. The rest of the "string tautening" operations are as described.

The protocol can be further extended to enable the interval between preallocated timeslots to change (by half or double), adapting to the varying wireless channel. This can be implemented by the receiver adding one more bit into the channel updates to indicate the interval until the next channel reporting timeslot. This allows the interval to quickly converge to the coherence time of the time-varying channel.

The protocol can also be extended to enable the real-time adjustment of the data deadlines,

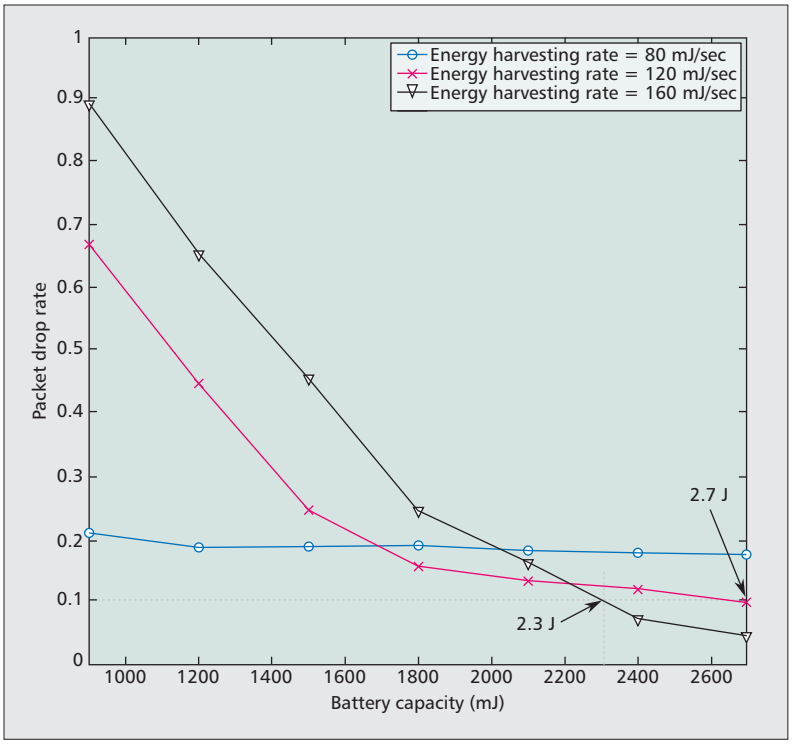

Figure 6. Battery capacity versus packet drop rate of our proposed "string tautening" scheme, where the data arrival rate is 1 packet/sec and the deadline requirement is 2.5 seconds.

adapting to the EH level and the packet drop rate. Specifically, the transmitter can decide to increase or decrease the deadline of every packet based on the current packet drop rate. This can be done by enabling the receiver to report the packet drop rate along with the channel updates. As a result of these extensions, the packet drop rate can be stabilized.

Extrapolation to Large-Scale Networks

To extrapolate EH powered transmissions to a real network, the key challenge is the design of the EH powered wireless receivers, because limited harvested energy can hardly afford continuous reception and decoding. To address the challenge, we can preallocate periodic timeslots to every pair of transmitter and receiver. The receiver only wakes up and receives during the timeslots, and is in sleep mode the rest of time. The transmitter schedules its transmissions only within the timeslots. In this case, the optimal transmission schedule can be generated in a similar way, as described previously. Specifically, the transmitter punctures out the time periods during which the receiver is in sleep mode, constructs a new timeline, and carries out the "dynamic string tautening" over the timeline. A careful design of the timeslot preallocation is important, which affects the QoS and packet drop rate, as well as the energy consumption of the transmitter and receiver.

Conclusions and Future Work

In this article we addressed the issue of providing QoS to EH powered wireless communications. A new "dynamic string tautening" method was developed to produce the most energy efficient schedule, adapting to the bursty arrival of traffic and energy. The method ensures that delay-sensitive data will be delivered by deadline, meanwhile substantially reducing complexity and circuit power consumption. Using the "string tautening" method, designs of EH powered transmitters were discussed. The EH rate, battery capacity, and deadline requirements can be jointly designed to leverage QoS and the energy requirement.

The future directions of the work include:

- Extrapolation of the optimal QoS schedule to large-scale EH powered networks, where careful designs of sleep mode are required for the transceivers.
- Upper layer adaptation to the EH powered transmissions, including routing, prioritization, and QoS provision over EH powered wireless networks.
- EH processes with multiple (discrete and/or continuous) power sources, and their impact on QoS.

References

[1] C. Knight, J. Davidson, and S. Behrens "Energy Options for Wireless Sensor Nodes," *Sensors*, no. 8, 2008, pp. 8037–66.

[2] M. Raju and M. Grazier, "Energy Harvesting ULP Meets Energy Harvesting: A Game-Changing Combination for Design Engineers," Texas Instruments, White Paper, Apr. 2010.

[3] D. Gunduz, K. Stamatiou, N. Michelusi, and M. Zorzi, "Designing Intelligent Energy Harvesting Communication Systems," *IEEE Commun. Mag.*, vol. 52, no. 1, Jan. 2014, pp. 210–16.

[4] K. Tutuncuoglu and A. Yener, "Optimum Transmission Policies for Battery Limited Energy Harvesting Nodes," *IEEE Trans. Wireless Commun.*, vol. 11, no. 3, Mar. 2012, pp. 1180–89.

[5] J. Yang and S. Ulukus, "Optimal Packet Scheduling in an Energy Harvesting Communication System," *IEEE Trans. Commun.*, vol. 60, no. 1, Jan. 2012, pp. 220–30.

[6] O. Ozel *et al.*, "Transmission with Energy Harvesting Nodes in Fading Wireless Channels: Optimal Policies," *IEEE JSAC*, vol. 29, 2011, pp. 1732–43.

[7] G. Miao, N. Himayat, and G. Li, "Energy-Efficient Link Adaptation in Frequency-Selective Channels," *IEEE Trans. Commun.*, vol. 58, no. 2, Feb. 2010, pp. 545–54.

[8] J. Xu and R. Zhang, "Throughput Optimal Policies for Energy Harvesting Wireless Transmitters with Non-Ideal Circuit Power," *IEEE JSAC*, vol. 32, no. 2, Feb. 2014, pp. 322–32.

[9] Q. Bai, J. Li, and J. Nossek, "Throughput Maximizing Transmission Strategy of Energy Harvesting Nodes," *Proc. IWCLD*, 2011.

[10] O. Orhan, D. Gunduz, and E. Erkip, "Throughput Maximization for an Energy Harvesting System with Processing Cost," *Proc. ITW*, 2012.

[11] X. Wang and Z. Li, "Energy-Efficient Transmissions of Bursty Data Packets with Strict Deadlines over Time-Varying Wireless Channels," *IEEE Trans. Wireless Commun.*, vol. 12, no. 5, May 2013, pp. 2533–43.

[12] Z. Nan, X. Wang, and W. Ni, "Energy-Efficient Transmission of Delay-Limited Bursty Data Packets under Non-Ideal Circuit Power Consumption," *Proc. ICC*, Sydney, Australia, June 10–14, 2014.

[13] X. Chen, X. Wang, and X. Zhou, "Energy-Harvesting Powered Transmissions of Delay-Limited Data Packets," *Proc. Globecom*, Austin, TX, Dec. 8–12, 2014.

[14] X. Chen, X. Wang, and Y. Sun, "Energy-Harvesting Powered Transmissions of Bursty Data Packets with Strict Deadlines," *Proc. ICC*, Sydney, Australia, June 10–14, 2014.

[15] M. Zafer and E. Modiano, "A Calculus Approach to Minimum Energy Transmission Policies with Quality Of Service Guarantees," *Proc. INFOCOM*, vol. 1, 2005, pp. 548–59.

BIOGRAPHIES

XIAOJING CHEN received her B.E. degree in communication engineering from Fudan University, China in 2013. Currently she is working toward her Ph.D. degree at Fudan University. Her research interests include wireless communications, energy-efficient communications, and stochastic network optimization.

WEI NI [M'09] received the B.E. and Ph.D. degrees in electronic engineering from Fudan University, China, in 2000 and 2005, respectively. Currently he is a senior scientist at Digital Productivity Flagship, CSIRO, Australia. Prior to this he was a research scientist and deputy project manager at the Bell Labs Research and Innovation Centre, Alcatel-Lucent (2005-2008), and a senior researcher at Devices R&D, Nokia (2008-2009). His research interests include multiuser MIMO, radio resource management, and software-defined networking. He has served as an editorial board member for *Hindawi Journal of Engineering* since 2012.

XIN WANG [SM'09] (xwang11@fudan.edu.cn) received the B.Sc. degree and the M.Sc. degree from Fudan University, Shanghai, China, in 1997 and 2000, respectively, and the Ph.D. degree from Auburn University, Auburn, AL, in 2004, all in electrical engineering. From September 2004 to August 2006 he was a postdoctoral research associate in the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis. In August 2006 he joined the Department of Computer & Electrical Engineering and Computer Science, Florida Atlantic University, Boca Raton, as an assistant professor, and then an associate professor from August 2010. He is now a professor in the Department of Communication Science and Engineering, Fudan University, China. His research interests include stochastic network optimization, energy-efficient communications, cross-layer design, and signal processing for communications.

YICHUANG SUN [M'90, SM'99] is a professor in the School of Engineering and Technology at the University of Hertfordshire, UK. His research interests are in wireless and mobile communications, and RF and mixed-signal circuits. He has published more than 260 papers and four books. He was a book series editor for the IET, an associate editor of *IEEE Transactions on Circuits and Systems-I*, and a guest editor of six special issues for *IET Communications*, *IET Signal Processing*, and *IET Circuits, Devices and Systems*.

Increasing Sustainability and Resiliency of Cellular Network Infrastructure by Harvesting Renewable Energy

Andres Kwasinski and Alexis Kwasinski

ABSTRACT

The carbon footprint of cellular base stations is a source of concern not only because of their power consumption, which accounts for more than half of all of the cellular infrastructure, but also because of the large rate of growth in their deployment. This article discusses how the use of harvested renewable energy can contribute to solving this problem. The article also addresses the challenges associated with harvesting wind and solar energy, namely the variability in available energy and the large physical footprint of energy harvesters. This article explains that these challenges can be better addressed by jointly considering the harvested energy availability and the dynamic characteristics of base station power consumption. A realization of this approach is the recently introduced idea of a "sustainable wireless area" that integrates energy harvesters and a group of base stations in a microgrid architecture. This architecture enables an integrated harvested energy-cellular traffic management technique that shapes the traffic serviced by a base station based on the predicted availability of renewable energy. As a result, longer periods of operation powered from renewable energy are achieved while degradation of the users' quality-of-experience (QoE) is minimal and occasional. This article also explains how harvested renewable energy also increases the resiliency of cellular networks because they do not depend on lifelines for operation.

INTRODUCTION

Multiple studies are providing strong support for the need to address the carbon footprint associated with the operation of a cellular network, [1, 2]. At a larger scale, the information and communications technology (ICT) infrastructure is contributing to two percent of the global carbon footprint (a value similar to the airline industry) [3], and is expected to grow yearly at a rate of four percent until 2020 [4]. Within the ICT infrastructure, the accelerating shift from wired to wireless networking results in cellular communication infrastructure growing in deployment and use (traffic) at a faster rate than other ICT infrastructures and, consequently, a faster carbon footprint growth rate.

For cellular network operators, a major center of attention is at the base stations because they account for more than half of the energy expenditure [1, 5]. As such, the carbon footprint of cellular networks can be reduced by developing sustainable approaches to powering base stations. One leading approach in this area is the use of harvested renewable energy.

It is commonly recognized that the use of harvested renewable energy reduces the carbon footprint of the cellular network and also allows for the deployment of cellular infrastructure in areas with limited or no electric power distribution infrastructure. Nevertheless, it is less recognized that the use of harvested renewable energy increases the resiliency of the cellular network, especially in the case of extreme disaster events [6]. Even without considering the effects of natural disasters, backup power systems based on harvested renewable energy could help address the reliability mismatch between power grids and communications equipments. Indeed, while power grids in the U.S. have an expected availability of what is called "3-nines" (operating 99.9 percent of the time) or total yearly outage time less than nine hours, communication systems require much shorter individual outage times and an overall availability of 5-nines. In this article we will expand on all these impacts that the use of harvested renewable energy could have on the operation of cellular infrastructure, and we will discuss techniques for a more effective use of harvested renewable energy.

ENERGY HARVESTING IN THE CELLULAR NETWORK INFRASTRUCTURE

The application of renewable energy to power cellular base stations needs to jointly consider the aspect of electric power generation and availability, together with the dynamic characteristics of the load. Consequently, in this section

Andres Kwasinski is with the Rochester Institute of Technology.

Alexis Kwasinski is with the University of Pittsburgh.

0163-6804/15/$25.00 © 2015 IEEE

we focus first on the electric power subsystem, and then we discuss issues related to the base station as an electric load.

The Electric Power Subsystem

Renewable energy can be obtained from multiple sources, such as wind, solar, hydropower, geothermal, and biomass [7]. In this article we will focus on solar and wind energy because harnessing some of the other sources requires a large infrastructure that cannot be considered as consistent with "harvesting" energy to power base stations, and also because wind and solar energy are available in the vast majority of geographical settings with practical interest. In addition, electric energy can be harvested from the sun and the wind without the need for a large energy distribution infrastructure, which is a key advantage that will be further discussed later.

Nevertheless, the use of solar and wind energy also introduces some technical challenges. The first challenge is their variability due to weather and, in the case of solar energy, the day-night cycle. This variability can be addressed through two approaches. The first approach consists of diversification by simultaneously harvesting energy from more than one source, which in this case takes the form of harvesting energy from both the wind and sun radiation. Figure 1 illustrates the rationale for this approach by showing how one energy source with low output may be complemented by another having at that time high power output (see days 96 and 97). The solar electric power data in Fig. 1 was collected from an array of 26 MX Solar USA Suncase MX60–240 panels installed at the Rochester Institute of Technology's Golisano Institute for Sustainability. Each panel has a size of 1.7 m^2 and is capable of generating a maximum of 240 W with an efficiency of 14.3 percent. The wind electric power data in Fig. 1 was collected from a 250 kW Fuhrlander FL250 wind turbine installed to partially power a plastic molding company near Rochester, NY. Note that harvesting renewable energy during nighttime is particularly important for powering base stations in residential areas because cellular traffic, and the base station electric power needs, tends to reach its maximum value in the early night hours, [8, 9]. The second approach to address renewable energy variability consists of the use of a battery bank to store renewable energy when there is surplus availability so that it can be used at a later time of deficit. The cost associated with battery banks is one of the main factors slowing down broader renewable energy adoption, but is less of a concern when powering cellular base stations because cell sites usually already have battery banks to provide backup power in cases of main electric grid outages. Even more, while in a traditional use as power backup, the battery banks are costly because they are typically needed on average a few hours a year, when used as renewable energy storage their utilization becomes much larger.

A second challenge encountered when harvesting wind or solar energy to power base stations is that these power sources present a relatively large physical footprint because their typical power density (e.g. about 200 W/m^2) is

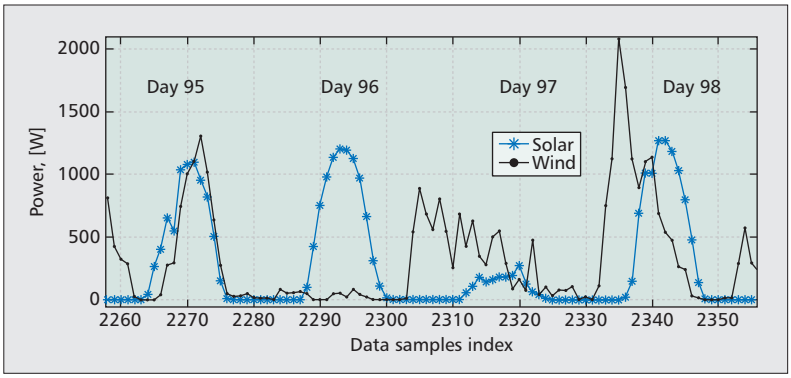

Figure 1. Wind and solar power output measured on days 95 through 98.

significantly smaller than that of base station loads (e.g. 5 kW/m^2). As a result, use of harvested renewable energy has been considered first to power cellular base stations in remote areas that are disconnected and far from an electric power distribution grid (or other energy distribution infrastructure), and for which the usual solution is the very costly and logistically challenging use of diesel generator sets (gensets). The requirements and feasibility of the practical realization of this application have been studied in many technical publications, e.g. [9, 10]. It was observed in [9] that powering a base station in a non-electrified area from renewable generators, while feasible, may need the acceptance of the generators' large physical footprint. Even then, the variability of the renewable energy sources results in a less reliable power system than what is expected for telecommunication systems. Of course, while a less reliable cellular service may be acceptable under conditions that otherwise would have no service at all, meeting typical telecommunication reliability goals requires costly solutions such as the use of gensets, deployment of large battery banks, or much larger wind or solar generators (as will be seen in the next subsection). Nevertheless, this application showcases a new paradigm for the electric grid where power generation becomes distributed and located close to the electric loads.

Part of the challenge presented by the large physical footprint of renewable electric power generators is due to an established thinking that considers the electric power system for each cell site independent of the others. This issue can be addressed by considering a new paradigm where power generation and energy storage resources are shared among a group of cell sites. In [11] we introduced the concept of a "sustainable wireless area" (SWA) as a self-contained and independently controlled power system that is formed by interconnecting a group of nearby cell sites in a common power distribution architecture. As shown in Fig. 2, an SWA is formed by a cluster of a few base stations with power obtained primarily from PV modules and wind turbines. Figure 2 also shows different wind turbines that could be used in the SWA. The horizontal axis wind turbine shown in Fig. 2 is the previously mentioned Fuhrlander FL250. Due to its size it would typically be found in a centralized location within the SWA, yet note that the

Microgrids are electric power systems that can operate powering their loads both when connected or disconnected from a large conventional grid by being self-contained electric power systems, with their own local power generators, controllers, loads, power distribution system, and, in most cases, energy storage devices.

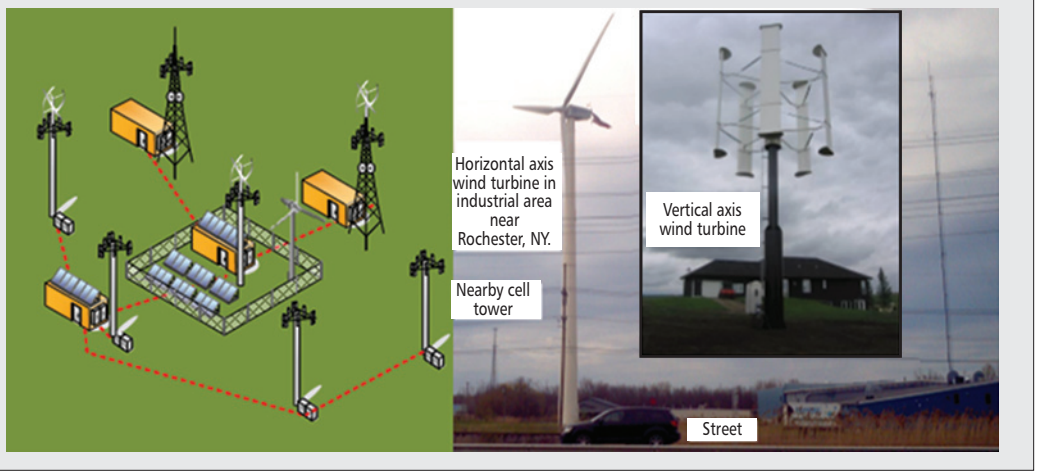

Figure 2. An example SWA with seven base stations and different wind turbines as could be installed in an SWA.

turbine height is similar to many cell site antenna towers, as can be seen in Fig. 2 by noticing the cell site tower located approximately 400 m apart. Smaller wind turbines, like the one with vertical axis shown in Fig. 2, which has a maximum power output of 10 kW and a diameter of 6 m, could be used at each cell site. Recall that the cell sites are already equipped with battery banks. From the perspective of an electric power system architecture, the SWA forms a "microgrid." Microgrids represent a paradigm change in electric power generation and distribution. Microgrids are electric power systems that can operate powering their loads both when connected or disconnected from a large conventional grid by being self-contained electric power systems, with their own local power generators, controllers, loads, power distribution system, and, in most cases, energy storage devices.

Cellular Base Station and Network Traffic as an Electric Load

We first look into the feasibility of powering base stations from harvested wind and solar energy. We considered for this an LTE base station (known as "eNBs") that belongs to an SWA. The eNB is a macro base station, transmitting at a power of 40 dBm, assumed to be separated 500 m from the closest eNBs and with a configuration typical for an urban environment, servicing three sectors and operating with an assumed two transmit and two receive MIMO configuration. We modeled the electric load (dynamic power consumed) from the eNB following [8]. According to this model, the load can be separated into a constant component and a dynamic component. The constant consumed power, equal to 65 W per transmitter unit in our model, accounts for the radio resource overhead (e.g. for pilot signals), cooling, processing, baseband interface, etc. In an LTE system, the dynamic power consumption component changes according to the intensity of cellular traffic through the eNB, increasing practically in a linear fashion following the fraction of the total resource blocks (RBs) in a frame that are allocated to transmit cellular traffic. In LTE, an RB is the minimum unit of time-frequency resource allocation and is composed by an array of 12 subcarriers by seven OFDM symbols. The number of RBs that needs to be allocated to a given call depends on the call's requirements and on the bit rate associated with the RB. The RB bit rate depends on the modulation (QPSK, 16-QAM, or 64-QAM) and error correcting settings chosen for the RB, which in turn depends on the signal-to-interference-plus-noise ratio (SINR) experienced during transmission of the RB. This channel was modeled considering path, penetration, and shadowing loss as well as delay spread (modeled with equal probability as either the Pedestrian B or the Vehicular A typical urban model from the Recommendation ITU-R M.1225 — "Guidelines for evaluation of radio transmission technologies for IMT-2000"). According to [8], the dynamic power consumption component adds at full cellular traffic load another 97 W per transmitter unit. The dynamic traffic intensity was modeled following the profiles in [8].

We considered three different configurations for harvesting wind and solar energy. The first configuration consisted of six Solar USA Suncase MX60–240 solar panels (occupying a total area of 10.2 m^2) and one 10 kW wind turbine at each cell site. The second configuration consisted of three MX60-240 solar panels and one 250 kW wind turbine shared by the seven eNBs in the SWA. The third configuration consisted of six MX60-240 solar panels and one 250 kW wind turbine shared by the seven eNBs in the SWA. In all configurations, each cell site counts with a battery bank dimensioned to provide energy to one eNB for eight hours at an average cellular traffic load of 80 percent. In the study we used actual measurements from solar and wind generation collected over a period of 100 days between May and August 2013. The power output for the 10 kW wind turbine was obtained by scaling the measurements from the Fuhrlander FL250 turbine, which is located in an area with average wind power capacity (wind zone class 5 in a scale from 1 to 7). In doing so, it was considered that the technology progress in the 12 years between

the older but larger FL250 turbine and the smaller 10 kW turbine resulted in a similar efficiency for both turbines. To quantify the feasibility of powering base stations from harvested wind and solar energy we measured the cumulative density function (CDF) of the battery bank state of charge (SoC). Since the batteries are only charged from harvested renewable energy, the event when the battery SoC equals zero represents the situation when the batteries have been discharged and the eNB cannot be powered any longer from the harvested renewable energy. The results shown in Fig. 3 indicate that all three configuration are able to operate powered from renewable sources a large proportion of time (68 percent, 81 percent, and 90 percent of time for the first, second and third configurations, respectively). These proportions of time would result in significant reductions in energy costs and base station carbon footprint. Yet they do not reach the typical expectation for telecommunications equipments to operate at least 99.999 percent of the time and, therefore, it is usually necessary to complement the harvested renewable energy with other sources of electric power (typically from a connection to the main grid). Nevertheless, these results underscore the value of techniques that save power at the base station to increase the proportion of time of operation powered from harvested renewable energy.

A straightforward technique to reduce the power consumed by a group of base stations is to set some of them to a "sleep" mode of operation where they consume very low power but they are unable to provide cellular service. As discussed in [9], this technique reduces the energy harvesting system requirements, making possible some implementations that would not be possible otherwise. Yet because the "sleeping" base stations do not provide cellular service, in the simplest implementation of this technique the base stations are set to sleep mode only when allowed by a small enough cellular traffic intensity.

The use of a sleep mode was improved in [12] by introducing a dynamic programming algorithm that determines, given the evolution of cellular traffic intensity and harvested energy over time, which base stations within a group are set to sleep mode and, for the active ones, the RBs utilized to carry cellular traffic. The algorithm considers the users' quality-of-service (QoS) requirements by minimizing the probability that an existing call cannot be allocated enough RBs or that a new call cannot be serviced. The algorithm achieves near-optimum performance, thus significantly reducing the base station power consumption compared to a setup that attempts no action to save power and extend the use of harvested renewable energy.

The work in [12] is representative of other recent published literature that aims at achieving a more effective use of renewable energy following approaches that reduce power consumption only when allowed by the given intensity of cellular traffic. Fewer works discuss techniques applicable to any traffic intensity conditions. The authors in [13] study the problem of reducing the energy consumed from the traditional grid in a cellular network that is harvesting energy during peak traffic hours. The problem is solved by addressing two interrelated sub-problems:

- Considering the temporal variations in traffic intensity and harvested energy, allocate the usage of harvested energy over a period of time divided into decision time slots.
- Control the coverage area of multiple base stations so that those with higher availability of harvested energy cover a larger area (and consume more power).

Heterogeneous cellular networks (a traditional cellular network equipped with macro base stations where there is an underlaid deployment of low power access nodes) present the opportunity for a more effective use of harvested energy by associating users to the base stations with higher harvested energy capacity. In [14] it was studied how to do this user association so as to avoid high traffic intensity and data forwarding delay in base stations with favorable harvested energy conditions.

Another technique that is applicable to any traffic intensity conditions takes advantage of the fact that in the microgrid within the SWA architecture, the generators, controllers, and loads are all located in the vicinity of each other, which allows for the control of cellular traffic intensity and the users' QoE based on the short-term prediction of harvested renewable energy availability, effectively adding an extra degree of freedom to the power management system. In [15] we presented an integrated harvested energy-cellular traffic management technique that shapes the traffic serviced by a base station based on the predicted availability of renewable energy. In this technique, traffic is shaped by controlling the RBs allocated to real-time video and data traffic because they account for almost all the traffic volume seen through a base station. The rationale for shaping traffic is based on the

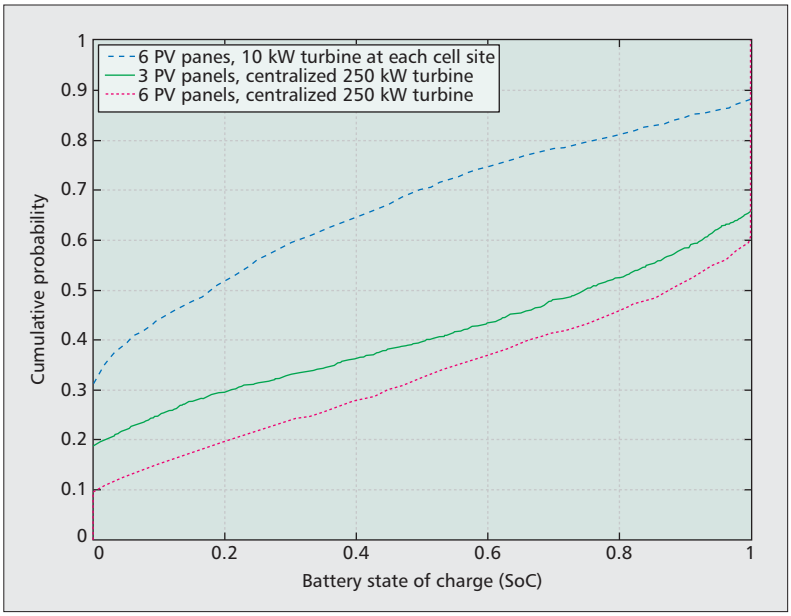

Figure 3. Cumulative distribution function (CDF) for the battery SoC for an eNB in an SWA using different configurations for wind and solar generators.

earlier observation that the power consumed at a base station changes linearly with the traffic load. In practice, real time video traffic can be controlled by leveraging the scalability features found in modern video codecs, and the data traffic is controlled by adapting the transmit throughput. As such, traffic is shaped through a controlled, smooth, and transient reduction of real-time video quality and increase in data delay. Then the integrated harvested energy-cellular traffic management technique presented in [15] computes for each control time period (chosen equal to one hour) a traffic shaping factor. This factor is a number between 0 and 1 that multiplies the number of RBs allocated without shaping, so as to smoothly reduce the service provided by the eNB in a controlled way and maintain the impact to the users at an acceptable level and little noticeable. The traffic shaping factors at each decision time are calculated by determining a target change in the energy stored in the batteries. This target change may be positive or negative and chosen based on the renewable energy predicted surplus or deficit with no traffic shaping. As illustrated in Fig. 4, the projected surplus or deficit is characterized by the net area S_A under the curve $S_{nc}[m] - S_{lim}$ over a chosen time period, where $S_{nc}[m]$ is the predicted energy stored in the batteries at time instant m without traffic shaping, and S_{lim} is a value indicative of a critically low energy stored in the batteries.

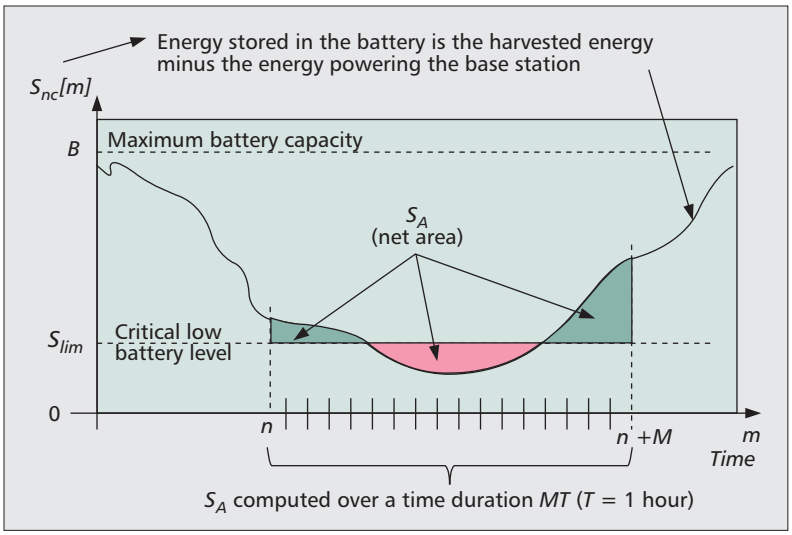

Figure 4. Illustration of the meaning of the magnitude S_A.

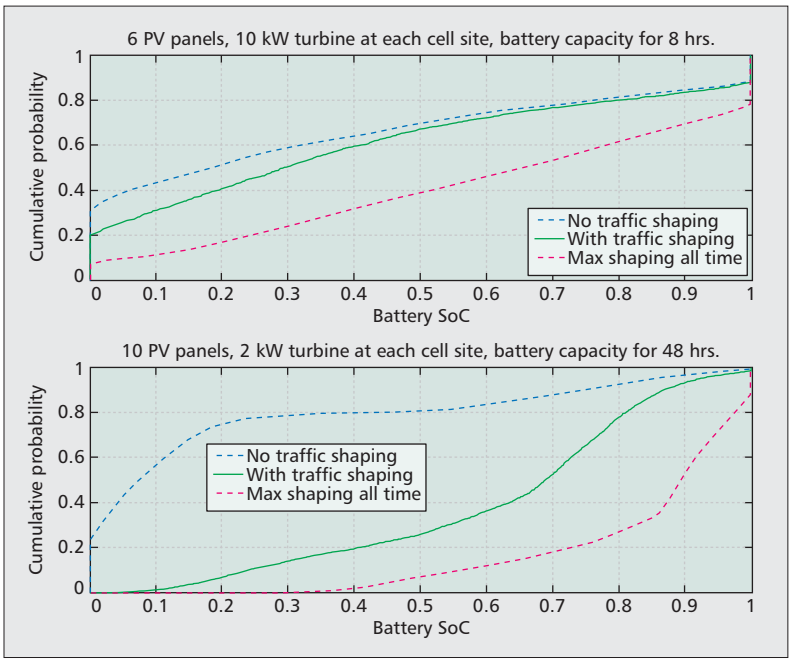

Figure 5. CDFs for the battery SoC under different operating settings.

Consequently, when the battery SoC nears 0 and/or S_A is close to a potential minimum (occurring if there is no renewable energy supply, e.g. during a windless night, and the eNB power consumption is maximum), traffic shaping shall be chosen to attempt recharging the batteries by choosing a large positive value for the target change in the energy stored in the batteries. This choice would likely imply shaping the traffic to be significantly reduced and should be avoided as much as possible to reduce the effects on video quality and data delay. When the SoC is much larger than 0 and/or S_A is close to its potential maximum value, there is less urgency to recharge the batteries and traffic shaping can be small or unnecessary. To consider all these requirements, it is explained in [15] how the target change in the energy stored in the batteries is calculated as a logistic function of the battery SoC and with the predicted value for S_A used as a parameter. After calculating the target change in the energy stored in the batteries, this result can be converted to the traffic shaping factors for the controlled time period.

Figure 5 compares the CDF of the battery bank SoC for a system with no traffic shaping, one with traffic shaped permanently, and one applying the integrated harvested energy-cellular traffic management technique described in the previous paragraphs. The results were obtained with the same LTE network setup described above and are shown for the first configuration in Fig. 3 and for a new configuration with 10 PV panels, a 2 kW wind turbine and a very large battery bank with capacity to power the base station for two days. As can be seen, when always applying traffic shaping, the proportion of time that the eNB is powered from harvested renewable energy increases from 68 percent to 91 percent for the first configuration, and from 75 percent to 100 percent for the new configuration. Because traffic is always shaped, these improvements come at the cost of permanently reducing the video quality to the minimum acceptable limit value and increasing average data delay to its maximum limit value always (both limit values are chosen during design). The traffic shaping technique described above only controls the traffic when it is needed due to a predicted deficit in harvested renewable energy. With this technique, the proportion of time that the eNB is powered from harvested renewable energy is 79 percent for the first configuration. Compared with no traffic shaping, this is a reduction of approximately 34 percent in the proportion of time that the eNB cannot be powered from renew-

able energy. While the improvements is less than when shaping traffic permanently, in this case the video quality is reduced only a minority of time (the video quality is reduced by more than 3 dBs in peak signal-to-noise ratio, PSNR, less than 30 percent of the time), and average data delay is increased almost imperceptibly (by the limit 30 percent less than 15 percent of the time). The new configuration was chosen to show a case where the traffic shaping technique achieves 100 percent powering from harvested renewable energy with little effect on video and data traffic. In this case, video quality is reduced by more than 3 dBs in PSNR less than 20 percent of the time, and data delay is increased by the limit 30 percent less than 25 percent of the time. For this new configuration, the dimensioning of the PV panels and wind turbines is such that powering 100 percent of the time from renewable energy without traffic shaping is unachievable in practice.

Increasing Cellular Network Resiliency Through Energy Harvesting

Cellular communication has become such an important part of today's society that the resiliency of its infrastructure is an important source of concern. Yet in many recent natural disasters cellular networks have been shown to be vulnerable even when direct damage affects a small percentage of sites (typically fewer than 10 percent) [6, 16]. During disasters, one of the main causes for loss of cellular service affecting base stations is long power outages that may last from a few days to a few weeks. To sustain operation during power outages, cell site infrastructures typically include batteries with capacity for a few hours (usually four hours but in some cases it may reach eight or more hours). However, batteries alone are not sufficient to keep base stations powered during the long power outages observed after disasters. Then the typical solution to reach longer backup times is to use gensets, either fixed or portable, with a fuel tank sufficient for about a day or two of operation. Nevertheless, gensets have a significant vulnerability that reduces wireless network resiliency: to maintain operation they all rely on lifelines, e.g. roads to deliver fuel or natural gas pipelines, that may likely also be affected by the natural disaster just as the power grid is. Figure 6 exemplifies this scenario by showing the condition of Lousiana's Highway 23 a few days after Hurricane Isaac in 2012. Naturally, in these conditions the power grid is inoperable and it becomes impossible to refuel gensets at cell sites that otherwise have all their hardware operable because they had avoided damage by being located on elevated platforms. Moreover, maintaining operation of gensets and/or deploying portable gensets in hundreds of cell sites requires a significant logistical effort and adds traffic to roads that are already in stress conditions due to the disaster.

The previously described SWA architecture prevents the vulnerability presented by gensets because harvested solar or wind energy have the key property that they do not depend on lifelines for operation [17]. Still, in view of the conditions depicted in Fig. 6, it is fair to question whether the renewable energy harvesters could also be damaged during the extreme event. Damage assessments performed during several recent earthquakes and hurricanes suggest that renewable energy sources' performance was good with minimal damage observed mainly in sites installed at ground level and affected by storm surges or tsunamis (which would also damage the cell sites). In these disasters, only one case was documented of damage to a PV solar farm located inland. No damage was observed to any wind turbine.

Figure 6. Louisiana's Highway 23 looking south near Jesuit Bend.

Conclusions

In this article we discussed the use of harvested renewable energy to power cellular base stations as a technique to reduce the carbon footprint of cellular infrastructure and to enable the deployment of cellular service in areas that lack electrification infrastructure. We also acknowledged the major challenges seen for this energy harvesting application, namely the variability in available energy and the large physical footprint of energy harvesters. We argued that to better address these challenges, it is necessary to jointly consider the aspect of electric power generation and availability, together with the dynamic characteristics of the load (base station power consumption). In this regard, we explained a simple approach to save power at cell sites consisting of setting some base stations to sleep mode when allowed by the cellular traffic intensity conditions. Furthermore, we explained how deploying renewable energy harvesters for a cluster of cell sites within a microgrid configuration allows for the tight integration of electric energy and cellular traffic management in what had been called a "sustainable wireless area" (SWA). This tight integration is leveraged in a traffic shaping technique that makes more effective use of renewable energy, enabling longer periods of operation powered from renewable energy with a minimal occasional degradation of the users' QoE. Importantly, the microgrid within the SWA takes advantage of the fact that solar and wind energy harvesters do not need lifelines for operation, resulting in an increase in cellular network resiliency.

This tight integration is leveraged in a traffic shaping technique that makes more effective use of renewable energy, enabling longer periods of operation powered from renewable energy with a minimal occasional degradation of the users' QoE.

ACKNOWLEDGMENT

This material is based on work supported by the National Science Foundation under awards No. CCF-1331788 and ECCS-0845828.

REFERENCES

[1] A. Fehske *et al.*, "The Global Footprint of Mobile Communications: The Ecological and Economic Perspective," *IEEE Commun. Mag.*, vol. 49, no. 8, Aug. 2011, pp. 55–62.

[2] W. Vereecken *et al.*, "Overall ICT Footprint and Green Communication Technologies," *Int'l. Symp. Commun., Control and Signal Processing (ISCCSP)*, Mar. 3–5, 2010, Limassol, Cyprus, pp. 1–6.

[3] M. Webb, "Smart 2020: Enabling the Low Carbon Economy in the Information Age," *The Climate Group Tech. Report*, 2008.

[4] A. P. Bianzinoe *et al.*, "A Survey of Green Networking Research," *IEEE Commun. Surveys and Tutorials*, vol. 14, no. 1, Jan. 2012, pp. 3–20.

[5] S. Vadgama and M. Hunukumbure, "Trends in Green Wireless Access Networks," *IEEE Int'l. Conf. Commun. Wksps. (ICC)*, 5–9 June, 2011, Kyoto, Japan.

[6] A. Kwasinski, "Lessons from Field Damage Assessments about Communication Networks Power Supply and Infrastructure Performance during Natural Disasters with a focus on Hurricane Sandy," *FCC Wksp. Network Resiliency*, Feb. 2013, New York City, NY, USA.

[7] NREL, "2012 Renewable Energy Data Book," http://www.nrel.gov/docs/fy14osti/60197.pdf..

[8] G. Auer *et al.*, "Energy Efficiency Analysis of the Reference Systems, Areas of Improvements and Target Breakdown," *Energy Aware Radio and neTwork tecHnologies (EARTH) Project deliverable D2.3*, document INFSO-ICT-247733 EARTH, (2010).

[9] M. A. Marsan *et al.*, "Towards Zero Grid Electricity Networking: Powering BSs with renewable Energy Sources," *IEEE Int'l. Conf. Commun. Wksps. (ICC)*, 9–13 June, 2013, Budapest, Hungary, pp. 59–601.

[10] Motorola, "Alternative Power for Mobile Telephony Base Stations," *Solutions Paper*, 2007, http://www.motorolasolutions.com/web/Business/Solutions/Technologies/WiMax/Access%20Services%20Network/_Documents/_Static%20Files/6682_MotDoc.pdf.

[11] A. Kwasinski and A. Kwasinski, "Operational Aspects and Power Architecture Design for a Microgrid to Increase the Use of Renewable Energy in Wireless Communication Networks," *IEEE the Int'l. Power Electronics Conf. (IPEC) ECCE-Asia*, 2014, Hiroshima, Japan.

[12] J. Gong *et al.*, "Base Station Sleeping and Resource Allocation in Renewable Energy Powered Cellular Networks," arXiv preprint: 1305.4996, 2013.

[13] T. Han and N. Ansari, "On Optimizing Green Energy Utilization for Cellular Networks with Hybrid Energy Supplies," *IEEE Trans. Wireless Commun.*, vol. 12, no. 8, Aug. 2013, pp. 3872–82.

[14] T. Han and N. Ansari, "Green-Energy Aware and Latency Aware User Association in Heterogeneous Cellular Networks," *IEEE Global Telecommun. Conf. (IEEE GLOBECOM)*, Dec. 2013, Atlanta, GA, USA, pp. 4946–51.

[15] A. Kwasinski and A. Kwasinski, "Traffic Management for Sustainable LTE Networks," *IEEE Global Telecommun. Conf. (IEEE GLOBECOM)*, 8–12 Dec. 2014, Austin, TX, USA, pp. 2520–25.

[16] A. Kwasinski, "Effects of Notable Natural Disasters from 2005 to 2011 on Telecommunications Infrastructure: Lessons from On-Site Damage Assessments," *Proc. INTELEC*, 9–13 Oct., 2011, Amsterdam, the Netherlands, 9 pages.

[17] A. Kwasinski *et al.*, "Availability Evaluation of Micro-Grids for Resistant Power Supply During Natural Disasters," *IEEE Trans. Smart Grid*, vol. 3, no. 4, Dec. 2012, pp. 2007–18.

BIOGRAPHIES

ANDRES KWASINSKI [M'92, SM'11] (axkeec@rit.edu) received the Ph.D. degree in electrical and computer engineering in 2004 from the University of Maryland, College Park, Maryland. He is currently an associate professor at the Department of Computer Engineering, Rochester Institute of Technology, Rochester, NY. He has co-authored more than fifty publications, including the books *Cooperative Communications and Networking* (Cambridge University Press, 2009) and *3D Visual Communications* (Wiley, 2013).

ALEXIS KWASINSKI [M'02, SM'14] (akwasins@pitt.edu) is an associate professor and R.K. Mellon Faculty Fellow in energy at the University of Pittsburgh. Previously he was a faculty member at The University of Texas at Austin, where he reached the rank of associate professor with tenure. He received a Ph.D. from the University of Illinois at Urbana-Champaign. He is specialized in power and communications resilience and leads groups in this area in both the IEEE and the ASCE.

Wireless Powered Communication: Opportunities and Challenges

Suzhi Bi, Chin Keong Ho, and Rui Zhang

ABSTRACT

The performance of wireless communication is fundamentally constrained by the limited battery life of wireless devices, the operations of which are frequently disrupted due to the need of manual battery replacement/recharging. The recent advance in RF-enabled wireless energy transfer (WET) technology provides an attractive solution named wireless powered communication (WPC), where the wireless devices are powered by dedicated wireless power transmitters to provide continuous and stable microwave energy over the air. As a key enabling technology for truly perpetual communications, WPC opens up the potential to build a network with larger throughput, higher robustness, and increased flexibility compared to its battery-powered counterpart. However, the combination of wireless energy and information transmissions also raises many new research problems and implementation issues that need to be addressed. In this article, we provide an overview of state-of-the-art RF-enabled WET technologies and their applications to wireless communications, highlighting the key design challenges, solutions, and opportunities ahead.

INTRODUCTION

Limited device battery life has always been a key consideration in the design of modern mobile wireless technologies. Frequent battery replacement/recharging is often costly due to the large number of wireless devices in use, and even infeasible in many critical applications (e.g., sensors embedded in structures and implanted medical devices). RF-enabled wireless energy transfer (WET) technology provides an attractive solution by powering wireless devices with continuous and stable energy over the air. By leveraging the far-field radiative properties of electromagnetic (EM) waves, wireless receivers could harvest energy remotely from RF signals radiated by an energy transmitter. RF-enabled WET enjoys many practical advantages, such as wide operating range, low production cost, small receiver form factor, and efficient energy multicasting thanks to the broadcast nature of EM waves.

One important application of RF-enabled WET is wireless powered communication (WPC), where wireless devices use harvested RF energy to transmit/decode information to/from other devices. Without being interrupted by energy depletion due to communication usage, WPC is expected to improve user experience and convenience, with higher and more sustainable throughput performance than conventional battery-powered communication. WPC can also be applied in sensors with much lower maintenance cost and enhanced flexibility in practical deployment. Due to the high attenuation of microwave energy over distance, RF-enabled WET is commonly used for supporting low-power devices, such as RFID tags and sensors. However, recent advances in antenna technologies and RF energy harvesting circuits have enabled much higher microwave power to be efficiently transferred and harvested by wireless devices [1]. Therefore, we envision in Fig. 1 that WPC will be an important building block of many popular commercial and industrial systems in the future, including the upcoming Internet of Things/Everything (IoT/IoE) systems consisting of billions of sensing/RFID devices as well as large-scale wireless sensor networks (WSNs). We also envision RF-enabled WET as a key component of the "last-mile" power delivery system, with the smart electrical power gird forming the backbone or core power network.

Before proceeding to the discussion of RF-enabled WET/WPC, it is worth pointing out its relation to another green communication technique, energy harvesting (EH), where wireless devices harness energy from energy sources in the environment not dedicated to powering wireless devices, such as solar power, wind power, and ambient EM radiation. Unlike RF-based EH from ambient transmitters, the energy source of WET is stable and, more importantly, fully controllable in its transmit power, waveforms, and occupied time/frequency dimensions to power the energy receivers. With a controllable energy source, a WPC network (WPCN) could be efficiently built to power multiple communication devices with different physical conditions and service requirements. Besides, with RF-enabled WET, information could also be jointly transmitted with energy using the same waveform. Such a design paradigm is referred to as simultaneous wireless information and power transfer (SWIPT), which has proved to be more efficient in spectrum usage than transmitting

The authors are with the National University of Singapore.

This work is supported in part by the National University of Singapore under Research Grant R-263-000-679-133.

With the continuing decrease of device operating power (as low as a few microwatts for some RFID tags) and the recent application of MIMO technology, which significantly enhances wireless energy transfer efficiency, we can expect more and more important applications of RF-enabled WET in the future.

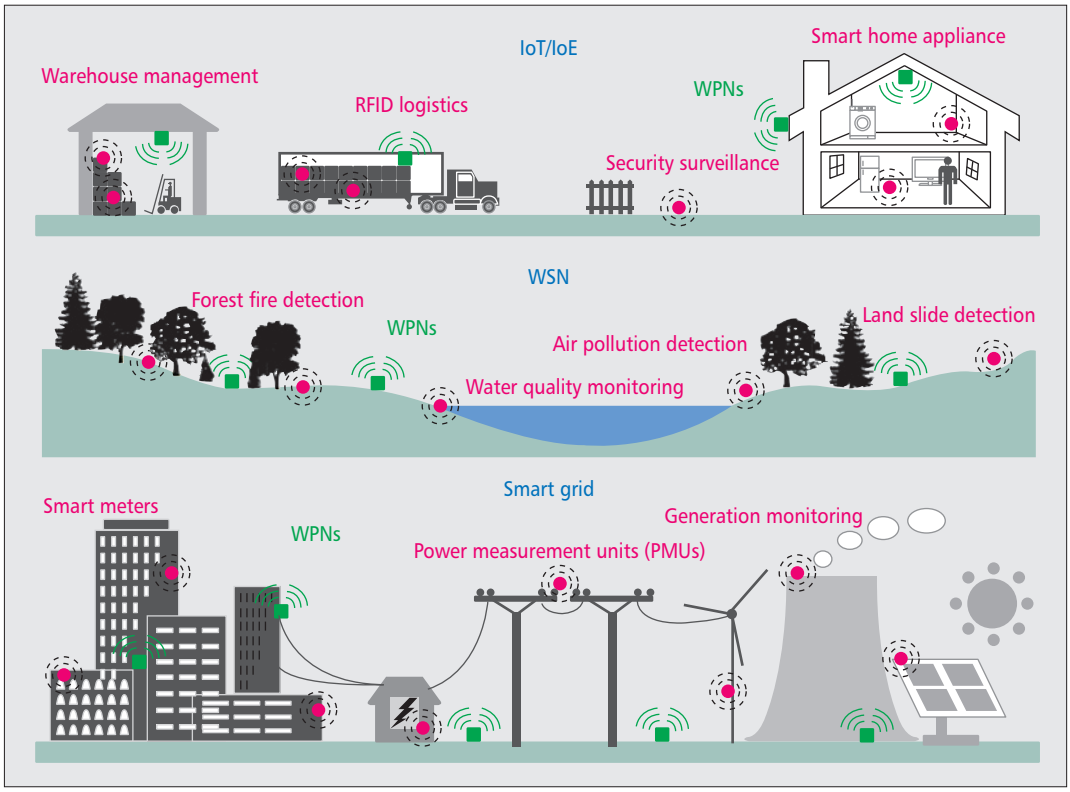

Figure 1. Example applications of WPC in IoT/IoE systems, WSNs for environment monitoring, and smart power grid. The green nodes denote wireless power nodes (WPNs), which transmit RF energy to wireless powered devices, denoted by red nodes in the figure.

information and energy in orthogonal time or frequency channels [2, 3].

In this article, we first provide a brief overview of state-of-the-art RF-enabled WET technologies. Then we focus on introducing RF-enabled WPC in the following three topics:

- The circuit model and advanced signal processing techniques used for WET
- The design trade-offs in joint energy and information transmission for SWIPT
- The design challenges and opportunities in WPCN

At last, we discuss future research directions for WPC and conclude the article.

The State of the Art of RF-Enabled WET

Although WET has gained popularity in recent years, it is in fact a technology that has been under development for more than a century (see [5] for the detailed historical development of WET). The existing WET technologies could be categorized into three classes based on the key physical mechanisms employed: inductive coupling, magnetic resonant coupling, and EM radiation. Among them, the first two types exploit the non-radiative near-field EM properties associated with an antenna for short-range high-power transfer. Currently, inductive coupling WET is well standardized, with applications such as charging mobile phones and implanted medical devices. However, due to the drastic drop in magnetic induction effect over distance, inductive coupling typically operates within a range of only several centimeters. The operating range of magnetic resonant coupling WET could be as large as a few meters. However, to maintain resonant coupling, the receiver may not be flexibly positioned as it is optimized for some fixed distance and circuit alignment settings. Besides, transmitting energy to multiple receivers is challenging as it requires careful tuning to avoid interference due to the mutual coupling effect.

On the other hand, RF-enabled WET exploits the far-field radiative properties of EM waves to power wireless devices over moderate to long distances. A typical RFID tag could be powered from 4 m away (with around 0.5 mW received RF power), and some RF energy harvesting chips have a maximum 12–14 m line-of-sight operating radius (with around 0.05 mW received RF power).[1] In practice, RF-enabled WET uses inexpensive RF energy receivers, which could be flexibly positioned and made very tiny to fit into commercial devices. Besides, transmitting energy to multiple receivers is easily achieved with the broadcasting property of microwaves. The major constraint on the application of RF-enabled WET is the high attenuation of microwave energy over distance. Nonetheless, with the continuing decrease in device operating power (as low as a few microwatts for some RFID tags), and the recent application of multiple-input multiple-output (MIMO) technology, which significantly enhances wireless energy transfer efficiency, we can expect more and more important applications of RF-enabled WET in the future.

[1] Please refer to the website of Powercast Corp. (http://www.powercastco.com) for detailed product specifications.

Network Model for Wireless Powered Communication

In Fig. 2, we present a network model to illustrate the basic concepts of WPC. In the downlink (DL), WET-enabled energy access points (APs) with stable power supply (e.g., AP2 in Fig. 2) transmit energy to a set of distributed wireless devices (WDs). Meanwhile, the WDs could use the harvested energy to transmit/receive information to/from the information APs (e.g., AP3 in Fig. 2) in the uplink (UL) and DL, respectively. Besides, energy and information APs could be integrated into a co-located energy/information AP (e.g., AP1 in Fig. 2), which both transmits energy and provides data access to the WDs. In particular, three canonical operating modes are specified as follows:

- WET: energy transfer in the DL only (e.g., AP1 to WD1 and AP2 to WD5)
- SWIPT: energy and information transfer in the DL (e.g., AP1 to WD4)
- WPCN: energy transfer in the DL and information transfer in the UL (e.g., AP1 to WD3)

Accordingly, the WDs all perform energy harvesting (EH) in the WET mode, with applications such as charging sensors for sensing operations. Additionally, the WDs perform EH in the DL transmission of WPCN mode while sending data in the UL with harvested energy, with applications such as sensor battery charging and data collection in a WSN [3]. For the SWIPT mode, the WDs perform both EH and information decoding (ID) in the DL with the same received signals, each using harvested energy to power its information decoder (e.g., in an energy-self-sustainable information broadcast network) [1, 2]. In practice, a WPC network could also include other general network models consisting of multiple co-located or separated information/energy transmitters, and receivers with heterogeneous operating modes. For instance, WD6 harvests energy from an energy AP (AP2) and transmits data to an information AP (AP3); AP1 transmits energy and information, respectively, to two separated energy and information receivers at the same time (i.e., WD1 and WD2). Besides, energy transmission could potentially generate interference with the information receivers operating in the same frequency band (the red dashed lines in Fig. 2), for which a joint design of energy/information transmissions is highly desired. In the following sections, we focus our discussions on the three canonical operating modes, WET, SWIPT, and WPCN, respectively.

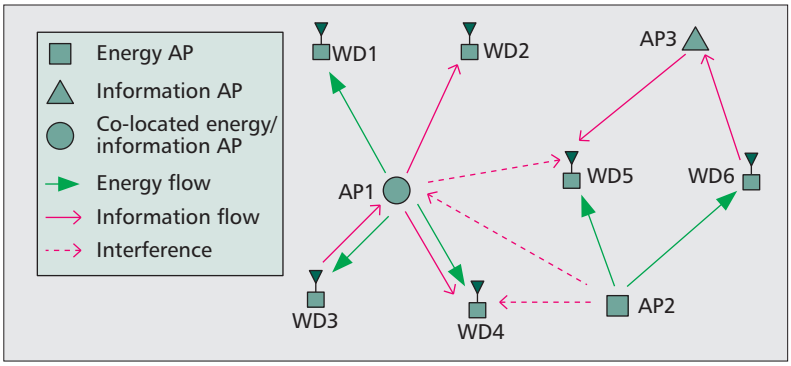

Figure 2. A network model for wireless powered communication.

Wireless Energy Transfer

In this section, we introduce the RF energy receiver structure and advanced signal processing methods to enhance the energy transfer efficiency of RF-enabled WET.

RF Energy Receiver Model

We consider in Fig. 3 an energy transmitter transferring RF energy to multiple energy receivers (ERs), where each transmitter/receiver is equipped with multiple antennas in general. The transmitted energy signals are in general modulated signals (e.g., pseudo-random signals) instead of an unmodulated sinusoid tone, which is used widely in practice. The signals could be designed to avoid a spike in the power spectral density (PSD), and to satisfy the PSD requirement for safety and interference considerations. Therefore, wireless energy transmission occupies a certain bandwidth similar to information transmission, determined by the modulated baseband signal.

The RF energy harvesting (EH) circuit model is also depicted for ER 3 in Fig. 3. The EH receiver is based on a *rectifying circuit*, consisting of a diode and a passive low-pass filter (LPF), which converts the received RF signal to a DC signal to charge the built-in battery, which stores the energy. By the law of energy conservation, the harvested energy per unit symbol time at an ER, denoted by Q, is proportional to the received RF power P_r; that is,

$$Q = \eta \lozenge P_r = \eta \lozenge P_t \lozenge D^{-\alpha} \lozenge G_A. \quad (1)$$

Here, $0 < \eta < 1$ denotes the overall receiver energy conversion efficiency, P_t denotes the transmit power, D denotes the distance (normalized with respect to a given reference distance) from the ER to the transmitter, $\alpha \geq 2$ denotes the path loss factor, and G_A denotes the combined antenna gain of the transmit and receive antennas. For instance, using two antennas at both the energy transmitter and the receiver, we could achieve a beamforming gain to increase the harvested energy by about four times (6 dB) compared to the case with a single-antenna transmitter and receiver with the same transmission power. This could be more cost effective in practice than the alternative approach of improving the energy conversion efficiency, η (say, from 25 to 99 percent) at the receiver with more sophisticated designs of rectifying circuits.

Energy Beamforming

Using multiple antennas not only provides antenna power gain as stated above, but also enables advanced energy beamforming (EB) techniques to focus the transmit power in a smaller region of space to bring significant improvement to energy transfer efficiency [1]. By carefully shaping the transmit waveform at each antenna, EB could control the collective behavior of the radi-

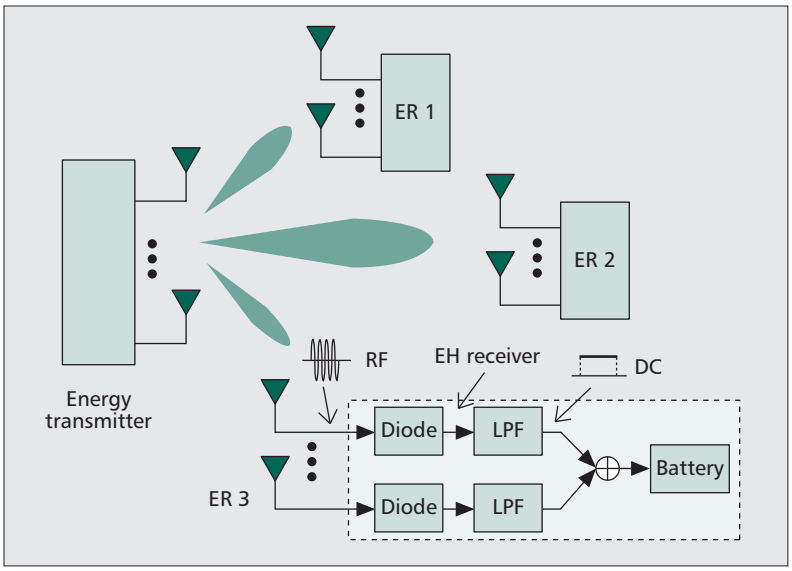

Figure 3. A wireless energy broadcast network and the energy receiver structure. At ER 3, the received RF signal is converted to a DC to charge a battery.

ated waveforms such that they are combined coherently at a specific receiver, but destructively at others. In general, the larger number of antennas installed at the energy transmitter, the sharper the energy beam that could be generated in a particular spatial direction. With only one ER, the transmitter could steer a single sharp beam to maximize the harvested energy. When there are multiple ERs as in Fig. 3, however, generating a single beam may result in severe unfairness among receivers, also known as the (energy) near-far problem, where users near the transmitter harvest much more energy than far users. In this case, the transmitter may need to generate multiple energy beams in different directions to balance the EH performance among the receivers [6].

An efficient EB design requires availability of accurate knowledge of the channel state information at the transmitter (CSIT). However, this is often difficult to achieve in practice. On one hand, many simple energy receivers have no baseband signal processing capability to perform channel estimation. On the other hand, accurate channel estimation consumes a significant amount of time and energy, which may offset the energy gain obtained from a refined EB. Besides, receiver mobility could cause time-varying channels, making channel tracking difficult.

Various efficient channel estimation methods have been proposed to perform EB under imperfect CSIT by exploiting the received energy levels over time and balancing the trade-off between energy consumption and EB gain [7, references therein]. Besides, robust EB design could be pursued to generate energy beams based on the statistical knowledge of CSIT. Another promising method is to perform EB using distributed antennas. This effectively reduces the amount of feedback signals for channel estimation, as a receiver harvests energy only from a small subset of nearby transmitting antennas. Besides, the deployment of distributed antennas also reduces the range between the energy transmitters and receivers, and thus is effective to solve the near-far problem caused by using a single energy transmitter. In this case, however, efficient coordination among the distributed antennas is needed.

Simultaneous Wireless Information and Power Transfer

When WET is applied to power communication devices, it inevitably occupies part of the spectrum used for communication purpose. To avoid the interference to communication, a simple but spectrally inefficient method is to transmit energy and information in orthogonal frequency channels. Alternatively, SWIPT designs seek to save spectrum by transmitting information and energy jointly using the same waveform. This is intuitively achievable, as any waveform for information transmission also carries energy to be harvested by the same or different receivers. However, an efficient SWIPT scheme involves a rate-energy trade-off in both the transmitter and receiver designs to balance the information decoding (ID) and EH performance.

Rate-Energy Tradeoff

In Fig. 4, we illustrate a SWIPT network with a multi-antenna hybrid access point (HAP) transmitting energy and information jointly to multiple receivers (Rxs). Some of the receivers only receive information (Rx 6) or harvest energy (Rx 7), while some do both simultaneously (Rx 1–4). It is worth pointing out that typical ID and EH receivers operate with rather different power sensitivities (e.g., –10 dBm for EH receivers vs. –60 dBm for ID receivers). Therefore, EH receivers are in general closer to the transmitter than ID receivers for effective energy reception.

At the transmitter side, the waveforms generated by the HAP directly determine the performance of information and energy transfer. In the extreme case, the HAP could ignore energy (information) receivers and optimize waveforms only to maximize information (energy) transmission efficiency. However, due to the fundamental difference in the optimal waveforms for information and energy transmissions, such an off-balance design may lead to poor performance of either information or energy transmission. In general, the waveform design needs to follow a *rate-energy trade-off* to achieve the best possible balance between the two objectives [2].

Meanwhile, the characterization of rate-energy trade-off is closely related to the receiver structure and the corresponding signal processing strategies [2, 3]. An ideal SWIPT receiver is assumed to be able to decode information and harvest energy from the same signal [8]; however, this cannot be realized by practical circuits. Some practical receiver structures are plotted in Fig. 4: time switching (TS), power splitting (PS), integrated ID/EH receiver (IntRx), and antenna switching (AS) [1, 2], and are specified later along with the respective rate-energy trade-off characterization.

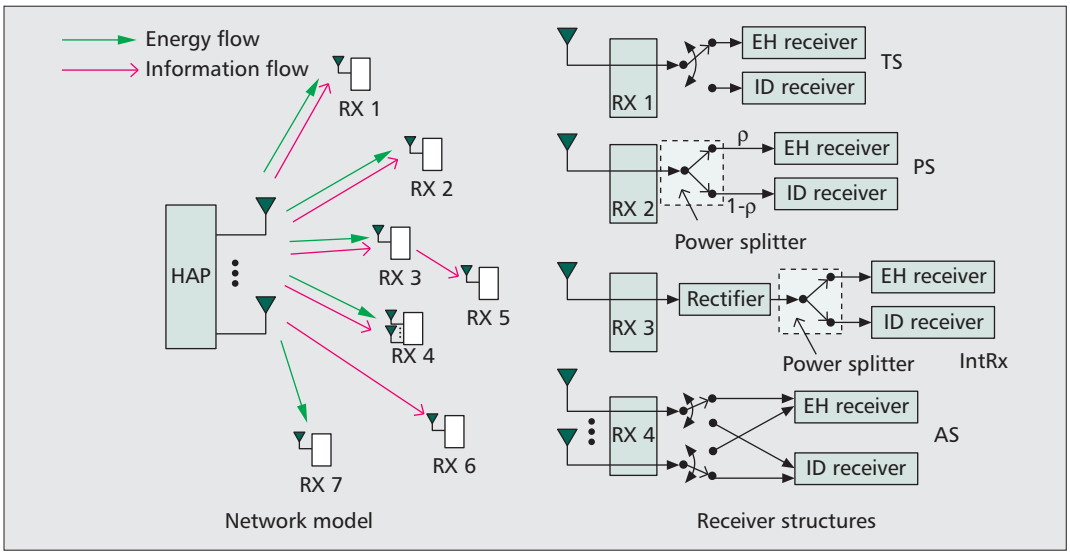

Figure 4. A SWIPT network model and the receiver structures.

When WET is applied to power communication devices, it will inevitably occupy part of the spectrum used for communication purposes. To avoid interference with communication, a simple but spectrally inefficient method is to transmit energy and information in orthogonal frequency channels.

Practical Receiver Structures

For the simplicity of illustration, we consider each transmitter and receiver pair separately to discuss the rate-energy trade-off in a point-to-point channel. In this case, the rate-energy trade-off is often characterized by the boundary of the *rate-energy region*, defined as the union of all the achievable rate-energy pairs by the receiver. Then any point on the boundary specifies the maximum achievable data rate given a harvested energy requirement.

Time Switching Receiver — This corresponds to Rx 1 in Fig. 4. A TS receiver consists of co-located ID and EH receivers, where the ID receiver is a conventional information decoder; the EH receiver's structure follows that in Fig. 3. In this case, the transmitter divides the transmission block into two orthogonal time slots, one for transferring power and the other for transmitting data. At each time slot, the transmitter could optimize its transmit waveforms for either energy or information transmission. Accordingly, the receiver switches its operations periodically between harvesting energy and decoding information between the two time slots. Then different R-E trade-offs could be achieved by varying the length of energy transfer slot.

Power Splitting Receiver — This corresponds to Rx 2 in Fig. 4. The EH and ID receiver components of a PS receiver are the same as those of a TS receiver. However, the HAP cannot optimize transmitted signals only for information or energy. Instead, the PS receiver splits the received signal into two streams, where one stream with power ratio $0 \le \rho \le 1$ is used for EH, and the other with power ratio $(1 - \rho)$ is used for ID. Different R-E trade-offs are achieved by adjusting the value of ρ.

Integrated Receiver — This corresponds to Rx 3 in Fig. 4. Unlike the TS and PS receivers that split the signal at the RF band, an IntRx combines the RF front-ends of ID and EH receivers, and splits the signal after converting it into DC current. Then the DC current is divided into two streams for battery charging and information decoding, respectively. IntRx uses a (passive) rectifier for RF-to-baseband conversion, which saves the circuit power consumed by the active mixer used in the information decoder of TS/PS receivers. However, the ID receiver of IntRx needs to perform noncoherent detection from the baseband signal (DC current). In this case, conventional phase amplitude modulation (PAM) must be replaced by *energy modulation*, where information is only encoded in the power of the input signal, resulting in a reduction of capacity [3]. However, IntRx is superior than PS/TS receivers when more harvested energy is required, because active frequency down conversion is not performed.

In Fig. 5, we give an example to illustrate the key characteristics of the rate-energy regions of the three practical receivers and the ideal receiver in a point-to-point additive white Gaussian noise (AWGN) channel. We can see that the rate-energy region of the ideal receiver is a box; thus, no design trade-off is involved in this case. A similar rate-energy region is also observed for IntRx. This is because the optimal strategy is to use an infinitesimally small amount of DC current for ID and the remaining DC current for EH. The rate-energy region of a TS receiver is a straight line connecting the two optimal operating points for EH and ID, respectively. Compared to a TS receiver, a PS receiver has a strictly larger rate-energy region. So far, the optimal EH-ID receiver is not known. It is unclear whether or not the nontrivial rate-energy region between the ideal and the introduced practical receivers could be achieved, which is left for future exploration likely involving different domains, such as physics, circuit theory, and information theory.

Design Challenges and Opportunities

Channel fading and interference are two major challenges to the transceiver design of wireless communication. However, they result in quite different performance degradations for ID and EH in SWIPT. While deep channel fading degrades the performance of both ID and EH, strong inter-

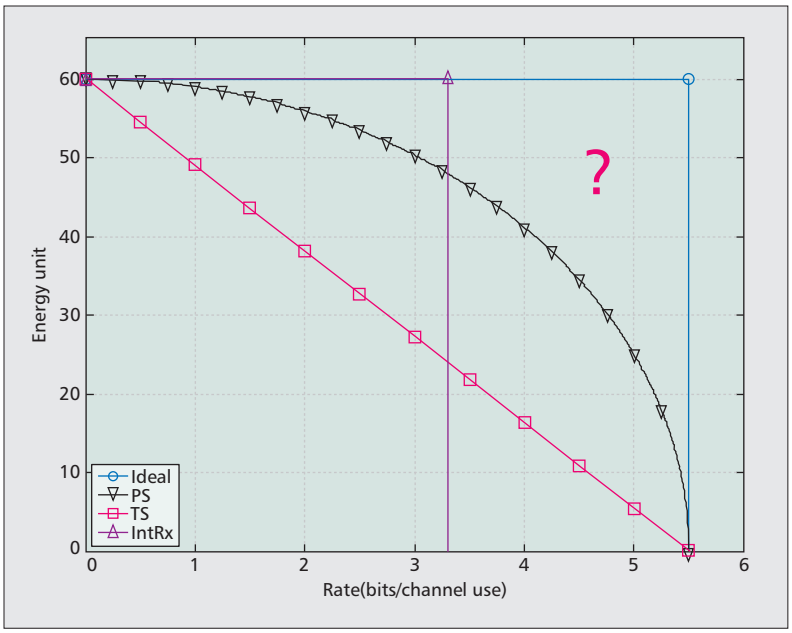

Figure 5. Comparison of rate-energy trade-offs of SWIPT receivers.

ference is only harmful to ID, but in fact could be helpful to increase the harvested energy for an EH receiver. To optimize EH and ID performance, the receiver could adapt its strategy to the channel conditions and interference level [9]. For instance, TS Rx 1 in Fig. 4 should switch to perform ID when the received signal (information and interference) is relatively weak and the signal-to-noise ratio (SNR) is sufficiently high, and EH otherwise. Similarly, PS Rx 2 should allocate more received power to an ID receiver when the channel is in poor condition, and more power to an EH receiver otherwise [10]. Intuitively, this is because the gain achieved by ID receiver is less than the gain achieved by an EH receiver when interference is strong (harmful for ID but helpful for EH), and when the channel is in good condition (logarithmic increase in rate for ID but linear increase in energy for EH). When CSIT is available, the transmitter could also adapt its transmit power to the channel state to achieve the maximum information and energy transfer efficiency (e.g., it does not waste energy to transmit in the case of deep channel fading) [9, 10].

The application of MIMO technology could significantly mitigate the effect of channel fading for both energy transmission (energy beamforming) and information transmission (spatial diversity and/or multiplexing) [1]. In a broadcast channel such as that shown in Fig. 4, a multi-antenna HAP could utilize the spatial degrees of freedom to focus the antenna radiation on specific locations, which not only enhances the harvested energy but also mitigates the interference with unintended information receivers. At the receiver side, the use of multiple antennas enables a low-complexity implementation for PS, *antenna switching* (Rx 4, Fig. 4), which uses a subset of antennas for EH ($\rho = 1$) and the rest for ID ($\rho = 0$). While PS requires a power splitter for each antenna, AS reduces the hardware complexity by simply connecting an antenna to either an ID receiver or an EH receiver with an inexpensive switch. The rate-energy region of AS approaches that of a PS receiver when the number of receive antennas is large enough [10].

SWIPT could also be extended to other useful application scenarios. In Fig. 4, for instance, the HAP could broadcast energy to the nearby receivers (Rx 7) and transmit information to the faraway receivers (Rx 6) simultaneously, while meeting the different sensitivities of EH and ID receivers. Besides, information secrecy could be achieved between the HAP and information receivers using physical layer secrecy coding techniques[6]. In addition, a relay node (Rx 3, Fig. 4) could harvest energy and receive the message dedicated to Rx 5, which is located further away from the HAP, and then forward the message to Rx 5 in another time slot to extend the coverage of the HAP.

Wireless Powered Communication Network

In SWIPT, wireless devices use harvested energy to decode the information sent to them. Here, we consider another scenario in which wireless devices use harvested energy to transmit information. This communication architecture is the WPCN, mentioned earlier [5]. In this section, we introduce the basic operations of a WPCN, the key design challenges and solutions, and interesting extensions to many practical network models.

Harvest-Then-Transmit Protocol

We consider a single-cell WPCN in Fig. 6, where a HAP broadcasts energy to multiple wireless devices in the DL, while the wireless devices communicate to the HAP in the UL using the energy harvested. Due to the half-duplex hardware constraint of HAP, the network operates under a two-phase *harvest-then-transmit* protocol within a transmission block of duration T. Specifically, the wireless devices harvest energy from DL WET in the first phase for a $\tau_0 T$ ($0 < \tau_0 < 1$) amount of time, and then transmit data in the second phase for the rest of the transmission block. This could easily be achieved by the time-switching circuit model shown for U3 in Fig. 6. Intuitively, with a larger τ_0, the UL data rate could be improved as the devices could harvest more energy in the first phase to transmit data. However, a larger τ_0 also decreases the data rate as it leaves a shorter data transmission time. In general, the optimal value τ_0 that results in the highest UL throughput is related to the users' wireless channel conditions. If the users are all close to the HAP, the optimal τ_0 is small, as each user could still harvest a sufficient amount of wireless energy within the short duration of DL WET. Otherwise, a larger τ_0 is required for far users to harvest sufficient energy before commencing reliable data transmission. In fact, it is shown in [4] that the optimal value of τ_0 that maximizes the sum-rate decreases as the sum of channel power gains of all users increases.

Doubly-Near-Far Problem

Besides setting the optimal duration for WET in the first phase, another important issue in the harvest-then-transmit protocol is to design an

efficient multiple access scheme in the second phase for coordinating UL information transmission of users. In a conventional wireless system, users far away from the base station in general achieve lower data rates than those in the vicinity. This fairness issue is even more critical and challenging in a WPCN. Due to significant signal attenuation, a user far away from the HAP (U2) harvests much lower wireless energy in the DL but consumes more to transmit data in the UL than a user near the HAP (U1). This coupled effect is referred to as the *doubly-near-far problem* [3], which could result in very low throughput for far users (e.g., 100 times less data rate than a nearby user) if the multiple access scheme is not properly designed.

When time-division multiple access (TDMA) is used in the second phase, the HAP could allocate a longer data transmission time to the far users in the second phase to tackle the doubly-near-far problem. On the other hand, when the HAP has multiple antennas, space-division multiple access (SDMA) could be applied in the UL. In this case, all the users transmit simultaneously to the HAP during the second phase, and the HAP jointly decodes the user messages using multi-user detection (MUD) techniques. SDMA in general achieves higher spectrum efficiency than the TDMA-based method. Besides, the doubly-near-far problem could be mitigated by user transmit power control in the UL and EB design in the DL. Specifically, the HAP uses EB to steer stronger energy beams toward the far users, and allows them to transmit at higher power than the near users to balance the throughput performance among all users [11].

Another effective method to tackle the doubly-near-far-problem is through user cooperation. In Fig. 6, after harvesting energy in the first phase, a nearby user, U4, uses part of its resources (transmit energy and time) in the second phase to forward a far-away user's (U3) messages to the HAP [12]. The more resources U4 consumes in helping U3, the more throughput improvement could be achieved for the far-away user. Due to the transmit time and energy constraints, relay node U4 needs to carefully allocate the resources for relaying the other's message and transmitting its own message. Interestingly, it is shown in [12] that both users could benefit from the cooperation. For the far-away user, the reason is obvious as the cooperation essentially increases the time and energy used for its message transmission. For the nearby user, its data rate loss due to cooperation could be made up by an overall longer data transmission time, because the gain from user cooperation allows the HAP to allocate more time for data transmission instead of WET.

Extensions

The efficient operation of a WPCN is highly dependent on accurate knowledge of channel state information (CSI) at the HAP, where both information decoding and resource allocation require accurate CSI. Similar to a conventional wireless network, the throughput performance of a WPCN would benefit from a longer channel estimation period with more accurate CSI estimation, but also suffer from a shorter energy/information transmission time. However, a WPCN-specific design trade-off arises due to the energy constraints at the wireless devices. This is because the wireless devices consume energy for channel estimation on decoding the pilot signals sent by the HAP, transmitting the CSI feedback, or sending pilot signals to the HAP in some channel estimation schemes that exploit the UL/DL channel reciprocity [7]. Evidently, more energy consumption on channel estimation would compromise the transmission rate (or reliability) because less energy is left for communication. However, this results in more accurate CSI estimation and hence more precise beamforming that both improves its transmission rate (or reliability) and increases the harvested energy in the following transmission blocks. The energy of wireless devices for channel estimation and information transmission should thus be carefully allocated to achieve optimal performance.

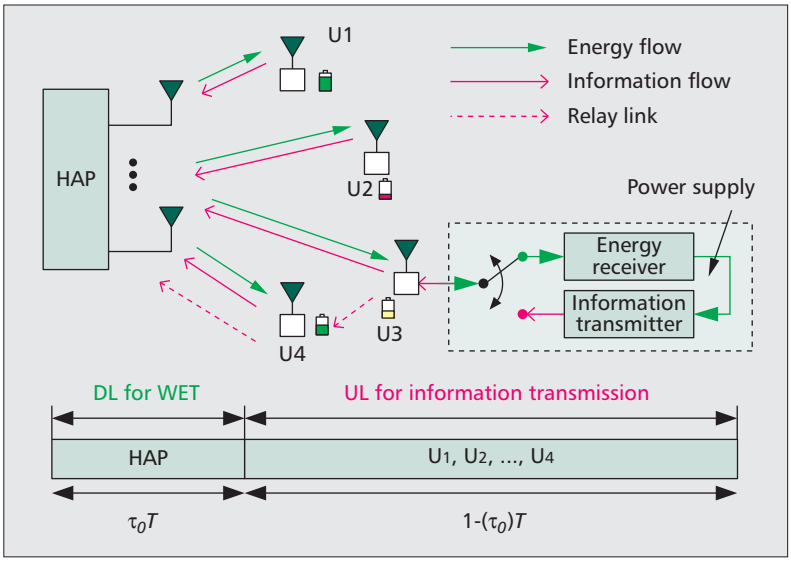

Figure 6. A WPCN model and the harvest-then-transmit protocol.

Another performance enhancing technique for a WPCN is massive MIMO, which employs a large number of antennas (tens to more than 100) at the HAP to exploit the high antenna array gain. On one hand, the large degree of freedom provided by massive MIMO enables spatial multiplexing to serve more mobile users for UL information transmission at the same time. On the other hand, the large antenna array also improves the EB performance in the DL by generation of very sharp beams to enhance the received signal power (e.g., more than 20 dB power gain). The application of massive MIMO to WPCN could result in multifold throughput improvement and also much longer operating range. Interestingly, the high number of antennas at a HAP does not necessarily translate to high processing complexity and cost. The property of asymptotic orthogonality of large user-to-HAP channels leads to largely simplified beamforming design, multiple access control, and power control solutions [13].

More generally, there could be a dedicated wireless energy network consisting of multiple

Hardware prototyping is urgently needed to evaluate the feasibility of WPC, and to test the applications of various technologies in joint energy and information transmissions, such as massive MIMO, millimeter wave, and distributed antenna systems.

power nodes that broadcast energy by means of WET. Moreover, multiple information receivers may decode the data transmitted by the wireless devices in the UL. In such a multi-cell WPCN (i.e., multiple energy transmitters and information receivers), it is complicated and often intractable to derive the network capacity using the resource allocation methods for single-cell capacity analysis. Instead, the method of *stochastic geometry* is widely used to study the scaling laws of network capacity as a function of system parameters. In the context of a cellular WPCN, [14] proposes to install wireless energy transferring nodes, called power beacons (PBs), to provide energy for mobile users to transmit data to some base stations (BSs). Based on a stochastic geometry model, it derives the functional relationships between the densities of BSs and PBs as well as their transmission power to achieve a prescribed communication outage probability. The application of WPCN is also exploited in a cognitive radio network in [15], where a secondary transmitter (e.g., sensor) could harvest energy from a primary transmitter (e.g., mobile phone) if they are close enough, and transmit to its intended secondary receiver if it is sufficiently far from any primary transmitter to avoid potential interference to the primary network.

Other than the extensions discussed above, a WPCN could also be applied to many wireless systems with energy constrained wireless devices; for instance, multihop communications with EH relays, systems with densely deployed HAPs using millimeter-wave technologies, and distributed antenna systems with coordinating energy/information beamforming.

Future Research Directions

WPC contains rich research problems of important applications yet to be studied. In this section. we highlight several interesting research topics we deem particularly worth investigating.

Energy and Information Transfer Coexistence

Due to the critical power constraints of wireless devices, the future wireless system is expected to be a mixture of wireless energy and communication networks. Under spectrum scarcity, it is likely that the two networks operate on overlapped spectrum. This raises problems in the coexistence of wireless energy and communication networks. Unlike the two-way interference in conventional multi-cell communication systems, the interference is one-way from an energy network to a communication network. Besides, the highly different sensitivity of information and energy receivers indicates that the interference due to WET is in general much stronger than information signals. There are many promising solutions to mitigate the interference, such as information/energy transfer scheduling, EB design, and opportunistic WET based on spectrum sensing. In particular, cognitive radio technology could be used to carry out effective spectrum sensing to minimize the interference from WET with a communication network.

Cross-Layer Design

So far we mainly focus on the physical (PHY) layer techniques to optimize the performance of WPC. In a practical system, medium access control (MAC) plays the key role in determining the fairness and efficiency of the system. An efficient wireless system design often takes a cross-layer approach, especially for the closely related PHY and MAC layers. In the context of WPC, an example of cross PHY-MAC design is for the HAP to steer the energy beam toward a user with a relatively strong wireless channel and many data packets backlogged in the queue, rather than considering the physical channel condition alone. Besides, efficient energy scheduling should also consider the residual battery life, wake up/sleep schedule, and expected energy consumption of all wireless devices.

Hardware Implementation

The current studies on WPC are mainly theoretical in nature. The achievable throughput performance using off-the-shelf EH and communication modules in a practical wireless environment is not known. Hardware prototyping is urgently needed to evaluate the feasibility of WPC, and to test the applications of various technologies in joint energy and information transmissions, such as massive MIMO, millimeter wave, and distributed antenna systems. An extensive testbed could also help identify the most suitable technology and proper application scenarios for WPC.

Health and Safety

With the potential use of massive MIMO and advanced beamforming technologies, the intensity of microwave in a particular area could be strong enough to harm human health and cause safety issues. In practice, the radiation power of any wireless device must satisfy the equivalent isotropically radiated power (EIRP) requirement on its operating frequency band; for example, the Federal Communications Commission (FCC) permits a maximum 36 dBm EIRP on 2.4 GHz band. One promising method to solve the safety problem is to use a distributed antenna system, such that for each antenna the radiation is omnidirectional and relatively weak (thus satisfying the given EIRP constraint), while the combined effect is constructive only at the destined location but destructive almost everywhere else. This will reduce the risk of "radiation burn" due to human blockage in a random direction. Besides, the distributed antenna system could be combined with advanced sensing technology to detect the presence of humans in real time, and cease energy transmission if it deems the transmission to be harmful.

Conclusions

In this article, we have provided an overview of state-of-the-art RF-enabled WET technologies and their applications to wireless communications. Promisingly, wireless powered communications could significantly improve on its battery-powered counterpart, and be practically achieved using simple and inexpensive transceiver structures. The opportunities and challenges

in the design of wireless powered communications were demonstrated by studying two new paradigms: SWIPT and WPCN. We hope that the design of wireless powered communications will spur research innovations in wireless technologies, as future wireless systems are expected to be a mixture of wireless information and energy transfer, with RF-enabled WET, SWIPT, and WPCN being important building blocks.

REFERENCES

[1] Y. H. Suh and K. Chang, "A High-Efficiency Dual-Frequency Rectenna for 2.45-and 5.8-GHz Wireless Power Transmission," *IEEE Trans. Microwave Theory and Techniques*, vol. 50, no. 7, July 2002, pp. 1784–89.

[2] R. Zhang and C. K. Ho, "MIMO Broadcasting for Simultaneous Wireless Information and Power Transfer," *IEEE Trans. Wireless Commun.*, vol. 12, no. 5, May 2013, pp. 1989–2001.

[3] X. Zhou, R. Zhang, and C. K. Ho, "Wireless Information and Power Transfer: Architecture Design and Rate-Energy Tradeoff," *IEEE Trans. Commun.*, vol. 61, no. 11, Nov. 2013, pp. 4754–67.

[4] H. Ju and R. Zhang, "Throughput Maximization in Wireless Powered Communication Networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 1, Jan. 2014.

[5] N. Shinohara, "Power Without Wires," *IEEE Microwave Mag.*, vol. 12, no. 7, Dec. 2011, pp. 564–73.

[6] L. Liu, R. Zhang, and K. C. Chua, "Secrecy Wireless Information and Power Transfer with MISO Beamforming," *IEEE Trans. Signal Proc.*, vol. 62, no. 7, Apr. 2014, pp. 1850–63.

[7] Y. Zeng and R. Zhang, "Optimized Training Design for Wireless Energy Transfer," *IEEE Trans. Commun.*, vol. 63, no. 2, Feb. 2015, pp. 536–50.

[8] L. R. Varshney, "Transporting Information and Energy Simultaneously," *Proc. IEEE ISIT*, July 2008, pp. 1612–16.

[9] L. Liu, R. Zhang, and K. C. Chua, "Wireless Information Transfer with Opportunistic Energy Harvesting," *IEEE Trans. Wireless Commun.*, vol. 12, no. 1, Jan. 2013, pp. 288–300.

[10] L. Liu, R. Zhang, and K. C. Chua, "Wireless Information and Power Transfer: A Dynamic Power Splitting approach," *IEEE Trans. Commun.*, vol. 61, no. 9, Sept. 2013, pp. 3990–4001.

[11] L. Liu, R. Zhang, and K. C. Chua, "Multi-Antenna Wireless Powered Communication with Energy Beamforming," *IEEE Trans. Commun.*, vol. 62, no. 12, Dec. 2014, pp. 4349–61

[12] H. Ju and R. Zhang, "User Cooperation in Wireless Powered Communication Networks," *Proc. IEEE GLOBECOM*, Dec. 2014.

[13] G. Yang *et al.*, "Throughput Optimization for Massive MIMO Systems Powered by Wireless Energy Transfer," to appear, *IEEE JSAC*, available online at arXiv:1403.3991

[14] K. Huang and V. K. N. Lau, "Enabling Wireless Power Transfer in Cellular Networks: Architecture, Modeling and Deployment," *IEEE Trans. Wireless Commun.*, vol. 13, no. 2, Feb. 2014, pp. 902–12.

[15] S. Lee, R. Zhang, and K. Huang, "Opportunistic Wireless Energy Harvesting in Cognitive Radio Networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 9, Sept. 2013, pp. 4788–99.

BIOGRAPHIES

SUZHI BI (bsz@nus.edu.sg) received his B.Eng. degree in communications engineering from Zhejiang University, Hangzhou, China, in 2009. He received his Ph.D degree in information engineering from the Chinese University of Hong Kong in 2013. He is currently a research fellow in the Department of Electrical and Computer Engineering at the National University of Singapore. His current research interests include wireless information and power transfer, medium access control in wireless networks, and smart power grid communications.

CHIN KEONG HO (hock@i2r.a-star.edu.sg) received B.Eng. and M. Eng degrees from the National University of Singapore. He received his Ph.D. degree from Eindhoven University of Technology, The Netherlands, where he concurrently conducted research work at Philips Research. He is the laboratory head of the Energy-Aware Communications Lab of the Institute for Infocomm Research, ASTAR. His research interest includes green wireless communications with focus on energy-efficient solutions and with energy harvesting constraints, and implementation aspects of multi-carrier and multi-antenna communications.

RUI ZHANG (elezhang@nus.edu.sg) received B.Eng. and M.Eng. degrees from the National University of Singapore and his Ph.D. degree from Stanford University, all in electrical engineering. He is now an assistant professor with the Department of Electrical and Computer Engineering at the National University of Singapore. His current research interests include multiuser MIMO, cognitive radio, energy-efficient and energy harvesting wireless communication, and wireless information and power transfer.

We hope that the design of wireless powered communications will spur research innovations in wireless technologies, as the future wireless systems is expected to be a mixture of wireless information and energy transfer, with RF-enabled WET, SWIPT and WPCN being important building blocks.

Fundamental Limits of Energy Harvesting Communications

Omur Ozel, Kaya Tutuncuoglu, Sennur Ulukus, and Aylin Yener

ABSTRACT

Wireless networks composed of energy harvesting devices will introduce several transformative changes in wireless networking as we know it: energy self-sufficient, energy self-sustaining, perpetual operation; reduced use of conventional energy and accompanying carbon footprint; untethered mobility; and an ability to deploy wireless networks in hard-to-reach places such as remote rural areas, within structures, and within the human body.

Energy harvesting brings new dimensions to the wireless communication problem in the form of intermittency and randomness of available energy, which necessitates a fresh look at wireless communication protocols at the physical, medium access, and networking layers. Scheduling and optimization aspects of energy harvesting communications in the medium access and networking layers have been relatively well-understood and surveyed in the recent paper [1]. This branch of literature takes a physical layer rate-power relationship that is valid in energy harvesting conditions under large-enough batteries and long-enough durations between energy harvests so that information-theoretic asymptotes are achieved, and optimizes the transmit power over time in order to maximize the throughput.

Another branch of recent literature aims to understand the fundamental capacity limits, i.e. information-theoretic capacities, of energy harvesting links under smaller scale dynamics, considering energy harvests at the channel use level. This branch necessitates a deeper look at the coding and transmission schemes in the physical layer, and ultimately aims to develop an information theory of energy harvesting communications, akin to Shannon's development of an information theory for average power constrained communications.

In this introductory article, we survey recent results in this branch and point to open problems that could be of interest to a broad set of researchers in the fields of communication theory, information theory, signal processing, and networking. In particular, we review capacities of energy harvesting links with infinite-sized, finite-sized, and no batteries at the transmitter.

Omur Ozel and Sennur Ulukus are with the University of Maryland.

Kaya Tutuncuoglu and Aylin Yener are with The Pennsylvania State University.

This work was supported by NSF Grants CNS 09-64632 / CNS 09-64364, and CCF 14-22111 / CCF 14-22347.

INTRODUCTION

Energy harvesting devices offer several significant advantages over conventional grid-powered and non-rechargeable battery-powered devices [2]. These advantages include energy self-sufficient, energy self-sustaining operation with lifetimes limited only by the lifetimes of their hardware. These devices utilize energy harvested from alternative natural resources such as solar, vibrational, electromagnetic, thermoelectric, etc. As such, their widespread adoption will reduce the use of conventional energy and the accompanying carbon footprint, and benefit the environment. In addition, these devices do not require conventional recharging, enabling untethered mobility, and deployment of these devices in hard-to-reach places such as remote rural areas, within structures, and within the human body.

The circuits and devices side of engineering has been contributing to the development of energy harvesting devices for decades. However, on the communications, networking, and systems side of engineering, the focus has been on energy-conscious communication system design, in the form of optimum average-power constrained communications, and energy-efficient networking. Such approaches have aimed to minimize total energy usage for a fixed performance, or maximize the performance for a fixed total energy availability. Only recently, communications subject to explicit energy harvesting conditions has garnered attention. In such systems, energy needed for data transmission becomes available (arrives at the transmitter) intermittently and in random amounts, due to the randomness in energy harvesting processes and sources. Then the goal becomes not only to minimize the overall energy consumption, but to maintain reliable communication under random and intermittent energy arrivals. Energy is an essential component of communication: in order to create a desired physical change at the receiving end of a channel, the transmitter needs to use (send) energy. In an energy harvesting system, there is uncertainty in this very basic component of communication, and the communication mechanisms need to be designed by explicitly accounting for these energy harvesting constraints.

Such explicit energy harvesting constraints

0163-6804/15/$25.00 © 2015 IEEE

have been introduced in the context of data scheduling and transmit power optimization in [3–6]; see also [1] for a review of recent results in this branch of literature. This branch of literature utilizes a concave rate-power relationship, such as the logarithmic Shannon capacity formula for the Gaussian channel, and optimizes the transmit power over time, subject to energy harvesting constraints, to maximize the throughput. Such an approach is valid under large-enough battery sizes and long-enough durations between energy harvests, which are valid abstractions in the medium access and networking layers. Such approaches are also valid approximations in the physical layer when the communication modulation and coding scheme is fixed to a potentially sub-optimum scheme. A departure from this approach and an important research direction is to determine the ultimate information-carrying capacity of an energy harvesting link under intermittent and random energy arrivals. While the capacity of an average power-constrained Gaussian channel is well known, what is the capacity of an energy harvesting link that harvests energy in the amount of E_i in channel use i with an average recharge rate of $E[E_i]$?

This question necessitates a deeper look at the explicit coding and transmission schemes in the physical layer. It also necessitates consideration of explicit interactions between energy arrivals, energy storage in the battery, and energy used by the codebook, for data transmission. In particular, the battery can be viewed as an energy queue where the incoming energy is saved for future use. The codebook, which transmits data, can be viewed as a server that serves energy out of this energy queue. On the other hand, energy arrives over time and the amount of available energy in the energy queue dictates the set of feasible codebooks. This introduces interactions between quantities that have been traditionally thought of separately, i.e. energy (which is physical) and data (which is cyber) interact explicitly at every time instant, as opposed to in some average sense as in the classical literature.

In an energy harvesting system, every transmitted symbol is instantaneously constrained by the energy available in the battery. This constraint, which is imposed at every channel use, is different than classical average-power and constant-amplitude constraints, and as such, results in a new form of channel input constraint. From an information-theoretic point of view, the amount of energy available in the battery may be viewed as a *state*, which is naturally causally known by the transmitter, but unknown by the receiver. Even when energy arrivals are independent and identically distributed (i.i.d.), code symbols cannot be generated independently due to the presence of a finite-sized battery and highly time-correlated nature of the battery state. In addition, the transmitter's own actions (past transmitted symbols) affect the future of the state. The goal of this article is to review this branch of literature at an introductory level and present open problems. In particular, we will see that the capacity of an energy harvesting link heavily depends on the size of the battery. We will review results for the cases of infinite-sized batteries, zero-sized batteries, and finite-sized batteries. We start with a brief summary of the results and developed insights in the scheduling and optimization side of the literature.

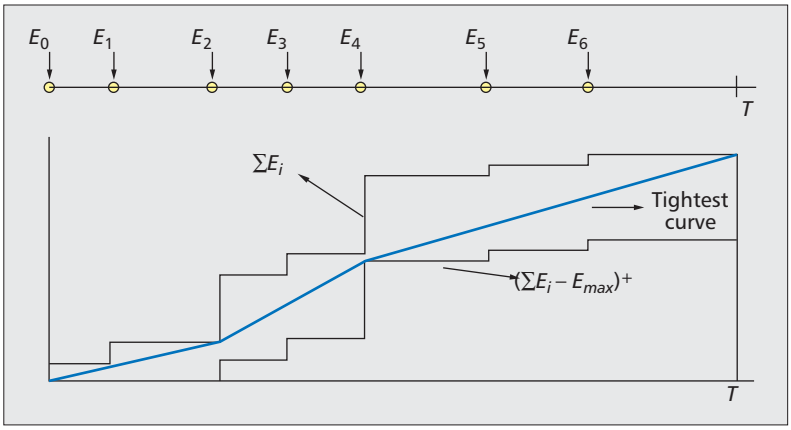

Figure 1. Optimal transmission policy is the *tightest curve* (or *shortest string*) in the *energy feasibility tunnel*.

Data Scheduling and Optimization in Energy Harvesting Communications

Data scheduling in energy harvesting communications is studied under two main categories: offline scheduling and online scheduling. Offline refers to the availability of knowledge of events, such as energy arrivals and channel fade level changes, prior to the start of data scheduling; online refers to the availability of this knowledge only causally over time, but not a priori. In this section we summarize offline scheduling; a more thorough review may be found in [1].

A key assumption in data scheduling literature [3–6] is that a rate-power relation, $r(p)$, determines the rate achieved when the transmission power is chosen as p; a typical example of $r(p)$ is the Shannon capacity formula for the Gaussian channel. Note that the Shannon capacity formula is a monotonically increasing concave function of p, and the development in this literature is valid for all monotonically increasing concave rate-power relationships. Energy arrives from an exogenous energy source over time and is saved in the battery before being utilized for data transmission. Since energy cannot be utilized before it arrives (is harvested), the transmission power has to obey the *energy causality* constraint [3], which states that the energy utilized until a time instant must be less than or equal to the energy harvested by that time instant. In addition, since the battery has a finite-size $E_{\max}$, energy may overflow if the battery does not have a sufficiently large space to accommodate it. Under offline knowledge of energy arrivals, the transmitter has to determine the transmit powers so that all incoming energies are accommodated in the battery, i.e. no energy is wasted; this constraint is called the *no-energy-overflow* constraint [4, 5].

The work in [3] and [4] determined the optimal transmission schemes for an energy harvest-

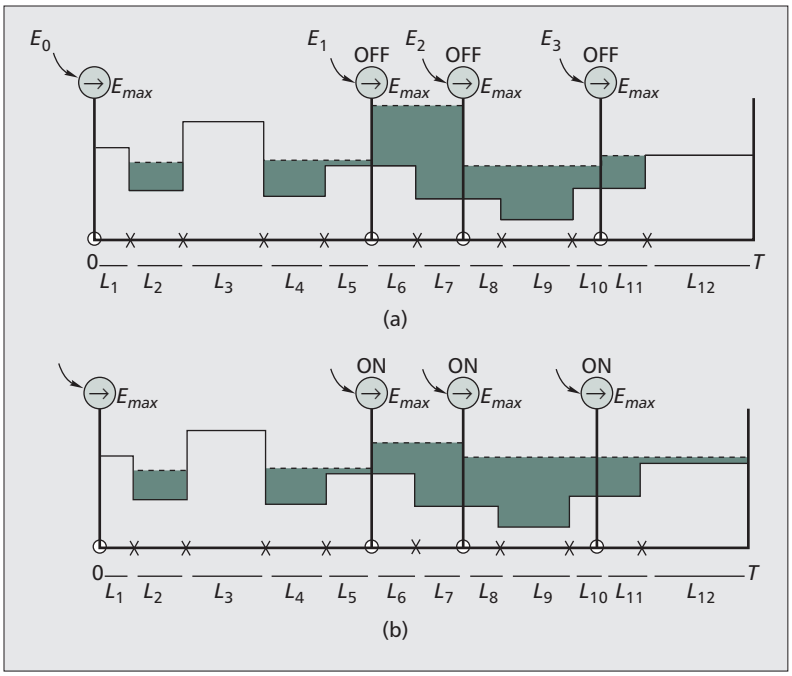

Figure 2. Directional water-filling algorithm: a) first, allocate arriving energy until next energy arrival; and b) then, allow energy to flow right until a balance is reached.

ing transmitter in a non-fading channel with infinite and finite batteries, respectively. The solution has a *shortest string* (or *tightest curve*) structure through an *energy feasibility tunnel*, as shown in Fig. 1, because the length of a curve is a convex function of its slope. The minimization of this convex function may be viewed as maximization of its negative, and the solution of the throughput maximization problem is invariant to the form of the objective function, so long as it is concave and monotone increasing. The upper curve that determines the energy feasibility tunnel represents the cumulative energy arrivals: the energy expenditure of any transmission policy should be below this curve due to *energy causality*. The lower curve in the energy feasibility tunnel is a vertically shifted version of the cumulative energy arrival curve by exactly $E_{\max}$, which is the size of the battery, since at most $E_{\max}$ units of energy can be saved in the battery. The energy expenditure of any transmission policy should be above this curve due to *no-energy-overflow*. Because of the concavity of the rate-power function, keeping the power as constant as possible subject to energy feasibility constraints is optimal. Therefore, the optimal transmit power policy has the *shortest string* structure: the power is kept constant to the extent possible subject to energy causality and no-energy-overflow conditions. The optimal transmission power sequence is the sequence of slopes of the shortest string that lies within the energy feasibility tunnel.

In [5] the solution of [3, 4] is extended to the case of fading channels. The solution is found by a *directional water-filling algorithm*, depicted in Fig. 2. This algorithm is based on a view of energy as water. The incoming energy (water) is allocated (poured) into the time interval till the next energy arrival, as shown in Fig. 2a. Then the water is allowed to flow through the taps from left to right (from the past to the future) due to energy causality (Fig. 2b). This signifies that energy can be saved and used in the future, but energy that has not arrived yet cannot be used. The extent of this energy flow is also limited by the battery limit $E_{\max}$: the amount of water that can flow through each tap is $E_{\max}$. The flow continues till the water levels are balanced. This is a generalization of the classical water-filling algorithm to the case of energy harvesting setting, which introduces a directionality (due to energy causality) and maximum-flow (due to no-energy-overflow) constraints. Note that the slopes in the tightest string interpretation for the static channel correspond to the water levels. See also [6] for a staircase water-filling algorithm.

The literature on data scheduling in energy harvesting communications, such as [3–6] and those that follow them, assumes a monotonically increasing concave rate-power relation $r(p)$ and optimizes the sequence of transmit powers subject to energy causality and no-energy-overflow constraints. However, exact capacities and the resulting rate-power relationships (if any) are not known in energy harvesting settings with dynamics occurring at the channel use level. In fact, we will see that the rate-power relationship highly depends on the size of the available battery $E_{\max}$. Starting in the next section, our goal will be to determine this relationship when possible. We will start by describing the detailed system model at the channel use level.

System Model

As shown in Fig. 3, the transmitter wishes to send a message W in n channel uses so that the receiver reliably decodes the message. $E_{\max}$ is the battery energy storage limit, i.e. the size of the battery. Channel input and output symbols may be discrete or continuous. The exogenous energy source supplies E_i units of energy in the ith channel use, where E_i is an i.i.d. sequence with an average value $E[E_i]$, which we denote as P_{avg}. When channel input X_i is transmitted in the ith channel use, the receiver gets Y_i. The underlying physical communication channel is memoryless.

S_i denotes the energy available in the battery at channel use i. The transmitter observes the available battery energy S_i and transmits a symbol X_i. At each channel use, the transmitter both harvests energy and transmits a symbol. The order of harvesting and transmission in a channel use could be in two different methods. The first method is *transmit first* where the energy of the channel symbol is constrained by the battery energy in that channel use. After sending the symbol, the transmitter harvests energy. Then incoming energy E_i is stored in the battery if there is sufficient space; otherwise, only $E_{\max}$ units of energy is stored. The second method of ordering the energy harvest is *simultaneous transmit*, where energy can be utilized for data transmission in the same channel use as it enters the system. Then remaining energy is stored in the battery if there is sufficient space; otherwise, only $E_{\max}$ units are stored. In general *transmit*

first and *simultaneous transmit* define two different systems and they yield different achievable rates. In both cases the next battery state S_{i+1} is updated as $\min\{S_i - X_i^2 + E_i, E_{\max}\}$. Note that the battery state information, S_i, is available at the transmitter only.

The constraints that force the current symbol to have energy less than or equal to the currently available energy are called *energy causality* constraints on the channel input at the channel use level. These constraints follow from the fact that energy cannot be utilized before it enters the system. Note that a non-trivial time correlation arises in the transmitted input sequence X_i due to the energy causality constraints. We also note that the energy arrivals E_i or the energy available at the battery S_i could be viewed as the state of the channel, which determines the set of feasible symbols that can be transmitted.

Channel Capacity with Infinite Energy Storage

In this section we focus on the case when $E_{\max}$ is set to infinity in a Gaussian channel. An immediate upper bound for the channel capacity C is the corresponding Shannon capacity with average power constrained to the average recharge rate P_{avg}. This bound holds due to the fact that each codeword satisfying the energy causality constraints automatically satisfies the average power constraint. The work in [7] established that this upper bound is, indeed, achievable and provided two different achievability schemes. We now consider these schemes.

Save-and-Transmit Scheme

In the save-and-transmit scheme, data transmission is performed in two phases, as shown in Fig. 4. The saving phase consists of the first $h(n)$ channel uses; the transmission phase consists of the following $n - h(n)$ channel uses. In the saving phase the transmit symbols are set to 0. Therefore, no energy is spent and the battery is fueled with harvested energy. In the transmission phase, information carrying code symbols are sent. We consider such $h(n)$ functions that grow to infinity sublinearly with n; that is, $h(n)$ is in $o(n)$ and grows to infinity as n grows. The remaining $n - h(n)$ code symbols are the information carrying symbols, which are selected from a Gaussian distribution with zero mean and average power smaller than but arbitrarily close to P_{avg}. Since $h(n)$ is in $o(n)$, the saving phase allows a sufficient number of channel uses for the data transmission phase so that the rate hit due to the saving phase is asymptotically zero. Moreover, since $h(n)$ grows to infinity, the energy saved in the saving phase is sufficient to send the designed code symbols without any energy shortages. This scheme achieves rates arbitrarily close to the Shannon capacity with average power constraint equal to the average recharge rate P_{avg}.

Best-Effort-Transmit Scheme

An alternative single-phase scheme that attains the capacity is the best-effort-transmit scheme. This scheme runs as follows. Let $X_1, X_2, \ldots, X_n$ be a codeword and select X_i as i.i.d. Gaussian

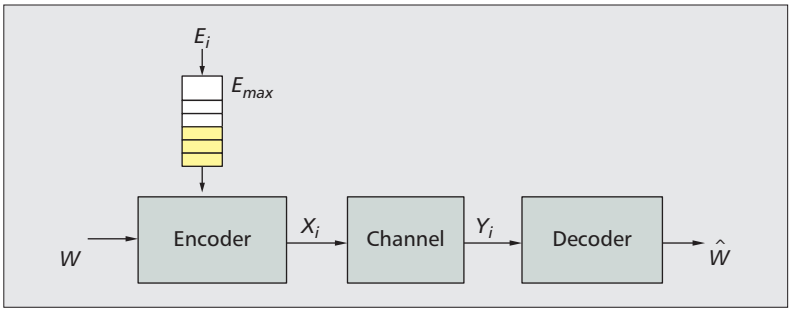

Figure 3. Information-theoretic model of point-to-point communications with an energy harvesting transmitter.

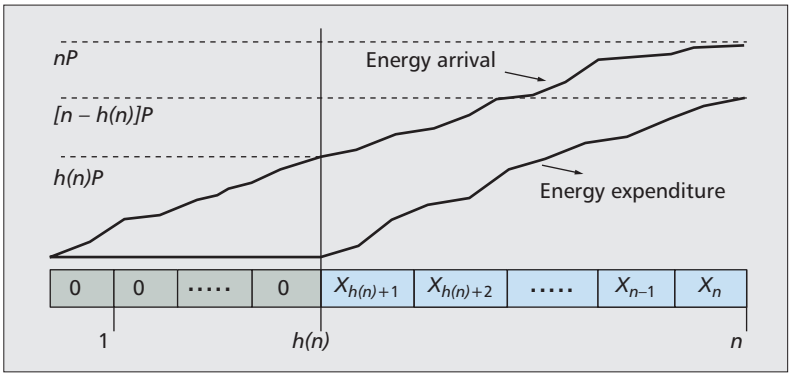

Figure 4. The code symbol energy expenditure and cumulative energy arrival curves in the save-and-transmit scheme.

with mean zero and variance smaller than but arbitrarily close to P_{avg}. We say that the symbol is infeasible if there is not sufficient energy to send it, i.e. if the battery energy S_i is less than X_i^2. In this case, the input to the channel is set to 0, otherwise it is X_i. That is, in the transmitted codeword, some of the symbols in the actual codeword in the codebook are replaced with zeros. This causes a mismatch between the encoder and the decoder. However, the resulting mismatch is proved to be negligible in [7], and communication with rates arbitrarily close to the Shannon capacity with average power constraint equal to the average recharge rate P_{avg} is possible.

Comments on the Schemes

In the save-and-transmit scheme, the available channel uses are divided into two phases. The saving phase duration $h(n)$ is selected to have a sublinear growth to infinity. For example, $h(n)$ can be selected as $\log(n)$ or $\sqrt{n}$. This, along with the unlimited sized battery, allows averaging out the uncertainty in the available energy. Remaining $n - h(n)$ channel uses are used for channel coding with an average power constraint equal to the average recharge rate. In contrast, in the best-effort-transmit scheme, transmission starts right away and code symbols are put to the channel only when it is feasible to do so. As long as the average energy expenditure of the codewords are below the average recharge rate, any infeasibilities become negligible asymptotically and reliable decoding is possible.

It is crucial to note that both save-and-transmit and best-effort-transmit schemes need unlim-

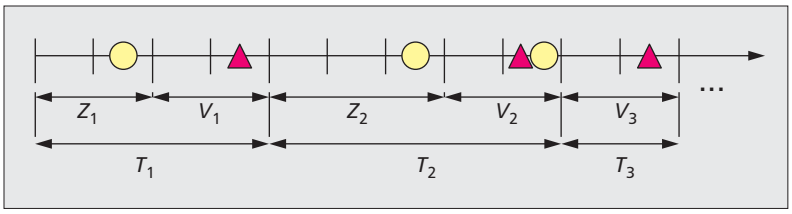

Figure 5. Graphical representation of the variables T_n, V_n and Z_n in the equivalent timing channel. An energy arrival is represented by a circle and a 1 channel symbol is represented by a triangle.

ited sized batteries. It is more obvious that the save-and-transmit scheme needs an unlimited sized battery, since the battery energy needs to go to infinity in the saving phase as the block length becomes large. The fact that the best-effort-transmit scheme also needs an unlimited sized battery is less obvious. While the best-effort-transmit scheme starts transmission right away, since average energy expenditure is less than average recharge rate, eventually the battery energy goes to infinity. In fact, this is the reason that energy shortages occur only in finitely many channel uses. Essentially, after a large enough channel use index, the battery has so much energy that no energy shortages occur.

We note that the capacity with an energy harvesting transmitter with an unlimited battery is invariant to the memory in the energy arrival process, so long as the energy arrival process is stationary and ergodic. That is, an i.i.d. energy arrival process and a non-i.i.d. energy arrival process with the same average arrival rate will yield the same capacity so long as the non-i.i.d. energy arrival process is stationary and ergodic.

Channel Capacity with Zero Energy Storage

In this section we consider the other extreme where there is no battery at the transmitter, i.e. $E_{\max}$ is zero, and at each channel use, the incoming energy is either used or it is lost. In this case, we assume *simultaneous transmit order*, as storing the incoming energy is not possible, which is required for the *transmit first order*. In this case, stochastic energy levels at the transmitter cause stochastic instantaneous amplitude constraints, which is a generalization of the static amplitude constrained Gaussian channel in [8]. Indeed, the channel could be viewed as a state-dependent channel with state information available at the transmitter only, where the state is the available energy [9]. In this case, Shannon strategies are the capacity achieving schemes. This is known due to Shannon's work [10] for the capacity of state-dependent channels with state information at the transmitter only. In particular, let us assume without loss of generality that the energy arrivals are selected from a finite alphabet with cardinality $|\mathcal{E}|$. By assigning an input T_i to be put as the channel input whenever E_i arrives, the channel capacity is the maximum mutual information between the extended channel input $T_1, \ldots, T_{|\mathcal{E}|}$ and the channel output; maximization is with respect to the joint distributions of $T_1, \ldots, T_{|\mathcal{E}|}$.

For the Gaussian channel, the optimal input distribution of the extended channel inputs $T_1, \ldots, T_{|\mathcal{E}|}$ are numerically verified to be discrete distributions in [9]. This observation is compatible with the discreteness result in [8]. However, the exact proof for the discreteness of the capacity achieving input distribution for the general case is still an open problem in the current literature.

Channel Capacity with Finite Energy Storage

We now focus on the practically relevant case where the battery size is neither zero nor infinity, but a finite number. Determining the channel capacity with a finite-sized battery, in general, is a difficult problem. In the infinite and zero energy storage cases, the channel capacity is achieved by transmission strategies that do not vary with the time index. However, this does not immediately extend to the finite-sized battery case, and the optimal transmission strategies may potentially need to use all past information of energy arrivals and signal transmissions. A complete answer to this question is still missing in the literature, but significant progress has been made. In the following we provide currently available results for this case.

General Case

In [11] Mao and Hassibi studied the channel capacity for the energy harvesting channel with finite energy storage, and found a general capacity formula in terms of Shannon strategies [10]. This formula involves maximization of the mutual information rate between a sequence of Shannon strategies and the sequence of corresponding channel outputs. Moreover, it is conjectured that the optimal sequence may be obtained by observing the current battery energy level only and ignoring the history of the battery level sequence. In addition, [11] presents a set of achievable schemes using Shannon strategies. These schemes have the desirable property that the corresponding achievable rates, i.e. the limit of the n-letter mutual information rates, can be calculated by a simple simulation-based algorithm.

Deterministic Energy Arrivals

In [12] Jog and Anantharam studied the capacity of the Gaussian channel with a finite-sized battery and deterministic energy arrivals. In this work the channel capacity is determined as the maximum n-letter mutual information rate between the input and output sequences. In contrast to the general capacity formula in [11], the n-letter mutual information rate does not involve Shannon strategies [10]. An implication of this result is that the need for Shannon strategies stems from the randomness in the energy arrivals, and the fact that the state information is available only at the transmitter. In addition, entropy-power-inequality based lower bounds are found for the capacity of this channel. Numerical results show that even with small battery sizes, the capacity approaches quickly to the capacity with an infinite-sized battery.

An Approximate Capacity Approach

In [13] Dong and Ozgur introduced a new approach to find an approximate capacity for the Gaussian channel with a finite-sized battery. The main result in this work is that the channel capacity with any finite-sized battery can be obtained within a constant gap from the channel capacity with infinite-sized battery. In order to establish this result, an achievable rate for a genie-aided receiver with energy arrival side information is derived. Then it is shown that this rate translates into a lower bound on the capacity by a novel evaluation of the side information.

Noiseless Binary Energy Harvesting Channel with Unit Energy Storage

Special cases may be helpful in the challenging problem of determining the channel capacity of the energy harvesting channel with a finite-sized battery. We now present the results of studying a noiseless binary channel with a unit-sized energy storage [14]. Note that even though the channel is noiseless, uncertainty of the receiver regarding the battery energy level at the transmitter side makes it challenging for the receiver to decode the messages of the transmitter.

In this binary energy harvesting channel, encoding and decoding can be equivalently performed over the time intervals between two consecutive 1s. This establishes an equivalence between this channel and a timing channel with additive geometric noise, where this noise state information is available at the transmitter only. In a timing channel, channel symbols are sent by the intervals between the time the packet is put in the queue and the time it is served out of the queue [15]. In our binary communication channel, a "packet" is replaced with a 1 symbol. The codebook used to transmit messages acts as a server to the 1 symbol in the energy queue, and determines its dynamics. The energy available at the energy queue is the state, which determines the set of feasible symbols that can be transmitted. Therefore, the transmitter causally observes the time a 1 symbol waits to be served in the queue.

The equivalent timing channel is represented by the input V_n, noise Z_n, and output T_n, as shown in Fig. 5. The output T_n is the time interval between two 1 symbols at the output, and the input V_n is the number of channel uses a 1 symbol is released from the queue as a channel symbol after Z_n units of waiting. Accordingly, T_n is equal to $V_n + Z_n$. The receiver observes T_n perfectly as the channel is noiseless.

Since Z_n is available at the transmitter before determining V_n, this is a state-dependent channel with state information available at the transmitter only, where the state is the noise Z_n. Hence, the channel capacity in this case can be expressed with a single-letter Shannon strategy at the transmitter [14]. In particular, the capacity is the maximum mutual information between an auxiliary random variable U (independent of the state Z) and the channel output T per average number of channel uses in the classical channel

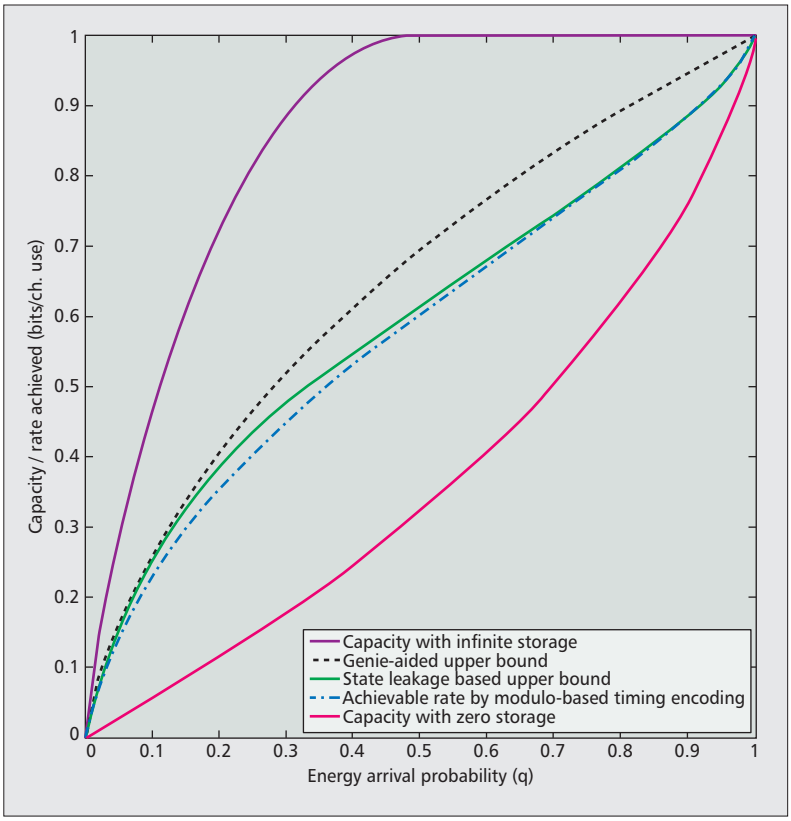

Figure 6. Comparison of upper bounds and achievable rates in [14].

$E[T]$. In this case the channel input is given as a deterministic function V of the auxiliary variable U and the noise Z. In general, a single-letter capacity expression is more desirable compared to an n-letter expression. However, the infinite cardinality requirement of the auxiliary random variable U and determination of the function V cause difficulties in evaluating the capacity.

The work in [14] proposes modulo-based timing encoding schemes. Moreover, [14] provides an upper bound for the channel capacity based on the state information leakage to the receiver side. Numerical studies show that this upper bound outperforms an older genie-aided upper bound and is quite tight. We present the comparison of upper bounds and achievable rates in Fig. 6.

Conclusions and Future Directions

In communications with energy harvesting devices, the energy source is intermittent and random. In addition, a battery is available to save the excess energy and use it later. The uncertainty of the available energy at the transmitter, which is also partially controlled by the transmitter's own actions, creates unprecedented constraints on the channel inputs, and this, together with the unavailability of the energy state information at the receiver side, makes the determination of the ultimate capacity limits of such channels a timely and challenging problem in modern wireless communication theory. In this article we have presented partial results for the general finite-sized battery case and com-

The uncertainty of the available energy at the transmitter, which is also partially controlled by the transmitter's own actions, creates unprecedented constraints on the channel inputs, and this, together with the unavailability of the energy state information at the receiver side, makes the determination of the ultimate capacity limits of such channels a timely and challenging problem.

plete results for the special cases of very large sized batteries and very small sized batteries.

We have observed that the capacity critically depends on the size of the battery. For instance, for the Gaussian channel, while for an infinite-sized battery case the capacity can be obtained by using classical Gaussian codebooks together with save-and-transmit and best-effort-transmit schemes, for the zero-sized battery case the capacity is achieved with time-varying discrete channel inputs. For the finite-sized battery case, we have presented results that give partial answers to the case of general channel and general energy arrivals, and lower and upper bounds for deterministic energy arrivals in Gaussian channels, and constant gap results for general energy arrivals in Gaussian channels. For a simple noiseless binary channel with unit energy storage, we have shown an interesting correspondence to the timing channel and an exact single-letter capacity expression whose evaluation required determination of an auxiliary random variable.

This literature is in its infancy and there are several open research directions. From an information theory point of view, the capacity for a general channel with a general finite-sized battery is still open. In addition, capacities of the channels with various side information availability conditions are also open. For instance, if the receiver also harvests energy from the same source or a similar source, then the receiver will have highly correlated side information regarding energy arrivals at the transmitter, and the capacity and achievable schemes in this case are open problems. From a coding theory point of view, explicit coding schemes are open problems. From a wireless communication theory point of view, practical adaptive coding and modulation schemes, and finite block length scenarios, are open directions. From a networking point of view, capacity regions of multi-user versions of these problems (multiple access, broadcast, interference, relay networks) are interesting future research directions. From a signal processing perspective, how the received signals should be processed, and how energy expenditure for processing and other essential components should be accounted for, are important open problems.

References

[1] D. Gunduz *et al.*, "Designing Intelligent Energy Harvesting Communication Systems," *IEEE Commun. Mag.*, vol. 52, Jan. 2014, pp. 210–16.

[2] V. Raghunathan, S. Ganeriwal, and M. Srivastava, "Emerging Techniques for Long-Lived Wireless Sensor Networks," *IEEE Commun. Mag.*, vol. 44, Apr. 2006, pp. 108–14.

[3] J. Yang and S. Ulukus, "Optimal Packet Scheduling in an Energy Harvesting Communication System," *IEEE Trans. Commun.*, vol. 60, Jan. 2012, pp. 220–30.

[4] K. Tutuncuoglu and A. Yener, "Optimum Transmission Policies for Battery Limited Energy Harvesting Nodes," *IEEE Trans. Wireless Commun.*, vol. 11, Mar. 2012, pp. 1180–89.

[5] O. Ozel *et al.*, "Transmission with Energy Harvesting Nodes in Fading Wireless Channels: Optimal Policies," *IEEE JSAC*, vol. 29, Sept. 2011, pp. 1732–43.

[6] C. K. Ho and R. Zhang, "Optimal Energy Allocation for Wireless Communications with Energy Harvesting Constraints," *IEEE Trans. Signal Proc.*, vol. 60, Sept. 2012, pp. 4808–18.

[7] O. Ozel and S. Ulukus, "Achieving AWGN Capacity under Stochastic Energy Harvesting," *IEEE Trans. Inf. Theory*, vol. 58, Oct. 2012, pp. 6471–83.

[8] J. G. Smith, "The Information Capacity of Amplitude and Variance-Constrained Scalar Gaussian Channels," *Information and Control*, vol. 18, Apr. 1971, pp. 203–19.

[9] O. Ozel and S. Ulukus, "AWGN Channel Under Time-Varying Amplitude Constraints with Causal Information at the Transmitter," *Asilomar Conf.*, Nov. 2011.

[10] C. E. Shannon, "Channels with Side Information at the Transmitter," *IBM J. Research and Development*, vol. 2, no. 4, 1958, pp. 289–93.

[11] W. Mao and B. Hassibi, "On the Capacity of a Communication System with Energy Harvesting and a Limited Battery," *IEEE ISIT*, July 2013.

[12] V. Jog and V. Anantharam, "An energy Harvesting AWGN Channel with a Finite Battery," *IEEE ISIT*, June 2014.

[13] Y. Dong and A. Ozgur, "Approximate Capacity of Energy Harvesting Communication with a Finite Battery," *IEEE ISIT*, June 2014.

[14] K. Tutuncuoglu *et al.*, "Improved Capacity Bounds for the Binary Energy Harvesting Channel," *IEEE ISIT*, June 2014.

[15] V. Anantharam and S. Verdu, "Bits Through Queues," *IEEE Trans. Info. Theory*, vol. 42, Jan. 1996, pp. 4–18.

Biographies

Omur Ozel [S'08, M'15] (ozel@berkeley.edu) received a Ph.D. degree in electrical and computer engineering (ECE) from the University of Maryland (UMD) College Park in 2014. He is currently a postdoctoral scholar in the Electrical Engineering and Computer Sciences Department at the University of California Berkeley. His doctoral research was awarded a Distinguished Dissertation Fellowship by the ECE Department at UMD. His research interests lie in the intersection of wireless communications, information theory, and networking.

Kaya Tutuncuoglu [S'08] (kaya@psu.edu) received two B.S. degrees in electrical engineering and in physics from the Middle East Technical University, Ankara, Turkey, in 2008 and 2009, respectively. He is currently pursuing the Ph.D. degree in the Department of Electrical Engineering, Pennsylvania State University, University Park, PA. His research interests include green communications, energy harvesting, and resource allocation for wireless networks. He received the AT&T Graduate Fellowship in 2012, and the IEEE Marconi Prize Paper Award in Wireless Communications in 2014.

Sennur Ulukus [S'90, M'98] (ulukus@umd.edu) is a professor of electrical and computer engineering at the University of Maryland at College Park, where she also holds a joint appointment with the Institute for Systems Research (ISR). Prior to joining UMD, she was a senior technical staff member at AT&T Labs-Research. She received her Ph.D. from WINLAB, Rutgers University, and B.S. and M.S. degrees from Bilkent University. Her research interests are in wireless communication theory and networking, and network information theory. She received the 2003 IEEE Marconi Prize Paper Award in Wireless Communications, a 2005 NSF CAREER Award, the 2010-2011 ISR Outstanding Systems Engineering Faculty Award, and the 2012 George Corcoran Education Award.

Aylin Yener [S'91, M'00, SM'13, F'14] (yener@ee.psu.edu) has been a professor of electrical engineering at The Pennsylvania State University, University Park, PA since 2010; she joined the faculty as an assistant professor in 2002. During the academic year 2008-2009 she was a visiting associate professor in the Department of Electrical Engineering, Stanford University, CA. Her research interests are in information theory, communication theory, and network science, with recent emphasis on energy harvesting wireless communications, information security, and foundations of complex networks.

Enhancing Wireless Information and Power Transfer by Exploiting Multi-Antenna Techniques

Xiaoming Chen, Zhaoyang Zhang, Hsiao-Hwa Chen, and Huazi Zhang

ABSTRACT

This article reviews an emerging wireless information and power transfer (WIPT) technique with an emphasis on its performance enhancement employing multi-antenna techniques. Compared to traditional wireless information transmission, WIPT faces numerous challenges. First, it is more susceptible to channel fading and path loss, resulting in a much shorter power transfer distance. Second, it gives rise to the issue of how to balance spectral efficiency for information transmission and energy efficiency for power transfer in order to obtain an optimal tradeoff. Third, there exists a security issue for information transmission in order to improve power transfer efficiency. In this context, multi-antenna techniques, e.g. energy beamforming, are introduced to solve these problems by exploiting spatial degree of freedom. This article provides a tutorial on various aspects of multi-antenna based WIPT techniques, with a focus on tackling the challenges by parameter optimization and protocol design. In particular, we investigate the WIPT tradeoffs based on two typical multi-antenna techniques: the limited feedback multi-antenna technique for short-distance transfer; and the large-scale multiple-input multiple-output (LS-MIMO, also known as massive MIMO) technique for long-distance transfer. Finally, simulation results validate the effectiveness of the proposed schemes.

INTRODUCTION

Wireless power transfer has attracted a lot of attention in the wireless research community, as it can effectively prolong the lifetime of a power-limited network in a relatively simple way, especially under extreme conditions, such as on the battlefield, underwater, and as part of body area networks [1]. For example, in medical care applications, devices implanted in the body send information to outside receivers with harvested power from the outside power source. Recently, wireless power transfer has been proposed for cellular systems, to provide mobiles practically infinitely long battery lives and eliminating the need for power cords and chargers. The radio frequency (RF) signal based wireless power transfer has attracted considerable attention in both academia and industry for the following two reasons [2–3]. First, it is a controllable and deterministic power transfer method. For example, it is possible to flexibly increase transmit power to enhance receive quality. Second, information and power can be simultaneously transferred in a form of RF signal. Then, the communications can be supported without external power sources.

In comparison with conventional wireless information transmission, wireless information and power transfer (WIPT) exhibits both similarities and differences. On one hand, both of them suffer from channel fading and path loss, resulting in performance loss. In particular, power transfer distance may be relatively short, since energy harvesting is more sensitive than information decoding [4]. Therefore, it is necessary to effectively combat the fading effects, so as to improve the efficiency and distance of power transfer. For traditional wireless information transmission, the multi-antenna technique is a powerful way to enhance the performance over fading channels. Through spatial beamforming, multi-antenna techniques can adapt the transmit signal to channel states, so that channel fading can be harnessed to improve the performance. Similarly, for wireless power transfer, the multi-antenna technique can also be used to align the the RF signal to a power receiver, thus improving energy efficiency. Therefore, it makes sense to exploit the benefits of the multi-antenna technique to enhance the performance of WIPT. On the other hand, WIPT has two performance metrics: spectral efficiency for information transmission, and energy efficiency for power transfer. In general, the two metrics are inconsistent and even contradictory, since information and power compete for the same RF signal and resources. Fortunately, it is convenient for multi-antenna techniques to achieve a good tradeoff between spectral and energy efficiencies by designing appropriate spatial beams for information and power transfer, respectively [5]. More importantly, multi-antenna techniques may concurrently

Xiaoming Chen is with Nanjing University of Aeronautics and Astronautics.

Zhaoyang Zhang and Huazi Zhang are with Zhejiang University.

Hsiao-Hwa Chen is with National Cheng Kung University.

This work was supported by the Natural Science Foundation of China (Nos. 61301102, 61371094, 61401388), the National Key Basic Research Program of China (No. 2012CB316104), the National Hi-Tech R&D Program (No. 2014AA01A702), Zhejiang Provincial Natural Science Foundation of China (No. LR12F01002), the Natural Science Foundation of Jiangsu Province (No. BK20130820), and Taiwan Ministry of Science & Technology (No. 102-2221-E-006-008-MY3).

WPT is not a new technology, although it has attracted considerable new interest recently. It was developed more than a century ago and its feasibility has been verified by many practical experiments. At the end of the 19th century, Nikola Tesla carried out the first WPT experiment.

support multiple streams of information and power transfer, and thus the efficiencies are improved significantly.

To exploit the benefits of multi-antenna techniques for WIPT, the transmitter requires full or partial channel state information (CSI), and then both information and power are transferred adaptively to the channel conditions. Specifically, based on the CSI, a transmitter selects the optimal transmit parameters, i.e. transmit beam, transmit power, and accessing users in order to maximize the efficiencies over fading channels. In [6] an optimal multiuser WIPT system was designed, assuming that full CSI is available at the transmitter. However, in multi-antenna systems, it is a nontrivial task to obtain instantaneous CSI at the transmitter, since the channel is a multi-dimensional time-varying random matrix. Generally, according to different duplex modes, there are two CSI acquisition methods in multi-antenna systems [7]. In frequency division duplex (FDD) systems, the CSI is usually conveyed from the information and power receivers to the transmitter by making use of quantization codebooks, so that the transmitter can obtain partial CSI. Note that a larger codebook size leads to more accurate CSI, but also increases feedback overheads. Therefore, it is possible to improve efficiencies by increasing the feedback amount. On the other hand, the CSI in time division duplex (TDD) systems can be estimated at a transmitter, directly making use of channel reciprocity. Compared to the CSI feedback in FDD systems, CSI estimation in TDD systems saves feedback resources, but may suffer from a performance loss due to transceiver hardware impairment. To solve the problem, robust beamforming for WIPT was proposed in [8] to guarantee high efficiencies even with imperfect CSI. Moreover, CSI can also be used to construct transmit beams. However, with respect to the beamforming based on instantaneous CSI, the one based on estimated CSI suffers an obvious performance degradation. Thus, adaptive multi-antenna transmission techniques via CSI feedback or estimation are effective ways to enhance performance for WIPT over fading channels.

For multi-antenna based WIPT techniques, there are a number of transmission frameworks proposed in the literature. First, for multi-antenna techniques, there are several different forms. For example, according to the number of antennas, there are traditional multi-antenna techniques and large-scale multiple-input multiple-output (LS-MIMO) techniques. Additionally, according to the number of accessing users, we have single-user and multi-user transmission techniques. Second, as mentioned earlier, there are two different CSI achievement methods: CSI feedback and CSI estimation. Third, according to the transmission protocols, WIPT can also be classified into two cases. In the first case, information and power are transferred simultaneously, namely simultaneous wireless information and power transfer (SWIPT) [9]. In the second case, power is first transferred, and then the harvested power is used to send information, namely wireless powered communication (WPC) or energy harvesting communication (EHC) [10]. Thus, combining the above three schemes, the multi-antenna based WIPT technique has a variety of forms, which are applicable to fit to different scenarios. In this article we intend to investigate various issues in the multi-antenna based WIPT technique from both theoretical and design perspectives. In particular, we analyze parameter optimization and protocol design for various multi-antenna based WIPT techniques. To facilitate understanding, we use the traditional multi-antenna technique and the LS-MIMO technique based WIPTs as two typical examples to instantiate the wireless information and power transfer tradeoff, and analyze the effect of CSI accuracy and the number of antennas on the tradeoff.

The rest of this article can be outlined as follows. First we introduce various multi-antenna based WIPT techniques, and then highlight the parameter optimization and protocol design. A discussion and comparison of two typical multi-antenna techniques for WIPT are given. Simulation results are illustrated to verify the tradeoff performance of the two typical multi-antenna based WIPT techniques, followed by the conclusions and discussions of several open issues.

Multi-Antenna Based WIPT Technique

WPT is not a new technology, although it has attracted considerable new interest recently. It was developed more than a century ago and its feasibility has been verified by many practical experiments. At the end of the 19th century, Nikola Tesla carried out the first WPT experiment, which tried to transmit approximately 300 KW power via 150 KHz radio waves. In the 1960s William C. Brown restarted WPT experiments with high-efficiency microwave technology, and the efficiency reached 50 percent at an output power of 4W DC. After the 1980s many experiments were carried out in Japan and the United States. In the 2000s, advances in microwave technologies pushed WPT back into consideration for wireless communications. Despite these advances, there are many challenging issues that remain for WIPT, because both information and power are carried by RF signals over wireless media, and they may suffer from attenuation, noise, interference, and interception. Thus, to effectively implement WIPT and evaluate the performance, several fundamental metrics are introduced as follows.

Transfer Efficiency: The RF signal will decay due to channel fading caused by reflection, scattering, and refraction in propagation processes. Thus, the received signal may be very weak, making it difficult to recover the transmit signal or harvest the signal energy. The problem becomes more prominent for wireless power transfer, since a power receiver is more sensitive to the magnitude of the RF signal. Hence, it is necessary to improve the efficiency of WIPT over wireless channels.

Transfer Distance: The attenuation of the RF signal is an increasing function of transfer distance. To guarantee a viable received power, wireless power transfer has a stringent limitation on transfer distance based on the current

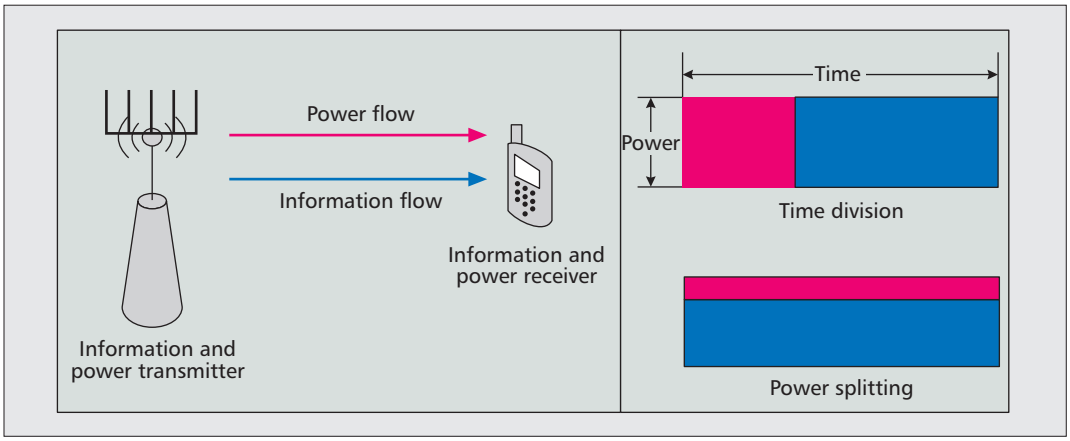

Figure 1. Model and protocol for combined case of SWIPT.

In order to realize efficient, reliable, secure, and long-distance WIPT, various advanced technologies have been identified recently, such as cooperative communication, resource allocation, and user scheduling. In particular, multi-antenna technique has a great potential due to its significant performance gain.

research. With an increasing demand on wireless power transmission, especially wireless powered communications, this limitation has become a major bottleneck in the development of wireless power transfer. Thus, it is imperative to increase the effective transfer distance.

Transfer Tradeoff: Limited power, spectrum, and time resources are shared by wireless information and power transfer, resulting in fundamental tradeoffs between the two. For example, in the SWIPT, the total transmit power is distributed for information and power transfer, while in wireless powered communications, each time slot is divided into information and power transfer durations. To balance the information and power transfer according to application requirements and enhance the overall performance, it is vital to analyze the optimal resource allocation between them.

Transfer Security: Due to the open nature of wireless media, information transmission is apt to be overheard. The security issue is even more serious in WIPT because the power receiver is usually placed closer to the transmitter than the information receiver. Traditionally encryption technology cannot fully solve the problem, because it requires a secure channel for exchanging private keys, becoming impractical in infrastructure-less or mobile networks.

In order to realize efficient, reliable, secure, and long-distance WIPT, various advanced technologies have been identified recently, such as cooperative communication, resource allocation, and user scheduling. In particular, the multi-antenna technique has great potential due to its significant performance gain. On one hand, multi-antenna diversity gain can be exploited to combat channel fading as an effort to improve transfer efficiency and increase transfer distance [12]. On the other hand, multi-antenna multiplexing gain can be leveraged to separate information and power transfer in space, so that the latter two metrics, i.e. transfer tradeoff and transfer security, can be realized simultaneously [4]. We now consider a simple example. If the information is transmitted in the null space of the channel for power transfer, then information security can be guaranteed with the help of physical layer security; even the power receiver is very close to the transmitter [4]. Due to these inherent advantages, the multi-antenna based WIPT technique is receiving considerable attention from both academia and industry. In what follows, we give a detailed investigation of multi-antenna WIPT. Due to space limitations, we only consider single-hop WIPT. In fact, the multi-hop transmission technology is also a powerful way of enhancing WIPT. For example, relay technology can shorten the transfer distance, and thus improve the performance [13]. The multi-hop cases will be studied in the future. According to the transfer model and protocol, WIPT can be further classified into SWIPT and WPC. Our investigation will cover a thorough case study of these two models and their integration.

Simultaneous Wireless Information and Power Transfer

As the name implies, SWIPT transmits information and power simultaneously. If the transmitter is equipped with multiple antennas, spatial beamforming adapted to the channel states can be used to improve the performance of WIPT. In this case, the information and power receivers can either be combined or separated. There are two subcases for SWIPT with different design principles.

Combined Case — In this case, a node plays the roles of both information and power receiver, as shown in the left-hand side of Fig. 1. The design of the transmitter is relatively simple. The core step is to perform spatial beamforming based on the CSI obtained through feedback in FDD systems or direct estimation in TDD systems. However, due to the dual roles, the receiver should be designed carefully. Note that the receiver cannot decode the information and harvest the energy simultaneously due to physical constraints. Then it is required to separate the information and power transfer by a certain protocol. Currently, there are mainly two protocols: the time division protocol [5] and the power splitting protocol [11]. Specifically, as shown in the right-hand side of Fig. 1, in the time division protocol each time slot is divided into information and power transfer durations. Then the roles of the receiver should switch between the two. Otherwise, in the power splitting protocol,

In the time division protocol, each time slot is divided into information and power transfer durations. Then, the roles of the receiver should switch between the two. Otherwise, in the power splitting protocol, the whole received signal is separated into two parts, one for information decoding and the other for power harvesting.

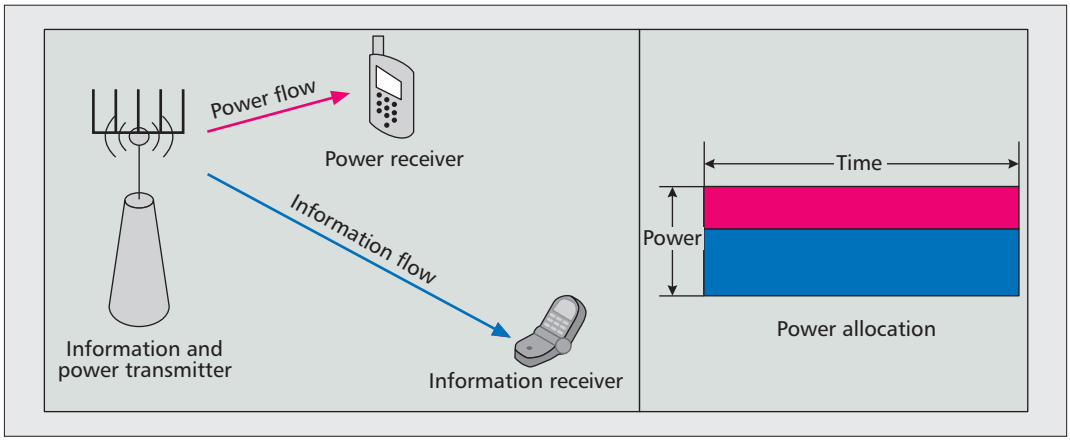

Figure 2. Model and protocol for separated case of SWIPT.

the whole received signal is separated into two parts, one for information decoding and the other for power harvesting.

Comparing the two protocols, we can find that time division requires two RF signal receive modules, since the signals for information decoding and power harvesting are separated at the RF side. In contrasting, power splitting only needs one RF signal receiver module, and the signals for information decoding and power harvesting are separated at the baseband. Note that there is a balance or tradeoff between information transmission and power transfer, since the time resource for time division and the power resource for power splitting are constrained and should be allocated to the two tasks according to a certain optimization objective. For example, the WIPT tradeoff can be formulated as an optimization problem of maximizing the information rate subject to a minimum harvested power, or maximizing the harvesting power subject to a minimum information rate.

Moreover, there may exist multiple receivers in the combined case. With respect to the single receiver case, there are more challenging problems to be solved. First, the receivers should be scheduled according to the urgency of information and power transfer. However, it is nontrivial to concurently determine the urgency of information and power transfer. Second, the WIPT tradeoff for each receiver may be distinct. In other words, each receiver may use different durations or powers for information decoding. Third, the beam design has contrasting goals for information and power transfer. For information transfer, the beams should be designed to mitigate inter-user interference; for power transfer, the inter-user interference can increase the received power. Fourth, one or more receivers may be an eavesdropper, which gives rise to security problems. A feasible method for multi-receiver SWIPT is to use time division multiplexing access (TDMA), such that each time slot is allocated to only one receiver. Then the multiple-receiver case is transformed to multiple single-receiver cases combined with receiver scheduling. However, the TDMA protocol may be suboptimal with respect to the space division multiplexing access (SDMA) protocol. It is still an open issue to design an optimal multiple access protocol.

Separated Case — In the case shown in Fig. 2, the information and power receivers are separated in different nodes. The transmitter is allowed to transmit RF signals for information and power transfer simultaneously in the same time and frequency resource block. As mentioned earlier, since the power receiver is more sensitive to the magnitude of the RF signal than the information receiver, it is usually placed closer to the transmitter, as shown in the left-hand side of Fig. 2. With respect to the combined case, the design focus of the separate case is on the transmitter, but not on the receiver. On one hand, the transmitter leverages the beamforming to separate the information and power transfer in space, in order to avoid the information leakage to the power receiver. On the other hand, the transmitter needs to allocate the transmit power to two beams, to achieve a tradeoff between information and power transfer. For example, the WIPT tradeoff can be formulated as an optimization problem of maximizing the secrecy rate subject to the minimum harvested power. It is worth pointing out that in the sense of maximizing the secrecy rate, the zero-forcing beamforming (ZFBF) and the use of maximum transmit power may not be optimal, since ZFBF and maximum transmit power may reduce the secrecy rate by decreasing the capacity of the legitimate channel from the transmitter to the information receiver and increasing information leakage to the eavesdropper, respectively. If we do not consider the security issues and only aim to maximize the information rate, the above tradeoff is reduced to a relatively simple optimization problem.

Similarly, the separated case may also comprise multiple information and power receivers. If a TDMA or OFDMA protocol is employed, the problem can be transformed to the case with one information receiver and multiple power receivers over each time slot or each subcarrier. In this subcase, if each power receiver, as an eavesdropper, overhears the information individually, the secrecy rate is determined by the power receiver with the strongest interception capability. Otherwise, if the power receivers cooperatively intercept the information, the secrecy rate is determined by the combined eavesdropper signal quality. Overall, by maximiz-

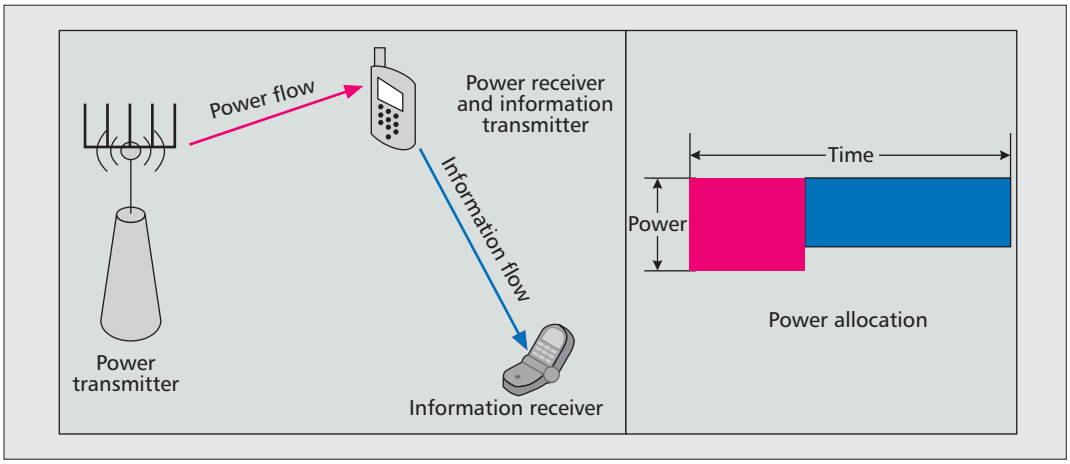

Figure 3. Models and protocols for WPC.

For the design of WPC, it is important to achieve the optimal tradeoff between information and power transfers. For example, based on the time division protocol, the tradeoff is to determine a switching point between power and information transfers.

ing the sum rate in all slots or subcarriers, it is possible to achieve the optimal receiver scheduling and spatial beamforming schemes. If a SDMA protocol is adopted, all information receivers are active over the same time-frequency resource block. Then the inter-user interference is inevitable, especially with imperfect CSI at the transmitter. Under such a circumstance, the design of transmit beams is more complicated, and is still an open issue.

Wireless Powered Communication

Different from SWIPT, WPC uses harvested power to transmit information, and thus it is also named energy harvesting communications. As a simple example, in medical care applications, the implanted equipment transmits the information it monitors to the instrument outside with the harvested power, as seen in the left-hand side of Fig. 3. With respect to SWIPT, WPC combines information and power transfer more closely, since the harvested power may also affect the information rate.

For the design of WPC, it is important to achieve the optimal tradeoff between information and power transfers. For example, based on the time division protocol, the tradeoff is to determine a switching point between power and information transfers, as shown at the right-hand side of Fig. 3. Since the power for information transmission comes solely from energy harvesting, the tradeoff based on the time division protocol can be formulated as an optimization problem maximizing the information rate with a given transmit power or minimizing the transmit power subject to a minimum rate.

More recently, several new technologies were introduced to further enhance the performance of multi-antenna based WIPT techniques. For example, the large-scale MIMO technique can generate high-resolution spatial beams by deploying tens or even hundreds of antennas. The benefit of large-scale MIMO technology for WPC is twofold [12]. First, the transfer efficiency and distance can be significantly improved by making use of its large array gain, so as to enable long-distance WPC with low power. Second, the high-resolution beam can reduce the information leakage to an unintended node to achieve information security. As the number of antennas increases, the performance gain becomes larger, which is a main advantage of LS-MIMO technique based WPC.

Integration of SWIPT and WPC

In fact, SWIPT and WPC can be integrated to give a more general WIPT scenario, described as follows. First, the transmitter sends information and power to one or multiple receivers, and then the power receivers send information to their next-hop receivers using the harvested power. The design of such a general WIPT can be considered as a concatenation of SWIPT and WPC. Its transmission protocol is also based on an integration of SWIPT and WPC components. In other words, each time slot is divided into two durations, one for SWIPT and the other for WPC.

If the time division protocol is adopted at the SWIPT stage, each time slot is partitioned into three non-overlapped durations. Typically, the power receiver will allocate constrained time duration to either receiving information from the power transmitter, or transmitting information to the next-hop receiver, which may lead to a low efficiency. Actually, it is possible to transmit and receive information simultaneously at the power receiver with recently introduced full-duplex technology [14]. For example, if the information transmitter at the SWIPT stage is also the information receiver at the WPC stage, the current full-duplex technology can be exploited to improve efficiency. A potential problem of the full-duplex technology is self interference from the information transmitter to the receiver. Fortunately, multi-antenna technology can be used to cancel self-interference by making use of the spatial degrees of freedom. Hence, the multi-antenna based WIPT technique combining full-duplex can significantly improve performance.

In all scenarios, the multi-antenna based WIPT technique can solve a series of challenging issues, making it an attractive solution to provide efficient, reliable, secure, and long-distance transfer. In Fig. 4 we give a summary of various multi-antenna based WIPT techniques together with their corresponding transfer protocols.

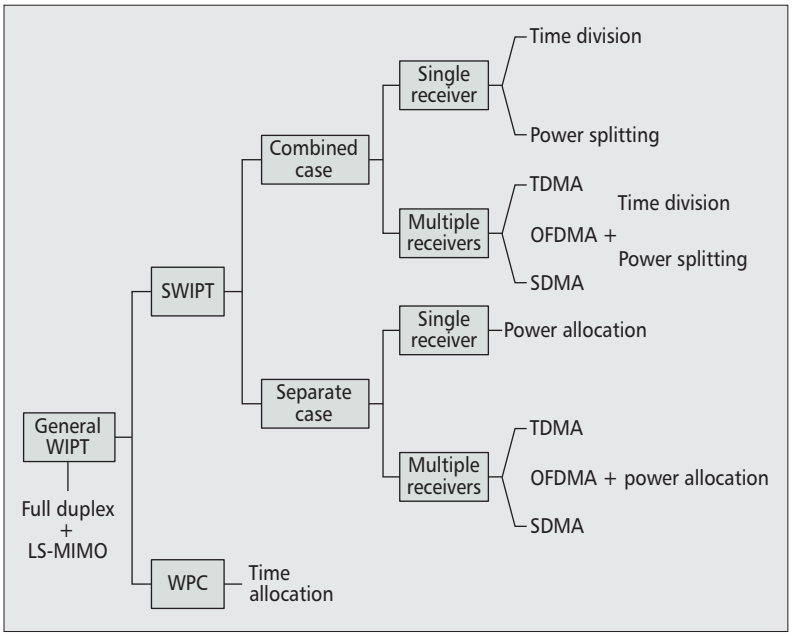

Figure 4. A summary of multi-antenna based WIPT.

Wireless Information and Power Transfer Trade-off

In this section we focus on the tradeoff or balance between wireless information and power transfer in single user multi-antenna systems. As discussed earlier, information and power transfers have different performance metrics. For example, information transmission mainly concerns rate, delay, and security, while power transfer emphasizes efficiency and distance. Intuitively, the goals for information and power transfers are inconsistent, and even contradictory. Thus, it is of importance to achieve an optimal tradeoff in the design of WIPT.

It is a common practice to achieve the performance objectives by optimizing the system parameters, e.g. transmit beam, transmit power, transfer duration, user scheduling, channel selection, and transfer protocol. For SWIPT, since the performance objectives are relatively independent, the tradeoff is usually formulated as three types of optimization problem. First, a multi-objective optimization scheme can be adopted in order to maximize the two performance indexes simultaneously. Second, the objective can be expressed as a general utility function. For example, it is reasonable to take a weighted sum of the efficiency as the objective. Third, the problem can be formulated by maximizing one performance index subject to a constraint on the other performance indices. For instance, a common problem in the existing related literature is to maximize the information rate subject to a minimum harvesting power constraint. Different from SWIPT, WPC relates the two performance metrics more closely, since the harvested power is used for information transmission. Therefore, the tradeoff for WPC has a direct and single formulation. In what follows, through two typical tradeoffs for multi-antenna technique based WPC, we present their protocol designs and parameter optimizations.

Information Rate Maximization in Traditional Multi-Antenna Systems

First let us consider a traditional multi-antenna based WPC technique, as shown in Fig. 3. A multi-antenna power transmitter charges a power receiver via RF signals at the beginning of each time slot, and then the power receiver sends information to an information receiver. Note that the power transmitter and the information receiver can be the same node in some cases. This is a typical application scenario in medial care (e.g. microchip implant) and underwater monitoring.

In order to improve the power transfer efficiency and thus maximize the information transmission rate, energy beamforming is conducted at the power transmitter. In practice, the multi-antenna power transmitter directs the RF signals to the receiver according to the current channel state, so as to overcome the negative effects of channel fading and propagation loss. Note that the performance of energy beamforming depends on the accuracy of CSI at the transmitter. As mentioned earlier, the CSI is obtained through feedback in FDD systems or direct estimation in TDD systems. Considering the fact that the power transmitter and the information receiver are in general separate, limited feedback based on a quantization codebook is a more practical choice. In such a system, the harvested power at the information transmitter can be considered as an increasing function of the CSI feedback amount, transfer duration, and transmit power, and at the same time a decreasing function of the transfer distance.

With the harvested power, the information transmitter sends information to the receiver in the remaining time of the slot. In general, the average transmit power for information transmission is equal to the quotient of the harvest energy and the duration left for information transmission. Therefore, according to the Shannon capacity equation, the average amount of information transmitted during a time slot can be expressed as a function of the average transmit power. Finally, the average information transmission rate can be derived through dividing the average amount of information transmitted during a time slot by the length of a time slot. Intuitively, it is a function of transmit power at the power transmitter, power transfer duration, and CSI feedback amount.

Taking the maximization of average information transmission rate as the optimization objective, we can derive the optimal transfer duration for a given CSI feedback amount, namely the switching point for power and information transfers. By adjusting the amount for CSI feedback, we can get different tradeoffs.

So far, we have only given a basic example. In fact, it can be extended to several more complex cases. First, when the information transmission has a certain quality of service (QoS) requirement, the above optimization problem should include a QoS constraint. It is worth pointing out that given the transmit power and feedback amount, there may be no feasible solutions for transfer duration. To solve it, we should increase transmit power or feedback amount. Additional-

ly, if the system is power-limited, we can formulate the problem as minimizing the transmit power, while satisfying the QoS requirement. Second, the basic model can also be naturally extended to the case of a general WIPT. Similarly, we need to add a minimum rate constraint for information transmission from the power transmitter to the power receiver. Meanwhile, if the time division protocol is adopted, an optimization variable of information transfer duration should be added. Otherwise, if the power splitting protocol is adopted, the added optimization variable should be the power splitting ratio instead. Third, in the case of an eavesdropper overhearing the information sent from the information transmitter, the above optimization problem is transformed to maximizing the secrecy rate.

Energy Efficiency Maximization in LS-MIMO Systems

The LS-MIMO technique can generate a high-resolution spatial beam through the deployment of a large number of antenna elements, and thus achieve substantial transfer efficiency and distance gains. In this case we consider a WPC system, where both the power transmitter and the information receiver are equipped with a large-scale antenna array.

To fully exploit the benefits of LS-MIMO techniques, the transmitter needs to know the exact CSI. However, due to a large amount of feedback (proportional to the number of antennas) in LS-MIMO systems, the CSI feedback scheme is practically infeasible. Thus, LS-MIMO systems usually work in TDD mode, and therefore the CSI can be estimated by making use of channel reciprocity. However, due to transceiver hardware impairment, the estimated CSI may be imperfect, resulting in certain performance loss. Hence, the CSI accuracy is also a decisive factor in determining the performance. Note that the number of antennas at the power transmitter and the information receiver is usually quite large (e.g. more than 100). According to the law of energy conservation, the harvested power at the power receiver (i.e. the information transmitter) is a function of transmit power at the power transmitter, power transfer duration, and CSI accuracy based on the TDD mode. In addition, due to channel hardening in LS-MIMO systems, it is also a deterministic function of the number of transmit antennas. Similarly, with the harvested power, the average information transmission rate can be expressed as a function of transmit power, transfer duration, CSI accuracy, the number of power transmit antennas, and the number of information receiver antennas by making use of the Shannon capacity expression.

Energy efficiency, defined as the bits transferred per Joule of energy, is a key performance metric for WIPT [15]. Therefore, we maximize energy efficiency to get the optimal tradeoff for such an LS-MIMO based WPC system. As pointed out earlier, the amount of information transferred during a time slot can be computed through multiplying the average information transmission rate by the length of a time slot, and the total energy consumption is the sum of the energy consumption in the power amplifier at the power transmitter and the constant energy consumption in the transmit filter, mixer, frequency synthesizer, and digital-to-analog converter (which are independent of the actual transmit power). Hence, by maximizing the ratio of the amount of information transmission and the total energy consumption, we can derive the optimal transfer duration. Similarly, we can add QoS and secrecy requirements on the basis of the above problem. Note that if the extended problem has no feasible solutions, we can make it feasible by simply adding more antennas at the power transmitter or the information receiver, which is a main advantage of the LS-MIMO based WIPT technique.

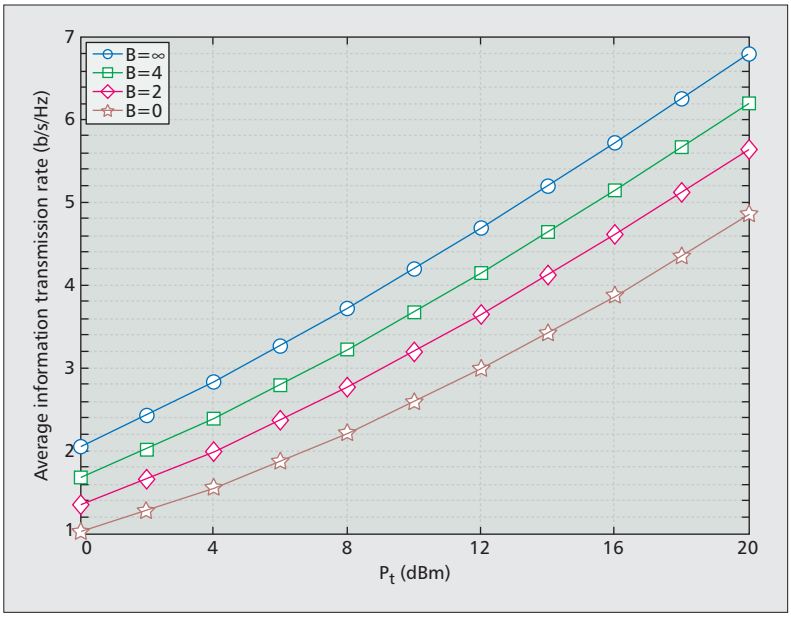

Figure 5. Information rate of traditional multi-antenna based WPC technique with different feedback amounts.

Performance Analysis and Simulations

In this section we present simulation results to validate the tradeoffs of the multi-antenna based WPC technique, where the power transmitter and the information receiver are integrated in one node. The parameters used are defined as follows. We set the length of a time slot as $T = 5$ ms, noise variance $\sigma^2 = -125$ dBm, energy conversion efficiency from RF signals to electric energy $\theta = 0.9$, constant power consumption $P_0 = 30$ dBm, and path loss for power transfer and information transmission $\alpha = \beta = 10^{-2}d^{-\nu}$, where d is the transfer distance and $\nu = 4$ is the path loss exponent. Note that in the given path loss model, a path loss of 20 dB is assumed at a reference distance of 1 meter. In addition, we use B and ρ to denote the feedback amount in traditional multi-antenna systems and the CSI accuracy in LS-MIMO systems, respectively.

First let us consider the tradeoff for a traditional multi-antenna based WPC technique with $N_t = N_r = 4$ and $d = 10$ m. As discussed earlier, we take the maximization of the average infor-

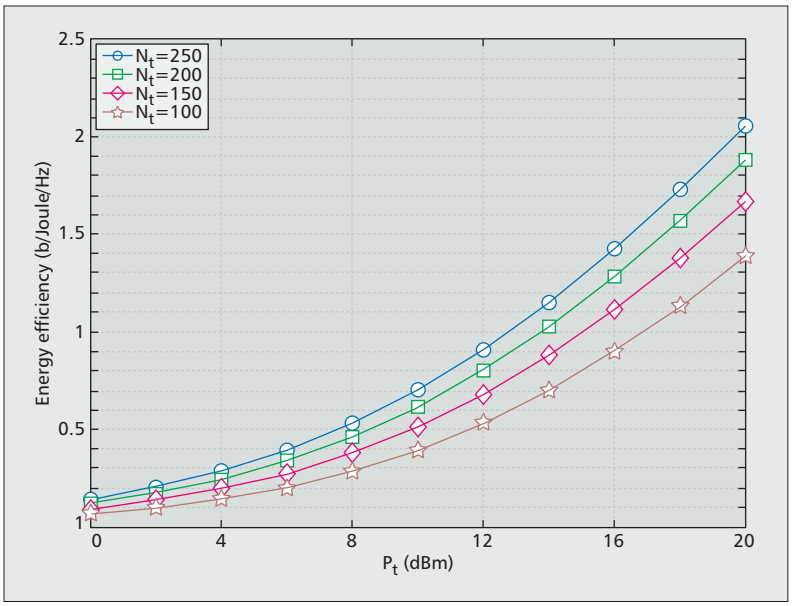

Figure 6. Energy efficiency of LS-MIMO based WPC technique with different numbers of antennas.

mation transmission rate as the optimization objective and adjust the transfer duration. It is shown in Fig. 5 that the feedback amount B has a great impact on the tradeoff, and thus affects the information rate. In comparison to the case without feedback, a small feedback amount, e.g. $B = 2$, can increase the information rate significantly. However, as the amount of feedback increases, the additional gain in terms of information rate diminishes. As seen from the results, with a finite feedback amount of $B = 4$, the performance gap to the ideal case (full feedback) is small. Thus the insight obtained here is that with even limited CSI feedback, the traditional multi-antenna technique can effectively enhance the performance of WIPT.

Second, let us examine the effect of the LS-MIMO technique on the tradeoff of WPC with $N_r = 100$, $\rho = 0.9$, and $d = 50$ m. This corresponds to a long-distance power transfer scenario. With respect to traditional multi-antenna techniques, only the LS-MIMO technique can support such a long transfer distance without consuming more transmit power, which is a very appealing feature. Taking energy efficiency as the optimization metric, we derive the optimal tradeoff of the LS-MIMO technique based WPC, as shown in Fig. 6. It is found that the number of antennas has a great impact on energy efficiency, which validates an antenna number versus energy efficiency tradeoff. By adding more antennas, energy efficiency can be improved further, which enables a high QoS for WPC with an affordable power even in the presence of imperfect CSI.

Conclusion and Future Works

This article reviewed the key technologies in WIPT and discussed several challenging issues, i.e. transfer efficiency, distance, tradeoff, and security. Through summarizing the existing work on multi-antenna based WIPT techniques, this article gives a comprehensive tutorial covering both parameter optimization and protocol design, and proposes to use full-duplex and LS-MIMO technologies to solve the challenges in various WIPT scenarios. In particular, a concept of WIPT tradeoff based on the multi-antenna technique is introduced and analyzed in detail. Finally, the tradeoffs are validated through simulations using the proposed schemes in two typical multi-antenna scenarios. It is worth pointing out that there are still many open issues for WIPT, especially for multiuser WIPT. First, the user scheduling schemes should be carefully designed to balance the QoS requirements, resource constraints, and information security. Second, transmit beams need to be elaborately constructed to achieve a proper tradeoff between information and power transfers, in particular with imperfect CSI. Third, the benefits of advanced multi-antenna techniques for WIPT should be further exploited. For instance, the self-interference of full-duplex techniques is adverse to information transmission, but can be harnessed to enhance power transfer. Hence, it is not optimal to cancel self-interference completely, and a more in-depth investigation is required.

References

[1] F. Zhang *et al.*, "Wireless Energy Transfer Platform for Medical Sensor and Implantable devices," *Proc. IEEE EMBS 31st Annual Int'l. Conf.*, Sept. 2009, pp. 1045–48.
[2] P. Grover and A. Sahai, "Shannon Meets Tesla: Wireless Information and Power Transfer," *Proc. IEEE Int'l. Symp. Info. Theory (ISIT)*, June 2010, pp. 2363–67.
[3] K. Huang and E. G. Larsson, "Simultaneous Information and Power Transfer for Broadband Wireless Systems," *IEEE Trans. Signal Proc.*, vol. 61, no. 23, Dec. 2013, pp. 5972–86.
[4] L. Liu, R. Zhang, and K.-C. Chua, "Secrecy Wireless Information and Power Transfer with MISO Beamforming," *IEEE Trans. Signal Proc.*, vol. 62, no. 7, Apr. 2014, pp. 1850–63.
[5] X. Chen, C. Yuen, and Z. Zhang, "Wireless Energy and Information Transfer Tradeoff for Limited Feedback Multiantenna Systems with Energy Beamforming," *IEEE Trans. Vehic. Tech.*, vol. 63, no. 1, pp. 407-412, Jan. 2014.
[6] R. Zhang and C. K. Ho, "MIMO Broadcasting for Simultaneous Wireless Information and Power Transfer," *IEEE Trans. Wirless Commun.*, vol. 12, no. 5, May 2013, pp. 3543–53.
[7] D. J. Love *et al.*, "An Overview of Limited Feedback in Wireless Communication Systems," *IEEE JSAC*, vol. 26, no. 8, Oct. 2008, pp. 1341–65.
[8] Z. Xiang, and M. Tao, "Robust Beamforming for Wireless Information and Power Transmission," *IEEE Wireless Commun. Lett.*, vol. 1, no. 4, Aug. 2012, pp. 372–75.
[9] I. Krikidis, "Simultaneous Information and Energy Transfer in Large-Scale Networks With/Without Relaying," *IEEE Trans. Commun.*, vol. 62, no. 3, Mar. 2014, pp. 900–12.
[10] X. Lu *et al.*, "Wireless Networks with RF Energy Harvesting, A Contemporary Survey," *IEEE Commun. Surveys & Tutorials*, 2015, http://arxiv.org/abs/1406.6470v1.
[11] L. Liu, R. Zhang, and K.-C. Chua, "Wireless Information and Power Transfer: A Dynamic Power Splitting Approach," *IEEE Trans. Commun.*, vol. 61, no. 9, pp. 3990-4001, Sept. 2013.
[12] X. Chen, X. Wang, and X. Chen, "Energy-Efficient Optimization for Wireless Information and Power Transfer in Large-Scale MIMO Systems Employing Energy Beamforming," *IEEE Wireless Commun. Lett.*, vol. 2, no. 6, Dec. 2013, pp. 667–70.
[13] D. S. Mishalogoulos, H. A. Suraweera, and R. Schober, "Relay Selection for Simultaneous Information Transmission and Wireless Energy Transfer: A Tradeoff Perspective," *IEEE JSAC*, vol. 33, Feb. 2015, http://arxiv.org/abs/1303.1647v1.

[14] C. Zhong *et al.*, "Wireless Information and Power Transfer with Full Duplex Relaying," *IEEE Trans. Commun.*, vol. 62, no. 10, Oct. 2014, pp. 3447–61.
[15] D. W. K. Ng, E. S. Lo, and R. Schober, "Energy-Efficient Resource Allocation in OFDMA Systems with Hybrid Energy Harvesting Base Station," *IEEE Trans. Wireless Commun.*, vol. 12, no. 7, July 2013, pp. 3412–27.

Biographies

XIAOMING CHEN [M'10, SM'14] received a B.Sc. degree from Hohai University in 2005, an M.Sc. degree from Nanjing University of Science and Technology in 2007, and a Ph.D. degree from Zhejiang University in 2011, all in electronic engineering. He is now with the College of Electronic and Information Engineering, Nanjing University of Aeronautics and Astronautics, Nanjing, China. His research interests mainly focus on cognitive radio, multi-antenna techniques, wireless security, interference networks, wireless power transfer, etc.

ZHAOYANG ZHANG [M'02] received a B.Sc. degree in radio technologies and a Ph.D. degree in communication and information systems from Zhejiang University, China, in 1994 and 1998, respectively. He is currently a full professor with the Department of Information Science and Electronic Engineering, Zhejiang University. His research interests are mainly focused on information theory and coding theory, signal processing for communications and in networks, and their applications in the next generation wireless mobile communication systems. He has co-authored more than 200 refereed papers in international journals and conferences, as well as two books in the above areas. He was a co-recipient of several conference Best Paper Awards/Best Student Paper Awards. He is currently serving as an editor for *IEEE Transactions on Communications*, *IET Communications*, and other international journals. He served as TPC co-chair or symposium co-chair for many international conferences such as WCSP 2013 and Globecom 2014 Wireless Communications Symposium, etc.

HSIAO-HWA CHEN [S'89, M'91-SM'00, F'10] is currently a distinguished professor in the Department of Engineering Science, National Cheng Kung University, Taiwan. He obtained BSc and MSc degrees from Zhejiang University, China, and a PhD degree from the University of Oulu, Finland, in 1982, 1985, and 1991, respectively. He is the founding editor-in-chief of Wiley's *Security and Communication Networks Journal* (http://www.interscience.wiley.com/security). Currently he is also serving as the editor-in-chief of *IEEE Wireless Communications*. He is a Fellow of the IEEE, a Fellow of IET, and an elected Member-at-Large of the IEEE Communications Society.

HUAZI ZHANG [S'11, M'13] received B.S. and Ph.D. degrees in electrical engineering from Zhejiang University, Hangzhou, China, in 2008 and 2013, respectively. From 2011 to 2013 he was a visiting scholar in the Department of Electrical and Computer Engineering, NC State University, Raleigh, NC. His research interests include OFDM and MIMO systems, clustering and compressive sensing, coding techniques, and cognitive radio networks.

GreenDelivery: Proactive Content Caching and Push with Energy-Harvesting-Based Small Cells

Sheng Zhou, Jie Gong, Zhenyu Zhou, Wei Chen, and Zhisheng Niu

ABSTRACT

The explosive growth of mobile multimedia traffic calls for scalable wireless access with high quality of service and low energy cost. Motivated by the emerging energy harvesting communications, and the trend of caching multimedia contents at the access edge and user terminals, we propose a paradigm shifting framework, GreenDelivery, enabling efficient content delivery with energy harvesting based small cells. To resolve the two-dimensional randomness of energy harvesting and content request arrivals, proactive caching and push are jointly optimized, with respect to the content popularity distribution and battery states. We thus develop a novel way of understanding the interplay between content and energy over time and space. Case studies are provided to show the substantial reduction of macro BS activities, and thus the related energy consumption from the power grid is reduced. Research issues of the proposed GreenDelivery framework are also discussed.

INTRODUCTION

Facing the rapidly growing multimedia traffic over the air and the concern regarding CO_2 emissions, it is crucial to innovate green wireless access. In particular, three emerging technologies have been demonstrated as effective, which are energy harvesting (EH) [1], traffic-aware service provisioning [2], and wireless multicasting [3].

More specifically, EH utilizes the energy from natural sources such as solar, wind, and kinetic activities, allowing wireless transmissions to consume less energy [5] or no energy [4] from the power grid. With EH-based access nodes, such as base stations (BSs) [6] and small cells (SCs), a more environment friendly network can be constructed. Traffic-aware service provisioning was proposed to match the wireless resources to the traffic demands, thereby achieving better energy efficiency (EE). For instance, one can exploit lazy scheduling (i.e., deliver the data with low rate to save power as long as a given deadline is met). Another example is optimizing BS sleeping based on the traffic demands and EH profile [6]. Finally, wireless multicast holds the promise of achieving significant EE gain via delivering commonly popular contents to multiple users simultaneously, which avoids retransmissions of the same content [3, 7].

However, there have been some barriers preventing these three methods from being practical. First, exploiting EH is limited by the state-of-the-art readiness for battery capacity. Due to the double randomness and temporal mismatch between energy arrivals and traffic arrivals, a large amount of harvested energy should be stored in batteries; otherwise, energy waste or shortage will occur. Second, the EE gain from on-demand service is also limited because of harsh and stringent quality of service (QoS) requirements of multimedia traffic such as video streaming, where many bits should be delivered before an urgent deadline. Finally, in current cellular infrastructures, wireless multicasting can only be enabled if and only if a number of users requires a common content concurrently. Otherwise, the transmitter has to delay the response to earlier demands for concurrence, which may severely damage the QoS of the earlier user demands.

To make the above three methods practical, we propose a paradigm shifting framework, GreenDelivery, where EH-based SCs provide content delivery services. Based on the EH status and content popularity distribution, the SCs proactively cache and push the contents before the actual arrival of user demands. In reward, the time duration in which the desired content can be delivered is greatly extended, so the delivery can flexibly match the EH process and enjoy low-rate transmission. The GreenDelivery framework is built on the recent trend of providing smart content service with the last-mile wireless access: contents can be cached at the SCs [8, 9] or relay nodes [10], with proactive caching schemes [11]. The benefits include reducing the core network overhead and enhancing user experience in terms of delay and rate thanks to shorter access distances. Correspondingly, there have been some initial developments of such technology in commercial products, such as HiWiFi,[1] Smart WiFi,[2] with large storage and advanced

Sheng Zhou, Jie Gong, Wei Chen, and Zhisheng Niu are with Tsinghua University.

Zhenyu Zhou is with North China Electric Power University.

This work is sponsored in part by the National Science Foundation of China (NSFC) under grant No. 61201191, the National Basic Research Program of China (973 Program: No. 2012CB316001 and No. 2013CB336600), the National Science Foundation of China (NSFC) under grant No. 61322111, No. 61321061, No. 61401250, and No. 61461136004, and Hitachi Ltd.

[1] http://www.hiwifi.com/j2

[2] http://www.linksys.com/en-us/smartwifi

0163-6804/15/$25.00 © 2015 IEEE

operating systems capable of running various applications to customize content caching schemes. To further reduce the energy consumption via multicast, proactive push [12] can be introduced on top of caching. The analysis to the network capacity gain provided by proactive push is presented in [13]. Practically, incorporating push into current mobile network is supported by the integration of broadcast and communication network [7]. Moreover, as the EH technology has been applied to access nodes [6], researchers are considering using such EH-based SCs to cache contents [14], getting the best of their deployment flexibility and low CO_2 emission.

In this article we provide a more general picture of the GreenDelivery framework, with the joint design of EH, push, and caching to provide twofold benefits of QoS and greenness. Enabling content caching nodes with push capabilities, specifically for EH-based SCs, can well match the multimedia traffic with random energy arrivals. The benefit will be reflected in the reduction of macro BS activities and thus the reduced energy consumption from the power grid. The next section will overview the framework, the intuitive ideas behind it, and its benefits. In subsequent sections, two case studies are illustrated, and related research challenges are discussed.

GreenDelivery: The Framework

The concept of GreenDelivery is illustrated in Fig. 1, where multiple GreenDelivery SCs cache popular contents and push them to users proactively. The design *objective* of the GreenDelivery framework is to minimize the number of user requests handled by the macro BS. The intuition of such a metric is twofold. First is energy saving. As the macro BS generally connects to the power grid, minimizing the activity of the macro BS reduces the grid power consumption, while the renewable energy used to power the GreenDelivery SCs can be regarded as free. Note that the backhaul link from the BS to SCs is generally good, and can enjoy low-power transmission. The second consideration is the user QoS. For those contents already pushed to users, users can get the contents with zero delay. Even for those requested contents that have not been pushed, as SCs are closer to users, unicast from an SC provides higher transmission rate and thus guarantees shorter delay.

Specifically, EH technology provides renewable energy for GreenDelivery SCs to:

- Fetch contents from a macro BS via the backhaul link. Since these SCs are EH-based, it is reasonable that they may only have wireless backhaul. As a result, the energy of fetch is not negligible, and accounts among major consumption portions of the harvested energy. Note that the wireline backhaul, as shown in Fig. 1, can also be considered as an option with less energy consumption but higher deployment cost.
- Cache the fetched content. The energy consumption depends on the storage method and the content storage volume. For GreenDelivery SCs serving limited numbers of users, the contents can be stored locally in the SC hardware, with negligible energy consumption for caching. When the cache size is large, additional hardware, like a cache server, is required, and the energy consumption cannot be ignored.
- Push the contents to users before the user might request it. Once contents have been fetched and cached at the GreenDelivery SC, they are selected to be pushed to users depending on their popularity and the battery status. As shown in Fig. 1, not only the users associated with the SC, but also those (receiving red signals) in the overlapping coverage of multiple SCs can enjoy the push service from one or more SCs. In other words, the caching and push can be coordinated among multiple SCs; an example of caching coordination can be found in [9]. Note that for commonly requested contents, multicast/broadcast is performed, while for private contents, unicast push is performed.
- Unicast the contents to users upon request. Users may request of their associated SC some content before it is pushed. If the SC has the content fetched and the battery has enough energy, it will unicast the content to the user upon request. Pushing the requested content can also be considered, but if the requested content is private or not popular, it is not beneficial considering the limited storage on user terminals.

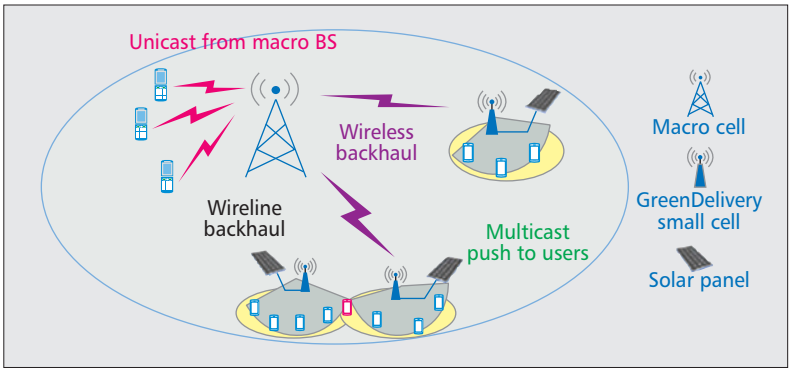

Figure 1. The concept of GreenDelivery. The small cells are energy-harvesting-based, while the macro BS is grid powered.

On the user side, if the upper layer application requests some content, the user will check its local storage first to see if the content has already been pushed. If not, it will request over the air from its associated SC. Note that in this article, the request is still counted even if it is satisfied by local storage. If the SC is not able to handle the request, the macro BS takes over and unicasts the content to the user.

The push mechanism can be realized by the existing broadcasting protocols without additional signaling overhead. Options include multimedia broadcast multicast services (MBMSs) proposed by the Third Generation Partnership Project (3GPP), or its new version, called broadcast and multicast service (BCMCS) [13]. The coordination of such a broadcast channel falls into the category of integrated communication and broadcast networks (ICBNs) [7]. If a specified broadcast channel does not exist, the GreenDelivery SC can reuse the unicast channel, but in this case users other than the default unicast receiver should be notified to receive via dedi-

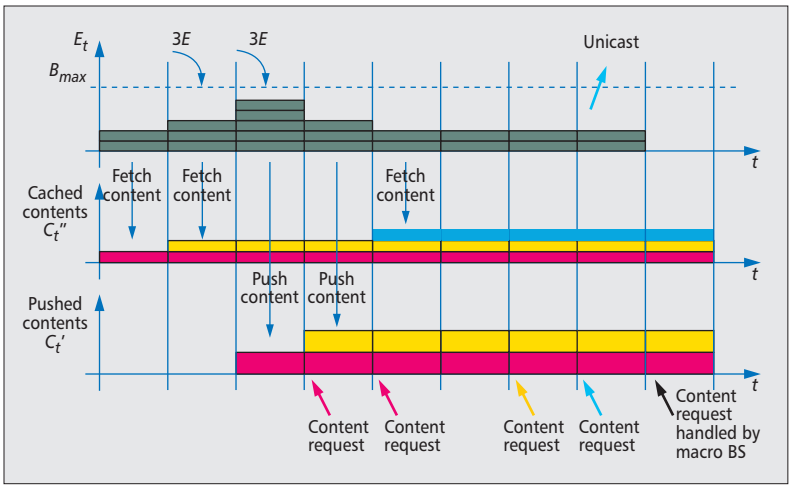

Figure 2. The behavior of a GreenDelivery SC: an example.

cated signaling. The signaling can be conveyed via the downlink control channel. Note that SCs should have the popularity distribution of the contents, which can be updated by the macro BS periodically. This overhead is proportional to the update rate of user interests, which is generally slow compared to the timescale of data transmission and EH.

Exploiting the Content and Energy Timeliness

The key idea of GreenDelivery is to exploit the timeliness of contents and energy via intelligent caching and push in order to match random energy arrivals and user requests over time and space. The timeliness of contents corresponds to their popularity and life span. The contents can only be interesting to users for a finite period of time, and the popularity ranking of the contents may change over time. The timeliness of the harvested energy comes with the causality of energy usage and limited battery capacity. The causality means that harvested energy cannot be used before its arrival. Moreover, limited battery capacity brings constraints on the delay in using such energy. In other words, if the arriving energy is not used in a timely manner, newly arrived energy will be wasted when the battery is full. This twofold randomness poses challenges in delivering contents efficiently. To solve this, based on the popularity and life span of the contents, together with the battery status, a GreenDelivery SC proactively fetches and caches popular contents and then pushes them to users. In this way, the delay constraint of using harvested energy is resolved, since the energy is used to provide stored contents at the users via push without waiting for requests. On the other hand, harvested energy can be regarded as being transferred from the time when the content is pushed into the future when the user actually requests the content. This is another way of information and power transfer over the hyper dimension of space (SC to users) and time (from the present to the future), which is different from the information and power transfer over space only [15].

An illustrative example of the aforementioned idea is shown in Fig. 2. The time horizon is divided into periods of equal length, which can be regarded as the broadcast frames on the broadcast channels in MBMS, and user requests arriving during some period are batched and responded to at the beginning of the next period. Energy is harvested and stored in the battery of the SC at the beginning of each period. In this example, the SC can fetch or transmit (push or unicast) at most one content in each period. The set of cached contents as of period t is denoted by $\mathcal{C}_t''$, and the set of pushed contents as of period t is denoted by $\mathcal{C}_t'$. Assume the length of the contents is the same, and the height of the contents in the figure represents the energy used to fetch or push it (i.e., unit energy E for fetching and caching a content, and $2E$ for pushing a content). As shown in Fig. 2, the energy arrives in the second and third periods, and the SC utilizes this energy (including the initial energy in the battery) to fetch and cache two contents. The SC then pushes the most popular content (the red one) in its cache to its users. Consequently, in the fourth and fifth periods, two requests for the red content arrive, and since it is pushed, the requests are instantaneously satisfied locally at user terminals. In the meantime, the SC can push and fetch more contents, satisfying the requests in the seventh and eighth periods. Note that the request for the blue content is served by unicast from the SC as it is not pushed yet. In the last period, since the SC is running out of energy, the user request is fulfilled by the macro BS. In the example, without proactive caching and push, the four requests require $11E$ ($2E \times 4 = 8E$ for content unicast, and $E \times 3 = 3E$ for fetching the three contents), which cannot be satisfied sincet the total energy budget is only $8E$. The battery capacity is $6E$, which means that if the energy is not used proactively, $2E$ of the harvested energy will be wasted.

Benefits of GreenDelivery

First, the temporal mismatch of content requests and energy arrivals can be resolved. Since the contents are cached, the SC may carry out content delivery whenever there is enough harvested energy in the battery. On the other hand, since the harvested energy can be effectively and promptly used, energy waste due to battery overflow can be avoided.

Second, SC push can greatly benefit from low-rate transmission, since there is no urgent delay constraint for proactive push, so the SC is allowed to transmit with reasonably low power. Note that there is a non-zero EE-optimal transmission time to balance the transmission and circuit power when a holistic power model is taken into account. Hence, the push holds the promise of achieving this EE-optimal transmission time in practice.

Finally, the joint use of push and caching enables more opportunities for wireless multicasting. During proactive push, all users who are potentially interested in these contents may overhear and decode the signal. In this case, wireless multicasting will not delay the service to any user because it is very flexible in aligning the time of push to different users. The only cost to be paid is the storage resource for caching, the

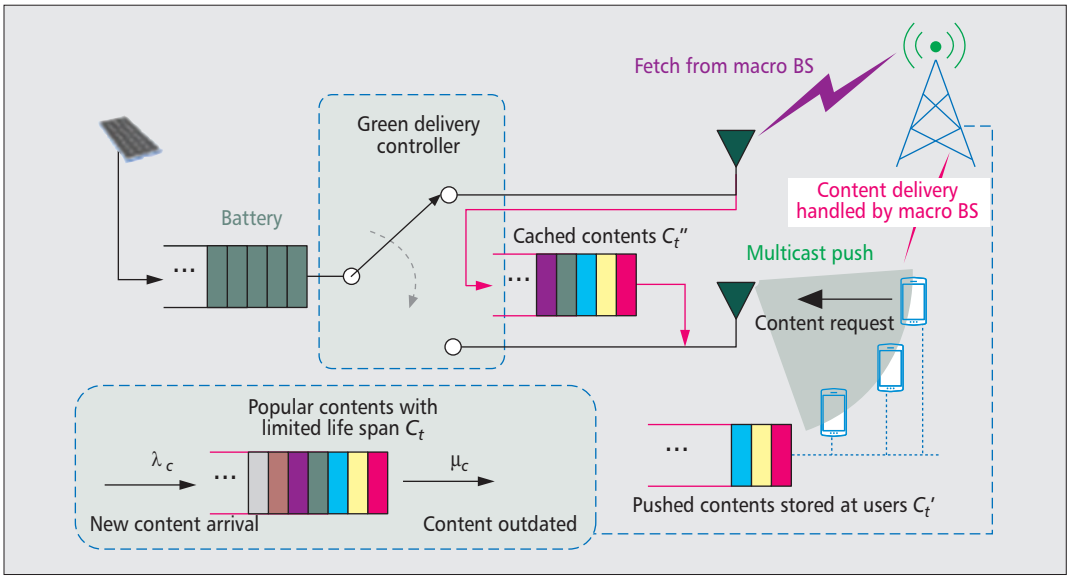

The key idea of GreenDelivery is to exploit the timeliness of the contents and energy via intelligent caching and push, so that to match random energy arrivals and user requests over time and space. The timeliness of the contents corresponds to their popularity and life span.

Figure 3. Block diagram of a GreenDelivery small cell.

price of which is dramatically dropping these days [8].

Case Study: Content Caching and Push under Dynamic Energy Arrivals

In this section, we express our idea through two case study examples. The considered model is presented in Fig. 3, where one EH-based GreenDelivery SC is illustrated. In each time period, the SC harvests a random amount of energy units, and stores them in a battery with finite capacity $E_{\max}$. Assume one or more contents can be fetched in a period, with energy consumption E_F, and one content can be transmitted, through either push or unicast, in a period with E_P units of energy. For example, in Fig. 2, $E_F = E$ and $E_P = 2E$. Note that here we simplify the channel to be static so that all transmissions are assumed successful. Specifically, the channel from the macro BS to the SC has a slight impact on the energy for fetch E_F as it is the receiving energy. The channels from SC to users are assumed to be static and identical, so E_P amount of push energy can always guarantee successful delivery.

The active content set $\mathcal{C}_t$ is time-varying; that is, new content comes into play over time, and contents can be outdated. This property is described as a birth-death process with birth rate λ_c and death rate μ_c, and the number of active contents at time period t is given by $|\mathcal{C}_t| = m$. The popularity of a content is defined as the probability that a user request corresponds to this content, denoted by f_m^i, which follows ZipF distribution [8]. By ranking the m contents with descending popularity, the popularity of the ith ranked content is

$$f_m^i = 1/(i^v \sum_{j=1}^{m} 1/j^v),$$

where $v > 0$ is the ZipF parameter, and larger v means that fewer contents account for the most popular ones.

The user request arrives at the beginning of a period with probability p_r, and we count both kinds of requests: those satisfied by the user local storage filled by proactive push, and the requests served over the air from the SC. In addition, when a content gets outdated and departs from the active content set $\mathcal{C}_t$, it will be removed from both $\mathcal{C}_t''$ and $\mathcal{C}_t'$ (if it is fetched and pushed to users), and the corresponding storage space is released.

At the beginning of each period, based on the current system state, including the active content set $\mathcal{C}_t$, the pushed content set $\mathcal{C}_t'$, the cached content set $\mathcal{C}_t''$, and the amount of energy units in the battery, the GreenDelivery Scheduler makes the action decision. The action set includes: fetch a content for caching; push a content; unicast a required content in $\mathcal{C}_t''$; do nothing. When the SC decides to do nothing, the user request, if it arrives, will be handled by the macro BS. As explained, our policy design objective is to minimize the ratio of user requests handled by the macro BS, denoted by η, subject to the energy causality constraint (i.e., the energy cannot be used before its arrival). In what follows, we first investigate the push behavior of a GreenDelivery SC, and then both fetch (for caching) and push are considered.

Energy-Harvesting-Based Proactive Push

To reflect the gain provided by proactive push, we first consider push only. Assume $\mathcal{C}_t'' = \mathcal{C}_t'$, which corresponds to the case in which the SC can get the content instantaneously via high-speed backhaul when it needs to push or unicast it. Therefore, the energy consumption of fetch and caching is ignored.

When the energy in the battery is sufficient and the SC decides to push, the most popular content in $\mathcal{C}_t''$ is pushed to users. We assume that the storage space of users is large enough to store the contents in set $\mathcal{C}_t'$.

We consider a simple energy-aware push

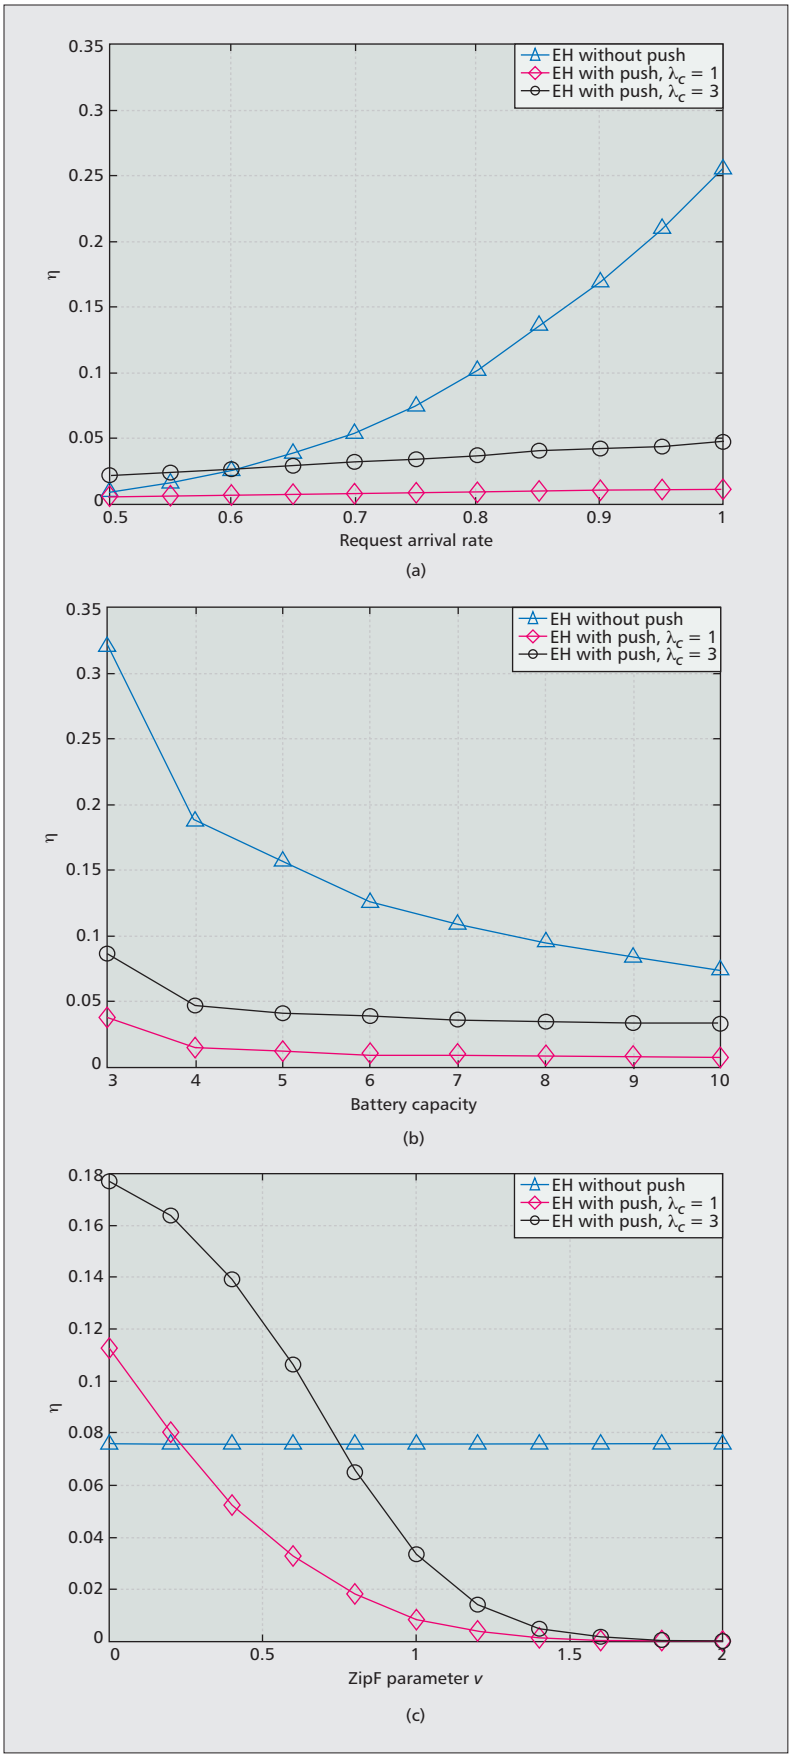

Figure 4. Evaluations of the ratio of the requests handled by the macro BS, with proactive push, $E_P = 2, E_H = 3$, where η denotes the the ratio of user requests handled by the macro BS: a) $E_{\max} = 10, v = 1$; b) $p_r = 0.75, v = 1$; c) $E_{\max} = 10, p_r = 0.75$.

scheme to see how proactive push can reduce the probability of handling user requests by the macro BS. Specifically, if there is no user request *over the air* in the current period and the battery energy is sufficient for pushing a content, the SC will push the most popular content in $\mathcal{C}_t''$ that has not been pushed. Otherwise, if a user requests a content, a unicast will be performed as long as the battery energy is sufficient to push a content. The user request is handled by the macro BS when the energy in the SC battery is not enough, and in this case the SC will do nothing in the current period.

Suppose the energy arrival follows Bernoulli distribution; that is, at the beginning of each period, the system can harvest E_H units of energy with probability p. We set $p = 0.5$, $\mu_c = 1 \times 10^{-3}$ per period, $v = 1$, $E_P = 2$, and $E_H = 3$. The ratio of requests handled by the macro BS is shown in Fig. 4. To compare, the policy without proactive push serves as the baseline, where the SC unicasts a required content as long as there is enough energy in the battery, or if no user request arrives, the SC does nothing. It can be seen from Fig. 4a that in the baseline scheme, the probability that the SC cannot provide service increases as the request arrival rate p_r increases. On the other hand, proactive push keeps this probability low and stable (i.e., almost irrelevant to p_r), hence greatly reducing the burden of the macro BS. One can also note that when p_r is low and the content refreshing rate is high (i.e., $\lambda_c = 3$ per period), proactive push does not bring any performance gain, because in this case user requests are diverse over different contents, and a pushed content has low probability of being requested by multiple users.

We also consider the influence of the battery capacity $E_{\max}$, since one major benefit of GreenDelivery is to solve the energy availability issue with limited battery. As shown in Fig. 4b, the ratio of requests handled by the macro BS decreases as $E_{\max}$ increases. Compared to the baseline, the reduction of the ratio is significant. In other words, to achieve the same performance, the required battery capacity with push is smaller than that without push.

In Fig. 4c, the impact of the content popularity distribution is depicted, where the ZipF distribution parameter v is varying from 0 to 2, that is, from a uniform distribution to a more skewed one. When the contents are uniformly distributed, it is better not to use proactive push, while the gain of proactive push increases with more skewed content distribution and lower content refreshing rate.

Energy-Harvesting-Based Caching and Push

We then take into account the cost of fetching the contents to cache at the GreenDelivery SC. Initially, the set of contents $\mathcal{C}_t$ is not available at the SC, and the SC has to first fetch the contents from the macro BS via the backhaul. It is reasonable to assume that the energy for fetching a content is less than that for pushing a content, and the SC can possibly fetch multiple contents in one period. A threshold-based fetch and push policy is proposed.

If the ratio of $|\mathcal{C}_t'|$ to $|\mathcal{C}_t''|$ is higher than the ratio of $|\mathcal{C}_t''|$ to $|\mathcal{C}_t|$, the number of cached contents is relatively small, and the SC needs to fetch more contents to avoid requests served by the macro BS. Then if the number of energy units in the battery is no less than a given threshold M_f, the SC fetches at most K contents and consume E_f units of energy. There is another threshold, M_p, for push. If the ratio of $|\mathcal{C}_t'|$ to $|\mathcal{C}_t''|$ is lower than that of $|\mathcal{C}_t''|$ to $|\mathcal{C}_t|$, the cached contents in the SC need to be pushed to reduce possible unicast events. In this case, the SC will push the most popular content in $\mathcal{C}_t''$ that has not been pushed, given the battery energy is no less than M_p. Finally, the user request is handled by the macro BS when the energy in the SC battery is not enough or the content has not been fetched, and in this case the SC will do nothing in the current period. Intuitively, larger thresholds M_f and M_p decrease the probability of battery outage when a request arrives, but on the other hand reduce the chances of push or fetch.

We set $\lambda_c = 1$ per period, $K = 3$, $E_F = 1$, and $E_{\max} = 10$, and other parameters are the same as those in the previous section. In Fig. 5, we compare the proposed algorithm with the case without proactive push. Similar to Fig. 4, it shows that push can significantly improve performance, and this also confirms the necessity of having push in GreenDelivery. The results also indicate that the thresholds should be carefully selected regarding the system parameters, especially the battery capacity as shown in Fig. 5b. From the figures, it is conjectured that having more aggressive fetch and caching (with smaller M_f) provides better performance, as contents to be pushed should already be cached. However, one should also note that using too much energy for fetch and caching leaves less energy for push, so these two sides of activities should be balanced.

Research Challenges

Several research challenges for releasing the benefits of the GreenDelivery framework are discussed as follows.

Intelligent Push under Random Energy Arrivals, Finite Battery, and Fading Channel

In practice, energy arrival and user requests cannot be predicted precisely. Therefore, online algorithms in charge of intelligent caching and push are required, with low computational complexity but close to optimal performance. Intuitively, more contents should be fetched and pushed if the energy arrival is sufficient and the battery is full, in order to avoid the battery overflow. On the other hand, when the harvested energy is not enough, it is better to reserve energy to handle randomly arriving user requests. Moreover, as in reality channel fading brings another dimension of randomness, the energy used for push is also varying. As a result, the algorithm design is not straightforward due to the threefold randomness of the EH process, user request arrivals, and channel fading.

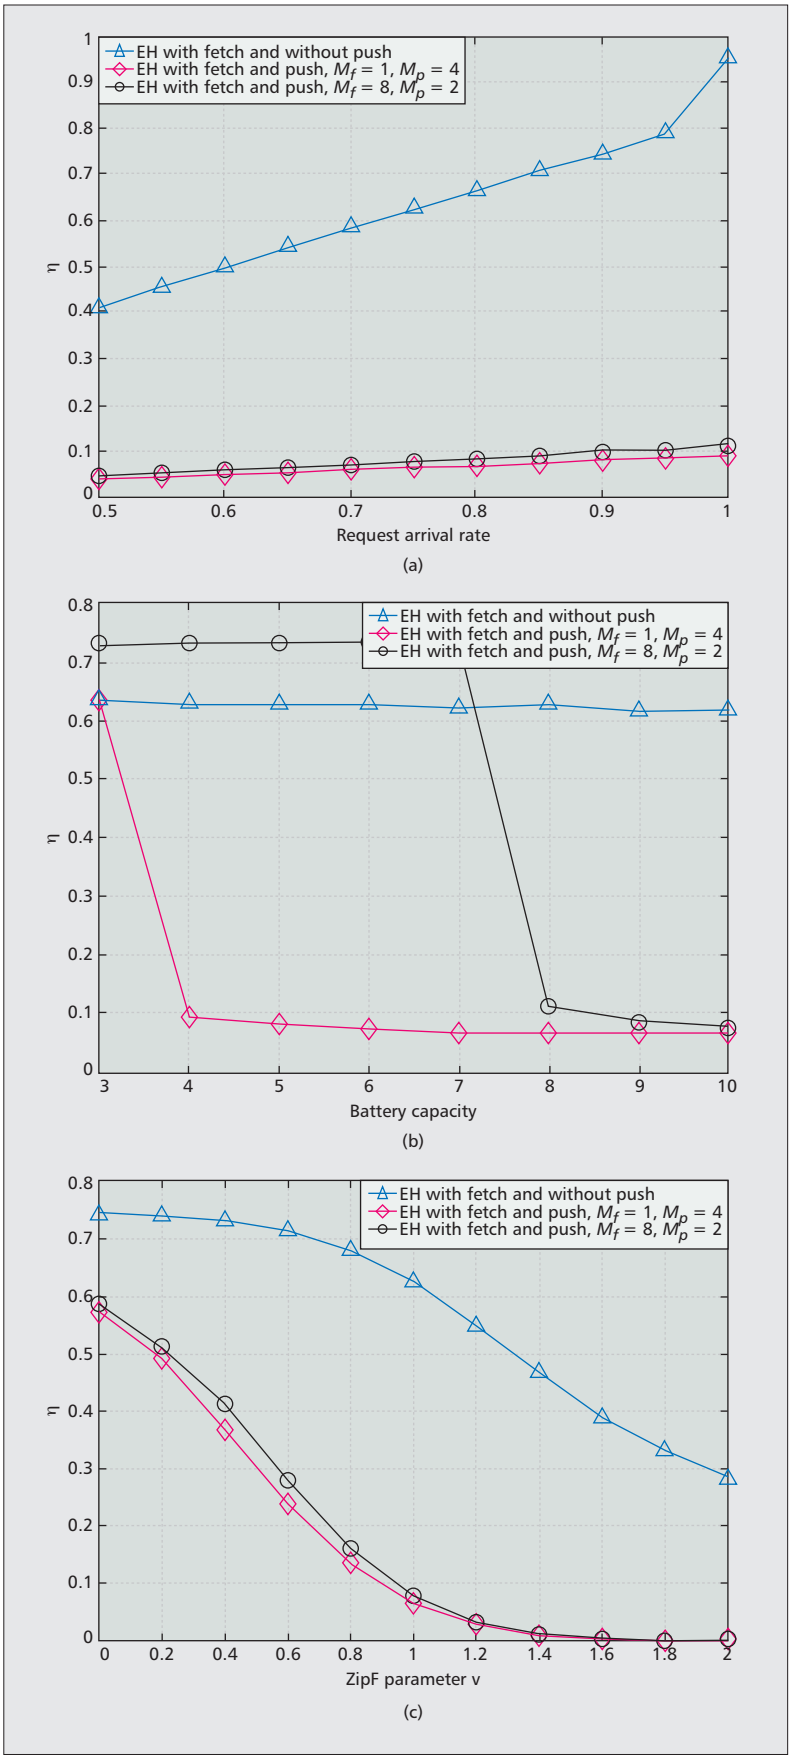

Figure 5. Evaluations of the ratio of the requests handled by the macro BS, with proactive fetch for caching and push, $E_F = 1$, $E_P = 2$, $E_H = 3$, where η denotes the the ratio of user requests handled by the macro BS: a) $E_{\max} = 10$, $v = 1$; b) $p_r = 0.75$, $v = 1$; c) $E_{\max} = 10$, $p_r = 0.75$.

We believe the idea of GreenDelivery is promising for delivering multimedia contents with densely deployed EH-based SCs, enjoying their deployment flexibility and energy scalability. GreenDelivery is also motivating future research.

Learning and Prediction of the Popularity and Energy Arrival Statistics

If the statistical information of either content popularity or the energy arrival process is unknown, the online algorithms should be able to learn and predict content popularity and energy arrival statistics. Due to the huge data volume and content variety, emerging big data learning technologies may be employed to provide both short-term and long-term popularity prediction. On the other hand, solar radiation models or wind speed models have been systematically studied for several decades. However, it is still an open problem to effectively combine popularity prediction and energy arrival prediction into the resource management for GreenDelivery SCs.

Trade-off between Benefits of Push and Content Storage Cost

In this article, we mainly consider the energy of fetch and push, while caching itself also introduces additional costs, including both energy consumption and storage occupancy. A large cache enhances the capacity of caching and push, but is expensive and energy consuming. Thus, the challenging problem is how to model and quantize the content storage cost, and how to achieve the optimal trade-off between the caching cost and the benefit of push. One could also improve the trade-off relation by caching and pushing the prefix of the multimedia contents rather than all of them, especially for videos, so the initial playout delay can still be reduced.

Cooperation among Multiple GreenDelivery Small Cells

GreenDelivery SCs can be densely deployed. Hence, it is possible to have cooperation and interaction among adjacent SCs to jointly optimize both customer-level QoS and system-level performance. For example, some SCs with larger request rate but low energy harvesting rate or small battery capacity need help from neighboring SCs. Multiple SCs can also form a cluster of coordinated transmission to combat channel fading. However, as the contents cached in different SCs are generally not the same, they may need to re-fetch the contents before push, which costs additional energy. Therefore, adjacent SCs should coordinate the re-fetch and push behaviors to efficiently handle content delivery, which poses design challenges, especially for large-scale heterogeneous networks.

Conclusion and Outlook

In this article, GreenDelivery as a new access network framework is proposed to enable efficient content delivery via EH-based SCs. Exploiting the content popularity information and battery status, proactive fetch/caching and push are implemented to match the random energy arrival and user content requests, and to provide more multicast opportunities. In this way, the limited harvested energy is wisely used, and the transmission cost of macro BSs is substantially reduced, which is illustrated via our case studies. We believe the idea of GreenDelivery is promising for delivering multimedia contents with densely deployed EH-based SCs, enjoying their deployment flexibility and energy scalability. GreenDelivery is also motivating future research directions including online policies for joint fetch-caching-push, learning from the EH and content statistics, and cooperation among multiple GreenDelivery SCs.

References

[1] D. Gunduz *et al.*, "Designing Intelligent Energy Harvesting Communication Systems," *IEEE Commun. Mag.*, vol. 52, no. 1, Jan. 2014, pp. 210–16.
[2] Z. Niu, "TANGO: Traffic-Aware Network Planning and Green Operation," *IEEE Wireless Commun.*, vol. 18, no. 5, Oct. 2011, pp. 25–29.
[3] J. Liu *et al.*, "A Utility Maximization Framework for Fair and Efficient Multicasting in Multicarrier Wireless Cellular Networks," *IEEE/ACM Trans. Net.*, vol. 21, no. 1, Feb. 2013, pp. 110–20.
[4] K. Tutuncuoglu and A. Yener, "Optimum Transmission Policies for Battery Limited Energy Harvesting Nodes," *IEEE Trans. Wireless Commun.*, vol. 11, no. 3, Mar. 2012, pp. 1180–89.
[5] J. Gong, S. Zhou, and Z. Niu, "Optimal Power Allocation for Energy Harvesting and Power Grid Coexisting Wireless Communication Systems," *IEEE Trans. Commun.*, vol. 61, no. 7, July 2013, pp. 3040–49.
[6] J. Gong *et al.*, "Base Station Sleeping and Resource Allocation in Renewable Energy Powered Cellular Networks," *IEEE Trans. Commun.*, vol. 62, no. 11, Nov. 2014, pp. 3801–13.
[7] Z. Niu *et al.*, "A New Paradigm for Mobile Multimedia Broadcasting Based on Integrated Communication and Broadcast Networks," *IEEE Commun. Mag.*, vol. 46, no. 7, July 2008, pp. 126–32.
[8] N. Golrezae *et al.*, "Femtocaching and Device-to-Device Collaboration: A New Architecture for Wireless Video Distribution," *IEEE Commun. Mag.*, vol. 51, no. 4, Apr. 2013, pp.142–49.
[9] K. Shanmugam *et al.*, "FemtoCaching: Wireless Content Delivery through Distributed Caching Helpers," *IEE Trans. Info. Theory*, vol. 59, no. 12, Dec. 2013, pp. 8402–13.
[10] X. Wang *et al.*, "On the Design of Relay Caching in Cellular Networks for Energy Efficiency," *Proc. IEEE INFOCOM '11 Wksp.*, 10–15 Apr. 2011, pp. 259–64.
[11] E. Bastug, M. Bennis, and M. Debbah, "Living on the Edge: The Role of Proactive Caching in 5G Wireless Networks," *IEEE Commun. Mag.*, vol. 52, no. 8, Aug. 2014, pp. 82–89.
[12] I. Podnar, M. Hauswirth, and M. Jazayeri, "Mobile Push: Delivering Content to Mobile Users," *Proc. ICDCSW '02*, 2002, pp. 563–68.
[13] K. Wang, Z. Chen, and H. Liu, "Push-Based Wireless Converged Networks for Massive Multimedia Content Delivery," *IEEE Trans. Wireless Commun.*, vol. 13, no. 5, May 2014, pp. 2894–2905.
[14] N. Sharma *et al.*, "GreenCache: Augmenting Off-the-Grid Cellular Towers with Multimedia Caches," *Proc. ACM MMsys '13*, Oslo, Norway, Feb. 2013.
[15] X. Zhou, R. Zhang, and C. K. Ho, "Wireless Information and Power Transfer: Architecture Design and Rate-Energy Tradeoff," *IEEE Trans. Commun.*, vol. 61, no. 11, Nov. 2013, pp. 4754–67.

Biographies

Sheng Zhou [M] (sheng.zhou@tsinghua.edu.cn) received his B.S. and Ph.D. degrees in electronic engineering from Tsinghua University, China, in 2005 and 2011, respectively. He is currently an assistant professor in the Electronic Engineering Department, Tsinghua University. From January to June 2010, he was a visiting student at the Wireless System Lab, Electrical Engineering Department, Stanford University, California. His research interests include cross-layer design for multiple antenna systems, cooperative transmission in cellular systems, and green wireless communications.

Jie Gong [M] (gongj13@tsinghua.edu.cn) received his B.S. and Ph.D. degrees from Tsinghua University in 2008 and

2013, respectively. He is currently a postdoctoral researcher with Tsinghua University. From July 2012 to January 2013, he visited the University of Edinburgh, United Kingdom. His research interests include base station cooperation in cellular networks, energy harvesting, and green communications.

ZHENYU ZHOU [M] (zhenyu_zhou@fuji.waseda.jp) received his M.E. and Ph.D degree from Waseda University, Tokyo, Japan, in 2008 and 2011 respectively. From April 2012 to March 2013, he was a chief researcher at the Department of Technology, KDDI, Tokyo. From March 2013 to now, he has been an associate professor at the School of Electrical and Electronic Engineering, North China Electric Power University. His research interests include energy-efficient wireless communications and smart grid communications.

WEI CHEN [SM] (wchen@tsinghua.edu.cn) received his B.S. degree in operations research and Ph.D. degree in electronic engineering (both with highest honors) from Tsinghua University in 2002 and 2007. He is a full professor and deputy head of the Electronic Engineering Department, Tsinghua University, as well as a National 973 Youth Project chief scientist and a winner of the National May 1st Medal. He has received the IEEE ComSoc APB Best Young Researcher Award and the IEEE Marconi Prize Paper Award.

ZHISHENG NIU [F] (niuzhs@tsinghua.edu.cn) graduated from Beijing Jiaotong University, China, in 1985, and got his M.E. and D.E. degrees from Toyohashi University of Technology, Japan, in 1989 and 1992, respectively. During 1992–1994, he worked for Fujitsu Laboratories Ltd., Japan, and in 1994 joined Tsinghua University, where he is now a professor in the Department of Electronic Engineering and deputy dean of the School of Information Science and Technology. He is also a guest chair professor of Shandong University, China. His major research interests include queueing theory, traffic engineering, mobile Internet, radio resource management of wireless networks, and green communication and networks. He is now a Fellow of IEICE, a distinguished lecturer (2012–2015) and Chair of the Emerging Technology Committee (2014–15) of IEEE Communication Society, and a Distinguished Lecturer (2014–2016) of the IEEE Vehicular Technologies Society.

NETWORK AND SERVICE VIRTUALIZATION: PART 2


Kostas Pentikousis


Catalin Meirosu


Diego R. Lopez


Spyros Denazis


Kohei Shiomoto


Fritz-Joachim Westphal

The second installment of the Feature Topic on network and service virtualization takes off from where Part I concluded: network function virtualization (NFV) deployment and operation. We are currently observing a transition from software-defined networking (SDN) focusing solely on programming packet-switching network elements based on lower-layer flow-oriented primitives toward the wider concept of infrastructure programmability. Such programmability extends vertically within the network stack to encompass optical layer capabilities as well as features located at upper layers of the traditional network stack. In addition, infrastructure programmability extends horizontally well beyond SDN to other resource stacks to encompass virtualized compute and storage resources.

In this issue we point our attention to technologies that unify the management of software-defined networks, clouds and application services in carrier-oriented environments. In particular, we look into the complex problems related to the integration of resource controllers and orchestrators with virtual network functions. A healthy dose of optical layer components, open source code, and security considerations takes the debate out of the abstract world of standardized interfaces and provides insights about the novel aspects of practical implementations.

The first article in this issue is "Optical Service Chaining for Network Function Virtualization." Ming Xia *et al.* explore service chaining to the optical domain. Service function chaining (SFC) is well known in the context of packet-level technologies, and there are ongoing standardization efforts in the Internet Engineering Task Force (IETF) on this topic. This article presents an integrated packet-optical architecture for data centers that enables steering large aggregated flows solely in the optical domain. Based on a combination of wavelength-selective switches, tunable transceivers, and shipping containers that house compute, packet network, and storage resources as an independent unit, the proposed architecture aims to significantly improve operational flexibility and scalability. Furthermore, the authors point to power efficiency benefits as more resource-intensive virtual network functions (VNFs) are deployed, which makes the architecture relevant to service providers or large enterprises that need to handle tens or hundreds of gigabits of traffic every second.

In "A Service-Oriented Hybrid Access Network and Cloud Architecture," Velasco *et al.* present an architectural solution that integrates distributed micro data centers and active optical nodes located close to the end users under a dynamic orchestration platform. The Service-Oriented Hybrid Access Network and Cloud Architecture (SHINE) leverages the IETF work on the application-based network operations (ABNO) framework for interacting with the transport network that interconnects the micro data centers and the active remote nodes. Four use cases based on scenarios defined by the European Telecommunications Standards Institute (ETSI) NFV ISG illustrate the flexibility and performance of the proposed architecture. The direct interface between VNFs (generally referred to as "applications" in the article) and the joint cloud and network resource orchestrator facilitates optimized placement of the instances as well as scale-in and scale-out operations.

In the third article in this Feature Topic, "A Service-Aware Virtualized Software-Defined Infrastructure," Mamatas *et al.* present a proposal to unify the management and control of networks and clouds while providing uniform virtualization abstractions for networks and applications. Virtual routers supported by the very lightweight software-driven network and services platform (VLSP) that allows deploying new network control and service components in support of advanced programmability features. The flexibility of VLSP is validated by testbed deployment and experimentation.The VLSP code has been open sourced, enabling interested readers to experiment and verify the conclusions of the authors.

Montero *et al.* present their work on trusted virtual domains in the access network, which can provide homogeneous security for a user regardless on the type of the

device (laptop, tablet, mobile phone) employed for accessing the network. As their article, "Virtualized Security at the Network Edge: A User-centric Approach," explains, shifting from a device-centric to a user-centric security model has significant advantages from the user perspective. That said, implementation raises a range of challenges that the authors detail in the article. Three domain-specific policy abstraction layers as well as translation services between them are described and discussed in the article, along with requirements on network functions and virtual infrastructure to support trusted virtual domains. This article is representative for the complexity of the technological backend that is required to simplify the life of end users of a modern telecommunications infrastructure, and the challenges related to security and privacy-preserving services.

The final article in this Feature Topic is "Toward an SDN-Enabled NFV architecture." Matias *et al.* vividly document their experience designing and implementing a VNF based on 802.1x access control functionality. The authors propose the decomposition of the VNF in a stateful component that can be executed on virtual compute resources and a stateless component that can be executed on a physical SDN switch. The architectural options for building the Flow-Based Network Access Control VNF are discussed in detail, along with the challenges and open research issues in this area. Resource isolation is considered one of the major challenges that must be addressed in order for the infrastructure to support widespread adoption of SDN-enabled VNFs.

Achieving a programmable software-defined infrastructure by jointly orchestrating resources at both the network and cloud layers, as presented in the articles included in this Feature Topic, is a common item on several research agendas. A large-scale initiative along these lines is the UNIFY project, funded in the Seventh Framework Programme by the European Union, which creates a software-defined platform that combines network and cloud resources to enable the delivery of new services at high velocity (see www.fp7-unify.eu for more details).

Future research directions call for extending the programmability to higher-layer virtualized network functions and services that have the potential to support more sophisticated infrastructure services related to network optimization and security, to name just a couple of relevant areas. Lower-layer functions such as switching and routing, which have been considered in the original SDN proposals, were stateless and allowed an immediately transparent distribution of functionality within the infrastructure. In contrast, many higher-layer VNFs are stateful, which further complicates the definition of appropriate programming primitives. Higher-layer services are often composed of several VNFs. Such composition adds to the challenge by introducing dependencies at the flow processing level that can no longer be solved within one self-contained processing pipeline, be it stateful or stateless. Infrastructure programmability in such environments will need to take inspiration, for example, from parallel and distributed computing technologies that evolved in the computer science field. This opens up exciting new opportunities for new protocols and programming languages that are optimized for high-throughput and low-latency packet processing.

It is our hope as Guest Editors that both practitioners and academics will find this Feature Topic a handy reference in the emergent area of network infrastructure and service virtualization. We conclude this introduction by thanking the numerous dedicated reviewers for their thorough analysis and constructive comments. We commend the work of the authors who submitted their articles to our Call for Papers and show particular appreciation for those who worked diligently to improve their manuscripts throughout the peer-review process. We gratefully acknowledge the magazine's Editor-in-Chief at the time this Feature Topic was being prepared, Dr. Sean Moore, and the Editorial Board for their continuous help. Finally, we say a big thank you to the ComSoc final production editors and staff for their professionalism, and in particular to Charis Scoggins for her guidance throughout the entire process.

BIOGRAPHIES

KOSTAS PENTIKOUSIS (k.pentikousis@eict.de) dedicates this Feature Topic to his father, Michael Pentikousis (1931–2015), who passed away unexpectedly as this Editorial was being prepared. He was a devoted father, and he is deeply missed by his family and all who knew him.

CATALIN MEIROSU (catalin.meirosu@ericsson.com) is a master researcher with Ericsson Research in Stockholm, Sweden, which he joined in 2007. He holds a Ph.D. in telecommunications (2005) from Politehnica University, Bucharest, Romania, and was a project associate of the ATLAS experiment at the Large Hadron Collider at CERN, Geneva, Switzerland. He has three granted patents and has co-authored over 30 scientific papers. He is currently working to develop DevOps capabilities for service providers in the FP7 UNIFY project.

DIEGO R. LOPEZ (diego@tid.es) joined Telefonica I+D in 2011 as a senior technology expert and is currently in charge of the Technology Exploration activities within the GCTO Unit. He is focused on network virtualization, infrastructural services, network management, new network architectures, and network security. He actively participates in the ETSI ISG on Network Function Virtualization (chairing its Technical Steering Committee), the ONF, and the IETF WGs connected to these activities.

SPYROS DENAZIS (sdena@upatras.gr) received his B.Sc. in mathematics from the University of Ioannina, Greece, in 1987, and in 1993 his Ph.D. in computer science from the University of Bradford, United Kingdom. He worked in European industry for eight years, and is now an associate professor in the Department of Electrical and Computer Engineering, University of Patras, Greece. His current research interests include P2P, SDN, and future Internet. He is currently technical manager of the STEER EU project. He has co-authored more than 50 papers and is a co-author of the book *Programmable Networks for IP Service Deployment*.

KOHEI SHIOMOTO [M] (shiomoto.kohei@lab.ntt.co.jp) is a senior manager at NTT Network Technology Laboratories. His research fields include IP/MPLS, IP+Optical, GMPLS/PCE, network virtualization, and traffic/QoE management. He has been engaged in standardization at the IETF and the organization of international conferences including WTC, MPLS, and iPOP. He is a Fellow of IEICE and a member of ACM. He co-authored *GMPLS Technologies: Broadband Backbone Networks and Systems*. He received his B.E., M.E., and Ph.D. degrees from Osaka University.

FRITZ-JOACHIM WESTPHAL (fritz-joachim.westphal@telekom.de) has over 25 years of experience in the field of telecommunications. For 15 years he has been a member of different research departments of Deutsche Telekom. His activities are focused on strategies of new network architectures and infrastructure-based services. In 2009 he joined Telekom Innovation Laboratories (T-Labs), and is currently responsible for the Network Architecture & Modelling research group. He has been involved in different projects on SDN and network virtualization, recently being active in the European FP7 project UNIFY. He has authored or coauthored more than 100 technical conference or journal papers.

Optical Service Chaining for Network Function Virtualization

Ming Xia, Meral Shirazipour, Ying Zhang, Howard Green, and Attila Takacs

ABSTRACT

This article presents an efficient optical service chaining architecture for network function virtualization in data centers. Service chaining (i.e., steering traffic through a sequence of network functions) is one emerging application of software-defined networking. However, existing schemes steer traffic solely in the packet domain, which is well suited for fine-grained (e.g., per-user level) flows carrying a relatively small volume of traffic. This article discusses how packet-based schemes do not yield sufficient efficiency for large/aggregated flows steered through high-capacity core network functions. It introduces an optical steering domain into the operator's data centers for NFV service chaining at a coarse-grained traffic level using wavelength switching. Performance evaluation shows that the optical steering domain can achieve significant power savings compared to using packet technologies as flow rates and the number of vNFs per service chain grow.

INTRODUCTION

To meet the increasing traffic demands while maintaining or improving average revenue per user (ARPU), operators are constantly seeking new ways to reduce their operational expenditure (OPEX) and capital expenditure (CAPEX). To this end, the concept of network function virtualization (NFV) was initiated within the European Telecommunications Standards Institute (ETSI) consortium [1]. Network function virtualization (NFV) calls for virtualization of network functions (NFs) currently provided by legacy middleboxes, such as firewalls, content filters, intrusion detection systems (IDSs), deep packet inspection (DPI), network address translation (NAT), as well as *high-capacity* gateways such as serving/gateway general packet radio service (GPRS) support node (SGSN/GGSN), broadband remote access server (BRAS), session border controller (SBC), provider edge (PE) routers, and so on. Using virtualization and cloud technologies, NFV allows legacy NFs offered by specialized equipment to run in software on generic hardware. Therefore, NFV makes it possible to deploy virtualized NFs (vNFs) in high-performance commodity servers in an operator's data center (DC), with great flexibility to spin on/off the vNFs on demand. In addition, by decoupling the NF software from the hardware, NFV facilitates a faster pace for innovations and results in shorter time to market for new services.

In parallel to NFV, operators have long struggled with the service chaining problem, which is the process of steering a traffic flow across a predefined set of (physical or now virtual) middleboxes (i.e., a service chain). Typically, a service chain is determined depending on the classification of traffic, service-level agreement (SLA), and operator's provisioning policies, and so on. These factors jointly dictate an ordered set of services for each traffic flow to go through. Traditionally, the NFs were implemented in proprietary middleboxes, and deployment of the middleboxes was static. Therefore, service chaining can be realized by static network configuration based on the locations and connectivity among the middleboxes. NFV puts the service chaining problem into a new context with different challenges. For example, NFV allows a vNF to be deployed when and where needed, which requires DC networks to support on-demand dynamic traffic steering. Another challenge is provisioning and energy efficiency. As high-capacity (HC) core NFs (e.g., BRAS) are being virtualized, these NFs are typically given more computing power to handle a large amount of traffic. To achieve high processing efficiency and reduce resource fragmentation, multiple fine-grained flows can be aggregated as a single entity for the service chaining process. The packet-based traffic steering techniques used today become inefficient due to high configuration complexity and energy consumption.

Optical communications have already enabled high-speed transmissions up to terabits per second. One representative technology is dense wavelength-division multiplexing (DWDM), which allows a single fiber to carry tens of wavelength channels simultaneously, offering huge transmission capacity and spectrum efficiency. On the other hand, reconfigurable wavelength switching devices have already been widely deployed in long-haul and metro transport net-

The authors are with Ericsson Research Silicon Valley.

0163-6804/15/$25.00 © 2015 IEEE

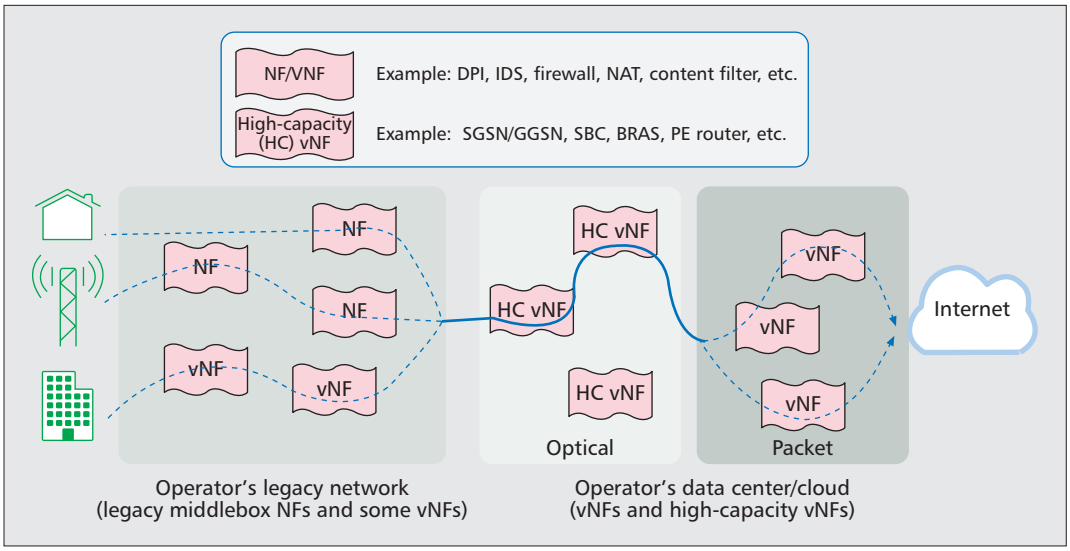

Figure 1. Illustration of service chaining. In the operator's DC/cloud, an aggregated flow first goes through the high-capacity core vNFs in the optical steering domain, and then continues to the fine-grained processing.

works to offer reconfiguration on layer-0 light-path topologies. While today's DCs have not yet fully leveraged optical technologies, it can be envisioned that optics could be enabled up to the end-of-rack (EoR) switches, the top-of-rack (ToR) switches, as well as the servers. Admittedly, switching in the optical domain offers less agility and flexibility than the packet-based approaches; however, it is suitable for the level of dynamicity required by service chains consisting of HC NFs and use of traffic aggregation.

This article shows that traffic steering for aggregated flows can elegantly be supported by optical technologies. In particular, this article presents a packet/optical hybrid DC architecture, which enables steering large aggregated flows in an optical steering domain. We detail the architecture, and describe how this proposal improves the performance of NFV service chaining in contrast to packet-based solutions. Our analysis shows a number of benefits including power savings and scalability.

This article is organized as follows. The next section reviews relevant related works. The subsequent section discusses the limitations of packet-based service chaining (Fig. 1). We then describe the proposed architecture. The following section discusses the benefits and conducts a performance evaluation on power consumption. Finally, we conclude this article.

Related Work

In this section, we briefly review some existing works and discuss how the proposed architecture differs from them. The potential for both CAPEX and OPEX savings has driven a variety of proposals for using optics in DC networks. Wang *et al.* [2] proposed a hybrid electrical-optical network called *c-Through*. Both packet and optical-circuit switches are used to connect ToR switches. Optical switches are employed to provide high-bandwidth connectivity between racks with intensive traffic. Farrington *et al.* [3] designed *Helios*, a hybrid DC architecture based on WDM. The architecture has two tiers: ToR switches are commodity electrical switches, and core switches are either electrical or optical. The optical switches are used for high-bandwidth slowly changing communications between the ToR switches. Chen *et al.* [4] proposed an optical switching architecture (OSA), which is claimed to achieve high configuration flexibility for redefining network topology and link capacity. One major benefit is high bisection bandwidth under dynamic traffic patterns. Our architecture introduces optical switching at the wavelength level, and employs a reconfigurable optical add/drop module (ROADM) as a dispatcher at the entry point of a DC, which schedules incoming traffic to either the optical steering domain or the packet steering domain. Our work is also complementary to the packet-based traffic steering schemes. Today's pure packet-based approaches include policy-based routing or flow steering based on software-defined networking (SDN) technologies such as OpenFlow. Typically, these approaches define rules at the packet flow level for each service chain. Our work also draws inspiration from a rich corpus of work in middlebox management. Sekar *et al.* [5] proposed to run software-centric middleboxes on general-purpose hardware platforms with open application programing interfaces (APIs). Sherry *et al.* [6] proposed a method to deploy middleboxes in the cloud. Joseph *et al.* [7] proposed a policy-aware switching layer for DCs, but it requires installing rules for each new flow, which may not scale. Our architecture builds on top of these packet-based traffic steering proposals. Table 1 makes a comparison between our work and the related works from the angle of using optics in DCs.

Limitations of Packet Switching

NFV traffic patterns are expected to fluctuate substantially compared to the legacy network services, because vNFs can be dynamically creat-

NFV traffic patterns are expected to fluctuate substantially compared to the legacy network services, because vNFs can be dynamically created and destroyed as needed. The configuration of matching rules can be complicated and error-prone as flow number goes large, and may lead to high operational complexity and cost.

	Type	Architecture	Traffic type	Building blocks	Novelty
This article	Hybrid	DC network	Intra-/Inter- DC	WSS, ROADM	Novel WSS interconnection for steering flexibility Wavelength switching for service chaining
C-Through [2]	Hybrid	DC network	Intra-DC	MEMS	Optical manager with algorithms for configuration optimization
Helios [3]	Hybrid	DC network	Intra-DC	MEMS, Mux	WDM enhancement from [3] Modules for circuit optimization
OSA [4]	Optical	DC network	Intra-DC	MEMS, WSS	Optimization for topology/ link capacity
Nagoya & NTT [8]	Optical	Optical switch	N/A	TWC, AWGR	Passive optical switch with high scalability

Table 1. A comparison between the proposed architecture and selected works using optics in DCs.

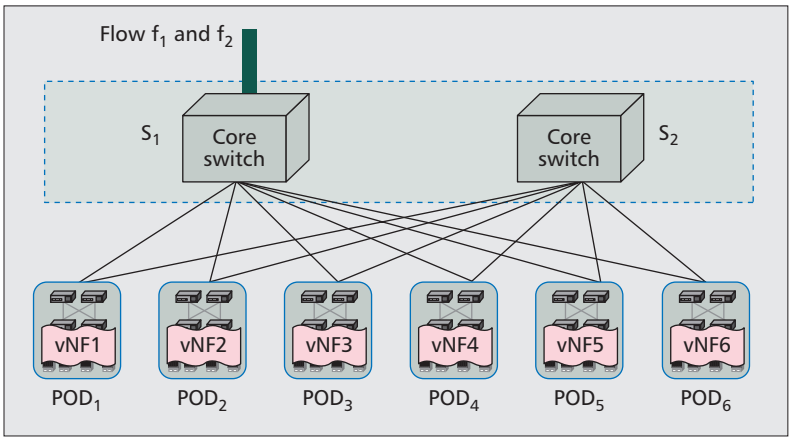

Figure 2. Packet-based traffic steering. (S_2 is added to accommodate the traffic growth of f_1 and f_2.)

ed, resized, moved, and destroyed as needed. The configuration of flow matching rules can be complicated and error-prone as the number of flows grows large, and may lead to high operational complexity and cost. Salability and power efficiency could also become major bottlenecks for performance. Figure 2 illustrates an example in which a DC with six *performance optimized data centers* (PODs) has to grow its network as *traffic rate* and the *number of vNFs perservice chain* increase. Service chaining within a POD is packet-based. Initially, only switch S_1 with 40 Gb/s capacity is deployed to handle two flows, f_1 and f_2, running at 5 Gb/s. It is assumed that the service chain for f_1 includes NF_1 and NF_2, and for f_2 includes NF_3 and NF_4. Two paths are configured as f_1:$S_1 \rightarrow POD_1 \rightarrow S_1 \rightarrow POD_2 \rightarrow S_1$ and f_2:$S_1 \rightarrow POD_3 \rightarrow S_1 \rightarrow POD_4 \rightarrow S_1$, respectively. In this scenario, S_1 has a throughput of 30 Gb/s, which is calculated as flow rate multiplying the times that flows traverse switches.

Next, we assume that both f_1 and f_2 increase to 10 Gb/s, and an additional NF is added to the two chains. Following a similar path configuration, S_1 alone will not be able to provide the needed throughput. When another switch, S_2, is considered, the path configuration for the two flows will be f_1:$S_1 \rightarrow POD_1 \rightarrow S_2 \rightarrow POD_2 \rightarrow S_2 \rightarrow POD_3 \rightarrow S_1$ and f_2:$S_1 \rightarrow POD_4 \rightarrow S_2 \rightarrow POD_5 \rightarrow S_2 \rightarrow POD_6 \rightarrow S_1$. In this case, both S_1 and S_2 have a throughput of 40 Gb/s. This example shows that more switching capacity is needed when flow rates are increased, and/or the number of vNFs on an NF chain is increased.

Architecture Description

Figure 3a illustrates an overview of the proposed architecture. The centralized orchestration, that is, the operations support system/business support system (OSS/BSS) module, interfaces with an SDN controller and a cloud/NFV manager. The SDN controller can be part of the cloud management subsystem or a separate entity. The OSS/BSS module specifies service chaining rules and performs an operator's policy enforcement. The SDN controller and cloud/NFV manager, on the other hand, perform resource provisioning. The southbound interface between the SDN controller and the network elements in the optical steering domain requires support of optical circuit switching. This interface can be realized by using proprietary interfaces provided by hardware vendors, or standardized protocols such as the OpenFlow protocol. OpenFlow v. 1.4 introduces extensions for optical circuit configuration, and the charter of the Open Networking Foundation (ONF) Optical Transport Working Group proposes further relevant extensions.

The cloud manager is responsible for cloud resource allocation and automating the provisioning of virtual machines (VMs) for vNFs. It also includes an NFV management module that handles instantiation of the required vNFs while ensuring correctness of configuration. The northbound interfaces of both the SDN controller and the cloud manager provide application programming interfaces (APIs) to the orchestration layer for injecting rules and parameters specified by the operator. If the SDN controller is a subsystem of the cloud management system, only the latter exposes APIs to the OSS/BSS system. Although a central/unified view of both network and cloud resources is assumed by this architecture, it is possible to delegate resource provisioning to the SDN controller and/or cloud manager for local decision and optimization. The end-to-end service chaining is performed similarly as is existing SDN-based service chaining. The OSS/BSS module needs to request vNFs and network resources, along with the

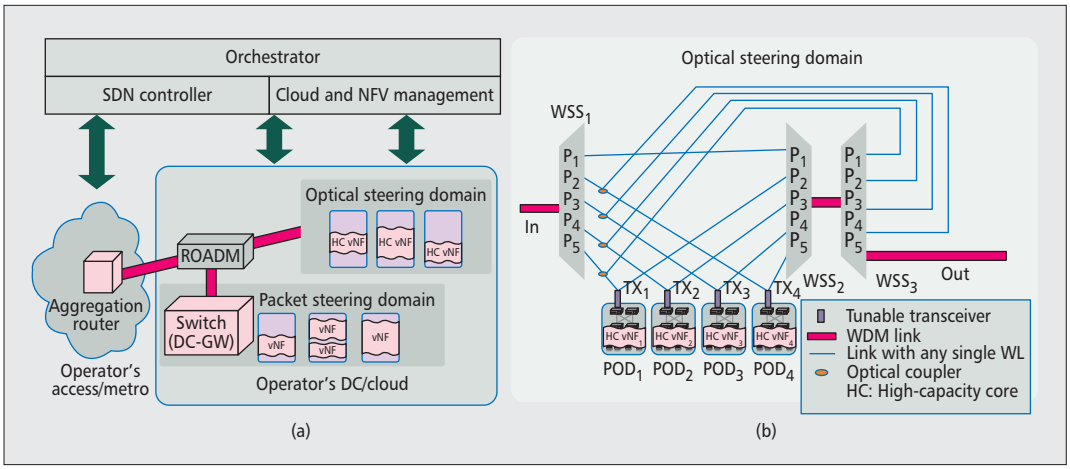

Figure 3. Proposed architecture for optical service chaining: a) overall architecture; b) optical steering domain.

The cloud manager is responsible for cloud resource allocation and for automating the provisioning of virtual machines for vNFs. It also includes an NFV management module that handles instantiation of the required vNFs while ensuring correctness of configuration.

policies to be applied during resource allocation. The SDN controller performs network-resource allocation by relying on a path computation entity that could be integrated with the SDN controller. In the proposed architecture, a service chain may partly reside in the optical domain and partly reside in the packet domain, as illustrated in Fig. 1.

In more detail, the forwarding plane consists of an operator's access/metro networks and an operator's DC/cloud. Traffic flows destined to the same set of high-capacity NFs are aggregated by the aggregation router at the edge of the operator's metro network. Aggregation can be done using multiprotocol label switching (MPLS) or an OpenFlow-based scheme with matching rules on one or a combination of packet field(s). The edge routers are equipped with optical modules to convert aggregate traffic into *wavelength flows*, which then are multiplexed into a fiber link to the DC.

The DC part is packet/optical hybrid, including an optical steering domain in addition to the conventional packet steering domain. The optical steering domain conducts traffic steering for large aggregated flows that otherwise have to be processed using packet technologies. The entry point of the DC is an ROADM. Configured by the SDN controller, the ROADM either forwards a wavelength flow to the optical domain or drops it to the packet domain for fine-grained processing. In the optical steering domain, when a wavelength flow has gone through the needed vNFs, it is steered back to the ROADM. Then the flow can either continue for additional fine-grained NF chaining or leave the DC. Small flows can be directly dropped by the ROADM to the packet domain, where low-capacity or infrequently used NFs are deployed for fine-grained processing. In this case, the flows completely bypass the optical domain. Although multiple factors have impact on NF placement such as inter-NF bandwidth and resource constraints, complementary to all these factors, we recommend that NFs handling large traffic aggregates should be deployed in the optical domain, where the flow rates match the wavelength line rate (e.g., 10 Gb/s or 40 Gb/s).

The optical steering domain serves as the backbone of the DC network to interconnect PODs that host high-capacity NFs. Figure 3b illustrates a reference implementation of the optical steering domain we proposed [9], which employs wavelength-selective switches (WSSs) as the building block for steering wavelength flows. A WSS is a $1 \times N$ active switching device, which has a single common port on one side and multiple tributary ports on the other side. By configuration, each wavelength entering the common port can be directed to one of the N tributary ports, independent of how all other wavelength channels are being switched. In this article, we use $P_z^{WSSy}:w_x$ to indicate that port z of WSS y is configured at wavelength x. There are other technologies that could facilitate the design of an optical steering domain, such as micro-electromechanical systems (MEMS) or liquid crystal on silicon (LCoS). In contrast to the typical design of DC, where the edge switches employ 10GbE SFP+ to interface with the optical steering domain, our design replaces them with tunable optical modules. While both cases need to perform optical/electrical (O/E) conversion, tunable optical modules support wavelength selection to allow more flexibility for wavelength switching. In the example shown in Fig. 3b, one tributary port of WSS_3 (P_5 in this example) is designated as an exit port to steer traffic back to the ROADM. The rest of the tributary ports, configured by the SDN controller, connect wavelength flows to the next vNFs in their service chains through fiber links and optical couplers. An optical coupler is a 2×1 passive device, allowing an optical signal to enter the device from either of the two input ports. However, it is not possible to combine two or more inputs of wavelength at the same frequency into one single-polarization output without significant losses. Therefore, the SDN controller needs to ensure no wavelength contention when flows are looped back. The numbers of WSS and PODs are determined at the dimensioning phase based on traffic volume and service-chaining demands. A potential limitation is the number of concurrent service chains that can be supported, which is limited

WSS_1	$p_1^{WSS_1} : w_3$	$p_2^{WSS_1} : w_2$	$p_5^{WSS_1} : w_1$	/	/
WSS_2	$p_1^{WSS_2} : w_3$	$p_2^{WSS_2} : w_4$	$p_4^{WSS_2} : w_1$	$p_4^{WSS_2} : w_6$	$p_5^{WSS_2} : w_5$
WSS_3*	$p_2^{WSS_3} : w_5$	$p_3^{WSS_3} : w_4$	$p_5^{WSS_3} : w_1$	$p_5^{WSS_3} : w_3$	$p_5^{WSS_3} : w_6$
*Tributary port 5 of WSS3 is assigned as the exit port.					

Table 2. WSS configuration for concurrent traffic steering.

by the number of wavelengths allowed by a fiber link, and by the port count of wavelength switching devices.

The proposed architecture offers concurrent traffic steering for multiple wavelength flows. Assume that there are three wavelength flows in Fig. 3b: flow 1 at wavelength w_1 for vNF_1 and vNF_3; flow 2 at wavelength w_2 for vNF_4 and vNF_2; and flow 3 at wavelength w_3 bypassing all the vNFs in the optical steering domain. In addition, wavelengths w_4, w_5, and w_6 are available for intermediate assignment. A possible configuration without incurring wavelength contention can be $w_1 \rightarrow vNF_1 \rightarrow w_4 \rightarrow vNF_3 \rightarrow w_1$ for flow 1, $w_2 \rightarrow vNF_4 \rightarrow w_5 \rightarrow vNF_2 \rightarrow w_6$ for flow 2, and $w_3 \rightarrow w_3$ for flow 3. This configuration also keeps the exiting wavelength of flow 1 the same as its entering wavelength, such that the optical steering domain appears transparent to flow 1. Table 2 shows the configuration of the WSSs, which is pushed down to the WSSs by the SDN controller.

Advantages and Power Consumption Analysis

This section discusses several key advantages of the proposed architecture.

Flexibility

The proposed architecture allows flexible configuration of wavelength paths as vNFs are dynamically created and destroyed. Compared to the time to instantiate a vNF (typically a few minutes on a standard VM), the tuning time of WSS is around a few hundred milliseconds, short enough for on-demand NFV service-chain provisioning. An aggregated flow has a relatively longer duration and more stable NF chain demand than a granular flow. Hence, the flexibility offered by wavelength switching matches the required level of dynamicity for service chaining through core NFs. Coordination between the SDN controller and the cloud/NFV manager further increases the flexibility for wavelength assignment, opening possibilities for resource optimization. For example, when wavelength contention blocks a service chain from going through a particular POD, the cloud manager can choose a different POD with both wavelength and NF available. In another scenario, when traffic volume is more than wavelength capacity, multiple wavelengths can be employed. A load balancing scheme can be used to distribute the wavelengths to multiple PODs with optimization of server utilization.

Scalability and Reduced Operational Complexity

The optical steering domain does not perform any packet forwarding above the edge switches of a POD, and hence saves a substantial amount of packet switching. Because of its analog nature, a wavelength is agnostic of the traffic it carries, which makes wavelength switching independent of transmission rate. Therefore, the same WSS can be used to transparently support flow-rate increases (e.g., from 2.5 Gb/s to 10 Gb/s). This feature makes the optical steering domain relatively insensitive to traffic growth. In contrast, traffic growth in the packet domain would require line cards and chassis (including switching fabric) to be upgraded to offer higher throughputs. The added scalability also results in less operational complexity. Compared to configuring service chains for a large number of flows, a wavelength flow only requires a single path in the optical steering domain. Without the need to define per-flow packet rules, the optical steering domain significantly reduces the complexity of networking configuration, and meanwhile saves effort in monitoring and error detection.

Power Efficiency

We conduct a simulation-based study to understand the power performance of the proposed architecture in various settings. The per-port power consumption at different rates is listed in Fig. 4a, where the numbers are obtained from measurement on a real DC system. The power consumption is based on the system level, that is, including all the components and peripheral circuits (transceivers, line cards, chassis, and etc.). In contrast to conventional DCs, which employ "gray light" transceivers (e.g., 10G SFP+)at edge switches, the proposed architecture uses DWDM 10G tunable transceivers with similar power consumption. For example, both the 10G DWDM laser and gray SFP+ have 1.5 W power consumption [10]. Currently, tunable transceivers at higher rates such as 40G and 100G are not yet commercially available for DCs. Therefore, we make an assumption that they have similar power consumption as their gray counterpart based on the 10G case. We also assume an economical realization of a 2-degree ROADM (the add/drop port is not counted as a degree) by using two WSSs.

Although a vNF may require multiple resources, such as CPU cycles, bandwidth, and storage, we consider a single resource type (e.g., CPU cycle) to simplify the resource constraint, so our study can focus on power consumption. Power consumption is calculated based on the

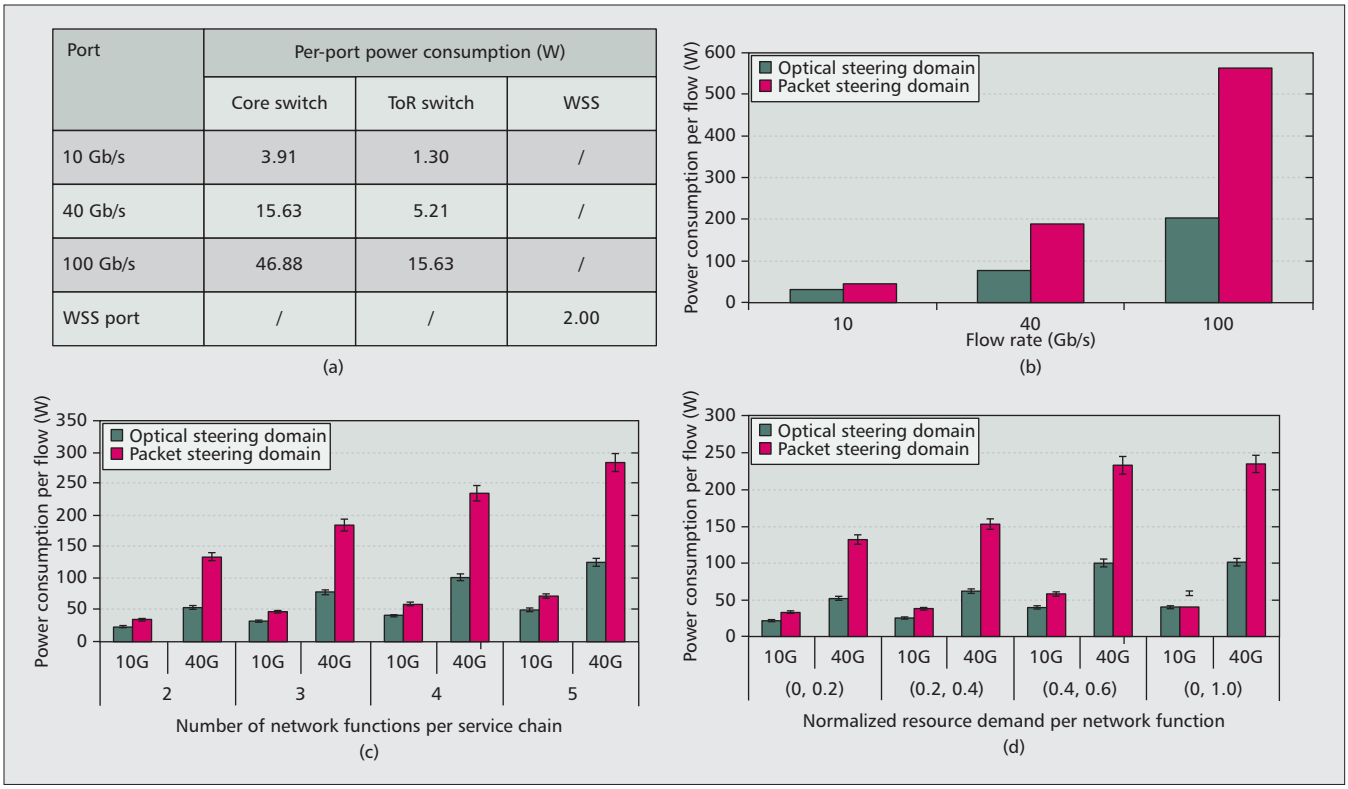

Figure 4. Per-port power consumption and performance evaluation: a) per-port power consumption (W); b) flow rate vs. power consumption per flow; c) number of vNFs vs. power consumption per flow; d) normalized resource demand per vNF vs. power consumption per flow. 95 percent confidence intervals in c) and d).

assumption that network equipment is fully utilized; otherwise, the numbers presented below will be higher due to more port power consumption. Figure 4b compares the per-flow power consumption in the packet and optical steering domains as flow rates increase. In this experiment, each vNF requires 1.0 unit of CPU resource, which is equal to what a POD can maximally offer. Therefore, a POD can host at most one vNF. Each flow needs to be steered through four vNFs. We can observe that the optical steering domain is always more power-efficient than using the packet steering domain. In addition, the gap between the two domains becomes larger as the traffic rate increases.

Next, we study the impact of the number of vNFs per service chain on power consumption. We choose10 Gb/s and 40 Gb/s as flow rates, and vary the number of vNFs per service chain from 2 to 5. For each vNF, the resource demand is uniformly generated between (0, 1.0]. We notice that if the sum of several vNFs' resource demands is less than the POD capacity, these vNFs can be placed within the same POD. This strategy reduces the inter-POD traffic as well as the throughput of the backbone steering domain. In this experiment, we employ a simple "first-hit" scheme to consolidate vNFs within fewer PODs. The first-hit scheme works as follows: for each vNF to be placed, always choose the first POD that has already hosted some vNFs and has enough resources. Figure 4c shows the per-flow power consumption as the number of vNFs grows. From this figure we can make several observations:

- Optical steering is always more power-efficient than packet steering in the given scenarios.
- The overall trend shows higher power consumption as the number of vNFs increases, which is due to more traversals in the backbone steering domain.
- The advantage of the optical steering domain becomes more significant as the number of vNFs grows.

Figure 4d shows the impact of resource demand per vNF on power consumption. We set the number of vNFs for each flow as 4, and use 10 Gb/s and 40 Gb/s as flow rates.The resource demand is uniformly distributed in (0, 0.2], (0.2, 0.4], and (0.4, 0.6], and we use(0, 1.0] as a reference. We do not include the range of (0.6, 1.0], as this would not leave any room for sharing a POD by multiple vNFs. Intuitively, when per-vNF resource demand is low, there is more chance to consolidate the vNFs of the same service chain into fewer PODs, and hence fewer traversals in the steering domain are needed. Figure 4d shows more benefits of using the optical steering domain as resource demands per vNF grow. Since advances in both hardware and software will eventually allow a vNF to perform high-volume processing, it will lead to more power savings using the optical steering domain. On the other hand, when packet-based steering has to be chosen for small-sized vNFs, its advantage of fine-grained flow control will be discounted by the low power efficiency of packet technologies. Finally, as in any long-

Our study shows that the optical steering domain achieves significant power savings compared to packet-based steering in various settings, as the flow rate and the number of vNFs per service chain increase.

term network operation, the CAPEX will gradually be surpassed by OPEX, employing an optical steering domain in a DC offers good potential in cutting energy costs.

Conclusion

Network function virtualization supports efficient resource utilization by dynamically spinning up/down virtual network functions (vNFs). Service chaining through high-capacity vNFs requires the efficient handling of large aggregated traffic flows. This article introduces an architecture for optical service chaining in data centers, which offloads large flows to an optical domain for steering across core vNFs at the wavelength level. By employing reconfigurable wavelength switching devices and a simple fiber-loopback scheme, this architecture offers high scalability and flexibility for on-demand wavelength service path configuration. It also supports fine-grained steering by dispatching traffic to a conventional packet steering domain. Our study shows that the optical steering domain achieves significant power savings compared to packet-based steering in various settings, as the flow rates and the number of vNFs per service chain increase.

References

[1] ETSI NFV ISG, "Network Functions Virtualization, An Introduction, Benefits, Enablers, Challenges & Call for Action," White Paper; http://portal.etsi.org/NFV/NFV_White_Paper.pdf, Oct. 2012.

[2] G. Wang *et al.*, "c-Through: Part-time Optics in Data Centers," *Proc. ACM SIGCOMM*, 2010.

[3] N. Farrington *et al.*, "Helios: A Hybrid Electrical/Optical Switch Architecture for Modular Data Centers," *Proc. ACM SIGCOMM 2010*, pp. 339–50.

[4] K. Chen *et al.*, "OSA: An Optical Switching Architecture for Data Center Networks with Unprecedented Flexibility," *Proc. USENIX/NSDI*, San Jose, CA, April 2012.

[5] V. Sekar *et al.*, "The Middlebox Manifesto: Enabling Innovation in Middlebox Deployment," *Proc. 10th ACM Wksp. Hot Topics in Networks*, 2011.

[6] J. Sherry *et al.*, "Making Middleboxes Someone Else's Problem: Network Processing as a Cloud Service," *Proc. ACM SIGCOMM*, 2012, pp. 13–24.

[7] D. Joseph et al., "A Policy-Aware Switching Layer for Data Centers," *Proc. ACM SIGCOMM*, 2008, pp. 51–62.

[8] K. Sato *et al.*, "A Large-Scale Wavelength Routing Optical Switch for Data Center Networks," *IEEE Commun. Mag.*, vol. 51, no. 9, Sept. 2013.

[9] M. Xia *et al.*, "SDN and Optical Flow Steering for Network Function Virtualization," Open Networking Summit research tracks, Santa Clara, CA, Mar. 3–5 2014.

[10] Cisco Transceiver Modules; http://www.cisco.com/c/en/us/products/collateral/interfaces-modules/dwdm-transceiver-modules/data_sheet_c78-711186.html.

Additional Reading

[1] N. Bitar *et al.*, "Technologies and Protocols for Data Center and Cloud Networking," *IEEE Commun. Mag.*, vol. 51, no. 9, Sept. 2013, pp. 24–31.

[2] Z. Ayyub Qazi *et al.*, "SIMPLE-Fying Middlebox Policy Enforcement using SDN," *Proc. ACM SIGCOMM*, 2013.

Biographies

Ming Xia received his Ph.D. degree in computer science from the University of California, Davis, in 2009. He is currently a senior research engineer at Ericsson Research, San Jose, California. He was an expert researcher at the National Institute of Information and Communications Technology, Japan. He serves as Associate Editor for the *Elsevier Journal of Telecommunication Systems*, and a Guest Editor of the *Journal of Computers & Electrical Engineering* Special Issue on Ubiquitous Computing and Communications. His research interests include computer networks and cloud/data centers.

Meral Shirazipour received a Bachelor's degree in computer engineering from Concordia University in 2002, and M.Sc.A and Ph.D. degrees in telecommunications from Ecole Polytechnique de Montreal in 2004 and 2010, respectively. She is currently part of the research staff at Ericsson Research Silicon Valley. Her research interests include distributed cloud infrastructures, network function virtualization, software-defined networking, measurement, and optical transport network control, and, more generally next generation networks, new networking architectures, protocols, and standardization.

Ying Zhang received her Ph.D. degree from the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, in 2009. She is a researcher in the IP and Transport Research group, Ericsson Research Silicon Valley. Her research interests are in networking and systems, including software-defined networking, cloud, Internet and cellular network management, Internet routing, and measurement and network security.

Howard Green joined Ericsson by way of the Marconi acquisition in 2006, and is currently working in broadband and transport research. He has been in the industry since 1980, working in research, development, and business strategy for many technologies, from public switching to SDH and photonics. He has a Ph.D. in mathematics and an M.B.A., both from Warwick University, United Kingdom.

Attila Takacs is a research manager of Ericsson Research in San Jose. He has been the technical lead of research projects on SDN, OpenFlow, GMPLS, traffic engineering, PCE, IP/MPLS, Ethernet, and OAM for transport networks. He is also an active contributor to standardization; in particular, he has worked in ONF, IETF, and IEEE. He holds more than 30 international patent applications granted and in progress. He holds an M.Sc. in computer science and a postgraduate degree in banking informatics, both from Budapest University of Technology and Economics, Hungary. He has an M.B.A. from the CEU Business School in Budapest.

A Service-Oriented Hybrid Access Network and Clouds Architecture

Luis Velasco, Luis Miguel Contreras, Giuseppe Ferraris, Alexandros Stavdas, Filippo Cugini, Manfred Wiegand, and Juan Pedro Fernández-Palacios

ABSTRACT

Many telecom operators are deploying their own cloud infrastructure with the two-fold objective of providing cloud services to their customers and enabling network function virtualization. In this article we present an architecture we call SHINE, which focuses on orchestrating cloud with heterogeneous access and core networks. In this architecture intra and inter DC connectivity is dynamically controlled, maximizing the overall performance in terms of throughput and latency while minimizing total costs. The main building blocks are: a future-proof network architecture that can scale to offer potentially unlimited bandwidth based on an active remote node (ARN) to interface end-users and the core network; an innovative distributed DC architecture consisting of micro-DCs placed in selected core locations to accelerate content delivery, reducing core network traffic, and ensuring very low latency; and dynamic orchestration of the distributed DC and access and core network segments. SHINE will provide unprecedented quality of experience, greatly reducing costs by coordinating network and cloud and facilitating service chaining by virtualizing network functions.

Luis Velasco is with Universitat Politecnica de Catalunya (UPC).

Luis Miguel Contreras and Juan Pedro Fernández-Palacios are with Telefonica Investigacion y Desarrollo (TID).

Giuseppe Ferraris is with Telecom Italia.

Alexandros Stavdas is with the University of Peloponnese (UoP).

Filippo Cugini is with Consorzio Nazionale Interuniversitario per le Telecomunicazioni (CNIT).

Manfred Wiegand is with Coriant.

INTRODUCTION

A revolution in access networks is underway. The revolution is driven by the continued transformation of cellular networks offering to portable devices bit-rates and quality of service (QoS) comparable to those traditionally made available only through fixed networks. Driven by demand for video and the proliferation of data centers (DC), more than 75 percent of that traffic will stay in access/metro networks by 2017, as compared to 57 percent today, as forecast in [1]. Accommodating the enormous traffic growth in a cost-effective and service-efficient way is essential for the viability of telecom operators and motivates a major network re-design. In fact, these shifting traffic patterns are the result of bringing content closer to the users to better manage quality of experience (QoE). For instance, the most popular video content can be cached and delivered to users locally over access/metro networks rather than being accessed from a central cache over the backbone network.

Coordinating these and new advanced services to be offered to a widely distributed number of customers requires building advanced service chains. Nevertheless, building service chains is very expensive and time consuming, since it requires, besides deploying dedicated hardware for each required network function, configuring each device using its proprietary command syntax, interfaces, and protocols. Moreover, since loads often change over time, building a new service chain typically requires estimating future demands and over-provisioning IT and network resources to support growth. This prevents operators from reducing the final price that users pay, undermining the average revenue per user (ARPU).

The continuous advances in computing hardware facilitate real-time processing to be performed on commodity hardware instead of specialized hardware. These advances enable network function virtualization (NFV) [2]. By eliminating specialized network processors, multiple heterogeneous workloads can be consolidated onto a single architecture, thus reducing complexity and simplifying operation, leading to total cost of ownership (TCO) reductions.

Cloud technology offers numerous benefits including economies of scale, cost-effectiveness, efficient hardware utilization, and TCO reductions, both in capital and in operational expenditures [3]. These benefits are all key objectives for telecom operators, so the appeal of cloud technologies is clear. In fact, as revealed in a recent survey [4], many telecom operators are deploying cloud infrastructures. Notwithstanding, deploying the telecom cloud presents a different set of challenges due to the industry's inherent requirements for availability (5-nines), very low latency, and complex networking (Ethernet, optical, wireless, etc.).

Scalability is also an issue since, in contrast to a small number of warehoused-sized DCs commonly used in public clouds, telecom cloud must support a large number of small, distributed DCs to reduce traffic in the core network. A distributed DC architecture brings many benefits for network operators. By encapsulating workloads in virtual machines (VM), a cloud resource manager can migrate workloads from one DC to another looking to improve the perceived quality of experience (QoE), reducing energy consump-

0163-6804/15/$25.00

SHINE includes an orchestrated management plane to provide elastic and resilient cloud and network resource provisioning, combining resources in geographically separated m-DCs. Dynamic network resource allocation will combine both, flexgrid core and access networks according to traffic needs.

tion [5], or even in response to situations such as network failures or high-demand events. In addition, placing DCs closer to end-users enables the development of services and applications that can take advantage of very low latency.

The efficient integration of cloud-based services under a distributed DC architecture, including the interconnecting network, is a challenging task due to the required performance and high availability guarantees. The answer from network providers to the increasing traffic dynamism is to migrate their networks to a cloud-ready transport network [6], as a platform able to handle dynamic traffic patterns and asymmetries. Although this approach enables a more elastic transport infrastructure, it has technical challenges on its own that must be addressed. In the recent work [7], the authors propose to use the dynamic connectivity provided by the flexgrid optical technology to improve resource utilization and save costs. The flexgrid technology enables a finer spectrum granularity adaptation and the ability to dynamically increase and decrease the amount of optical resources assigned to connections. The availability of flexgrid ready spectrum selective switches enables building bandwidth-variable optical cross-connects (OXC), whereas the advent of sliceable bandwidth-variable optical transponders (SBVT), able to deal with several flows in parallel, adds even more flexibility and reduces costs [8].

In the access, higher speeds together with multiple data plane interfaces will drive the evolution of aggregation elements to multi-service nodes, abstracting capabilities from data plane specificities. The necessary support of legacy services and interfaces and the multi-service scope for those devices motivates the definition of programmable control, adapting the generic conception of the node to the specific need.

That control has to consider both service and transport characteristics to orchestrate resources end-to-end. The advent of software-defined networking (SDN) is fueling the deployment of programmable control methods. In fact, several initiatives are currently under way to define architectural frameworks for centralized control elements, such as the OpenDayLight project[1] or the Application Based Network Operations (ABNO) architecture [9]. The OpenFlow protocol[2] is well suited to handle transmission specifics and intra-DC connectivity [10]; for example, extensions to OpenFlow can be defined to configure SBVTs. In contrast, some telecom operators might prefer using ABNO to control interconnection networks since it is based on working functional elements and facilitates network re-configuration [11].

In this article we present a Service-oriented HybrId access Network and Cloud ArchitecturE (SHINE) that orchestrates cloud with heterogeneous access and core networks, dynamically controlling intra and inter DC connectivity. A number of challenges associated with end-to-end coordination and the migration from the existing networking framework need to be faced. Separation of service and transport oriented tasks are key to allow a scalable orchestration, facilitating its independent evolution; clear interfaces and taxonomy of functions is required.

A set of NFV use cases have been recently identified by the NFV group within ETSI [12] and several initiatives are being developed in that field, with a relevant number of proofs-of-concept in place.[3] In addition, the recently launched mobile-edge computing (MEC) initiative aims at adding cloud-computing capabilities at the edge of the mobile network.[4] Notwithstanding, because of its versatile and adaptable architecture, SHINE offers a common infrastructure to deploy many different NFV scenarios.

SHINE Architecture

SHINE proposes a new optical architecture capable of fulfilling the requirements in terms of capacity and dynamicity of future access networks bypassing metro aggregation layers currently deployed (Fig. 1a). An active remote node (ARN) serves as a gateway for a number of heterogeneous networks and uses transmission and multiplexing to incorporate traffic from large geographic areas (rural and urban) directly to the core network. The ARN directly interfaces OXCs in the core network by means of point-to-point connections through dedicated links exploiting adaptive modulation formats to capitalize on their distance adaptive transmission properties (Fig. 1b).

A number of μ-DCs are placed in some core locations to accelerate content access times and to reduce core network traffic. μ-DCs are geographically distributed and connected through a flexgrid core network to behave as one single large DC. Large DCs can also co-exist to feed μ-DCs with contents.

SHINE also includes an orchestrated management plane to provide elastic and resilient cloud and network resource provisioning, combining resources in geographically separated μ-DCs. Dynamic network resource allocation will combine both flexgrid core and access networks according to traffic needs. The architecture of the SHINE's ARN and μ-DC is illustrated in Fig. 2.

Access Network System

The SHINE's ARN works as a protocol termination point where frame aggregation is implemented using either IP/MPLS, Ethernet, or OTN platforms (Fig. 2a). The conceived functionality offers the potential for a service transparent solution whenever this is needed (e.g. for mobile front/back hauling). The main building blocks of the ARN include:

- Access interface implementation by means of transceiver modules. 10 GbE modules serving point-to-point connections from a multitude of services are envisaged, as well as 10 GbE PON for residential access.
- Upstream interface implementation by means of SBVTs allowing access to a number of client signals.
- OpenFlow switching and programmable network processing backplane.

μ–DCs

The SHINE's μ-DC architecture aims at creating an energy-efficient cloud infrastructure while keeping latency ultra-low. To that end, one sin-

[1] OpenDayLight: http://www.opendaylight.org/

[2] Open Networking Foundation: https://www.opennetworking.org/

[3] http://nfvwiki.etsi.org/index.php?title=Ongoing_PoCs

[4] https://portal.etsi.org/Portals/0/TBpages/MEC/Docs/MEC%20Executive%20Brief%20v1%2028-09-14.pdf

gle multi-granular switch is used to connect every server in the μ-DC, using 100 Gb/s optical interfaces, to the interconnection network (Fig. 2b).

The multi-granular switch is able to switch packets, flows, and/or optical signals from/to the SBVTs, enabling configurable multiplexing toward the SBVT front-end cards. Based on the modulation formats supported by the SBVTs, the multi-granular switch can be configured to aggregate heterogeneous client lower granularity packet flows (by performing full electronic packet processing) or entire optical flows (by performing optical port-to-port forwarding) to a given tributary signal having a certain destination (e.g. a remote μ-DC). Such flexibility, orchestrated by the local SDN controller, allows the adoption of energy-efficient grooming strategies, aiming at reducing the impact of electronic processing only where and when needed. As an example, if enough traffic is generated by a local server toward a remote μ-DC, such traffic can be assigned to single or multiple flows. Such flows can be optically switched directly toward one or more tributary lines of the SBVT that reaches the remote μ-DC, thus bypassing electronic processing. A local SDN controller is in charge of such optimization by automatically and dynamically configuring the flow entries of the multi-granular switch.

Orchestrated Service Management and Control

SHINE considers the deployment of an orchestrated service management and control architecture spanning along the μ-DCs (Fig. 3); this architecture leverages the ABNO framework for the interaction with the transport infrastructure. A parent module is in charge of the overall coordination of cloud and networking resources, being the common entry point for services. Specific management and control modules at the μ-DC level are in charge of the resources internal to a given μ-DC, whereas ABNO coordinates both optical nodes in the core and ARNs in the access network. See [13] for details on the iteration between components.

Components of *service management and control* are:

- The scheduler, which assigns VMs to servers seeking to use resources effectively and achieve the target QoE. In addition, energy efficiency can result from energy-aware VM scheduling and server consolidation.
- The QoE estimation module estimates parameters related to the QoE experienced by end-users, mainly delay.
- The statistics and monitoring module gathers information regarding the use of resources and the performance of services to be used to predict likely scenarios.

Components of *network management and control* are:

- The network control module issues commands to μ-DC level modules and the ABNO looking to create virtual networks among VMs running in one or more μ-DCs.
- The SDN controller is in charge of intra μ-DC network resources and controls both the multi-granular switch and the SBVTs installed on it.

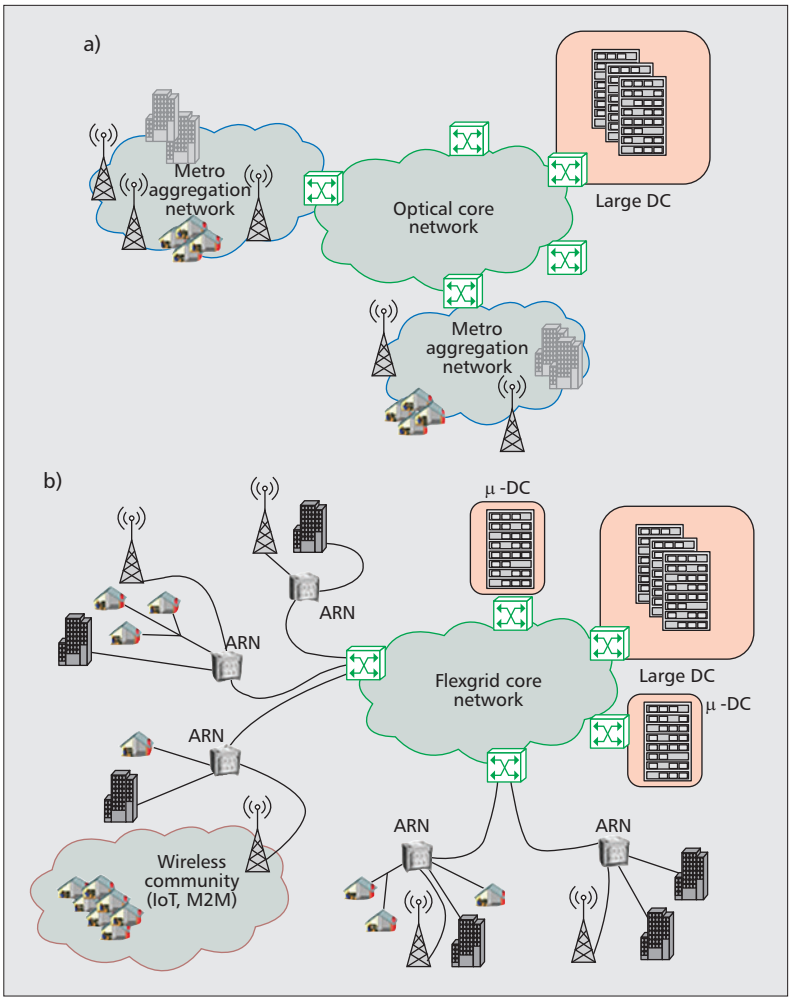

Figure 1. a) Current and b) SHINE network scenario. SHINE combines hybrid access and a distributed DC connected through a flexgrid core network.

- The ABNO module is in charge of the connections among DCs and from them to ARNs in the access network.

SHINE IaaS in Support to NFV

IaaS-based cloud services can be offered on top of the resulting SHINE architecture, were VMs can run on the servers available in the μ-DCs and large DCs, while connectivity can be created to connect VMs belonging to the same client, disregarding their placement. The same infrastructure can be shared for NFV applications, where in addition to servers in DCs, VMs supporting NFV applications can run in ARNs. This creates three levels with different characteristics:

- ARNs can host those functions that require proximity to the end-users because of latency or to aggregate data from a reduced number of sources.
- μ-DCs offer a good trade-off between latency and end-user proximity.
- Large DCs offer economies of scale and can be used for those delay tolerant services.

The service manager module is in charge of managing dynamic optimal VM placement and, once decided, optimal virtual networks are created or reconfigured using online optimization

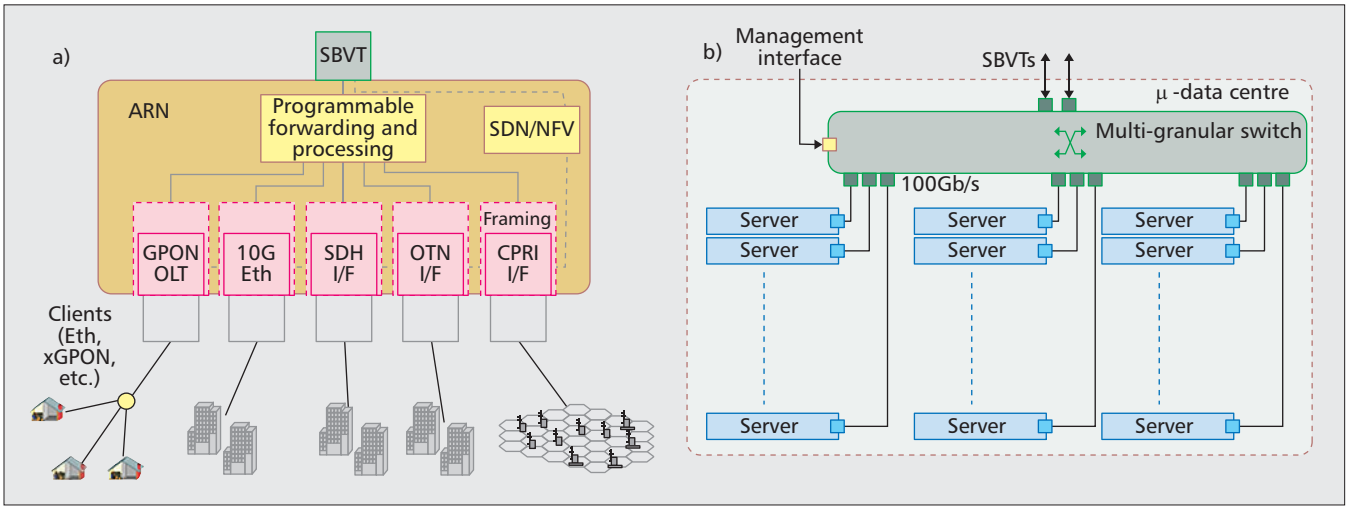

Figure 2. SHINE's ARN (a) and μ-DC (b) architecture.

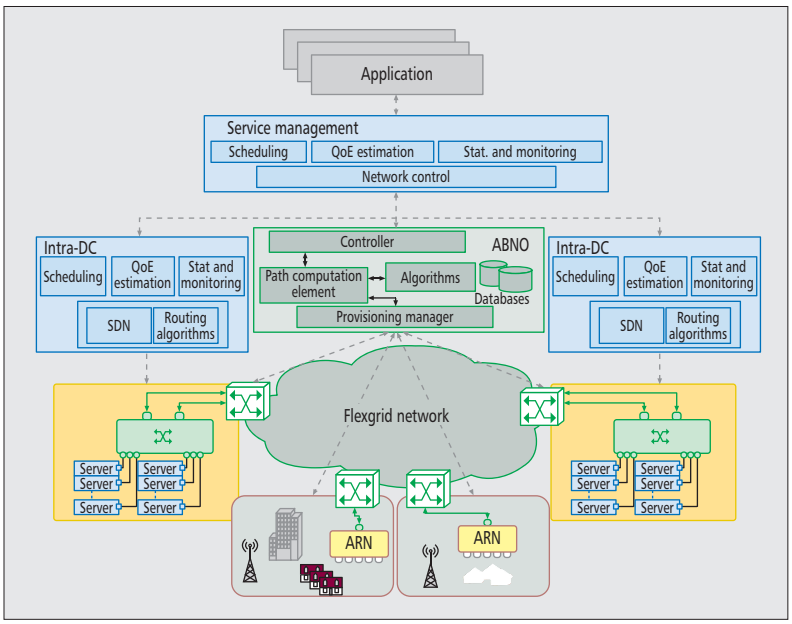

Figure 3. SHINE's management architecture.

algorithms. This *self-management* allows applications to be deployed based on SLA agreements, including QoE parameters (e.g. max user delay, max average delay, etc.). For scalability, application administrators can request elastic operations being applied to *scale-out* by adding more VMs, *scale-up* adding more resources to a VM, and *scale-down* their cloud services.

All the above can be used to support NFV applications. Next we present several use cases to illustrate the automated composition and allocation of computing and network resources and the interaction with the rest of the elements in the SHINE architecture.

Use Cases

This section presents potential use cases addressing different markets of interest for any network operator, namely content delivery, and business, mobile, and fixed broadband access. For that goal, we extract the corresponding use cases from the set defined in [12], aiming to briefly describe how these NFV scenarios fit into the proposed SHINE architecture.

Use Case I: Content Delivery Network

Content delivery networks (CDN) incorporate a number of components, e.g. cache nodes, that are orchestrated by a controller. The CDN controller is a centralized component that selects a cache node to serve an end-user request, and then redirect the end-user to the selected cache node; selecting nodes closer to the end user reduces traffic in the core network and enables delivering higher quality multimedia flows. CDN cache nodes are distributed within the network and are currently deployed as dedicated physical appliances or software running on dedicated hardware.

CDN cache nodes can be virtualized to run on VMs placed in μ-DCs and/or ARNs, whereas the CDN controller can run in large DCs. Based on SLA agreements, VMs encapsulating CDN nodes are autonomously placed by SHINE's service manager to meet QoE parameters. In addition, the performance and load of the CDN nodes need to be monitored by their own CDN service administrator so as to elastically adapt the deployed nodes to the current service needs. In case the load of some cache nodes reaches an upper threshold, elastic operations to scale-up specific VMs or to scale-out to add new VMs encapsulating cache nodes can be requested from SHINE's service manager. In contrast, when the load decreases, opportunities appear to reduce the resources (CPU or memory) available to some VMs or to consolidate workload in a few VMs. As before, the CDN service should detect these opportunities to request the proper configuration from the SHINE's service manager. Finally, it is worth noting that each CDN service runs isolated from other services, so several CDNs can be deployed, where cache nodes share the underlying infrastructure.

Finally, one of the main characteristics of the SHINE service manager is its ability to reactively

reconfigure deployed services in the event of QoE degradation and even proactively reconfigure them to improve overall performance. As an example, let us imagine that a failure in a link in the core network has triggered restoration, and the length of the restoration path suddenly causes the measured delay from the users to the serving application to increase past a given threshold. In that case, SHINE's service manager re-computes optimal VM placement to meet the committed QoE, which might result in decreasing the traffic through the restoration path, thus saving resources in the core network.

Use Case II: Business Access

The business connectivity market is typically characterized by the provision of isolated virtual private networks (VPN) to a variety of geographically dispersed access points. Even more, business services are demanding connectivity to some form of cloud networking, accessing either private, public, or hybrid clouds.

Specific network elements are deployed at customer premises, but also at the access point of presence (PoP) to collect such traffic, ideally aggregating a high number of enterprises demanding similar services. At the customer side, not only the customer edge (CE) equipment, but also some other devices like firewalls, could be in place. On the network side, what is required at a minimum is the deployment of a provider edge (PE) router for customer access. The new trends in NFV can facilitate the virtualization of such network elements by instantiation of network functions. This can have a direct impact on service savings, since for managed services, both the CE's cost and its operation are entirely allocated to the enterprise customer.

On the other hand, the cost of a PE can be shared among the customers connected to it. However, in practice, overprovisioning is required to provide the needed service coverage, thus it is necessary to deploy a huge number of PE equipment for enterprise access, which in reality exceeds the connectivity demand per area. This complex trade-off between service footprint and adequate platform dimensioning could be highly optimized by rolling-out virtualized PE functionality.

Considering the SHINE architecture, CEs can be deployed at the ARNs, together with additional functions if needed (e.g. firewalls, as mentioned before). As for the PE function, it can be located deeper in the network, at the core level, or it can even be distributed to the ARNs in case of scalability concerns. In the centralized case, the underlying flexgrid transport network can guarantee the required SLAs for the service. Regarding the cloud resources complementing the business service, they can be placed and moved among DCs (large and μ) according to actual service needs. Fig. 4 shows a potential deployment scenario.

From the operation viewpoint, a VPN application should be placed on top of the SHINE's service management and control module for *programming* forwarding rules among virtual CE functions residing in ARNs across the network, with the necessary isolation among customers. The intelligence needed for routing among customer branches or some other rich functions, for example, network address translation (NAT), will be part of the VPN application, which interacts properly with the central ABNO controller for accomplishing end-to-end services. Additionally, in the case of hybrid cloud services, the needed orchestration with the DC infrastructure can be managed from the service management and control module in a transparent way to the customer. This interaction would allow for elastic cloud services and isolation between customer and operator cloud management, for example, when moving VMs among DCs [14, 15].

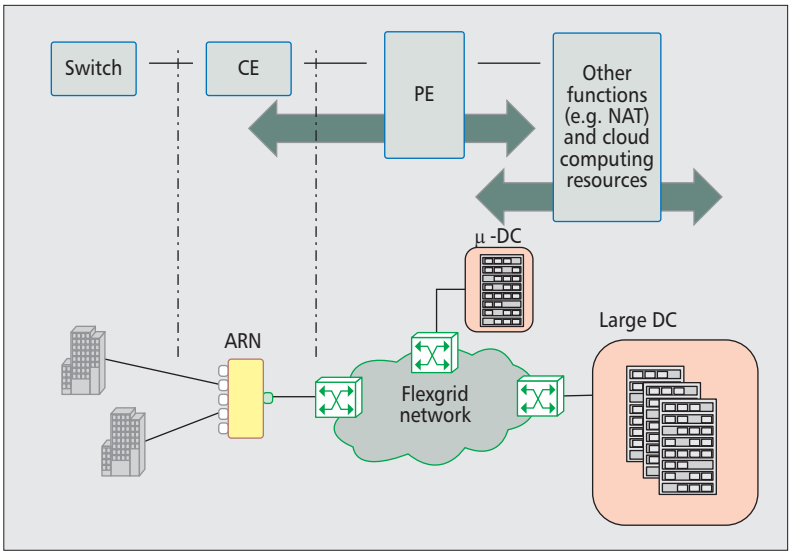

Figure 4. Business access scenario.

Use Case III: Mobile Broadband Access

Mobile access networks are of particular interest to network operators because of the high capacity and capillarity they require to satisfy end-user expectations; this will become even more evident with the advent of 5G wireless networks. This scenario forces operators to explore new ways of deploying the necessary infrastructure to fulfill end-user requirements in a cost effective way.

One of the recent trends in the mobile industry is the centralization of some functions of the radio access network (RAN), named the centralized-RAN (C-RAN) approach. C-RAN proposes allocating common radio-access processing resources, base-band units (BBUs), currently deployed in mobile stations, in a central node, while just keeping remote only the infrastructure strictly needed to provide the wireless connectivity, that is, the radio remote units (RRUs).

The flexibility of C-RAN can be further extended by virtualizing the BBU functionality. Fig. 5 suggests a mapping of the C-RAN approach to the SHINE architecture. In this case, the BBU is deployed inside the ARNs and connected to the RRUs in the coverage area defined for this service. Such connection is implemented by means of high-speed common public radio interface line cards, supported by the ARN architecture.

Communication between two mobile stations allocated to the same BBU is performed through

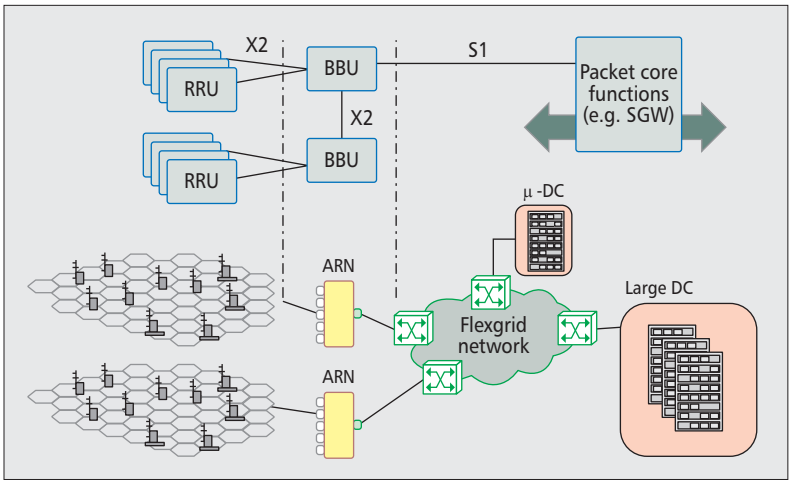

Figure 5. Mobile broadband access scenario.

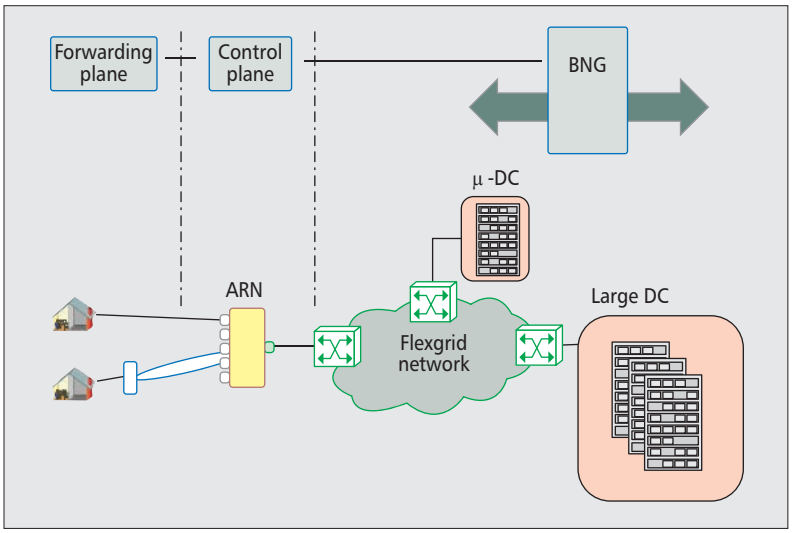

Figure 6. Fixed broadband access scenario.

an X2 interface implemented directly in the ARN. When mobile stations are attached to distinct ARNs, connectivity is performed through the flexgrid core network available in the SHINE architecture, thus minimizing latency and guaranteeing quality indicators.

The SHINE's service management and control module will be responsible for handling all the necessary connections in the access to ensure service provision. This involves not only the connections for the X2 interface, but also the connections needed for the S1 interface that allows communication with the mobile packet core elements placed deep in the network. In fact, such core elements, for example, the mobility management entity (MME), the serving gateway (SGW), or the packet data network gateway (PGW) in the LTE architecture, can also be deployed in the form of virtualized functions to run either in the large DC or the μ-DCs, depending on the required scalability.

An application running on top of the SHINE's service management and control module requests function deployment and the overall connectivity to the system.

Use Case IV: Fixed Broadband Access

Similar to mobile access, network operators are considering the viability of centralizing certain fixed broadband processing capabilities in access nodes, looking to simplify those network elements that provide connectivity to the end users. This simplification, ideally implementing just programmable forwarding capabilities, will clearly impact cost savings and service flexibility.

By doing so, access service provisioning can be highly simplified; end user connectivity will only be a matter of how much bandwidth the available infrastructure (e.g. xDSL, FTTx) provides, while centralizing all service logic. Aspects like QoS configuration, traffic filtering and prioritization, and so on, can be governed independently of the underlying technology. New service creation would only imply updating service logic in a central point, whereas connectivity upgrades would become just a question of migrating to an access transport technology supporting more capacity.

In addition, this approach would simplify the way access infrastructure is shared among operators; the flexibility introduced by separating forwarding and control planes in the access allows slicing of network assets, facilitating the control of dedicated portions of the network to different operators. Thus, infrastructure deployed by different operators can be shared, reducing costs and time to market.

Apart from the access, network functions like the broadband network gateway (BNG) can be virtualized and instantiated on cloud infrastructures, scaling according to real needs. Nowadays, monolithic BNGs are rolled-out per PoP considering a maximum user forecast. The reality is that the load of those BNGs is far lower than its maximum capacity, with the constraint that the vacant resources cannot attend (at least in an easy and optimal way) the demand in other PoPs, while consuming scarce resources, like IPv4 addressing. By deploying virtualized instances of BNGs, the right scale for accommodating the actual demand, which changes along the day, can be ensured.

Figure 6 presents how the SHINE architecture can address this scenario. ARNs host control plane capabilities of fixed access network nodes, with the local instantiation for handling a variety of access technologies (G.PON, Ethernet, etc). Each ARN supports client interfaces for all "last drop" technologies, abstracting control from data-plane characteristics.

On the other hand, the BNG function is deployed in the large DC or the μ-DCs, where the optimal placement depends on the actual demand. The SHINE architecture facilitates two ways of scaling:

- When a BNG is attending a huge number of dispersed customers across the network (i.e. accessing from different PoPs) the BNG can be deployed in the large DC and its network functionality scaled-out/down, producing a bigger/smaller BNG;
- In the case where a punctual high demand arises in a specific area, a new instance of the BNG can be deployed in a nearby μ-DC.

As in previous use cases, an application running on top of the SHINE's service management and control module requests function deployment.

SUMMARY

In this article the SHINE architecture has been presented; it orchestrates cloud with heterogeneous access and core networks to provide cloud services, being the base to support NFV. The SHINE approach incorporates: ARN nodes to interface end-users directly to the core; a set of μ-DC placed close to the access to reduce network traffic while ensuring low latency; and a service management and control module to dynamically orchestrate cloud and network.

Four use cases addressing different markets of interest for any network operator have been used to illustrate how the SHINE architecture can be used to facilitate virtualizing network functions and orchestrating services.

The major challenges in the deployment of SHINE are associated with end-to-end coordination and the migration from the existing networking framework. In that regard, further specific studies are needed and migration steps need to be considered.

The major challenges in the deployment of SHINE are associated with end-to-end coordination and the migration from the existing networking framework. In that regard, further specific studies are needed and migration steps need to be considered.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement no. 317999 IDEALIST project, and from the Spanish MINECO SYNERGY project (TEC2014-59995-R).

REFERENCES

[1] Alcatel Lucent white paper, "Bell Labs Metro Network Traffic Growth: An Architecture Impact Study," 2013.
[2] NFV white paper, "Network Functions Virtualisation," (http://portal.etsi.org/NFV/NFV_White_Paper.pdf), 2012.
[3] Ericsson white paper, "The Telecom Cloud Opportunity: How Telecom Operators Can Leverage Their Unique Advantages in the Emerging Cloud Market," 2012.
[4] IBM Institute for Business Value, "The Natural Fit of Cloud with Telecommunications," 2012.
[5] L. Velasco *et al.*, "Elastic Operations in Federated Datacenters for Performance and Cost Optimization," *Elsevier Computer Commun.*, vol. 50, 2014, pp. 142–51.
[6] L. Contreras *et al.*, "Towards Cloud-Ready Transport Networks," *IEEE Commun. Mag.*, vol. 50, 2012, pp. 48–55.
[7] L. Velasco *et al.*, "Cross-Stratum Orchestration and Flexgrid Optical Networks for Datacenter Federations," *IEEE Network*, vol. 27, 2013, pp. 23–30.
[8] M. Jinno *et al.*, "Multiflow Optical Transponder for Efficient Multilayer Optical Networking," *IEEE Commun. Mag.*, vol. 50, 2012, pp. 56–65.
[9] D. King and A. Farrel, "A PCE-Based Architecture for Application-Based Network Operations," RFC 7491, March 2015.
[10] ONF Solution Brief, "OpenFlow-enable Transport SDN," May 2014.
[11] L. Velasco *et al.*, "In-Operation Network Planning," *IEEE Commun. Mag.*, vol. 52, 2014, pp. 52–60.
[12] ETSI GS NFV 001, "Network Functions Virtualisation (NFV): Use Cases," V1.1.1, Oct. 2013.
[13] Ll. Gifre *et al.*, "Experimental Assessment of ABNO-driven Multicast Connectivity in Flexgrid Networks," *IEEE J. Lightw. Technol.* (JLT), vol. 33, pp. 1-8, 2015.
[14] M. Mishra *et al.*, "Dynamic Resource Management using Virtual Machine Migrations," *IEEE Commun. Mag.*, vol. 50, 2012, pp. 34–40.
[15] J. Barrera, M. Ruiz, and L. Velasco, "Orchestrating Virtual Machine Migrations in Telecom Clouds," *Proc. OFC*, 2015.

BIOGRAPHIES

LUIS VELASCO received the M.Sc. degree in physics from the Universidad Complutense de Madrid (UCM) in 1993, and the Ph.D. degree from the Universitat Politecnica de Catalunya (UPC) in 2009. In 1989 he joined Telefonica of Spain and was involved with the specifications and first office application of SDH transport network. In 2004 he joined UPC, where he is currently an associate professor in the Department of Computer Architecture (DAC) and senior researcher at the Advanced Broadband Communications Center (CCABA). His interests include planning, routing, and resilience mechanisms in multilayer networks and software defined environments.

LUIS MIGUEL CONTRERAS received the M.Sc. degree in telecommunications engineering from the Universidad Politecnica de Madrid (1997) and the M.Sc. in telematics from the Universidad Carlos III de Madrid (UC3M) (2010). In 1997 he joined Alcatel Spain, working in both wireless and fixed networks. In 2006 he joined the Network Planning department of Orange Spain (France Telcom group), working on IP backbone planning. Since 2011 he has been with Telefonica Global CTO, working on scalable networks and their interaction with cloud and distributed services, and participating on several FP7 projects. He is currently working toward the Ph.D. degree in the Telematics Department of UC3M.

GIUSEPPE FERRARIS received a degree in electronic engineering from Politecnico di Torino (Italy) in 1988. He then joined Telecom Italia Lab (TILab), the Corporate R&D center of Telecom Italia. He is currently a senior project manager in the Optical Network and Planning Department of TILab. He was involved in research activities related to the evolution of the transport network, both in the long distance and in the metropolitan area. He is currently active in the area of advanced optical networking and on the interworking with the IP layer. He was active for several years in ITU and ETSI on the standardization of SDH and WDM based transport networks. He was also involved in the ISP Project MOON, in the IST Project NOBEL (project coordinator) and in several EURESCOM Projects on transport networks.

ALEXANDROS STAVDAS received the B.Sc. in physics (University of Athens), the M.Sc. in optoelectronics and laser devices (Heriot-Watt /St-Andrews University), and the Ph.D. (University College London) in the field of wavelength routed WDM networks. Currently he is a professor of optical networks in the Department of Telecommunications Science & Technology of UoP. He is the author or co-author of more than 80 journal publications and conference articles. He has also served as the technical program committee chairman and member of the Technical Program Committees in various international conferences. His current research interests include physical layer modelling of optical networks, ultra-high capacity end-to-end optical networks, OXC architectures, optical packet/burst switching, and DWDM access networks.

FILIPPO CUGINI received the M.Sc. degree in telecommunication engineering from the University of Parma, Italy. Since 2001 he has been with the National Laboratory of Photonic Networks, Consorzio Nazionale Interuniversitario per le Telecomunicazioni (CNIT), Pisa, Italy. His research interests include theoretical and experimental studies in the field of optical communications. In particular, the focus is on Ethernet, GMPLS, and PCE protocols and architectures, survivability and traffic engineering in IP over optical networks, multi-rate flexible optical networks, and software defined networking (SDN). He served as co-chair of several international symposia and as an editorial board member of the Elsevier journal *Optical Switching and Networking*. He is the co-author of twelve international patents and more than one hundred IEEE publications.

MANFRED WIEGAND has more than 25 years of international experience in telecommunications. He is now with Coriant (previously the ON business of Nokia Siemens Networks) and focuses on business development for optical transport networks. Prior to this he has held various management positions in sales, project management and R&D in Siemens Telecommunications Group and Nokia Siemens Networks, in Germany and overseas. He received a Dr.rer.nat degree in physics in 1981 with a thesis on quantum optics.

JUAN PEDRO FERNANDEZ-PALACIOS received the M.Sc. in telecommunications engineering from the Universidad Politecnica de Valencia in 2000. In 2000 he joined Telefonica I+D, where he is currently leading the Core Network Evolution unit. He has been involved in several European projects such as NOBEL, NOBEL-2, STRONGEST, and MAINS, as well as in the design of core network architectures in the Telefonica Group. Currently he is coordinating the FP7 project IDEALIST and the standardization activities within the CAON cluster.

A Service-Aware Virtualized Software-Defined Infrastructure

Lefteris Mamatas, Stuart Clayman, and Alex Galis

ABSTRACT

The Internet infrastructure is gradually improving its flexibility and adaptability due to the incorporation of new promising technologies, such as the software-defined networks and the network function virtualization. The main goal is to meet the diverse communication needs of the users, while the global system operation satisfies the business and societal goals of the infrastructure and service providers. This calls for solutions that consider both local and global network viewpoints and provide sophisticated system control in a stable and predictable way, while being service-aware.

We propose a fully integrated solution along these lines: the VLSP, a service-aware software-defined infrastructure for networks and clouds. The VLSP consists of three main distributed systems: a facility performing uniformly logically-centralized management and control of the infrastructure, called the virtual infrastructure management; an information management infrastructure able to maintain an accurate view of the infrastructure environment at both the local and system levels, called the virtual infrastructure information service; and a lightweight virtualization hypervisor able to perform configuration changes in the infrastructure resources, called the lightweight network hypervisor. We discuss representative use-case scenarios, while we demonstrate how VLSP tunes performance trade-offs for particular service demands.

INTRODUCTION AND MOTIVATION

The Internet is a global infrastructure that accommodates a wide range of applications with diverse quality of service requirements. A primary goal is to utilize the physical resources according to the needs of the deployed network services, while at the same time Internet stakeholders, such as the network and service providers, target the realization of services in a manner consistent with their business plans. This diversity calls for flexible resource control and management, with programmability and improved determinism in the system behavior.

Along these lines, programmable networks [1], software-defined networks (SDNs), and OpenFlow [2] were introduced. SDNs are characterized by:

- Decoupled network control from forwarding, with control embedded in a logically-centralized component.
- Programmability via software functions interacting with the network.
- Appropriate abstractions that allow SDN applications and services to be network-aware.

OpenFlow is a de jure standardized way to control flow tables in switches and routers. It allows a logically-centralized software application that has a global viewpoint of the network, called an SDN controller, to interact with the network equipment and make changes in the flow tables of the network equipment. Higher-level SDN applications interact with one or more SDN controllers in order to manipulate their general behavior, and to achieve significant performance improvements [3]. Two survey papers that cover the area of SDNs are [4] and [5].

The network function virtualization (NFV) architectural concept [6] brings networks closer to IT domains and their related operations. NFV targets both flexibility in service provisioning and reduction of cost expenditure. SDNs and NFV, although independent, can be mutually beneficial and may co-exist in the same network environment. SDNs offer flexibility at the network control level, while virtualization is a good candidate technology for hiding network device heterogeneity.

Based on the above, we have envisaged the following *four research challenges* for the future evolution of SDN/NFV technologies with the aim of improving service-awareness in the infrastructure.

Challenge 1: The enablement of new SDN applications, beyond centralized traffic engineering, without being constrained by existing hardware characteristics. Virtualization can hide the heterogeneity at the hardware level and can serve as a migration path toward adopting SDN-like technologies.

Challenge 2: The introduction of new abstractions for realizing sophisticated management features on top of flexible and programmable network control components. Such technologies will bridge the gap between the local viewpoints (i.e. solutions handling network control issues) and the global viewpoints (i.e. higher-level management features).

Challenge 3: The fast and simple deployment

Lefteris Mamatas is with the University of Macedonia, Greece. This work was performed when he was with University College London

Stuart Clayman, and Alex Galis are with University College London.

0163-6804/15/$25.00 © 2015 IEEE

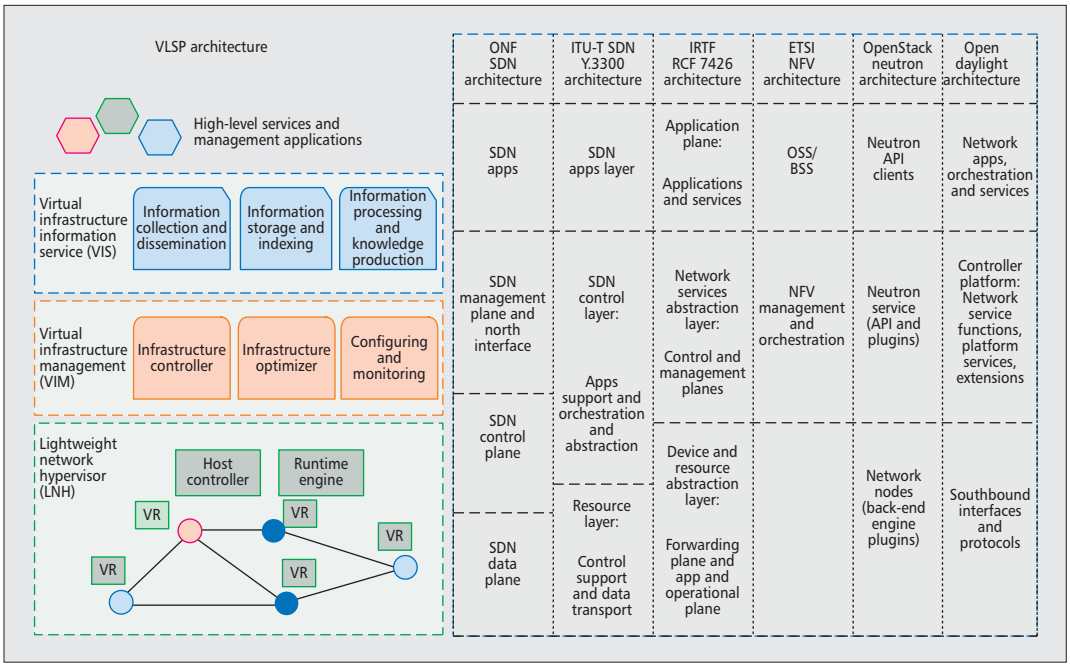

Figure 1. VLSP architecture, components and relation to ONF SDN, ITU-T SDN, IRTF SDN, ETSI NFV, OpenStack Neutron and OpenDaylight architectures.

of network resources that collectively support a large number of services, from the level of a single node up to a large-scale network.

Challenge 4: The maintenance of a global picture at both the network level and domain level, using logically-centralized intelligence, programmable techniques, and an abstracted design.

Consequently, a high-level unification and integration of the virtualization, network function virtualization, SDN, programmability, and management would need to be achieved. This calls for new infrastructure architectures, standardized interfaces, and management applications. As such, we foresee that the research focus would move toward the incorporation of SDN with NFV, whereby novel network management approaches can bring services and networks closer.

Targeting the above challenges, we present the Very Lightweight Software-Driven Network and Services Platform (VLSP), a fully integrated open-source software-defined infrastructure and architecture for networks and clouds that we have designed and implemented from the scratch. VLSP is differentiated from SDN architectures [7] and open solution initiatives [8] by having the following architectural features:

- A deeper integration of NFVs and SDNs by introducing uniform virtualization of networks and applications.
- Unified management and control for both networks and clouds, reducing management cost and complexity.
- Suitability for reliability and scalability evaluation through the use of a lightweight virtualization hypervisor.

Additionally, VLSP supports the following novel features:

- A software execution environment within the virtual routers, allowing the deployment, at run-time, of diverse network control and service components in order to enable programmability capabilities.
- A focus on elastic, adaptable, and autonomic service provisioning and management. Supporting lightweight application components being realized as virtual resources deployed and managed uniformly by the same environment.
- Hierarchical and distributed control components that offer scalable and logically-centralized network control and management, without overloading centralized software nodes. This enables the autonomic network and service management logic to be designed and operated on a global network view.
- An integrated and abstracted state information manipulation facility, building the global-picture for the system, while supporting local level and domain level views, yet adaptable to the diverse requirements of the involved entities.

In Fig. 1 we show the VLSP architecture, components and their relation [7, 8] to the ONF SDN ITU-T Y.3300 SDN, IRFT RFC 7426, ETSI NFV, OpenStack Neutron, and OpenDaylight architectures. ONF is working on OpenFlow standardization aspects, and ETSI is studying the architecture of NFV from the network operators' perspective. Neutron augments OpenStack clouds with networking as a service capabilities, and is based on a loosely-coupled architecture, whereby service and virtual network device plugins realize the targeted behavior and are hidden behind a common API. All available plugins present different performance trade-offs, scalability, manageability, and compatibility aspects, which cloud operators should weigh-up for themselves.

Several SDN/NFV efforts have appeared within the IETF in the form of working groups

This calls for new infrastructure architectures, standardized interfaces, and management applications. As such, we foresee that the research focus would move towards the incorporation of SDN with NFV, whereby novel network management approaches can bring services and networks closer.

In the IRTF, the Software-Defined Networking Research Group identifies future research challenges, including scalability, abstractions, security, and programming languages for SDN environments. Other SDN initiatives have a stronger focus on combining SDNs with NFV, including the IEEE SDN technical community.

or working group initiatives. The data plane oriented approaches are enabling higher-level control (e.g. the Interface to the Routing System Group and the Abstraction and Control of Transport Networks Group), with some of them introducing application-awareness (e.g. the Application-Layer Traffic Optimization Group). The Virtual Network Function Pools Group and the Network Function Virtualization Configuration Group propose design details for a more efficient usage of NFVs. The Service Function Chaining Working Group is studying the deployment of service functions in large-scale environments.

The ITU's Telecommunication Standardization Sector (ITU-T) hosts study groups on SDNs. SG13 focuses on future networks (e.g. clouds, mobile, virtual networks) and SG11 on relevant network protocols. In the IRTF, the Software-Defined Networking Research Group identifies future research challenges, including scalability, abstractions, security, and programming languages for SDN environments. Other SDN initiatives have a stronger focus on combining SDNs with NFV, including the IEEE SDN technical community.

Many researchers have pointed out the importance of both reliability and scalability in SDN environments. Shalimov *et al*. [9] carried out an extensive study of seven SDN/OpenFlow controllers. They claim, based on their results, that current controllers do not scale well over the network cores and that they are not able to meet the increasing demands in communication. In our case, VLSP considers scalability as a basic design requirement.

In another case, Google documented an outage incident in its SDN WAN deployment [3], which could have been avoided if latency sensitive operations had received higher priority, compared to the throughput-intensive operations. The outage problem was detected at a very late stage, due to the lack of enough performance profiling and reporting. Levin *et al.* assess how inconsistency of SDN control state information significantly degrades performance of logically-centralized control applications [10]. In their work and many others, state information management is integrated within a corresponding SDN application or SDN controller.

From our point of view, the state information, its manipulation, and the exchange capabilities can be abstracted away and realized through a separate component in the SDN architecture. The SDN applications and controllers should be able to communicate information based on their own diverse requirements and constraints. This requires not only supporting alternative methods to create the network-wide state, but also a flexible way to choose the most appropriate configuration each time.

In contrast to the related initiatives, we target our identified four research challenges, which we foresee as main research trends in the near future. We focus on aspects such as service-awareness, a better and uniform control, all of which are beyond performance issues, network protocol issues, and the existing constraints of deployed infrastructures. VLSP is based on a lightweight virtual router implementation, suitable for scalability and reliability evaluations. As we highlight in Fig. 1, the VLSP components could be integrated into existing deployed infrastructures at their equivalent architectural blocks, since they use similar design strategies (e.g. RESTful communication).

The following section discusses the VLSP design details and architecture. After that we provide representative use-case scenarios of our platform. Then we highlight our experimental methodology and provide experimental results. Following that we discuss the lessons we have learned during its design and implementation. Finally, we conclude this article.

Platform Design and Architecture

The VLSP provides reusable management and control facilities that are utilized by additional software entities, called management and control entities (MCEs) in this article. These MCEs, when combined with the VLSP, enable logically-centralized management and control of the system. A categorization of the different types of MCE, from a deployment related perspective, follows:

- High-level services and management applications, which are responsible for the efficient operation of the whole system, at both the network and service levels. They take optimization decisions based on the global picture.
- MCEs deployed at the physical hosts, controlling a part of the network, such as the SDN controllers.
- MCEs deployed at the virtual routers, which are responsible for resource-facing operations at the virtualization hypervisor level.

Figure 1 gives a high-level overview of the VLSP. An earlier version of the VIM and the LNH layers are presented in [11]. We have done further design and implemented an integrated working version of the VLSP and released it as open-source software.[1] In the following subsections, we describe all three of the VLSP layers and their main functions.

The Virtual Infrastructure Information Service

The VIS offers abstracted and logically-centralized information manipulation across all of the deployed software entities. In Fig. 2, we present a high-level view of the VIS architecture and its basic interactions. The VIS uses two separate interfaces for communication with the MCEs, namely:

- The information management interface, which handles the configuration of information manipulation, including the MCE's registration to the VIS, the management of internal VIS information manipulation functions, and the establishment, operation, and optimization of information flows;
- The information exchange interface, which offers the actual information transfer and exchange capability to the deployed MCEs.

The VIS has the following core functions.

[1] The VIS, the VIM, and the LNH components, the relevant documentation, and research papers can all be downloaded from: http://clayfour.ee.ucl.ac.uk.

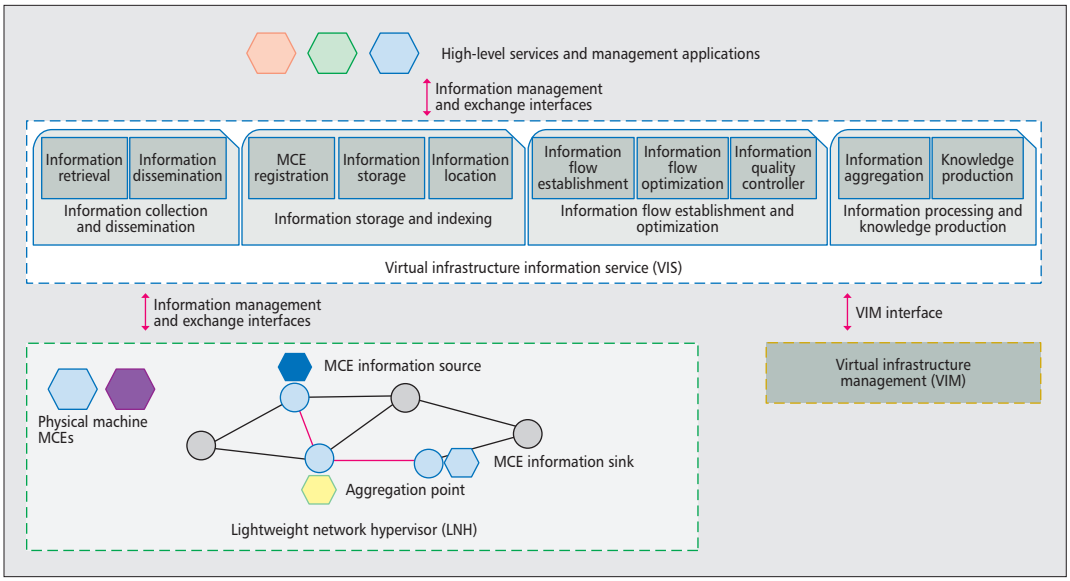

Figure 2. The Virtual Infrastructure Information Service.

The Virtual Infrastructure Management (VIM) provides high-level control and management of the virtual infrastructure. It is responsible for the manipulation and lifecycle of virtual topologies and the service and management software running on top of them, to ensure continued operation and consistency.

Information Collection and Dissemination (ICD): The ICD is responsible for organizing the communication of information, including the optimization of information flows. It offers facilities for *information retrieval* and *information dissemination*, plus the *information flow controller*. Alternate selectable communication methods are supported, such as the push/pull, publish/subscribe, and direct communication methods. The *information flow controller* oversees such functions, including controlling the information flow establishment, operation, and relevant optimization aspects;

Information Storage and Indexing (ISI): The ISI provides storage and indexing functionalities to the VIS. The *MCE registration* module allows the MCEs to express their information handling requirements and capabilities. The ISI function maintains an MCE registry, storing specifications for the available information to be retrieved or disseminated. The *information storage* module offers alternative storage options, according to requirements and characteristics, specified during an MCE's registration phase.

Information Processing and Knowledge Production (IPKP): The IPKP augments VIS with information processing, aggregation, and global-picture information production capabilities. The *information aggregation* module applies aggregation functions (e.g. MAX, MIN, AVERAGE) to the collected data before they are stored or disseminated. The data may be filtered at the aggregation level for optimization purposes. The *knowledge production* module generates global-picture information through processing and/or aggregating information.

The Virtual Infrastructure Management

The virtual infrastructure management (VIM) provides high-level control and management of the virtual infrastructure. In Fig. 3 we show the VIM with its basic functions and interfaces. It is responsible for the manipulation and lifecycle of virtual topologies and the service and management software running on top of them, to ensure continued operation and consistency. The VIM interacts with the other two VLSP layers through the *virtual infrastructure management interface*. The core VIM architectural components are as follows.

Infrastructure Controller: The *infrastructure controller* acts as a control point for managing the virtual elements and the applications they support. It accepts all of its input via the *VIM interface* from high-level management applications (i.e. for service orchestration aspects) and from the VIS (i.e. for information management related activities). The *service orchestrator* performs the automatic allocation and full lifecycle management of distributed application nodes which run on the virtual routers and realize particular network services. The *infrastructure controller* is also responsible for the allocation and efficient operation of the virtual resources, through the *resources orchestrator*. This module handles the optimal placement of the virtual routers, the decisions to add or remove new resources, the manipulation of the corresponding virtual links, and so on. It manages the full lifecycle of the virtual resources, aligned to the availability of physical resources and the service requirements.

Infrastructure Optimizer: This function handles VLSP optimization aspects for efficient allocation of virtual resources and the corresponding data paths. It is responsive to the changes in network conditions or requirements, triggering appropriate optimization processes or mitigating stability problems (such as the relocation of service nodes or virtual routers). The *placement engine* performs the actual placement of all VLSP entities according to the current topology and the virtual network elements load. The *flow controller* is responsible for performing logically-centralized data flow allocation. It makes decisions on the establishment of particular data

The virtual router is a software element that offers basic network protocol functionality (e.g. routing and transport), an isolated execution environment, and a virtual sockets API. Our focus here was to provide a lightweight foundation for a scalable infrastructure (i.e. the VRs can bootstrap and shutdown in less than a second).

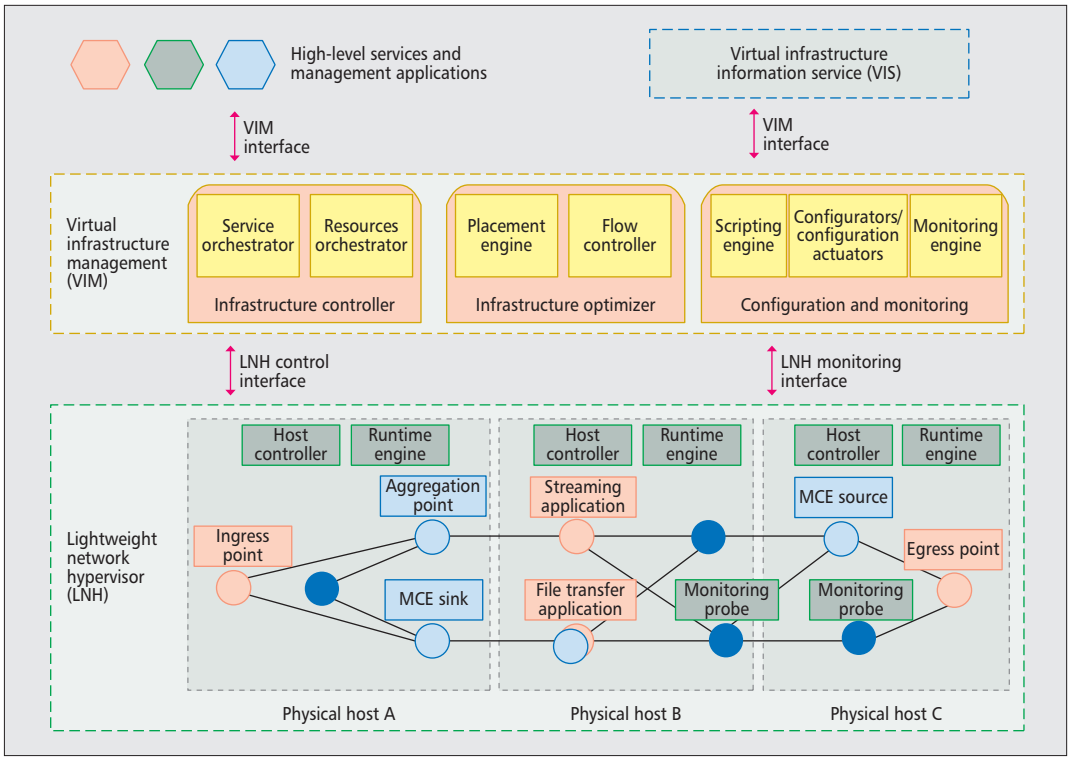

Figure 3. The Virtual Infrastructure Management and the Lightweight Network Hypervisor.

flows. The same component realizes the information flows being handled from the VIS.

Configuration and Monitoring: The configuration and monitoring function supports the VIM interactions with the high-level applications and the LNH. The *configurators and configuration actuators* handle operations related to the translation of certain management and control commands to configuration settings at the level of LNH. It provides alternative configuration options to the management entities for a targeted network behavior. The *monitoring engine* is responsible for low-level monitoring activities, such as the manipulation of monitoring probes in the virtual routers. It collects information to be further processed from the VIS (either aggregated or global-picture information) and used from the VIM as an input for optimization decisions.

The Lightweight Network Hypervisor

The lightweight network hypervisor (LNH), as shown in Fig. 3, includes a fully operational lightweight virtual router (VR) combined with virtual network connectivity and its associated low-level control components. The LNH communicates with the other two VLSP layers for network control activities through the *LNH control interface*. This communication involves enforcement of the VIM decisions regarding the deployment and configuration of virtual networks and network services. Furthermore, it handles the establishment and efficient parameterization of the information flows being overseen from the VIS. The VIM *monitoring engine* communicates with a number of deployed monitoring probes in the VRs, in order to collect real-time state information via the *LNH monitoring interface*. The LNH consists of the following.

The Virtual Router: The virtual router is a software element that offers basic network protocol functionality (e.g. routing and transport), an isolated execution environment, and a virtual sockets API. Our focus here was to provide a lightweight foundation for a scalable infrastructure (i.e. the VRs can bootstrap and shutdown in less than a second). The routers support core IP functionality, such as the distance vector routing protocol, the time-to-live option, and so on. The router offers an application layer interface enabling the deployment of Java software applications at runtime. These act as the service elements within the platform by using a virtual sockets API which can send and receive datagrams or packets.

The Host Controller: The VRs are controlled by distributed components residing at each physical host, the *host controllers (HCs)*. The HCs directly control the VRs within a host by passing commands from the VIM, such as to manipulate virtual routers, links, application nodes, and their configuration.

The Runtime Engine: The *runtime engine* complements the *host controller* with functionalities related to the runtime operation of the deployed virtual networks and network services. An example is to enforce periodical network activities, such as network maintenance processes and timely detection of failures.

The Monitoring Probes: Each virtual router is instrumented with the VLSP monitoring system, based on Lattice [12], in order to gather data on

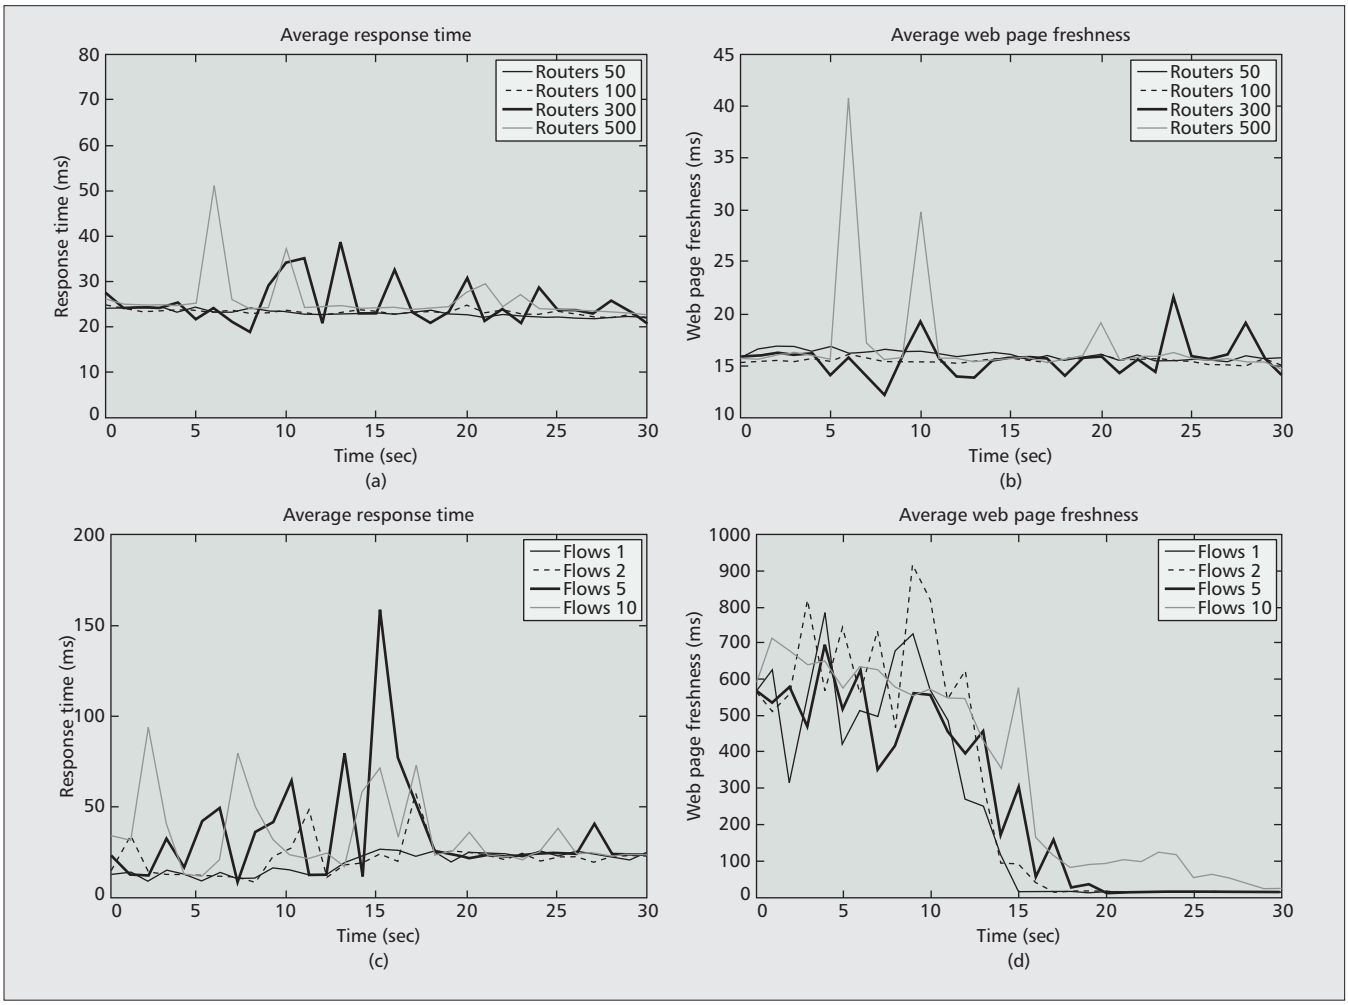

Figure 4. VLSP validation results: a) average response time (scalability); b) average web page freshness (scalability); c) average response time (handling jitter issue); d) web page freshness (handling jitter issue).

the virtual network or the hosted applications via modules called *monitoring probes*. The monitoring system collects the raw data and passes in for further processing onto the VIS.

Use Case Scenarios

Here we present an example workflow involving all three of the VLSP layers. Assume a high-level application expressing particular service requirements: a request to deploy a web-based application to be used by a given number of users. The application also requests global-picture information on the web application performance, such as the average response time. The service deployment request is directed to the VIM through the *VIM interface* and the *information flows establishment request* to the VIS through the *information management interface*.

The VIM invokes the *infrastructure controller*, which decides to deploy a minimum number of required web-servers hosting the web application and a number of web-proxies that may be used to improve communication performance. Furthermore, the same component, after a request from the VIS, decides to deploy a number of monitoring probes sampling web application performance, and an aggregation point to calculate the average values. All of the above are handled by the *service orchestrator*.

The *resource orchestrator* organizes the deployment of virtual machines on a number of physical servers. The *infrastructure optimizer* uses the *placement engine* to decide the most efficient locations of the involved application nodes (i.e. the web-servers, web-proxies, and the aggregation point), in order to utilize physical resources in an optimal way, as in [11]. The *flow controller* deploys the data flows between the web-servers, web-proxies, and the clients, and the information flows that communicate and process information regarding the web application performance. All flows and node placements are optimized according to an active global performance goal (for the minimum energy consumption, as an example).

The monitoring probes are handled by the *monitoring engine*. The *scripting engine* and the *configurators and configuration actuators* modules handle the appropriate representations and configuration settings. Those settings are communicated to the LNH, so it can enforce all the above network and service decisions using the appropriate *host controllers* and *runtime engines* which are handling the required virtual routers. After this, the deployed network service is ready to offer the web-based application to users.

We have created small and simple data applications resembling Web Clients, Web Servers, and Web Proxies. The VLSP determines the most appropriate data paths for the data flows by having the global network view as an input and using the dynamic node selection algorithms presented in the paper.

In practice, all of the above configuration decisions may be revoked by the high-level application, in the case that the active performance goal changes, or by the VIM, if it foresees issues that may cause faults or performance problems. In the experimental results section, we study a similar scenario in which the VIM requests, after some time, a change in the global performance goal; it enforces a general guideline to stabilize the average response time.

Some of the use-case examples that have benefited from the VLSP platform are:

- The VLSP was used as a soft core network connected to hardware based wireless and mobile network elements in order to realize network services on top of heterogeneous network paths. The VLSP operated efficient data flows and enabled allocation of resources near the service consumers, based on the unified local and global network views and the high-level service requirements. We demonstrated a fully working and integrated system in the context of UniverSELF project.[2]
- The VLSP has been used to implement an information-centric networking (ICN) testbed. Specifically, it has been utilized in the implementation of the CURLING ICN algorithm [13], which employs a hop-by-hop hierarchical content-based publish-subscribe paradigm to content distribution.
- Within the Dolfin project,[3] the aim is to optimize energy consumption within the limits of a single data center and in a group of data centers, based on system virtualization techniques and the optimal distribution and placement of virtual machines. To realize such a system, we augmented the VLSP with energy-aware monitoring, modeling, and resource allocation.

Experimental Results

Here we present our experiments with the main use-case scenario discussed in the previous section. We particularly focus on how a change in the general performance goal by the VIM is realized, because of scalability and stability problems. We next detail our experimental methodology and then present our experimental results.

In our testbed we used 11 servers each with four CPU cores and 32GB of physical memory. The VLSP platform software itself consists of more than 600 Java classes and more than 100,000 lines of code. On a desktop machine it is possible to deploy topologies of approximately 70 virtual routers. We have tested a deployment of 700 communicating virtual routers in complex multipath topologies on 11 dedicated physical servers. The underlying monitoring framework used by the testbed on the routers, known as Lattice, is described in [12] and is available as open-source software.[4]

Every experimental run includes the creation of a new network topology. Such a process involves interactions between the VIM and the corresponding *host controllers* deployed at all physical servers. We have created small and simple data applications resembling web clients, web servers, and web proxies. The VLSP determines the most appropriate data paths for the data flows by having the global network view as an input and using the dynamic node selection algorithms presented in [14]. This activity involves a number of distributed nodes being deployed over the virtual network (i.e. the web servers, the web proxies, and the aggregation point).

As the next step, the VIM *placement engine* assigns all of the web clients and web servers to the most appropriate web proxies (namely the web proxy that is closest). Then, after a warm-up period, the communication begins. The web clients periodically transmit performance measurements to the VLSP over the negotiated data flows using the following metrics:

- Average response time. The average time taken from the request of a web page from a web client to the point that it is received.
- Web page freshness The time taken from a web page update to the point it reaches the requesting web client.

In our experiments we have stress tested our infrastructure with large topologies. The main goal here is to investigate its behavior in terms of scalability and stability and how resource exhaustion can be tackled by changing the global performance goal. As shown in Figs. 4a and 4b, large scales can be reached (in this case we go up to 500 virtual routers).

After some time during the runs, the VLSP detects stability problems (in this case, increased jitter in the average response time). When this occurs the *service orchestrator* enforces a change in the global performance goal that bypasses the use of web proxies so that the web clients communicate with the web servers directly. As a consequence of this change the jitter decreases, the average response time increases, and we observe in Fig. 4d a significant improvement in the web page freshness. We can see in Figs. 4c and 4d that the VLSP can trade the increasing response time jitter for a slight increase in the average response time. Such strategies can be associated with control loops that detect and tackle systematic stability problems.

Lessons Learned

Since we have been continuously designing and implementing VLSP features for a number of years, we have faced a number of issues that we document in this section. In a nutshell, we observed the following:

- It is difficult to experiment with medium to large-scale topologies without having scalability and stability as cornerstone requirements. In our case, we started building on top of XEN virtual machines which hosted a very small Linux distribution and routing software. However, as this was still too heavy for our needs, we ended up building our own virtual router implementation supporting the essential features only. Lightweight communication protocols are important as well, so we implemented basic transport and routing protocols. All of our management-level protocols run on top of REST.
- Efficient resource consumption calls for a clear identification and handling of the involved performance trade-offs. For exam-

[2] UniverSELF project website: http://www.univerself-project.eu.

[3] Dolfin project website: http://www.dolfin-fp7.eu.

[4] Lattice home page: http://clayfour.ee.ucl.ac.uk/lattice/.

ple, a physical server may run out of memory while the other resources can be minimally utilized. The logically centralized view and control give unique capabilities along these lines. Furthermore, many OS related configuration parameters need to be carefully adjusted, such as the heap memory allocation.

- The architectural aspects and relevant design abstractions need to avoid replicating functionalities, such as overlapping communication protocol features.
- There is a need to focus on design and integration aspects beyond simple performance improvements. The lesson from operating system design is that flexibility and adaptability may nominally reduce performance but can enable a greater number of services that were not possible before.

The reader may benefit from our experience for their own software design, development, and implementations of large scale distributed systems.

Conclusions

We have presented VLSP, a complete design and implementation of a novel service-aware virtualized software-defined infrastructure that exhibits the main aspects of autonomic service provisioning and network/service management. We highlighted its relation to well known SDN, NFV architectures, and relevant open-source initiatives. VLSP includes:

- The VIS, which provides an accurate, timely, and complete view of the network.
- The VIM, which brings sophisticated high-level network management for virtual network topologies.
- The LNH, which abstracts resources and enforces decisions taken.

The above three distributed components inter-operate in order to provide a number of control loops that realize the management and operations that actually utilize the physical resources according to the needs of the deployed network services.

In summary, the novel qualities of our proposal are as follows:

- We designed and implemented a uniform software-defined infrastructure from scratch, bringing together the NFV with SDN technologies.
- We introduced a lightweight virtualization hypervisor, suitable for quick and scalable deployment of network infrastructures, having both computation and connectivity nodes, as a suitable facility to experiment with our ideas.
- We proposed a new infrastructure bringing context-awareness in such environments, associating global-picture information with local views, requirements, and resource constraints.

Last but not least, we demonstrated how these ideas work together to tune performance tradeoffs toward planned network behavior.

We have shown how VLSP has been used for various network evaluations, and in the future we plan to explore using the VLSP for flexible service creation and deployment focusing on distributed networked cloud infrastructures and network service chaining. This involves extensions of the VLSP with methods and processes elaborating and using [15] for continuous dynamic operation of services, service composition, aggregation and management of service blocks, including service delivery based on orchestration, programmability [1] and automatic (re)deployment.

We have shown how VLSP has been used for various network evaluations, and in the future we plan to explore using the VLSP for flexible service creation and deployment focusing on distributed networked cloud infrastructures and network service chaining.

Acknowledgment

This work is partially supported by the European Union DOLFIN (http://www.dolfin-fp7.eu) and UniverSELF (http://www.univerself-project.eu) projects.

References

[1] A. Galis *et al.*, *Programmable Networks for IP Service Deployment*, Artech House, 2004.
[2] N. McKeown *et al.*, "Openflow: Enabling Innovation in Campus Networks," *ACM SIGCOMM Computer Commun. Rev.*, vol. 38, no. 2, 2008, pp. 69–74.
[3] S. Jain *et al.*, "B4: Experience with a Globally-Deployed Software Defined WAN," *Proc. ACM SIGCOMM 2013 Conf. SIGCOMM*, ACM, 2013, pp. 3–14.
[4] D. Kreutz *et al.*, "Software-Defined Networking: A Comprehensive Survey," *CoRR*, vol. abs/1406.0440, 2014.
[5] F. Hu, Q. Hao, and K. Bao, "A Survey on Software Defined Networking (SDN) and Openflow: From Concept to Implementation," *IEEE Commun. Surveys Tutorials*, vol. 16, no. 4, 2014, pp. 2181–2206.
[6] M. Chiosi *et al.*, "Network Functions Virtualisation," white paper at the SDN and OpenFlow World Congress, ETSI, Tech. Rep., 2012.
[7] SDN De-Facto/Dejure Standards, "IRTF, "Software-Defined Networking (SDN): Layers and Architecture Terminology," Tech. Rep., Jan. 2015; available: https://tools.ietf.org/html/rfc7426; ONF, "Software-Defined Networking: The New Norm for Networks," Open Network Foundation, Tech. Rep., 2012. ; ITU-T, "Recommendation y.3300 (2014) — Framework of Software-Defined Networking, https://www.itu.int/rec/t-rec-y.3300201406-i/en; Recommendation y.3001 (2012) — Future Networks: Objectives and Design Goals, http://www.itu.int/rec/trec-y.3001-201105-i; Recommendation y.3011 (2012) — Framework of Network Virtualization for Future Networks, https://www.itu.int/rec/t-rec-y.3011-201201-i/en," Tech. Rep. ; ETSI NFV: M. Chiosi et al., "Network Functions Virtualisation," White paper at the SDN and OpenFlow World Congress, ETSI, Tech. Rep., 2012," Tech. Rep.
[8] Open Source initiatives, "OPNFV, "Open Platform for NFV, https://www.opnfv.org/.; OpenDaylight, "SDN and NFV platform that enables network control and programmability," Tech. Rep., http://www.opendaylight.org/.; OpenStack, "OpenStack Networking (Neutron)," Tech. Rep., https://wiki.openstack.org/wiki/Neutron," Tech. Rep.
[9] A. Shalimov *et al.*, "Advanced Study of SDN/Openflow Controllers," *Proc. 9th Central & Eastern European Software Engineering Conf. Russia*, ACM, 2013, p. 1.
[10] D. Levin *et al.*, "Logically Centralized? State Distribution Trade-Offs in Software Defined Networks," *Proc. first workshop on Hot topics in software defined networks*, ACM, 2012, pp. 1—6.
[11] S. Clayman *et al.*, "The Dynamic Placement of Virtual Network Functions," *Proc. 1st IEEE/IFIP Int'l Workshop SDN Management and Orchestration*, 2014.
[12] S. Clayman, A. Galis, and L. Mamatas, "Monitoring Virtual Networks with Lattice," *Proc. Network Operations and Management Symposium Workshops (NOMS Wksps)*, 2010 IEEE/IFIP, 2010, pp. 239–46.
[13] W. K. Chai *et al.*, "Curling: Content-Ubiquitous Resolution and Delivery Infra*structure for Next-Generation Services," IEEE Commun. Mag.*, vol. 49, no. 3, 2011, pp. 112–20.
[14] R. G. Clegg *et al.*, "On the Selection of Management/Monitoring Nodes in Highly Dynamic Networks," *IEEE Trans. Computers*, vol. 62, no. 6, 2013, pp. 1207–20.
[15] B. Rochwerger *et al.*, "The Reservoir Model and Architecture for Open Federated Cloud Computing," IBM Journal of Research and Development, vol. 53, no. 4, 2009, pp. 4—1.

BIOGRAPHIES

LEFTERIS MAMATAS (https://sites.google.com/site/emamatas/) is a lecturer in the Department of Applied Informatics, University of Macedonia, Greece. Before that he was a senior researcher at University College London, Space Internetworking Center/Democritus University of Thrace, and DoCoMo Eurolabs in Munich. He received his Ph.D. from the Department of Electrical and Computer Engineering, Democritus University of Thrace in Greece. His research interests are in the areas of software-defined networks, network management, opportunistic networks, and energy efficient communication. He participated in several international research projects, such as Dolfin (FP7), Autonomic Internet (FP7), UniverSELF (FP7), Ambient Networks (FP6), and others. He has published more than 40 papers in international journals and conferences. He served as a TPC chair for the WWIC 2012 and E-DTN 2009 conferences and as a guest editor for the Elsevier *Ad Hoc Networks Journal*.

STUART CLAYMAN received his Ph.D. in computer science from University College London in 1994. He has worked as a research lecturer at Kingston University and at UCL. He is currently a senior research fellow at the UCL EEE department. He has co-authored more than 30 conference and journal papers. His research interests and expertise lie in the areas of software engineering and programming paradigms; distributed systems; virtualized compute and network systems, network and systems management; networked media; and knowledge-based systems. He has been involved in several European research projects since 1994. He also has extensive experience in the commercial arena undertaking architecture and development for software engineering, distributed systems, and networking systems. He has run his own technology start-up in the area of NoSQL databases, sensors, and digital media.

ALEX GALIS (a.galis@ucl.ac.uk, www.ee.ucl.ac.uk/ agalis) is a professorial research fellow in networked and service systems at University College London. He has co-authored 10 research books and more that 200 publications in the Future Internet areas: system management, networks and services, networking clouds, virtualization and programmability. In 2004 he co-authored *Programmable Networks for IP Service Deployment*, one of the first textbooks on management and programmbility of virtualized network and computing resources (http://en.wikipedia.org/wiki/Active_Networking). He was a vice chair of the ITU-T SG13/Group on Future Networking (www.itu.int/ITU-T/focusgroups/fn/index.html), where five ITU-T SDN related recommendations were developed. He is currently the co-chair of the IEEE SDN publication committee (http://sdn.ieee.org) and TPC co-chair of IEEE Network Softwarization 2015 (NetSoft 2015 http://sites.ieee.org/netsoft/).

CALL FOR PAPERS
IEEE COMMUNICATIONS MAGAZINE
COMMUNICATIONS STANDARDS SUPPLEMENT

BACKGROUND

Communications Standards enable the global marketplace to offer interoperable products and services at affordable cost. Standards Development Organizations (SDOs) bring together stake holders to develop consensus standards for use by a global industry. The importance of standards to the work and careers of communications practitioners has motivated the creation of a new publication on standards that meets the needs of a broad range of individuals including: industrial researchers, industry practitioners, business entrepreneurs, marketing managers, compliance/interoperability specialists, social scientists, regulators, intellectual property managers, and end users. This new publication will be incubated as a Communications Standards Supplement in *IEEE Communications Magazine*, which, if successful, will transition into a full-fledged new magazine. It is a platform for presenting and discussing standards-related topics in the areas of communications, networking and related disciplines. Contributions are also encouraged from relevant disciplines of computer science, information systems, management, business studies, social sciences, economics, engineering, political science, public policy, sociology, and human factors/usability.

SCOPE OF CONTRIBUTIONS

Submissions are solicited on topics related to the areas of communications and networking standards and standardization research, in at least the following topical areas:

Analysis of new topic areas for standardization, either enhancements to existing standards, or of a new area. The standards activity may be just starting or nearing completion. For example, current topics of interest include:
- 5G radio access
- Wireless LAN
- SDN
- Ethernet
- Media codecs
- Cloud computing

Tutorials on, analysis of, and comparisons of IEEE and non-IEEE standards. For example, possible topics of interest include:
- Optical transport
- Radio access
- Power line carrier

The relationship between innovation and standardization, including, but not limited to:
- Patent policies, intellectual property rights, and antitrust law
- Examples and case studies of different kinds of innovation processes, analytical models of innovation, and new innovation methods

Technology governance aspects of standards focusing on both the socio-economic impact as well as the policies that guide it. This would include, but are not limited to:
- The national, regional, and global impacts of standards on industry, society, and economies
- The processes and organizations for creation and diffusion of standards, including the roles of organizations such as IEEE and IEEE-SA
- National and international policies and regulation for standards
- Standards and developing countries

The history of standardization, including, but not limited to:
- The cultures of different SDOs
- Standards education and its impact
- Corporate standards strategies
- The impact of Open Source on standards
- The impact of technology development and convergence on standards

Research-to-Standards, including standards-oriented research, standards-related research, research on standards

Compatibility and interoperability, including testing methodologies and certification to standards

Tools and services related to any or all aspects of the standardization lifecycle

Proposals are also solicited for Feature Topic issues of the Communications Standards Supplement.

Articles should be submitted to the IEEE Communications Magazine submissions site at

http://mc.manuscriptcentral.com/commag-ieee

Select "Standards Supplement" from the drop down menu of submission options.

Virtualized Security at the Network Edge: A User-Centric Approach

Diego Montero, Marcelo Yannuzzi, Adrian Shaw, Ludovic Jacquin, Antonio Pastor, René Serral-Gracià, Antonio Lioy, Fulvio Risso, Cataldo Basile, Roberto Sassu, Mario Nemirovsky, Francesco Ciaccia, Michael Georgiades, Savvas Charalambides, Jarkko Kuusijärvi, and Francesca Bosco

ABSTRACT

The current device-centric protection model against security threats has serious limitations. On one hand, the proliferation of user terminals such as smartphones, tablets, notebooks, smart TVs, game consoles, and desktop computers makes it extremely difficult to achieve the same level of protection regardless of the device used. On the other hand, when various users share devices (e.g., parents and kids using the same devices at home), the setup of distinct security profiles, policies, and protection rules for the different users of a terminal is far from trivial. In light of this, this article advocates for a paradigm shift in user protection. In our model, protection is decoupled from users' terminals, and it is provided by the access network through a trusted virtual domain. Each trusted virtual domain provides unified and homogeneous security for a single user irrespective of the terminal employed. We describe a user-centric model where non-technically savvy users can define their own profiles and protection rules in an intuitive way. We show that our model can harness the virtualization power offered by next-generation access networks, especially from network functions virtualization in the points of presence at the edge of telecom operators. We also analyze the distinctive features of our model, and the challenges faced based on the experience gained in the development of a proof of concept.

Diego Montero, Marcelo Yannuzzi, and René Serral-Gracià are with Technical University of Catalonia (UPC).

Adrian Shaw and Ludovic Jacquin are with Hewlett-Packard Laboratories, United Kingdom.

Antonio Pastor is with Telefónica I+D.

Antonio Lioy, Fulvio Risso, Cataldo Basile, and Roberto Sassu are with Politecnico di Torino.

Francesco Ciaccia is with Barcelona Supercomputing Center (BSC).

Mario Nemirovsky is with ICREA Researcher Professor at BSC.

Michael Georgiades and Savvas Charalambides are with PrimeTel PLC.

Jarkko Kuusijärvi is with VTT Technical Research Centre of Finland Ltd.

Francesca Bosco is with United Nations Interregional Crime and justice Research Institute.

INTRODUCTION

The protection of users' terminals against Internet threats is largely dominated by a device-centric model. This basically consists of installing a set of security applications on each terminal, such as anti-virus software and a personal firewall. An average user nowadays has multiple terminals, including a smartphone, a smart TV, and a notebook, and in many cases also a tablet, a desktop computer, and even a game console. These devices usually have different architectures (e.g., Intel or ARM) as well as different capabilities and operating systems (e.g., Android, Windows, or Linux), so the appropriate protection tools may not be available for all platforms. As a result, the most common practice is to install different security applications on the various terminals — or simply rely on the default protection means provided by the operating systems. Let us assume for a moment that users would like to have the same security policy and exactly the same protection level enforced on all of their devices. In the context of this article, we will call this the "uniform security aim." To achieve this goal, the user typically needs to understand the configuration details of each device, which typically involves the setup of different security applications on different platforms. For non-technically savvy people, this turns out to be an impossible hurdle to overcome. As a result, most Internet users suffer from wide variations in their protection levels, and this problem is exacerbated as the number of devices per user grows.

In this article, we propose a paradigm shift from device-centric protection to a user-centric model. The latter specifically addresses the two main drawbacks of the former: the need for dissimilar installations of security applications in different devices due to their different platforms, and the problem of non-uniform protection due to the difficulties in the configurations needed.

To cope with the first problem, we propose a model in which the protection and security policies are now unified and remain homogeneous for each user, independent of the terminal used. This is achieved by means of a user-specific trusted virtual domain (TVD), which is dynamically instantiated at a secure place in the network edge. As we shall show, the TVD can be instantiated either on the user's side (e.g., on a home gateway) or on the provider's side (e.g., on a next-generation broadband access server handling the user's connections).

To cope with the second problem identified above, we propose a user-defined security model that aims at ease of use by design. We discuss the importance of exposing the selection of high-level protection policies to the average user, and the necessity to enforce the configurations required transparently to the latter. This simple strategy detaches the definition of the protection policies from their corresponding configurations, thus allowing tailored protection even by non-

0163-6804/15/$25.00 © 2015 IEEE

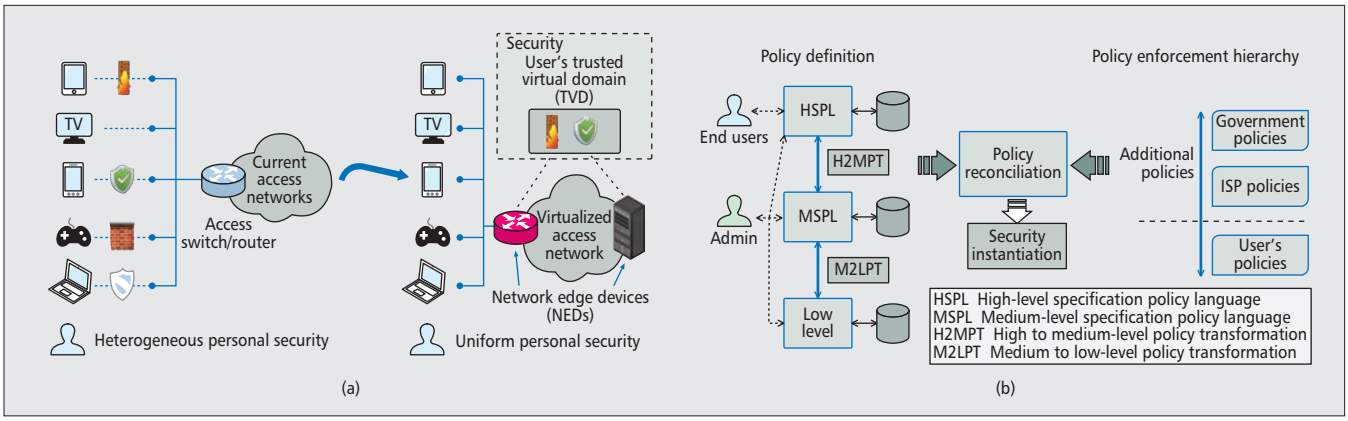

Figure 1. The two main objectives of our user-centric model, that is, uniform protection and ease of configuration: a) offloading security to the virtualized access network; b) policy definition and the policy enforcement hierarchy.

technically savvy users. It is worth highlighting that the virtualized security model described in this article can be applied both to residential and corporate scenarios. We describe its application in the form of a multi-tenant platform, considering the main stakeholders involved (i.e., service providers, infrastructure providers, security application developers, and users).

The remainder of the article is structured as follows. First, we outline the essentials of the paradigm proposed, including the new protection model and the security policy approach. Next, we introduce the general architecture and its main components. After that, we analyze the distinctive factors of our model, and outline some of the main conclusions that can be drawn from our prototype implementation. Finally, we conclude the article.

Toward a New Protection Paradigm

Figure 1a depicts the basic concepts, showing the evolution from device-specific security to a common security framework for all devices hosted in the access network. In our model, security applications that are commonplace today (anti-virus, firewalls, content inspection tools, etc.) shall be called personal security applications (PSAs). Observe that under the current protection model, the heterogeneity of devices and platforms requires the installation of various PSAs with similar roles and functions; actually, four PSAs are required in the example shown in Fig. 1a. Also observe that some devices may remain completely unprotected, as in the case of smart TVs.

Under our paradigm, the heterogeneous set of PSAs protecting the different devices is now moved and consolidated into a TVD. Each TVD only needs to host the minimum set of complementary PSAs required by the user (e.g., an anti-virus and a firewall in the example). A TVD is a "logical container" that is instantiated *per user*, and is composed of the following elements:

- The execution environments hosting the user's PSAs
- The required data, control, and management plane interconnectivity in order to guarantee the isolation between different users' TVDs (we delve into this later; Fig. 3).

The right side of Fig. 1a shows that a user TVD can be instantiated at either end of the access link. Indeed, as a logical container, a TVD may run entirely within a single network edge device (NED), or in a distributed way involving several NEDs. In our terminology, a NED is a device with virtualization capabilities that supports the instantiation of TVDs in a multi-tenant fashion. If the TVD is placed in a user's premises, the NED could be either an enhanced home gateway or customer premises equipment (CPE). Those devices may need additional compute, storage, and networking resources, and could be managed by the Internet service provider (ISP). If the TVD is placed in the ISP premises, as will be the case with the upcoming network functions virtualization (NFV) based access networks [1], a pool of nodes belonging to the NFV infrastructure could be the NEDs devoted to host our TVDs. Note that this second deployment strategy leverages the virtualization and processing power of commodity hardware, and the unquestionable trend toward its ubiquity at the network edge — although it does not exclude the adoption of the first deployment strategy as well. It is worth highlighting that our model has a remarkable advantage over cloud-based protection [2]. Whereas in the latter case the virtualized resources supporting the users' security are rarely on the path that would naturally be followed by user traffic, in our model, the TVD is always instantiated on the natural path. In other words, our model avoids routing detours, which would occur if the NEDs were located off the path between the user terminal and its traffic destinations (e.g., in the cloud).

As its name indicates, the TVD must be trusted, since it will execute security applications on behalf of the user on one or more nodes that are typically owned and managed by a third party. Appropriate techniques, such as remote attestations [3] or contractual agreements, must be put in place to guarantee the appropriate level of trust according to the security needs of a specific user. Also, observe that the NEDs must be secure, since they will host the applications of several users that could potentially affect each other. As we shall show, the NED must be connected with a secure channel to the user termi-

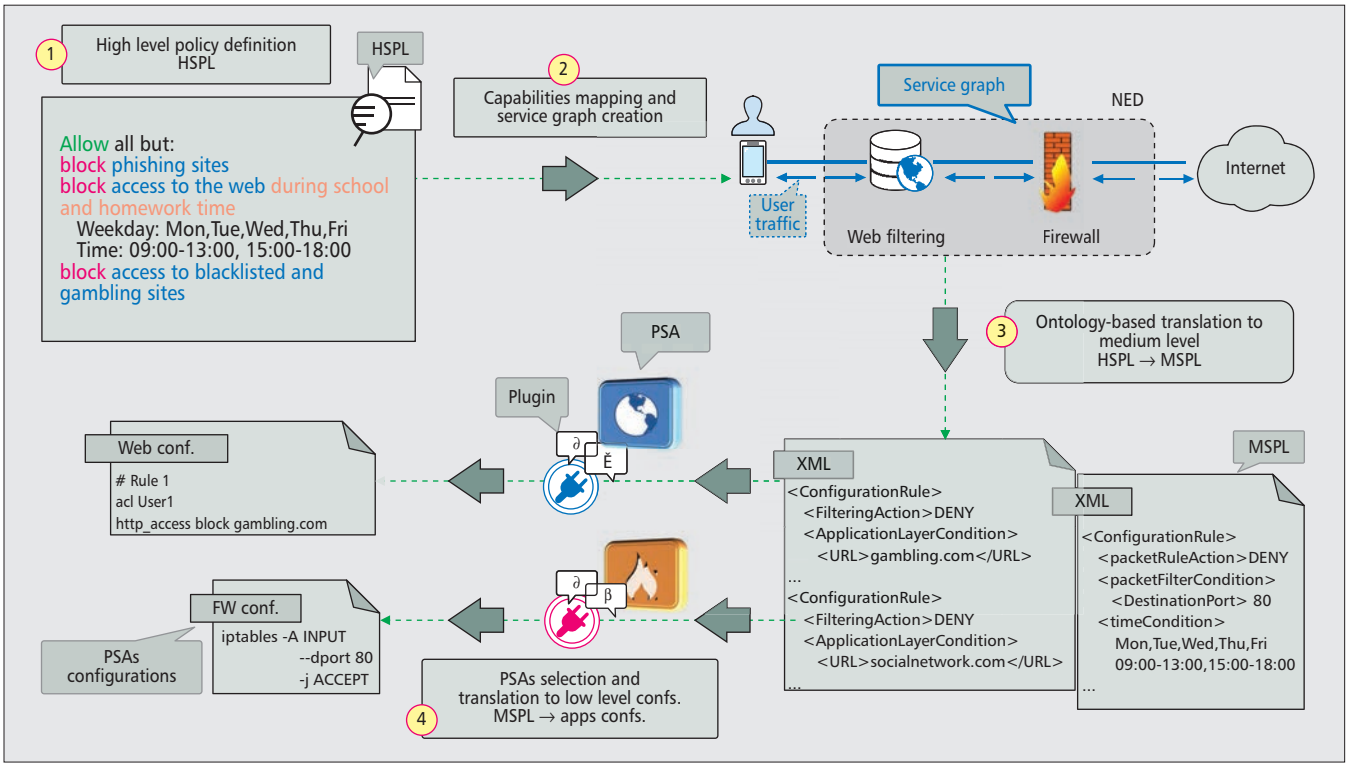

Figure 2. Example of policy definition and enforcement, going from HSPL to MSPL and then to low-level configurations.

nal, because this path may be subject to attacks that could try to bypass the security controls performed at the NED.

Each PSA within a TVD implements one (or possibly more) security controls that need to be configured according to the needs of a specific user. However, the configuration of security applications is often complex and not well understood by the majority of users. To simplify this task, we propose the model shown in Fig. 1b. The rationale behind it is that to build a real user-centric model, it is mandatory to allow users to specify their own security requirements (i.e., their *security policy*) in a straightforward way. Our design principle aims to meet the expectations of both non-technically savvy users and experts in the field, such as security administrators. For the former, the goal is to allow them to specify their security policy without needing to deal with the technicalities. For the latter, the goal is to allow them to fine-tune their policies while simplifying the configuration of the security applications under their administration.

To achieve these goals, our model is composed of three policy abstraction layers, and two translation services between them (see the left side of Fig. 1b). The first abstraction layer is supplied by the High-Level Security Policy Language (HSPL), a user-oriented authorization language suitable for expressing concepts related to user protection. HSPL allows users to express general protection requirements by means of sentences that are very close to natural language, such as "do not permit access to war content," "block my son from accessing gambling sites," or "allow email scanning." In our model, HSPL policies can be selected from a set of candidate policies that can be then customized and grouped (e.g., "block access to gaming sites" + "only during weekdays"). The policy sentences are internally mapped to a *subject-verb-object-attribute* authorization language that is currently under definition as an XACML profile [4]. For instance, the policy "block my son from accessing gambling sites" is interpreted as "block" (verb) "my son" (subject) "from accessing gambling sites" (the object). Predefined lists of subjects, verbs, and objects are made ready for the users, so they can easily compose their own sentences. Available attributes depend on the verb-object pair. Moreover, users can extend the predefined fields without being experts. The specific details of HSPL are out of the scope of this article, so for additional information the reader is referred to [5].

The lowest layer in the policy abstraction stack is what we call the "low level" in Fig. 1b, as it is the one that deals with the configuration details of the PSAs. This configuration procedure is clearly application-specific, and hence is not under our control.

With the aim of abstracting the specific configuration procedures while meeting the experts' needs, we have created an intermediate abstraction layer that allows the specification of PSA configurations using a PSA-independent format. The security policies in this abstraction layer are specified by means of the Medium-Level Security Policy Language (MSPL). The effort in the definition of the MSPL is not trivial. Indeed, depending on the heterogeneity of the different security control languages, the mappings can be arbitrarily complex. We address this complexity by means of an MSPL model that defines the main concepts (e.g., policies, rules, conditions, and actions), and is organized by *capabilities*. In

this context, capabilities are defined as basic features that can be configured to enforce a security policy (channel protection, filtering, anti-virus, parental control, etc.). Our approach also allows families of languages with similar concepts to be grouped (e.g., attributes, actions, or condition types), which can be captured by specific submodels built by analyzing several languages of controls sharing the same capability. For instance, through MSPL it is possible to write the configuration of a general packet filter or to configure the options of a general anti-virus. An illustrative example of MSPL outlining the translation from HSPL up to low-level configurations is sketched in Fig. 2.

Overall, writing policies in MSPL demands the same security awareness and level of expertise as specifying the configurations directly in the PSAs. The advantage, however, is that MSPL spares experts the burden of mastering several semantically equivalent security controls and syntaxes. Observe that PSA developers will need to provide their plug-ins jointly with their PSAs, in the form of a medium- to low-level policy translation (M2LPT) service (Fig. 2). Also note that the complexity mainly resides in the language definition, so these translators fundamentally perform syntax adaptation. Thanks to this approach, a security policy written in MSPL can be embodied by different PSAs, provided that the candidate PSAs offer the capabilities required by the user. In addition, the PSAs can be replaced without impacting the security policy specified by the user (e.g., replacing a Cisco packet filtering application by one provided by Checkpoint). For further details on MSPL, the reader is referred to [5].

As shown in Fig. 1b, the binding between HSPL and MSPL is supplied by the high- to medium-level policy translation (H2MPT) service. Different from the M2LPT translation, which is provided by the PSA developer, H2MPT represents a translation service that is natively provided by our architecture. H2MPT uses formal ontologies to provide the semantics implied by the high-level policy statements. Our ontology is based on [6], and it models the high-level concepts (subjects, objects, verbs, and attributes) as well as the medium-level concepts (rules, conditions, actions, resolution strategies) and the capabilities. The ontology also contains information on how predefined HSPL concepts are expanded into useful information for building MSPL rules. The translation process first identifies a set of applications that can enforce the security policies (e.g., a web filter and a firewall), and then generates the MSPL for the selected applications. The HSPL verb-object pairs are used to match the capabilities needed for policy enforcement, while the capabilities per se are used to determine the PSAs and their interactions.

Moreover, a meta-model defines how HSPL sentences are mapped into MSPL concepts, and how these concepts must be assembled to build valid rules. This meta-model is used by a set of enrichment modules and by a standard ontology reasoner to gather all the information needed to create MSPL policies that enforce the HSPL policy [5, 6]. Finally, an H2MPT component combines this information into MSPL policies. This translation is done transparently for nontechnically savvy users (i.e., for those users specifying their policies through HSPL). We contend that by having a high-level policy specification language, our model provides far more flexibility and expressiveness than approaches based on profiles or templates. This is because these latter basically wrap under a common name a set of low-level settings, which are basically applied for a fixed set of security controls.

In the model we conceive, the PSAs can be selected by the users themselves or by a provider. If the user only specifies the HSPL, the PSAs are automatically selected from a catalog of available applications based on the PSAs that meet the functionality required by the policies. In our model, the capabilities of a PSA are specified through a "PSA manifest." In this context, the selection may be straightforward — when only one PSA is available with the required capabilities — or it may be based on various criteria if multiple PSAs could offer those capabilities (on the PSA reputation or its cost, etc.).

Another important aspect is that according to recent studies, human mistakes are the major cause of breaches and vulnerabilities [7]. Thus, our model provides analytics that help reduce the likelihood of such mistakes. These include contradictions among policies in different PSAs, policy contradictions within a PSA, or cases leading to suboptimal performance (e.g., rules that are never matched and simply increase the processing time). Our model identifies these types of anomalies by means of state-of-the-art techniques [8]. We represent clauses as hyperrectangles so that anomalies can be detected by using geometric intersections. Anomalies are classified by evaluating geometric relations among conditions (e.g., inclusion, intersecting conditions but no one includes the other), as well as relations between actions (e.g., same action, equivalent actions, conflicting actions). The resolution is dealt with by formally modeled strategies, which cover a set of existing security control resolution mechanisms. Upon detection, we provide hints on how to resolve them and notify the effects of each decision.

Moreover, the model we envision should support multiple actors, which could simultaneously operate on the same traffic (see the right side of Fig. 1b). Each of these actors may possibly impose its potentially conflicting security policy. For instance, a user can decide the level of protection needed, but the ISP may impose other limitations in order to guarantee the integrity of its network. In turn, the government may impose additional restrictions. In the case of conflict between the different policies in the hierarchy, our approach is to automatically resolve such anomalies, and inform the user about the issue and its outcome.

In order to resolve such conflicts, a "reconciliation" [9] process is performed. The latter takes the policies of the different actors that must be reconciled, and obtains a single MSPL policy to be enforced by the user's PSAs. The core of this process is the resolution of contradictions among rules from different policies. Priorities and hierarchies are some of the simplest ways to resolve

A derived requirement posed by multi-tenancy is network isolation. The SECURED architecture must ensure the isolation of traffic among different users. More precisely, each tenant will be configured with a dedicated and private virtual network.

contradictions (i.e., rules from higher-priority policies/actors prevail), and they typically map well to contractual frameworks. However, custom reconciliation strategies can be defined. The reconciliation process copies non-conflicting rules in the reconciled policy, while each resolved contradiction generates a new rule. The latter have higher priority than the original ones, and the correct action is decided by the selected reconciliation strategy. More details on our reconciliation approach can be found in [10].

Observe that actors may decide not to disclose their policies to other actors. In that case, reconciliation strategies that require full access to the policy set are not possible. An alternative approach is to use *policy chaining*. This consists of redirecting the output of one set of PSAs in an administration domain (e.g., the user PSAs) to a set of PSAs in another domain (e.g., the ISP PSAs). The user must not necessarily own the PSAs in other domains when chaining is performed. This is useful when more sophisticated controls are required by the entities that specify the higher policies in the stack.

An Example of Policy Translation and Enforcement

To better describe our new paradigm, we present an example that illustrates the step-by-step process, starting from the definition of high-level policies up to the configurations made to guarantee their enforcement. Figure 2 depicts a simplified but complete example of the policy definition process for a non-technically savvy user. It comprises four basic steps. First, the user is requested to define its policies using HSPL. This user-oriented authorization language allows a set of general security rules to be expressed and customized by means of sentences that are very close to natural language (e.g., *block phishing sites*).

Next, the HSPL policy sentences are mapped to a subject-verb-object-attribute authorization language aiming to extract the different security capabilities required by the user (Fig. 2, step 2). As a result, a service graph is built, where the nodes represent generic applications (PSAs) capable of fulfilling the security requirements. Observe that two applications are required in the example, web filtering and a firewall. The selection of PSAs is based on the manifest provided along with each PSA, which indicates its specific capabilities. Third, by using the ontology and the service graph information, the security policies are translated into MSPL, obtaining the application-independent definition of policies requested by the user (Fig. 2, step 3). The representation of MSPL policies is stored and managed in XML format. Fourth, specific PSAs are selected satisfying the capabilities and requirements of the user. As mentioned above, the specific PSAs can be selected by either the user or the provider. For each PSA, the configurations are created using an application-specific translation plugin. These plugins convert the generic MSPL rules to application-specific configurations (Fig. 2, step 4). These configurations will be the inputs once the PSAs are instantiated and linked.

Finally, once the PSA configurations are created, an orchestration system instantiates each PSA and enforces its particular configuration, hence providing the security policies defined by the user.

The SECURED Architecture

This section introduces the envisioned architecture, which we call SECURED [5]. As explained above, SECURED provides a system where users can offload their PSAs to their nearest compatible NED. The architecture is specifically devised to be heavily multi-tenanted and flexible enough to be used in scale-out systems. From a use case point of view, it can be expanded and deployed in a variety of ways, ranging from small set-top boxes or home gateways up to deployments on a much larger scale in a distributed environment (e.g., in localized data centers at the edge of ISP networks). Our focus in this section is on the main architectural components.

General Overview

The architecture must support the dynamic allocation and instantiation of users' security. The security functionality of each user can be comprised of different PSAs in a defined arrangement through service chaining, and these PSAs can be deployed within the same physical host or in a distributed manner. As a result, two general requirements are imposed on the architecture:

- Massive multi-tenancy, which implies isolation of users, their applications, and network traffic
- A secure and verifiable infrastructure and environment, which users can trust to host their security applications

A general view of the basic architecture is depicted in Fig. 3. The figure shows a generic deployment (e.g., on an NFV POP of an ISP). It is worth noting that in simpler deployments (e.g., when the NED is a home gateway), the functionality provided by some of the systems at the top of Fig. 3 could be simplified and embedded in the NED itself, or might not be needed, such as the case of the NFV orchestrator.

Overall, the first requirement is to guarantee complete isolation between different users. In light of this, the TVD was designed as an isolated environment that will hold the security applications of a user and in turn process the user traffic. A TVD comprises one or more execution environments (EEs). An EE is a lightweight and heavily controlled environment that contains and executes one or more user PSAs,each operating on the principle of least privilege. Thus, within SECURED, two levels of isolation are defined (Fig. 3):

- The *compartmentalization layer*, which is mainly responsible for the isolation between user TVDs
- The *containment layer*, which handles isolation between PSAs within an EE

Thus, an EE could be either a compartment or containment layer, respectively.

A derived requirement posed by multi-tenancy is network isolation. The SECURED architecture must ensure the isolation of traffic among different users. More precisely, each tenant will be configured with a dedicated and private virtu-

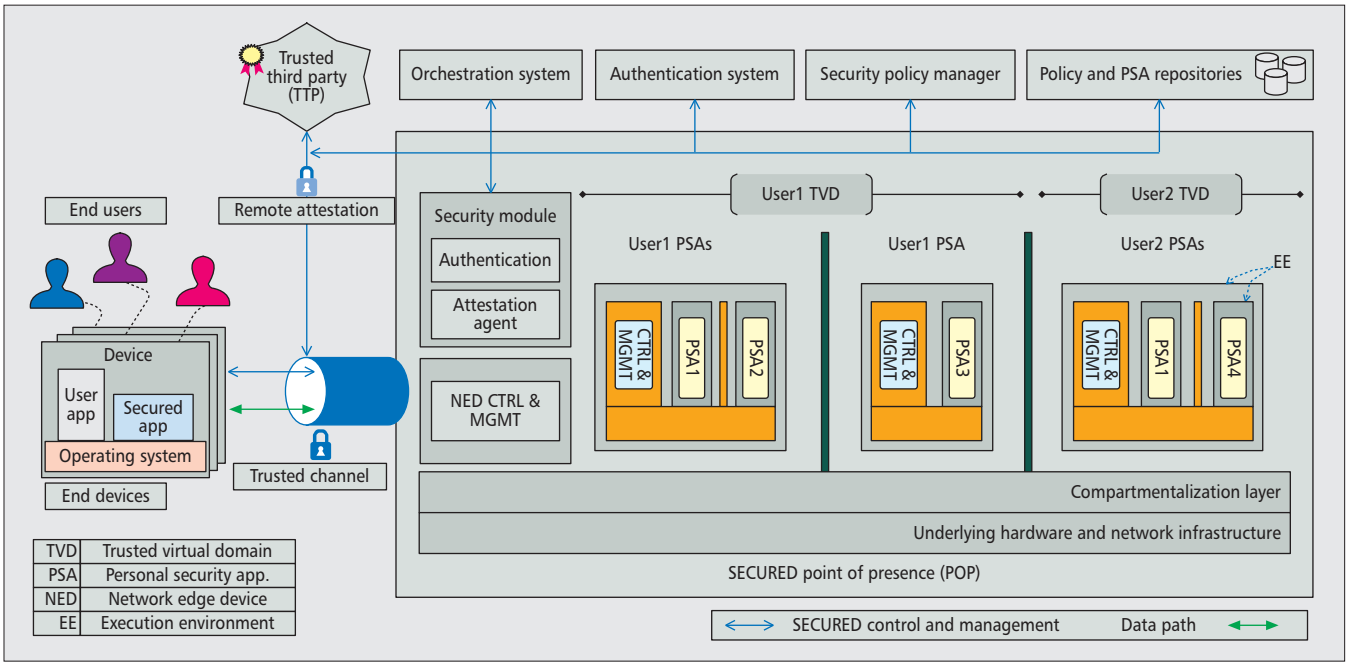

Figure 3. The basic SECURED architecture showing a multi-tenant scheme on a point of presence (POP).

al network. This network connects the different PSAs with the end user on one side and the Internet on the other. Furthermore, the architecture defines a private management network that sets up, controls, and manages the different TVDs. Both the compartmentalization and containment layers have a *control and management* component, which aims to establish separation between the technology-independent part and the implementation-dependent technology.

The second requirement is related to the establishment of trust between the end user and SECURED. This requirement is vital, since users would like to establish a certain level of trust with SECURED prior to requesting the instantiation of security applications and sending their traffic. We address this requirement by using the concept of *remote attestation* (RA). SECURED leverages trusted computing mechanisms to measure the system software upon component startup, where resulting measurement digests are held by a secure root of trust, such as a hardware device like a trusted platform module (TPM) [11]. These measurements can be cryptographically signed by the device and sent to the users whenever they send an attestation request. The process of RA poses a major challenge for SECURED, and preliminary insight on a proof-of-concept implementation is described later.

Main Components

Security Module — This module is the front-end, which is contacted during connection establishment. It comprises two elements, the *attestation agent* and the *authentication module*. Prior to authenticating, the end user first contacts SECURED in an attempt to establish a secure connection while also performing the remote attestation protocol. To this end, the SECURED system receives a challenge request to perform an attestation of its software configuration. A mutually trusted third party (TTP) system is involved in the attestation process. The TTP is responsible to keep a copy of known-good measurements, and provide a secure verification service to the user for verifying remote attestation responses. After a successful check, a secure channel is created, and the user safely sends her credentials to the *authentication module.*

Authentication System — The authentication of users is a key component of SECURED. This can be implemented either using a local (standalone) authentication system or relying on an existing external authentication infrastructure (e.g., an AAA+ system). The result of the authentication process is to obtain tokens allowing the interplay between the main components within a NED, and external subsystems such as PSA repositories. Once the user is authenticated, the instantiation of his security must be enforced.

NED Control and Management — Once the user is authenticated, this module retrieves the user policies and metadata related to the composition of the required security applications. After that, the control and management module drives the instantiation of the user TVD, including its applications and setup of the virtual network. More specifically, this module determines the resources required for the user TVD, and commands the instantiations required as well as the deployment and interconnection of the PSAs. This computation encompasses an analysis of the required compartments, containments, and virtual networks to be allocated in order to instantiate the security applications. This analysis considers the PSA requirements along with the availability of resources, and the required configuration of the network (physical and virtual). In addition, this module also manages the extension

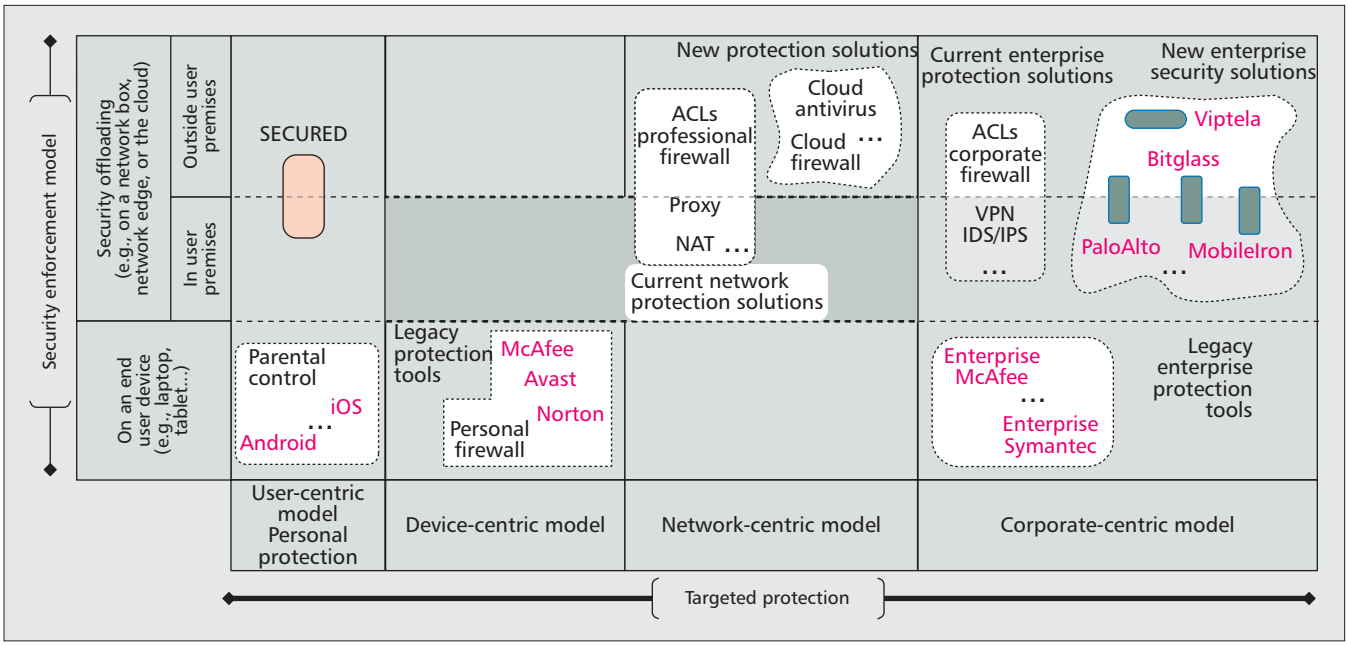

Figure 4. Positioning SECURED considering some of the most common tools as well as some of the most recent and compelling solutions in the area.

of the user data path to connect the user's device to the newly created TVD.

Orchestration System — In the case of an NFV POP, the NED control and management module will be assisted by the NFV orchestration system. However, in simpler scenarios, the former could entirely handle all the configurations required. In other words, when the NED is embodied in the home gateway of a residential user, the orchestrations needed will be handled locally without requiring any external orchestrator. In general terms, the orchestration system should deal with the instantiations and configurations in large distributed systems (e.g., an NFV POP), preferably in a "technology-agnostic" way. The technology-dependent part could be managed by the control and management module embedded in the NED. In our model, the attestation agent keeps track of the different components during the instantiation phase (i.e., compartments, containments, and PSAs), and manages the corresponding measurements in order to present an attestation proof back to the user concerning her TVD.

Security Policy Manager — This module is in charge of handling the user's policies and the reconciliation process prior to performing the configuration of the user's PSAs.

PSA Repositories — The applications are retrieved from these repositories with their respective MSPL plugins, which then need to be loaded into one or more TVD containments.

SECURED App — This is the only application that needs to be installed in a user device. Its role is basically to support the secure communications with the NED, and handle the remote attestations and their outcomes.

Overall, the architecture introduced in this section allows the dynamic creation of trusted and virtualized execution environments throughout the access network. In this framework, several actors such as users, corporate information and communications technology (ICT) managers, infrastructure providers, security service providers, and software developers can interplay and benefit from our user-centric protection model. An important remark about the proposed architecture is its alignment with the emerging NFV technology. NFV is an enabler for SECURED, and will be essential for guaranteeing its scalability.

Analysis of SECURED

The security model proposed in this article has several distinctive factors that make it unique. To show this, we position SECURED in the current spectrum of protection techniques and highlight its main differences with state-of-the-art solutions. In addition, we present and discuss our initial evaluations of a proof-of-concept implementation, with special focus on performance aspects related to the security, trust, and service verification offered by SECURED.

Positioning SECURED within the Security Panorama

The spectrum of solutions designed to counter security threats is really broad. The solutions available today can be reasonably categorized according to the table shown in Fig. 4. As can be observed, there are solutions that are focused on protecting the end-user device, while others propose different forms of security offloading. Moreover, current protection schemes can be classified based on whether they are user-centric, device-centric, network-centric, or corporate-centric. In a nutshell, Fig. 4 presents a high-level comparison of different security protection

schemes according to two general criteria: the targeted protection model and where the security is enforced.

To the best of our knowledge, SECURED is the only solution available nowadays that proposes a true user-centric model which specifically addresses the need for device-independent security. As described previously, the user-centric approach is achieved thanks to the HSPL and MSPL languages, and the H2MPT and M2LPT translation services between the three abstraction layers involved. This allows users and even experts in the field to focus on their security policies rather than on the configuration details of specific security applications. Another important aspect is that, in contrast to many of the offloading solutions available today, which are typically deployed in the cloud, our solution admits a rich variety of deployments on either edge of an access link. Cloud-based solutions provide compelling protection schemes while avoiding several of the overheads for end users (e.g., corporate customers). The downside, however, is that they require routing detours, are not really user-centric (at least not yet), do not provide essential trust means such as remote attestation, and do not support advanced features such as anomaly verification and policy reconciliation techniques. These latter two are a couple of distinctive aspects in SECURED, and therefore are the center of our assessment and analysis at this stage. We proceed to provide insight foresee based on a proof-of-concept implementation.

Remote Attestations

Trust establishment between an end user and the protection platform is a critical step toward security offloading. In our model, we use remote attestations (RAs) and verification techniques for the trust establishment process. Let us assume the following scenario: A user connects through an insecure channel and requests protection from SECURED. Prior to starting traffic exchange, the user is requested to create a trusted channel toward a NED. A trusted channel is an instance of a secure channel (e.g., a virtual private network, VPN), where the endpoints are attested before any data exchange. In SECURED, the trusted channel protects users against a potentially compromised NED. However, enabling these security countermeasures introduces overhead. On one hand, users may experience delay during the establishment of the connection with the NED. This is due to the integrity check needed, which is issued only once per user during the connection. Likewise, administrators may face scalability problems, since a portion of the network and the computational resources will be dedicated to the security checks as users connect. Normally, solutions offering this feature use a cryptographic chip — the trusted platform module (TPM) [11] — that may pose a performance bottleneck while issuing the required verifications. SECURED overcomes this issue by introducing a trusted third party (TTP) system (Fig. 3). This is an entity that is trusted by users and infrastructure administrators, which asynchronously attests a set of controlled NEDs in a configurable time interval. The advantage of this approach is twofold. First, the workload for the attestation process does not increase with the number of connecting users, since the NED is common to all users. Second, end users will get a response regarding the integrity of the NED almost immediately.

We have developed a prototype that uses *strongSwan* [12] for the creation of a trusted channel with IPsec. To this end, strongSwan has been adapted to generate RA requests to the TTP, and either continue or drop the connection depending on the result of the integrity verification. The TTP has been implemented with *OpenAttestation* [13], a framework for attesting large infrastructures. Our initial results show that the establishment of an IPsec connection without attestation is very fast (around 76 ms), and the asynchronous attestation with the TTP in the same setting does not introduce noticeable delays (around 217 ms). Unlike our solution, performing synchronous RA adds a significant delay on the creation time of the tunnel (around 4.119 s).

Another source of overhead is due to the size of the integrity reports. Figure 5 shows the size of the reports exchanged between the NED and the TTP. The results were obtained over a 10-min period, where a user repeatedly connected to the NED. While the first report generated is near 300 kB, subsequent reports are very small (between 4 and 8 kB) due to the fact that *OpenAttestation* only sends new integrity measurements, which are performed on the NED with the Integrity Measurement Architecture (IMA) [14] software. Note that the first report contains all the measurements performed at boot time. Furthermore, new reports will be generated only if new measurements are produced on the NED (i.e., when new software is executed).

These initial results show that smartly performing RA does not incur a noticeable overhead for the end user, as all the heavy lifting is asynchronously performed behind the scenes. The analysis also sheds light on the feasibility of enabling end users to remotely verify the status of a NED. It is worth highlighting that the interval between two consecutive attestations can be configured, thereby offering the possibility of defining convenient trade-offs depending on the case. So far, we have seen that the RA of a single NED will introduce negligible overhead.

However, performing the RA over a distributed infrastructure poses complex challenges and remains an open problem. These challenges increase when we also include multi-domain scenarios or requirements such as user mobility and roaming. Furthermore, the assessment of time bounds for dynamic service deployment, as well as the appraisal of the multi-tenant isolation model, will need to be deeply analyzed in the near future. We plan to develop a comprehensive prototype that will address these issues. Our research and future evaluations will prioritize the following aspects: security and isolation, ease of use, deployment and service provisioning in relatively short timescales, and, related to the latter, support for user mobility.

User-Centric Policy Framework

Our policy-based framework also needs an in-depth performance assessment to evaluate if the policy services can actually be used in real sce-

These initial results show that smartly performing RA does not incur a noticeable overhead for the end user, as all the heavy lifting is asynchronously performed behind the scenes. The analysis also sheds light on the feasibility of enabling end users to remotely verify the status of a NED.

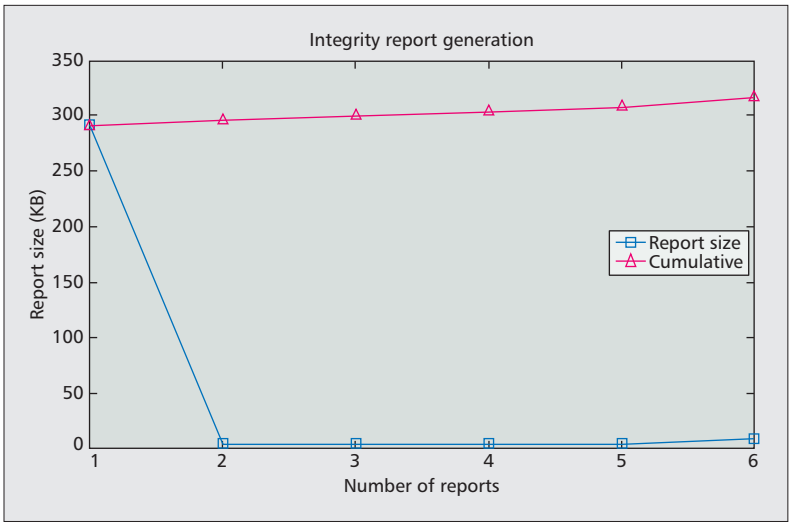

Figure 5. Size of the integrity reports generated with OpenAttestation.

narios. To this purpose, we tested the performance of the reconciliation, anomaly analysis, and translation with an off-the-shelf computer equipped with an Intel processor i7-3630QM (2.4 GHz), with 16 GB of RAM, running OpenJDK RE 1.7.0 55 on top of a Linux operating system. We performed two different rule processing experiments: an average case with a realistic amount of rules, and a higher bound worst case scenario with thousands of rules. In both cases, we considered two types of filtering within the PSAs: a *packet filter* and an *L7 filter*. During the experiments, we measured the time required to process and validate the filtering rules. As discussed earlier, such validation is composed of three parts: anomaly analysis, reconciliation, and M2L translation.

The first tests evaluate the performance of a small/medium scenario, where the number of rules per user are on average in the range of tens or hundreds. This estimation was derived from a use case with four actors, where policies included 10 to 50 rules for each PSA, amounting to an average of 100 rules to be processed. We consider that these numbers per user are representative of a reasonable average, since in a user-centric approach, the size of the rule set will not raise to thousands — which is typically the case found on border firewalls of large companies. As reported in the first row of Table 1, all three measured policy-related tasks were completed in less than 1 ms.

The second experiment aims to assess the scalability for large-scale policy scenarios. This means scenarios that, as stated on [8], statistically satisfy significant parameters of the policies that can be found in practice. This experiment provides two different results. On one hand, we compute the necessary processing time for a very large amount of rules. On the other hand, we compute the amount of rules that can be processed in 1 s — a amount for interactive purposes. Both results are reported in Table 1. We observe that for the anomaly analysis, our prototype can process 5000 rules in 12 s for the packet filtering case. In contrast, L7 filtering requires 90 s to perform the same task, due to the massive usage of regular expressions. In terms of the number of rules processed in less than a second, we obtained 2000 rules for the packet filter case and 1000 for L7 filtering. Regarding the reconciliation part, we were able to process 1500 packet filter policies and 1000 L7 filter policies in less than 1 s. However, the worst cases for the 5000 rules considered yielded reconciliation times of 74 s and 364 s for the packet filter and L7 filter, respectively. Finally, the translation of MSPL into low-level configurations is a linear problem that took approximately 1 s with 5000 rules with both an XSLT-based approach and a SAX-based Java program. All these results are summarized in Table 1.

Given that these computations are performed at infrastructure elements, wherein computational power can be adjusted as needed, we consider that our approach can reasonably scale in several real scenarios. For instance, the average cases are representative of residential scenarios, and all computations can be resolved online. We also consider that the processing of 5000 rules is quite representative of a corporate user case (e.g., an SME), and the worst cases are highly unlikely to occur in practice. Anyway, the bounds found indicate that there are cases in which the reconciliation cannot be handled online, and therefore, this analysis serves as a starting point for investigating new strategies and optimizations.

Conclusion

In this article, we have argued that for the large majority of Internet users, the current protection model against security threats is broken. Users typically have multiple devices, but achieving the same level of protection irrespective of the device used has become "mission impossible." We have proposed a paradigm shift in user protection through a user-centric model that also decouples security from user terminals. The protection model we envision is based on the setup of a trusted virtual domain per user, placed in the access network. Our approach facilitates security policy configuration, and enables uniform protection independent of the terminal used. We have also shown that the trust and security verification mechanisms offered by a prototype implementation can be applied in many practical scenarios, such as the case of residential users.

In spite of this, several of the issues addressed in this article require significant efforts in terms of research. The list is large, and includes aspects such as remotely attesting distributed systems, multi-domain scenarios (i.e., the interplay among different ISPs), user mobility and roaming scenarios, scalability analysis, assessment of upper bounds for dynamic service deployment, isolation assessment, development of a comprehensive threat model, constraints and deeper analysis of corporate scenarios, and more.

Acknowledgment

The research described in this article is part of the SECURED project [5], co-funded by the European Commission under the ICT theme of FP7 (grant agreement no. 611458).

	Filtering level	Anomaly analysis	Reconciliation	M2L translation
Average case (time to process 100 rules)	Packet filter	< 1 ms	< 1 ms	< 1 ms
	L7 filter	< 1 ms	< 1 ms	< 1 ms
Worst case (time to process 5000 rules)	Packet filter	12 s	74 s	< 1 s
	L7 filter	90 s	364 s	< 1 s
Number of rules processed in 1 s	Packet filter	2000	1500	> 5000
	L7 filter	1000	1000	> 5000

Table 1. Results of the tests for policy-based tasks.

The protection model that we envision is based on the setup of a trusted virtual domain per-user, placed in the access network. Our approach facilitates security policy configuration, and enables uniform protection independently of the terminal used.

References

[1] ETSI, "Network Functions Virtualisation (NFV) Architectural Framework," http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf, December 2014.

[2] J. Sherry *et al.*, "Making Middleboxes Someone Else's Problem: Network Processing as a Cloud Service," *Proc. ACM SIGCOMM 2012 Conf. Applications, Technologies, Architectures, and Protocols for Computer .*, 2012, pp. 13–24.

[3] K. Goldman, R. Perez, and R. Sailer, "Linking Remote Attestation to Secure Tunnel Endpoints," *Proc. 1st ACM Wksp. on Scalable Trusted Computing*, 2006, pp. 21–24.

[4] OASIS, "eXtensible Access Control Markup Language (XACML) Version 3.0," http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.pdf, Jan. 2013.

[5] "Security at the Network Edge (SECURED)," http://www.secured-fp7.eu/.

[6] C. Basile *et al.*, "Ontology-Based Security Policy Translation," *J. Info. Assurance and Security*, vol. 5, no. 1, 2010, pp. 437–45.

[7] IBM Global Technology Services, "IBM Security Services 2014 Cyber Security Intelligence Index," http://media.scmagazine.com/documents/82/ibm_cyber_security_intelligenc_20450.pdf, June 2014.

[8] C. Basile, A. Cappadonia, and A. Lioy, "Network-Level Access Control Policy Analysis and Transformation," *IEEE/ACM Trans. Networking*, vol. 20, no. 4, 2012, pp. 985–98.

[9] P. McDaniel and A. Prakash, "Methods and Limitations of Security Policy Reconciliation," *ACM Trans. Info. System Security*, vol. 9, no. 3, Aug. 2006, pp. 259–91.

[10] C. Basile *et al.*, "A Formal Model of Policy Reconciliation," *Proc. 23th Euromicro Int'l. Conf. Parallel, Distributed, and Network-Based Processing*, March 4–6, 2015.

[11] H. Zhang, Z. Qin, and Q. Yang, "Design and Implementation of the TPM chip J3210," *Proc. 2008 Third Asia-Pacific Trusted Infrastructure Technologies Conf.*, 2008, pp. 72–78.

[12] A. Steffen, "strongSwan — IPsec for Linux," https://www.strongswan.org.

[13] Intel, "OpenAttestation SDK: A SDK for Remote Attestation," https://github.com/OpenAttestation/OpenAttestation.

[14] R. Sailer *et al.*, "Design and Implementation of a TCG-based Integrity Measurement Architecture," *Proc. 13th Conf. USENIX Security Symp.*, 2004, pp. 223–38.

Biographies

Diego Montero (dmontero@ac.upc.edu) received his B.Sc. in computer engineering from the University of Cuenca, Ecuador. He completed his M.Sc. in computer architecture, networks and systems from the Technical University of Catalonia (UPC), Spain. He is currently a Ph.D. candidate at the Networking and Information Technology Lab (NetITLab), where his research interests include network security, SDN, network virtualization, and mobility.

Marcelo Yannuzzi (mayannuz@cisco.com) received a degree in electrical engineering from the University of the Republic, Uruguay, and MSc. and Ph.D. degrees in computer science from the Department of Computer Architecture, UPC, Spain. He is with the Corporate Technology Group at Cisco Systems International, Switzerland. Before that, he was head of NetITLab, as well as the Advanced Network Architectures (ANA) research group at UPC. He has led several projects in close interaction with European and U.S. companies and research centers. His interests lie in the areas of fog computing, IoE, security, NFV, orchestration and management, and mobility.

Adrian Shaw (adrian.shaw@hp.com) is a research scientist at HP Labs, Bristol, United Kingdom, where he is a member of the Embedded Control Points (ECP) group in the Security and Cloud Lab. His research covers the areas of operating systems, virtualization, and trusted computing. He received his M.Sc. in computer security from the University of Birmingham, United Kingdom.

Ludovic Jacquin (ludovic.jacquin@hp.com) is a research scientist at HP Labs in the Security and Cloud Laboratory at Bristol. He received his Ph.D. in computer science from Grenoble University in 2013, with a thesis titled *Performance/Security Trade-off for High-Bandwidth Internet VPN Gateways* supervised by Vincent Roca and Jean-Louis Roch. His main research currently focuses on infrastructure security, especially attestation of the network devices in the new "softwarized" paradigm.

Antonio Pastor (apastor@tid.es) is a technology expert in security on networks working for the network virtualization group in the GCTO unit within Telefónica I+D. Since 2006 he has been working as an expert in IP network security designs and services, and holds several certifications from ISACA and GIAC in this area. Currently, he is working on SDN and NFV technologies oriented toward security, including applied research projects and close-to-market services.

René Serral-Gracià (rserral@ac.upc.edu) received his degree in computer science (2003) and a Ph.D. (2009) from UPC. He is the R&D head of NetITLab at UPC, where he is leading different research initiatives, including projects under the European FP7 Research Framework as well as with industry. He is also an associate professor in the Department of Computer Architecture at UPC. His research interests are focused on SDNs, overlay networks, network security, routing optimization, and QoE assessment of multimedia traffic.

Antonio Lioy (lioy@polito.it) (M.Sc. in electronic engineering and Ph.D. in computer engineering) is a full professor at the Politecnico di Torino, Italy, where he leads the TORSEC group. His current research interests are network security (especially optimization and automatic configuration), PKI applications (e-identity and digital workflows), and policy-based protection of ICT systems. He is the coordinator of the SECURED project, and frequently acts as a cybersecurity expert for the Italian government and the European Commission.

Fulvio Risso (fulvio.risso@polito.it) received his Ph.D. degree in computer and system engineering from Politecnico di Torino in 2000. He is currently an assistant professor with the Department of Control and Computer Engineering, Politecnico di Torino. His current research activities focus on efficient packet processing, traffic analysis, and programmable networks.

Cataldo Basile (cataldo.basile@polito.it) received his Ph.D. degree in computer and system engineering from Politecnico di Torino in 2005. He is currently a research assistant at Politec-

nico di Torino. His research is concerned with policy-based management of security in networked environments, policy refinement, general models for detection, resolution and reconciliation of specification conflicts, and software security.

ROBERTO SASSU (rsassu@suse.de) is a senior engineer at SUSE Linux GmbH. Previously, he worked as a research assistant at Politecnico di Torino. His research activity focused on sensitive data protection, platform integrity evaluation, and cloud computing. He received his M.Sc. in computer engineering from Politecnico di Torino.

MARIO NEMIROVSKY (mario.nemirovsky@bsc.es) received a Ph.D. degree in electrical and computer engineering from the University of California, Santa Barbara in 1990. He was an adjunct professor at the same university from 1991 to 1998. After being chief architect at companies such as Apple, Inc., National Semiconductors, and General Motors (GM), he founded several renowned startups including XStream Logic, FlowStorm Networks, ConSentry Networks, and Miraveo. In 2007, he became an ICREA Senior Research Professor with the Barcelona Supercomputing Center (BSC), Spain. He holds more than 60 issued patents. His current research interests include multithreaded multicore systems, high-performance systems, IoT, big data, and network processors.

FRANCESCO CIACCIA (francesco.ciaccia@bsc.es) is currently a researcher at BSC. He is part of the Unconventional Computer Architecure and Networks research group. He received his M.S. degree in computer engineering from Politecnico di Torino. His main reaserch interests include network security, SDNs, virtualization, and the Internet of Things.

MICHAEL GEORGIADES (michaelg@prime-tel.com) is the R&D manager at PrimeTel PLC and an adjunct faculty member at the Open University of Cyprus. He received a B.Eng. from King's College London (2000), an M.Sc. from University College London (2001). and a Ph.D. from the University of Surrey (2008) in telecommunications. He has been involved in more than 10 EU ICT projects, and published more than 40 journals, book chapters, and conference publications. He received the Nokia Prize of Research Excellence for a Patent in 2004.

SAVVAS CHARALAMBIDES (savvasch@prime-tel.com) is an R&D research engineer at PrimeTel PLC. He received his B.Sc. in computer science and engineering from the University of Patras, Greece, in 2012 and his M.Phil. in Advanced Computer Science from the Computer Laboratory of the University of Cambridge, United Kingdom, in 2013. He has five publications in international journals and conferences, and is a co-author of a book chapter in the area of computer network simulations.

JARKKO KUUSIJÄRVI (jarkko.kuusijarvi@vtt.fi) received B.Sc. (Tech.) and M.Sc. (Tech.) degrees from the University of Oulu, Finland, in 2008 and 2010, respectively. Since 2010, he has been a research scientist at VTT. His current areas of research interests include mobile applications, security visualization, and cybersecurity.

FRANCESCA BOSCO (bosco@unicri.it) is a UN project officer working in UNICRI responsible for developing research and capacity building activities on misuse of technology and technology-enabled crimes. She is a member of the Advisory Groups on Gender and Secure Societies in the framework of Horizon2020 and of the Internet Security Expert Group of the EC3. She is a co-founder of the Tech and Law Center and a member of the Centre for Internet & Human Rights of European University Viadrina.

Toward an SDN-Enabled NFV Architecture

Jon Matias, Jokin Garay, Nerea Toledo, Juanjo Unzilla, and Eduardo Jacob

ABSTRACT

This article presents the progressive evolution of NFV from the initial SDN-agnostic initiative to a fully SDN-enabled NFV solution, where SDN is not only used as infrastructure support but also influences how virtual network functions (VNFs) are designed. In the latest approach, when possible, stateless processing in the VNF shifts from the computing element to the networking element. To support these claims, the article presents the implementation of a flow-based network access control solution, with an SDN-enabled VNF built on IEEE 802.1x, which establishes services as sets of flow definitions that are authorized as the result of an end user authentication process. Enforcing the access to the network is done at the network element, while the authentication and authorization state is maintained at the compute element. The application of this proposal allows the performance to be enhanced, while traffic in the control channel is reduced to a minimum. The SDN-enabled NFV approach sets the foundation to increase the areas of application of NFV, in particular in those areas where massive stateless processing of packets is expected.

INTRODUCTION

The design, management, and operation of network infrastructure have evolved during the last few years, leveraging on innovative technologies and architectures. On top of this innovation, the service delivery has also progressed. Software-defined networking (SDN) and network functions virtualization (NFV) are key enablers for this evolution.

SDN has been one of the pillars of innovation in network infrastructures, allowing the decoupling of the control and data planes through an open and standard interface that enables the programmability of the network. OpenFlow, ForCES, and I2RS are some examples of SDN technology. SDN has also contributed to the virtualization of the network infrastructure, providing the foundation to isolate, abstract, and share the network resources.

Service provisioning is often based on proprietary hardware appliances, which imposes some restrictions when trying to deploy new network services, such as capacity and availability. In this scenario, the network infrastructure is not flexible enough to accommodate new services or migrate them to other locations due to its dependence on the physical appliances.

NFV has been proposed to innovate in the service delivery arena by using standard computing virtualization technology to consolidate in commodity hardware (i.e., standard high volume servers, storage, and switches) the functions previously performed by specific hardware appliances. Virtualized network functions (VNFs), which compose the service chain, are the basic elements to achieve the complete virtualization of service delivery and are commonly based on computing resources. The interconnection of VNFs, or traffic steering, is a challenging goal for the underlying network infrastructure. The migration of VNFs and dynamic composition of services make this task even harder than in legacy networks.

SDN and NFV are complementary technologies, and each one can leverage off the other to improve the flexibility and simplicity of networks and service delivery over them. For this aim, new architectures and interfaces between them are needed, and several proposals are emerging. In this article we try to clarify the evolution of their relationship in the context of service provisioning. We explain the evolution of the NFV architectures from an SDN-agnostic approach to a fully SDN-enabled architecture, a step forward in this evolution that exposes the programmability of the underlying network to build VNFs, enriching their architecture and improving their performance.

Adapting the use of resources to the actual demand is one of the main outcomes from a virtualized infrastructure, providing elasticity of resources instead of overprovisioning. However, the trade-off between flexibility and performance must be considered when designing the solution. The evolution of the NFV architecture presented in this article contributes to overcoming this limitation.

A real example, FlowNAC [1], is presented to illustrate the applicability of NFV using different architectures and the advantages offered by the SDN-enabled NFV approach in service provisioning scenarios. The main idea behind FlowNAC is to achieve fine-grained control of which traffic from the user is granted access to the network. The users are authenticated and authorized on a per service basis. All the incoming traffic from the user is independently evaluated and categorized in a specific service or none at the data plane. Then a basic allow or deny decision is enforced for each frame

The authors are with the University of the Basque Country (UPV/EHU).

0163-6804/15/$25.00

SDN and NFV are complementary technologies, and each can leverage off the other to improve the flexibility and simplicity of networks and service delivery over them. For this aim, new architectures and interfaces between them are needed, and several proposals are emerging.

depending on the associated service being authorized or not.

The depicted scenario is challenging in several aspects, and some topics are still open. The article concludes presenting the main challenges introduced by the SDN-enabled NFV architecture and relates them to the current efforts to bring the NFV proposal to fulfill the expected benefits.

Related Work

Major standardization efforts of the emerging NFV technology are being led by the European Telecommunications Standards Institute (ETSI), where the NFV Industry Specification Group (ISG) has recently published 11 NFV specifications, including the NFV architecture [2]. The defined architecture focuses on the aspects unique to virtualization, such as the transformation of the management and orchestration of VNFs, rather than common challenges to both physical and virtualized NFs, such as the control and operation of the end-to-end network service. Moreover, the NFV ISG is also coordinating and promoting public demonstrations of proofs of concept (PoCs) [3] that illustrate key aspects of NFVs, such as scalability, multi-tenancy, and migration issues.

The Open Networking Foundation (ONF) has also been active in the NFV arena and has proposed "A Flexible NFV Networking Solution" [4], outlining the benefits for the NFV deployment of an OpenFlow-enabled SDN approach to deal with the dynamic provisioning of networking services.

Recently, the Internet Research Task Force has published RFC 7426 [5], which proposes a common terminology for SDN layering and architecture based on significant related work from the SDN research community. In this regard, the work presented in this article mainly focuses on the separation between the control and forwarding planes, and is compatible with this RFC, but narrows down the scope to multi-tenancy and NFV-related implementation aspects.

The work done by ETSI and ONF is relevant as a reference to the two first steps of the NFV architecture evolution, presented later.

As a contribution to progress in NFV research, European Commission funded projects like T-NOVA [6] and UNIFY [7]. The T-NOVA project has the goal of designing a framework for providing NFs as a service for all the stakeholders, while the UNIFY project seeks to open up all the potential of virtualization and automation across the whole networking and cloud infrastructure. The work presented in this article is related to the latter.

Some other works are also related to the evolution of the NFV architecture and FlowNAC proposal. Reference [8] demonstrates the applicability of the ETSI architecture to deploy VNFs in a resilient and scalable manner. The SIP-PBX service proposed is similar to FlowNAC in the sense that the control and data planes are clearly separated. Reference [9] introduces an ETSI-based architecture that ensures high availability and scalability for a virtual session border controller and proposes its combination with SDN. Similar to FlowNAC, there are also other efforts that aim at demonstrating the integration of NFV and SDN. Reference [10] presents the implementation of a routing function virtualization based on NFV concepts that leverages on OpenFlow, although from the architectural standpoint this work is not related to the architecture proposed by ETSI. Other SDN technologies, like ForCES, have also positioned themselves to demonstrate their applicability to NFV [11] and how the networking resources could be exposed to the VNFs.

From the related work, we can conclude that currently there is a clear momentum for exploiting networking innovation in the light of SDN and NFV. Moreover, open issues arise as a step forward is achieved; thus, work should continue to bring us closer to a dynamic and flexible networking infrastructure.

Toward an SDN-Enabled Network Function Virtualization Architecture

The scenario to enable the dynamic deployment of VNFs is challenging from the networking point of view. It must support multi-tenancy, multiple service chains sharing the same physical resources, and traffic steering between the VNFs to develop the service chain. In this context, the traffic must be isolated not only among service chains but between the NFs that compose the service as well. SDN is a perfect complement to deal with these requirements and with the dynamicity imposed to the network resources. Although the interaction between NFV and SDN is complementary, there is space for innovation in this area, and their relationship can evolve beyond providing a network infrastructure with enhanced capabilities to significantly improve how the VNFs are designed. Next, we present a vision on the evolution of the NFV architecture from an SDN-agnostic approach to a fully SDN-enabled architecture, also represented in Fig. 1.

SDN-Agnostic NFV Architecture

Before the appearance of NFV, NFs were built as a closed combination of software and hardware from vendors. NFV is a step forward for the provisioning of network functions and enables the decoupling of software from hardware. This decomposition relies on the virtualization layer, which exposes virtual resources (i.e., computing, storage, and network) that become the building blocks for NFs. Moreover, the deployment of NFs becomes more flexible as they are based on software and not attached to specific hardware. As the NF turns into instantiable software (VNF), it provides more flexibility to scale up/down with finer granularity according to the actual traffic and NF performance.

The NFV architecture promoted by ETSI leverages on compute, storage, and network virtual resources, and NFs are virtualized and encapsulated as a software package, like a virtu-

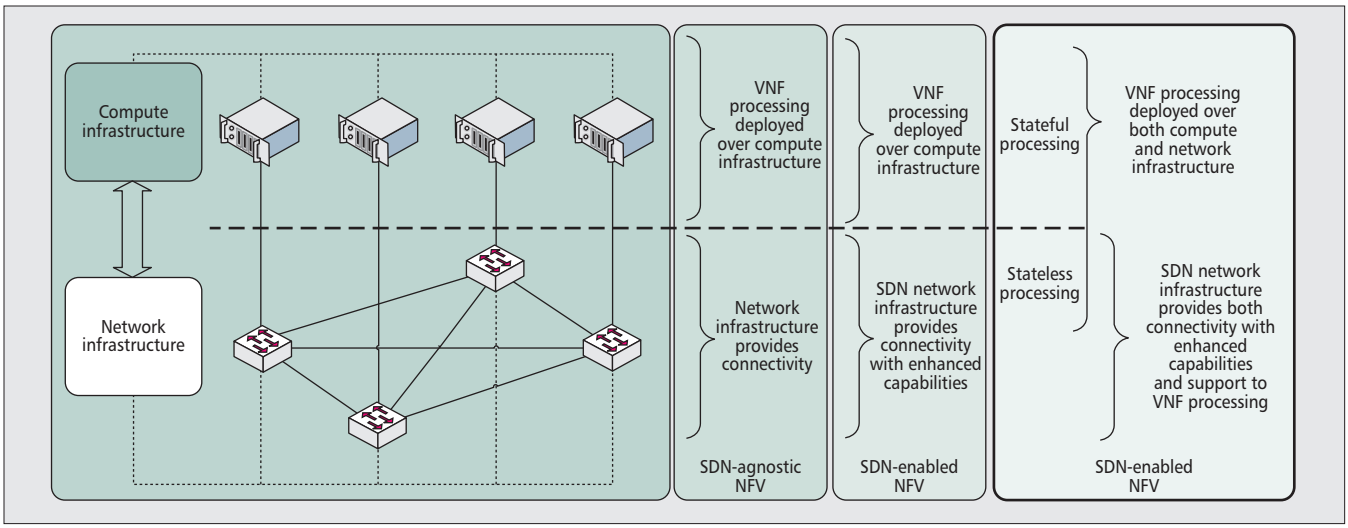

Figure 1. NFV architecture evolution.

al machine (VM), that relies on these components. The underlying network infrastructure, included in the NFV infrastructure (NFVI), is abstracted to realize virtualized network paths that provide connectivity to support the interconnection between VNFs and with the endpoints [2]. Therefore, the VNFs are software boxes running on commodity servers to process the frames coming from the underlying network, and consequently, compute resources are the main architectural component to build the NFs. In this regard, virtual network resources are limited to providing an interface to the underlying network resources, which are mostly considered just for providing connectivity service.

Some examples proposed by ETSI that can benefit from this approach are vBNG, vCPE, vRouter, CG-NAT, software-based DPI, and mobile network nodes [12].

SDN-Aware NFV Architecture

Since the initial proposal of the NFV concept [12], its relationship with SDN was argued to be complementary and potentially of added value when both technologies are combined. The separation of data forwarding from the control plane improves the flexibility of the network and simplifies the dynamic deployment and operation of resources. In addition, the usage of commodity servers and switches, avoiding specific hardware-based components provided by vendors, is a shared objective between NFV and SDN. Moreover, some of the networking challenges of the NFV architecture to be addressed match the design goals of SDN, such as dynamic control and configuration of network nodes and automated management of the network. Others, like elastic and fine-grained scalability adapted to the actual needs, seamless mobility of resources, and efficient multi-tenancy support, can be built on SDN capabilities.

Thus, even if the first ETSI architecture [2] did not explicitly mention SDN as part of the NFVI, the ONF [4] quickly published the envisioned scenario of cooperation between SDN and NFV, which contributed to simplify the integration of both physical and virtual networking infrastructures by using a common interface. As part of this contribution, the ONF depicts a possible interaction between the NFV Orchestration component and the OpenFlow Controller, which is based on a northbound interface exposed by the latter.

Further pursuing this approach, the latest documents from ETSI [13] integrate SDN with the defined architecture and reference points for NFVI. However, the contributions of SDN remain in the infrastructure network domain of the NFVI, focused on providing connectivity services; and despite being a perfect complement for NFV, it does not tackle the compute-based design of VNFs. Since the main difference from the previous approach lies only in the network infrastructure, the same examples apply here. The main idea is that the compute resources (e.g., the CG-NAT) must process all the data traffic. The network infrastructure still only provides connectivity services, albeit more dynamic and programmable.

SDN-Enabled NFV Architecture

This last step in the evolution of NFV toward a fully SDN-enabled architecture means a valuable advance in the way the VNFs are designed and implemented. As explained in the previous section, the synergies between both technologies advocate for deploying NFV over an SDN network infrastructure. Based on this programmability already in place, the proposal is to explore the possibility of exploiting the network infrastructure layer to implement part of the VNF functionality. At this point, it must be highlighted that current SDN datapath implementations are mostly stateless, since there is no (or limited) state associated with the flow entries. In general, previously matched frames do not affect subsequent frames, meaning that no state is associated (e.g., a stateless firewall or load balancer). However, some limited or lightweight state can be kept in the data path, such as flow-level counters, timers for flow expiration, and queue-level counters for QoS support. As a consequence, the networking devices, which are supposed to be optimized for data plane pro-

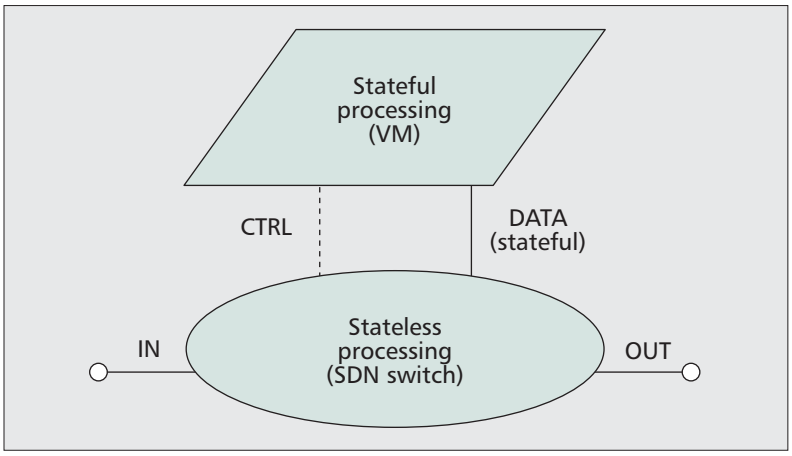

Figure 2. SDN-based component decomposition of the VNF.

cessing, can be used to perform the stateless processing of data traffic.

The aforementioned proposal has an impact on the way VNFs are designed, and also on the NFV architecture, enabling a new framework for further innovation. There are several factors that justify this approach, including data traffic performance improvements; a wider placement scope (e.g., a VNF functionality to be performed in the first element in the datapath); better resource utilization, to avoid traffic going up to and down from a VM to process the frames; and so on. The network devices are the basic substrate for building the service chains, which means that they must be crossed anyway, and their availability as a resource is more widespread. On the contrary, the commodity servers can be limited to specific locations, such as data center infrastructures, thus requiring more complicated traffic steering and possibly imposing some delay due to the longer path the traffic must traverse.

VNF Design Considerations

The proposed evolution for the NFV architecture and the VNF design approach distinguishes between two components: the stateful network function component and the stateless data path processing component, as shown in Fig. 2. The stateful component is based on virtual compute resources (e.g., VMs) to keep the state associated with the VNF, since software packages are good at managing state machines. The stateless component makes use of SDN datapath resources to perform data traffic processing efficiently. The separation between these two components is inspired by SDN decoupling of the data and control planes. The main idea is to keep data processing in hardware as much as possible, and only forward the data traffic to the stateful component when the processing is also stateful.

An example to illustrate this separation is the Evolved Packet Core (EPC), where the packet gateway (PGW) component can be implemented following this design approach. In this case, the GTP-U tunneling endpoint performing the (de)encapsulation (stateless) and gathering associated statistics (lightweight state) can be implemented in the SDN datapath, whereas the creation of the GTP-U tunnels and statistics collection (and further actions) can be implemented on the compute resources (stateful). Reference [11] proposes a ForCES-based implementation of the PGW.

The interfaces between the stateful and stateless components must be clearly defined to describe the new architecture. There are two different planes that can be used to interact between the mentioned components: the control plane and data plane. The control plane interface is used to configure and update the behavior of the stateless datapath processing component as a result of logic executed in the stateful component. The data plane interface is used when some portion of the traffic needs stateful processing and must be redirected to the stateful network function component.

Considering this new approach, the proposed NFV architecture opens up the supported design alternatives for VNFs: solely with compute resources (option 1); solely with network resources (option 2); or splitting the design following the aforementioned separation of stateful and stateless components (option 3). Each of the options is best suited for a different scenario, and the selection criteria would include the requirements of the service related to data processing (low vs. high speed) and the complexity of the state to be kept (stateless, lightweight state, or stateful), as represented in Table 1 and Fig. 3.

The main benefits from the decoupled VNF design are the efficient data processing achieved by optimized hardware, optimization of resources by means of avoiding data traffic going up/down to/from a VM, and independent scalability of each component. This independence is fundamental when the stateless and stateful processing is unbalanced. In this case, the most demanding processing can increase its assigned resources independently and optimize the resource utilization.

FlowNAC: A Real Experience Exploring Novel NFV Architectures

This section illustrates the applicability of NFV architectures described above with a real example, FlowNAC, a flow-based network access control (NAC) [1]. This solution has already been deployed over the OpenFlow-based EHU-OEF infrastructure [14] and demonstrated at several conferences like IEEE GLOBECOM 2013. The following analysis performed over the FlowNAC example focuses on one hand on SDN-agnostic and SDN-aware NFV architectures, as both share the same compute-based approach for VNF design, and on the other hand SDN-enabled NFV architecture, which are compared in Fig. 4.

FlowNAC is a NAC system for services that can be identified as flows, inspired by the IEEE 802.1X port-based NAC, which is a basic NAC solution enforcing the access control at the port level (i.e., the physical port of the network node). The main difference between both approaches is the granularity of each solution: port vs. flow.

The traffic originated in the user can be classified in three categories:

a-type: Authentication and authorization (AA) traffic that must be processed by the FlowNAC VNF to keep the state associated with each AA process. Being related to the AA process, it will be limited in time and volume.

b-type: Data traffic for the authorized services that must be granted access to the network. It will depend on the actual service provided, which can be bandwidth-intensive (e.g., multimedia services) or time-sensitive. This traffic must be evaluated by the FlowNAC VNF to enforce the access control functionality.

c-type: Data traffic for non-authorized services (the remaining traffic generated by the user), which must be denied access to the network. This traffic completely depends on the user and its operating system, running applications, viruses, and so on. Similar to b-type traffic, the FlowNAC VNF must enforce the access control, but in this case to prevent these frames accessing the resources.

Following the compute-based design of the first two approaches to the NFV architecture (SDN-agnostic and SDN-aware NFV architectures), FlowNAC can be implemented as a software package running on a single VM. It could also be decomposed into different lower-level functions running on separated VMs, but for simplicity we assume just one element.

In this approach, all the traffic from users must be redirected to the VM by the network infrastructure. The AA traffic is processed, authorized data traffic is allowed back into the network to reach the authorized service, and non-authorized data traffic is dropped, as shown on the left side of Fig. 4.

This approach has several benefits like its ease of implementation and deployment, as the FlowNAC VNF relies solely on computing resources and is deployed by simply instantiating one VM and redirecting all the traffic from the user to it. Migration of the FlowNAC VNF is also straightforward as it is based on legacy computing virtualization technology with well-known interfaces, and the functions are completely isolated from the underlying infrastructures.

On the other hand, this architecture also has several drawbacks. As with any other VNF, the scalability is improved by virtualization techniques, which allow more or fewer computation resources to adapt the processing capacity to the actual demand. However, the scalability would be limited by the availability of computing resources.

The performance of data processing is one of the main limitations of this approach. Although there are some benchmarks that claim high throughput when traffic crosses the VM, the underlying hardware devices are expected to achieve better performance [15]. As a consequence, the data traffic processing in virtual compute resources is worse than in hardware-based networking devices, which impacts authorized service data traffic that must be processed by the VM.

Data performance requirements	Complexity of State	Best alternative
High-speed processing	Stateful	Option 3: VNF with split design
High-speed processing	Stateless/ lightweight state	Option 2: VNF designed with network resources
Medium-/low-speed processing	Stateful	Option 1: VNF designed with compute resources
Medium-/low-speed processing	Stateless/ lightweight state	Option 1 or 2

Table 1. VNF design alternative selection criteria.

The overall resource utilization also becomes an issue. All the traffic from the user must be redirected to a location with computing resources. Then all the traffic must be processed by the VM, and finally, the authorized data traffic must be redirected to the following step toward the service, as shown in Fig. 4. This approach has a negative impact on the bandwidth occupation as all the aggregate traffic entering the network must effectively reach a VM.

Following the proposed approach (SDN-enabled NFV architecture), a FlowNAC VNF is designed separated into two functional blocks: the AA block, which keeps the state of the AA processes currently executed; and the access control enforcing block, which limits the access only to already authorized services and does not require any state. Regarding its implementation, the AA process relies on computing resources, and the access control is enforced by the networking devices. This separation permits redirecting each type of traffic to the appropriate resource, as shown on the right side of Fig. 4. Only the AA traffic is sent to the VM to be processed, and depending on the result of the AA process, the SDN switch is configured to allow the authorized services. In our implementation, this interface is realized using OpenFlow, and the FlowNAC VNF includes an SDN controller that sends the appropriate instructions to the SDN switch to allow the flows associated with the authorized service. Authorized and unauthorized data traffic is processed by the SDN switch, and only the authorized traffic is allowed, while the unauthorized data traffic is dropped.

A FlowNAC VNF benefits significantly from this approach. First, the scalability requirements of the AA and access control enforcing blocks are very different. The computing resources used for the AA block could scale slower as AA traffic is expected to be less demanding, and the offloading of processing would make them lighter and easier to move and scale. The networking resources used for the access control enforcing block would require more data processing capacity as the traffic associated with the services increases. Moreover, it must also be considered that unauthorized data traffic, which must be dropped, could also demand increasing the data processing capacity.

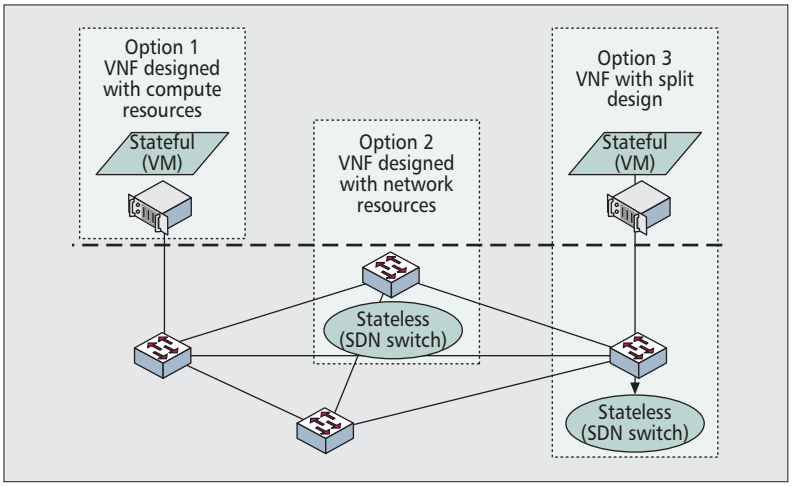

Figure 3. VNF design alternatives in SDN-enabled NFV architecture.

Second, the data processing performance of the VNF is improved, as specialized hardware processes authorized and unauthorized data traffic, which is expected to be the most intensive and bandwidth consuming. The AA traffic, which is less intensive and requires stateful processing, is redirected to the computing resources.

Finally, the availability of networking resources, as part of the underlying connectivity substrate, provides the possibility to ingress only the authorized traffic data and drop the unauthorized data at the first node, reducing general network utilization. Moreover, it also opens the way to better utilization of resources, as the network must be traversed anyway, and the network node does not process the authorized data traffic twice, to and from the VM.

Issues and Challenges

The proposed approach to an SDN-based NFV architecture, discussed previously, brings up a series of challenges stemming from the two major changes depicted in Fig. 1: on one hand, the VNF must be designed splitting the components to be deployed over compute and network resources; on the other hand, the network infrastructure must support a dual role for traffic steering and VNF processing.

In order to achieve the performance gain the split design allows, the services must be carefully redesigned. Thus, each service must be analyzed to determine if the performance gain overcomes the effort involved in the redesign process. Also, the functionalities are supported by the SDN data path condition in which part of the VNF processing can be offloaded to the network elements. As services are designed and deployed over NFV architectures, emerging best practices in design of VNFs should provide guidance in these issues.

VNFs designed also employing network elements add complexity to the optimal placement decision. The NFV framework must now orchestrate an additional type of resource with its own constraints (e.g., the rule insertion performance or flow table size). Research on orchestrators is a hot topic in the NFV field, so significant improvement is expected in this area, which should also allow limiting the impact of the increased complexity.

The effective use of compute and network virtual resources requires the virtualization layer to provide the necessary interfaces and allow equivalent flexibility and dynamicity in deployment and migration over both of them, as well as enforcing the required isolation among the virtual resources of different tenants sharing the same physical resources. Progress in this area is also expected in order to provide a full and high-performance carrier grade virtualization of service delivery.

Dual use of the infrastructure, for traffic steering and stateless NF processing, implies that the underlying network infrastructure must guarantee isolation between both functionalities, and also between the processing of different NFs, not only in the data plane but also in the control plane, as the VNFs influence the behavior of the SDN data path. Performance isolation is also required to avoid VNF functionality hindering the correct behavior of the NFV architecture. In any case, in the future any expectations will undoubtedly be met.

Conclusions

The NFV concept is probably one of the major recent revolutions on the information technology landscape. From a chronological point of view it was more or less coetaneous with SDN, but at that time, NFV just did not take into account this last technology or possible relationship.

As SDN technology matures and NFV becomes a real technological trend, the convergence of both technologies was something to be expected. And this is happening: both ONF and ETSI NFV ISG are already proposing a SDN-

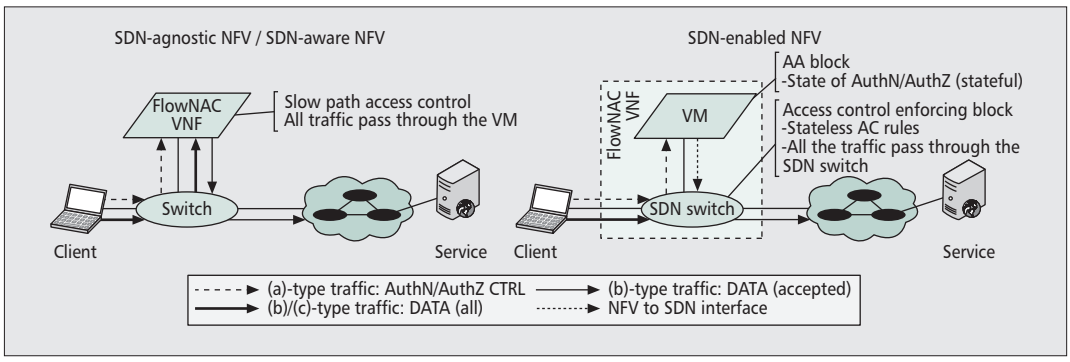

Figure 4. FlowNAC VNF alternatives.

aware NFV, which offers the network the dynamicity to support new network-aware service provisioning paradigms.

In this article we propose a taxonomy of the evolution of the NFV/SDN relationship. In the latest approach, the processing of network packets is partially offloaded to the network element (the SDN switch) while maintaining the stateful processing of the VNF on the compute element. This effectively means that the VNF logically extends to the networking element.

To demonstrate the validity of the concept, a real implemented use case that shows the suitability is presented: an access control virtualized network function (FlowNAC VNF) using FlowNAC, our own flow-based network access control.

Several challenges still need to be addressed to fully attain the benefits of this approach. But overall, we consider that there are many reasons to believe that SDN-enabled NFV will boost NFV deployment to support new efficient and cost-effective services.

Several challenges still need to be addressed to fully attain the benefits of this approach. But overall, we consider that there are many reasons to believe that an SDN-enabled NFV will boost the NFV deployment to support new efficient and cost effective services.

Acknowledgments

This research was partly funded by the Spanish Ministry of Economy and Competitiveness under the "Secure deployment of services over SDN and NFV based networks" project S&N-SEC TEC2013-47960-C4-3-P and by the European Commission under the FP7 UNIFY (Unifying Cloud and Carrier Networks) project CNECT-ICT-619609. This has been produced within the Training and Research Unit UFI11/16 supported by the UPV/EHU.

References

[1] J. Matias *et al.*, "FlowNAC: Flow-Based Network Access Control," *Proc. Euro. Wksp. Software Defined Networks*, Budapest, Hungary, 2014.
[2] ETSI ISG for NFV, "ETSI GS NFV 002: Network Functions Virtualisation (NFV); Architectural Framework," http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf; accessed 15 January 2015.
[3] ETSI ISG for NFV, "NFV Proofs of Concept," http://www.etsi.org/technologies-clusters/technologies/nfv/nfv-poc; accessed 15 Jan. 2015.
[4] ONF, "OpenFlow-Enabled SDN and Network Functions Virtualization," https://www.opennetworking.org/images/stories/downloads/sdn-resources/solution-briefs/sb-sdn-nvf-solution.pdf; accessed 15 Jan. 2015.
[5] E. Haleplidis *et al.*, "Software-Defined Networking (SDN): Layers and Architecture Terminology," RFC 7426, Jan. 2015; http://tools.ietf.org/search/rfc7426; accessed 01 Feb. 2015.
[6] G. Xilouris et al., "T-NOVA: A Marketplace for Viortualized Network Functions," *Proc. Euro. Conf. Networks and Commun.*, Bologna, Italy, 2014.
[7] P. Skoldstrom *et al.*, "Towards Unified Programmability of Cloud and Carrier Networks," *Proc. Euro. Wksp. Software Defined Networks*, Budapest, Hungary, 2014.
[8] M. Schöller *et al.*, "Resilient Deployment of Virtual Network Functions," *Proc. Int'l. Congress on Ultra Modern Telecommunications and Control Systems and Wksps. '13*, Almaty, Kazajistan, 10–13 Sept.
[9] G. Monteleone and P. Paglierani, "Session Border Controller Virtualization towards Service-Defined Networks Based on NFV and SDN," *Proc. Int'l. Conf. IEEE SDN for Future Networks and Services '13*, Trento, Italy, 11–13 Nov.
[10] J. Batall *et al.*, "On the Implementation of NFV over an OpenFlow infrastructure: Routing Function Virtualization," *Proc. Int'l Conf. IEEE SDN for Future Networks and Services '13*, Trento, Italy, 11–13 Nov.
[11] E. Haleplidis *et al.*, "ForCES Applicability to SDN-enhanced NFV," *Proc. Euro. Wksp. Software Defined Networks*, Budapest, Hungary, 2014.
[12] ETSI ISG for NFV, NFV White paper: "Network Functions Virtualisation, An Introduction, Benefits, Enablers, Challenges & Call for Action. Issue 1," http://portal.etsi.org/NFV/NFV_White_Paper.pdf; accessed 15 Jan. 2015.
[13] ETSI ISG for NFV, ETSI GS NFV-INF 001, "Network Functions Virtualisation (NFV); Infrastructure Overview;" http://www.etsi.org/deliver/etsi_gs/NFV-INF/001_099/001/01.01.01_60/gs_nfv-inf001v010101p.pdf; accessed 15 Jan. 2015.
[14] J. Matias *et al.*, "The EHU-OEF: An OpenFlow-Based Layer-2 Experimental Facility," *Computer Networks*, vol. 63, Special Issue on Future Internet Testbeds, 2013, pp. 101–27.
[15] C. Rotsos *et al.*, "OFLOPS: An Open Framework for OpenFlow Switch Evaluation," *Proc. Passive and Active Measurement*, Vienna, Austria, 2012.

Biographies

Jon Matias received his B.S. and M.S. degrees in telecommunication engineering from the University of the Basque Country (UPV/EHU) in 2003. He currently works as a researcher at the Communications Engineering Department of the same university and is pursuing a Ph.D. degree focused on access networks and security. His research interests include software defined networking, network functions virtualization, broadband access networks, and security.

Jokin Garay received his B.S. and M.S. degrees in telecommunication engineering in 2003 from UPV/EHU. After a period in the private sector, he came back to the university to pursue a Ph.D.. His research interests include software defined networking, network functions virtualisation, and cloud computing.

Nerea Toledo received her B.Sc. and M.Sc. degrees in telecommunications engineering in 2007 from UPV/EHU and her Ph.D. degree from the same university in 2012. Since 2008 she has been an assistant professor at UPV/EHU and as a researcher in the I2T lab (http://i2t.ehu.es). She has been a visiting researcher at Institut Telecom-Telecom Bretagne. Her current research interests include SDN, wireless networking, and security.

Juanjo Unzilla holds B.S., M.S. (1990), and Ph.D. (1999) degrees in communications engineering, and is a professor in the Communications Engineering Department at UPV/EHU where he teaches subjects related to telecommunications networks and services. He is part of the I2T Research Group, where he participates in several national and European R&D projects. His research interests include SDN and NFV, network security, and techno-economic models for access networks. Other interests are models and metrics for knowledge transfer from universities to enterprises.

Eduardo Jacob [SM], after spending a few years in the private sector, first as a network manager and an R&D project leader later, came back to UPV/EHU. He is a professor and leads a research group at his university that is participating in several national and European R&D projects. Other interests are industrial applications of SDN and NFV for resiliency, experimental network infrastructures, and cyber physical systems.

INTEGRATED CIRCUITS FOR COMMUNICATIONS

Charles Chien *Zhiwei Xu*

In this issue of Topics in Integrated Circuits for Communications, we have selected three papers that mark recent progress in the communications semiconductor industry which is enabling emerging short-range millimeter-wave (mmWave) communications and power-efficient transmissions.

With the proliferation of wireless communications in recent years, frequency spectrum has become extremely congested, especially for frequencies below 6 GHz. The spectral congestion is aggravated by a steep rise in demand for higher data rates in excess of 1 Gb/s. To support the rise in the number of users as well as data rates, wider available bandwidth or more complex signal processing is required to increase spectral efficiency. For mid-range to long-range systems, such as IEEE 802.11n/ac and LTE-A, the communications industry is leveraging high-order modulation and multiple-antenna technology to implement diversity, and multiple-input multiple-output (MIMO) to boost channel capacity while supporting high-bit-rate services. In addition, carrier aggregation is employed to group channels in existing allocated bands to provide effectively wider transmission spectrum.

In contrast, millimeter-wave and terahertz frequencies offer alternative means to achieve high data rate because of the availability of abundant spectrum, as well as high propagation loss and short wavelengths. Given large available bandwidths (e.g., 7 GHz), less spectrally efficient modulations could be used to achieve data rates > 10 Gb/s at short distances. The high propagation loss facilitates frequency reuse at short distances, while short wavelengths lead to small antenna aperture, making it possible to integrate large numbers of antennas in a phased array radio. The ability to achieve short-range focused radio beams enables high-resolution medical imaging and radar ranging. Thus, mmWave devices are poised to usher in portable devices for medical imaging and onboard radar systems for self-driving cars of the future.

Advanced silicon technologies have played a key role in materializing integrated mmWave radios due to their high density for digital integration and high effective cutoff frequency. In particular, SiGe BiCMOS technologies deliver better power handling capability than silicon CMOS, and may bring unique advantages in implementing integrated mmWave and terahertz transceivers. However, the manufacturability and cost of large phased arrays becomes challenging due to low yield. In the article "W-Band Scalable Phased Arrays for Imaging and Communications," the authors propose a modular approach to address the conflicting requirements between high directivity and manufacturing cost for a phased array. After a brief overview, the authors discuss both circuit and architecture trade-offs leading to a 64-element phased array that contains four 16 dual-polarized RF phase-shifting front-ends together with 64 dual-polarized antennas housed in a single package, occupying 16.2 mm × 16.2 mm × 0.7 mm. Their work has demonstrated the feasibility of achieving large arrays by tiling smaller array modules using low-cost silicon-based technology.

Over the past two decades, we have come to embrace the notion of the "last mile," portraying the reach of communication infrastructure to every home. Now with the realization of mmWave communications, we can talk about perhaps an even more challenging problem — the "last centimeter," the ability to connect large numbers of high-speed integrated circuits or backplanes. With the abundant bandwidths available in the terahertz range (i.e., 50–300 GHz), terahertz interconnect can potentially meet the bandwidth density (Gb/s/mm^2) and energy efficiency (Joules per bit) required for future high-speed computing devices and network backplanes. In the article "THz Interconnect: the Last Centimeter Communication," the author gives a detailed overview of the trade-offs that exist for designing THz circuits to meet the bandwidth density and energy efficiency required for high-speed inter-chip communications. To overcome the bandwidth limitations in process technology, the article describes an interesting circuit technique to introduce negative resistance into the signal path such that resistive losses can be reduced, effectively lifting the bandwidth limitation imposed by process technology. Moreover, to achieve data rates in excess

of 100 Gb/s, the author proposes a micro-machined dielectric waveguide in a silicon process to address the significant line loss exhibited at such high frequency. This article highlights the potential of realizing mmWave circuits to address the increasing demand for IO bandwidth using low-cost silicon-based technologies.

The last article in this issue of our Series addresses the age-old problem of low power efficiency in wireless transmissions. In a mobile device, more than half of the power is being dissipated by the power amplifier, which makes it the dominant factor in determining battery lifetime. Likewise, in a base station, the power amplifier is the dominant source of heat leading to high air-conditioning cost. Making the power amplifier more efficient will greatly enhance the user experience with longer battery lifetime, as well as reduce the maintenance cost of base stations by reducing heat generation. Several articles from past issues of this Series have described digital enhancement techniques to improve the efficiency of power amplifiers based on RF digital-to-analog (DAC) architecture. In the last article, "Outphasing Transmitters, Enabling Digital-Like Amplifier Operation with High-Efficiency and Spectral Purity," the authors give an in-depth overview of another widely known PA architecture based on outphasing. They have shown its potential to provide greater than 80 percent power efficiency in contrast to the 40–50 percent typically achieved in conventional PAs. Moreover, a 70 W power amplifier has been demonstrated using a CMOS-driven GaN outphasing amplifier. Their work shows that outphasing is a strong contender in solving this age-old problem in transmission systems, although reaching toward 100 percent efficiency remains a challenging goal for future circuit designers.

We would like to take this opportunity to thank all the authors and reviewers for their contributions to this Series. Future issues will continue to cover circuit technologies that are enabling new emerging communication systems. If readers are interested in submitting a paper to this Series, please send your paper title and an abstract to either of the Series Editors for consideration.

BIOGRAPHIES

CHARLES CHIEN (charles.chien@creonexsystems.com) is the president and CTO of CreoNex Systems, which focuses on technology development for next generation communication systems. Previously he held various key roles at Conexant Systems, SST Communications, and Rockwell. In his career, he has architected several key products including a CMOS/SiGe chip-set for multimedia over coax (MoCA), an IEEE 802.11abg WLAN RF CMOS transceiver and GaAs PA/RF switches, a wireless audio CMOS chip-set for home theatre in a box, digitally-assisted cellular transceivers, and low-power wireless networked sensors. He was also an assistant adjunct professor at the University of California, Los Angeles (UCLA) from 1998 to 2009. His interests focus mainly on the design of system-on-chip solutions for wireless multimedia and networking applications. He has published in various journals and conferences, and has authored a book. Digital Radio Systems on a Chip. He received his B.S.E.E. from UC Berkeley, and his M.S. and Ph.D. from UCLA. He served as a member of the Technical Program Committee of ISSCC from 1998 to 2006.

ZHIWEI XU (xuzhw@yahoo.com) received B.S. and M.S. degrees from Fudan University, Shanghai, China, and his Ph.D. from UCLA, all in electrical engineering. He held industry positions with G-Plus Inc., SST Communications, Conexant Systems, and NXP Inc., where he did development for wireless LAN and SoC solutions for proprietary wireless multimedia systems, CMOS cellular transceivers, MoCA systems, and TV tuners. He is currently with HRL Laboratories, working on software defined radios, high-speed ADC, and analog VLSI. He has published in various journals and conferences, made one contribution to the *Encyclopedia of Wireless and Mobile Communications*, and has five granted patents.

W-Band Scalable Phased Arrays for Imaging and Communications

Xiaoxiong Gu, Alberto Valdes-Garcia, Arun Natarajan, Bodhisatwa Sadhu, Duixian Liu, and Scott K. Reynolds

ABSTRACT

This article discusses the benefits and challenges associated with the design of multi-function scalable phased arrays at millimeter wave frequencies. First, applications for phased arrays with tens to hundreds of elements are discussed. Existing solutions for scaling silicon-based phased arrays from microwave to terahertz frequencies are reviewed. The challenges and trade-offs associated with multiple integration options for W-band phased arrays are analyzed, with special consideration given to packaging and antenna performance. Finally, a solution based on SiGe ICs and organic packages for a 64-element dual-polarized 94 GHz phased array is described, along with associated measurement results.

INTRODUCTION

W-band frequencies, which range from 75–110 GHz, have been garnering significant attention recently, specifically in the areas of automotive radar, backhaul communications, millimeter-wave (mmWave) radar, and imaging. In the area of communications and imaging, the availability of a large bandwidth at these frequencies, as well as the presence of a low-absorption atmospheric window, makes W-band frequencies particularly attractive. A number of imaging and point-to-point wireless link applications require highly directional transceivers, the ability to rapidly scan in two dimensions, and support for dual-polarized operation to meet performance needs. For such applications, phased-array-antenna-based solutions, with their beamforming and electronic beam-steering capabilities, higher range, and higher signal-to-noise ratio (SNR), are particularly suitable. Moreover, many of these applications require significant beam control capability and functional versatility, while demanding relatively little radiated power per element. These requirements make silicon-based implementations, which are inherently friendly to multi-function integration, high degrees of digital programmability, and built-in calibration, particularly attractive [1].

While advanced complementary metal oxide semiconductor (CMOS) nodes provide higher density for digital integration, the effective cut-off frequency f_{MAX} (including wiring to device terminals) and output power delivery capabilities of CMOS devices are limited. In contrast, SiGe BiCMOS technologies provide higher effective f_{MAX} and power handling capability, as well better potential for further f_{MAX} increase while still providing CMOS technology for digital functions [2]. However, for either advanced CMOS or SiGe BiCMOS technology, the cost benefits of silicon integrated circuit (IC) manufacture only become compelling when production is scaled to more than tens or hundreds of thousands of units. This fact presents particular challenges for low-volume phased array applications that require large numbers of elements for high directionality, translating to large, expensive ICs if all elements are integrated on a single IC. In addition, the maximum IC size has manufacturability constrains. A modular approach consisting of scalable phased arrays, in which an arbitrary number of repeated unit cells operate in unison, is therefore an attractive alternative from many perspectives, including cost, yield, robustness, and ease of testing.

In this article, we discuss the advantages and challenges of designing such scalable, multi-functional, W-band phased arrays in silicon. The next section discusses the impact of scalability on W-band communications. A review of existing scalable phased arrays is covered after that. Trade-offs associated with various scalability options are then discussed; measurement results from a prototype scalable 64-element phased array at 94GHz are presented in the following section, and some concluding remarks are presented in the final section.

THE IMPACT OF PHASED ARRAY SCALABILITY ON W-BAND COMMUNICATIONS

Millimeter-wave links are now an integral part of the wireless backhaul infrastructure, particularly at E-band frequencies (71–76 GHz, 81–86 GHz, and 92–94 GHz) [3]. Note that these frequencies are a subset of the W-band frequency range. E-band links are currently implemented by a com-

Xiaoxiong Gu, Alberto Valdes-Garcia, Bodhisatwa Sadhu, Duixian Liu, and Scott K. Reynolds are with IBM Thomas J. Watson Research Center.

Arun Natarajan was with IBM Thomas J. Watson Research Center. He is now with Oregon State University.

0163-6804/15/$25.00 © 2015 IEEE

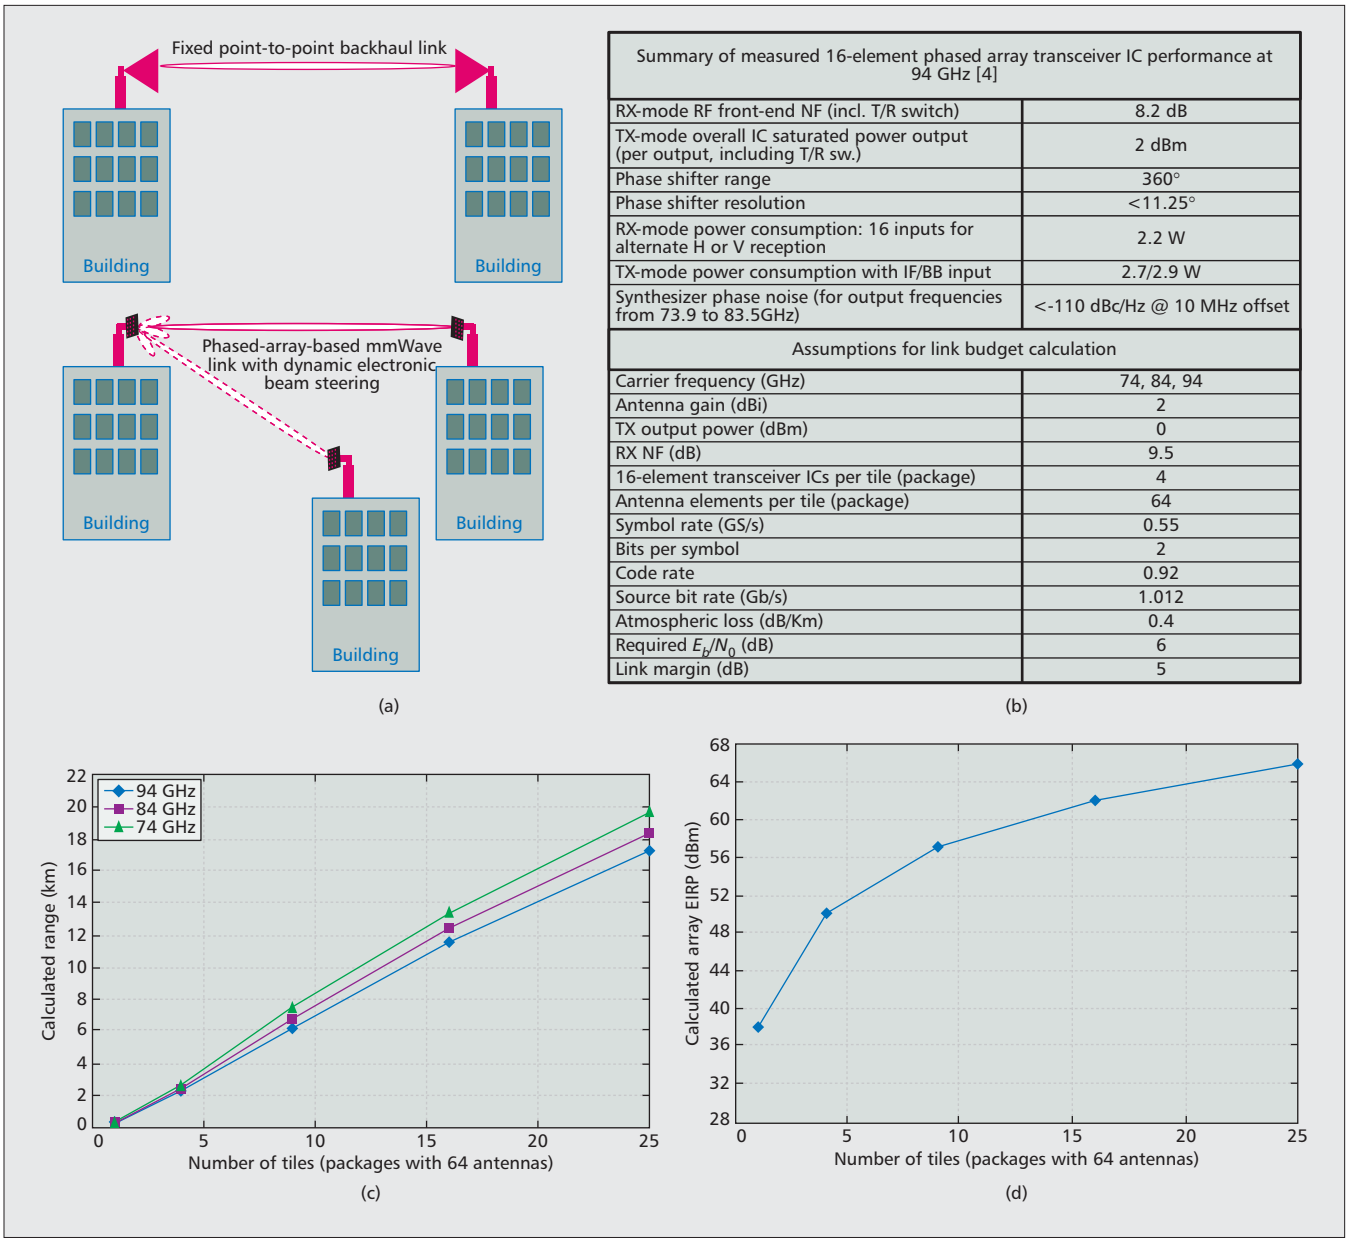

Figure 1. a) Illustration of two different types of mmWave backhaul link, with fixed high-gain antennas (top) and with a phased array (bottom); b) summary table of measured transceiver IC performance [4] and assumptions for link budget calculation; c) estimated range for a 1 Gb/s terrestrial data link at 74, 84, and 94 GHz, constructed using scalable phased arrays of various sizes. The estimates assume QPSK (2b/symbol) in 800 MHz RF bandwidth, LDPC code (1369, 1260) with code rate $R = 0.92$, and BER = 10^{-7} with 5dB implementation loss and 0.4 dB/km atmospheric loss; d) calculated EIRP as a function of the number of tiles.

bination of single-element transceivers and antennas with high gain (i.e., > 30 dB) and consequently occupy large form factors. Mechanical alignment is required for these antennas; as such, links are established and maintained in a single fixed direction, as shown in the top example in Fig. 1a.

Monolithic phased arrays intended for indoor applications at 60 GHz have already demonstrated links with data rates in excess of 5 Gb/s at distances of ~10 m employing only 16 antennas [1]. These and other similar results have motivated research on highly integrated phased arrays at higher frequencies and with a larger number of elements.

Silicon-based scalable phased arrays at W-band offer the possibility of attaining similar coverage range as current fixed-beam solutions with the additional advantage of dynamic steerability, at the cost of complexity and power consumption. Electronic steerability would not only eliminate the need for mechanical alignment, but would also open the possibility of dynamic backhaul networking, as shown in the bottom example in Fig. 1a. The scalability of a unit cell array with a moderate number of elements is key, since different links and usage scenarios may require a different number of elements.

To illustrate the potential of a scalable phased array at W-band, the table in Fig. 1b introduces the link budget considerations for a 1 Gb/s link formed with tiles of 64 antenna elements sup-

It can be observed that a 10+ km range is potentially achievable with a 4 × 4 array of 64-element tiles (1024 elements), and such an array would occupy an area smaller than 70 mm^2.

ported by four 16-element phased array ICs. This link budget calculation considers the measured phased array IC performance reported in [4], which is also shown in Fig. 1b. It should be noted that the NF and output power performance used in the link budget calculation is slightly worse (1.5–2 dB) than that reported for an IC at room temperature to account for performance degradation at higher temperatures. Figure 1c shows the potential link range as a function of the number of tiles for three different E-band frequencies. It can be observed that a 10+ km range is potentially achievable with a 4 × 4 array of 64-element tiles (1024 elements), and such an array would occupy an area smaller than 70 mm^2. The calculated array equivalent isotropic radiation power (EIRP) as a function of the number of tiles is plotted in Fig. 1d.

Review of Scalable Phased Arrays

Implementing a single multifunctional or reconfigurable RX and/or TX element in an IC, which is similar to approaches using discrete components, is one possible approach that can be taken to enable realization of scalable arrays [5, 6]. Such an approach provides high levels of flexibility by maximizing the granularities at which arrays can be created, but fails to take advantage of a key aspect of silicon technologies: the availability of multiple well-controlled interconnect layers. The single-unit approach places the significant complexity and cost burden associated with routing large numbers of RF, IF, and/or baseband signals on the PCB and packaging instead of fully leveraging the wiring capabilities of silicon to enable complex array realization. The single-unit approach also leads to high power consumption since impedance-matched drivers are required on every IC. At low RF frequencies (< 10 GHz), large antenna spacing is offset by low packaging losses, making single-element unit cells feasible. At these frequencies, the interface to the antenna can also be considered independent of the IC unit cell due to the relatively flexible packaging requirement. However, at mmWave frequencies (> 30 GHz), physically short antenna spacings ($\sim\lambda/2 < 5$ mm), packaging losses, and manufacturing challenges with impedance-controlled multi-layer packaging interconnects make multi-element unit ICs more attractive. Trade-offs with respect to system packaging and antenna integration are discussed in detail in the next section.

In the case of the array TX (RX), a unit cell that contains N elements must distribute (combine) the input signal to (from) each of the N elements while providing variable phase-shift and variable-gain functionality in each element. The unit cell may or may not include frequency translation. Note that intermediate frequency (IF) signal distribution is preferable to RF signal distribution in the package; however, frequency translation implies that multiple unit cells need to be phase locked. Therefore, in addition to the IF signal, a frequency reference at the local oscillator (LO) frequency or a lower frequency must also be distributed. As mentioned earlier, at mmWave frequencies, the N-element unit cell must be envisioned while also considering the interface between the IC and antennas.

Research and development efforts focused on RFICs have led to scalable integrated phased array architectures based on RF and/or LO-path phase shifting at frequencies from 6 GHz to beyond 100 GHz [7, 8]. The scalable low-IF 6–18 GHz array receiver in [7] incorporates two receivers, with each receiver capable of providing two outputs with independent variable phase shift and variable gain for multi-beam arrays. The IC includes PLLs that operate from a 50 MHz reference enabling phase locking between multiple ICs. The sub-100 MHz IF and reference frequencies simplify multi-IC packaging; however, an N-element array requires N such ICs, increasing packaging complexity. In [9], a scalable Q-band array is presented that leverages a 16-element phased array T/R phase shifting and combining front-end [10] along with a 4 × 4 array of wide-scan patch antennas. This approach also addresses the challenges of RF and DC interconnect within the unit cell with a micromachined silicon interposer for signal routing. Including the antennas and ICs in the unit cell simplifies subsequent packaging, but the signal distribution is still at RF, and additional ICs are required for signal combining and distribution to the phased array unit cell. In [8], a scalable CMOS transmit phased array element is developed at 140 GHz, with each element containing a phase locked loop (PLL) capable of LO-path phase shifting, digital-to-RF upconversion, and antenna-on-PCB that eliminates the need for mmWave signal distribution when multiple elements are tiled.

In general, digital beamforming arrays that incorporate analog-to-digital (A/D) and D/A conversion in each unit cell considerably reduce RF, IF, or analog baseband signal distribution and reduce sensitivity to packaging when the array is scaled to larger numbers of elements. Furthermore, digitization is an approach that is well suited for advanced CMOS technologies. The potential for such digital-intensive scalable arrays has been demonstrated in X-band using monolithic microwave ICs (MMICs) and commercial off-the-shelf (COTS) components in [11]. It must be noted, however, that digital IO can lead to high power consumption in wideband arrays, limiting array size (8 bits I and Q at 1 GS/s, implies 16 Gb/s, translating to 160 mW/IC assuming 10 pJ/b serial link efficiency). Therefore, an attractive approach to realizing wideband large-scale arrays is a combination of N-element RF-combined unit cells and digital IO at the sub-array level. Hybrid analog and digital beamforming has also been explored in [12].

Scalability Options and Trade-offs

Each N-element unit cell forming the scalable array can be designed to use RF-path, LO-path, or IF-path phase shifting. Among these options, although LO- and IF-path phase shifting are relatively easier to implement, RF-path phase shifting offers significant hardware and performance

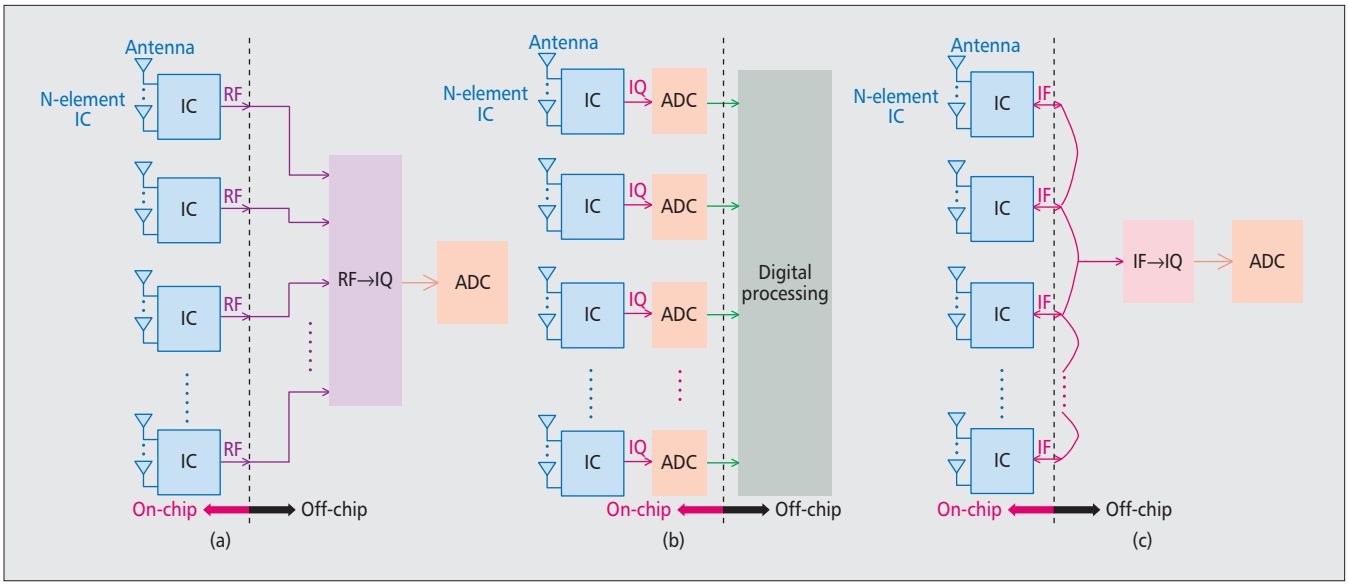

Figure 2. Scalable array system architecture concepts: a) distributing RF signals from IC output; b) digital beamforming; c) distributing IF signals between daisy-chained unit cells.

advantages. An RF-path phase shifting architecture uses the minimum amount of hardware and achieves the lowest power consumption. In the receiver, since all the interferers are nulled out at RF, the linearity requirements of the IF baseband stage are reduced [13].

Figure 2 illustrates different architecture concepts for combining and distributing the input signals for each of the *N*-element unit cells to form a large scalable array. In Fig. 2a, the combined RF signals are output from each unit cell and are distributed on the package or board level before being down-converted to baseband. Each *N*-element RFIC in this case uses RF-path combining. The relative simplicity and flexibility of the RFICs come at the expense of complex packaging and board design. In Fig. 2b, A/D and D/A conversion is incorporated in each unit cell. For each *N*-element IC, any combination of RF-path, LO-path, or IF-path can be used. The incorporation of A/D and D/A inside the unit cell enables significantly reduced signal routing for package and board implementation; however, repeating the A/D/A in each unit cell results in significantly higher power consumption. In Fig. 2c, an intermediate balanced approach is outlined where IF signals from each unit cell are daisy chained and further down-converted to baseband on the package or PCB level. Again, for the option in Fig. 2c, each *N*-element RFIC can use any of the three combining options. While this replicates the on-IC IF hardware, as opposed to the option in Fig. 2a, the distribution of IF signals instead of RF signals significantly simplifies packaging and board design.

For silicon-based scalable phased array at W-band frequencies (75–110 GHz), the λ/2 antenna pitch ranges from 1.35 to 2 mm, which requires tightly integrated antenna solutions for system packaging. Figure 3 illustrates three categories of packaging options for antenna integration at these frequencies.

The first of these options involves implementing antennas directly on the application PCB. The RFICs are flip-chip bonded to the board on the opposite side from the direction of radiation. Note that wire bonding of the ICs would not meet the tight pitch demands for the array that are imposed by scalability considerations. In this approach, the major challenges comprise designing antennas with sufficient gain and efficiency, as well as controlling the antenna variation arising from the PCB manufacturing tolerances. For example, the registration (i.e., lateral movement of the position) of vertical vias used for carrying mmWave feed signals degrades as the number of layers grows, which needs to be accounted for in the antenna and board design.

The second option shown in Fig. 3 involves implementing antennas on the first-level package. Together with the flip-chip bonded ICs, the package forms a unit tile module. The modules can then be attached to a second-level PCB through ball grid arrays (BGAs) to form a larger array. Depending on the application, there are many substrate technologies that have appropriate properties for implementing the package with embedded antennas, including but not limited to low-temperature co-fired ceramic (LTCC), glass, multi-layer organic polymer, and embedded wafer level BGA (eWLB). The structure tolerance on the package is typically significantly improved over that of the PCB process, which allows antennas to be built with better uniformity at a tight pitch. Each assembled package can also be prescreened and tested by checking digital functions of all elements before the package is mounted on the board. On the other hand, this approach does have a disadvantage, which is the higher complexity of the system assembly.

The third option considered in Fig. 3 is implementing antennas directly on the RFIC. In [14], it has been demonstrated that by stacking a glass (quartz) substrate with a metal patch on top of the RFIC with antenna feed, the peak gain of the superstrate antenna can be boosted to 4 dBi at 110 GHz with 50 percent efficiency. Refer-

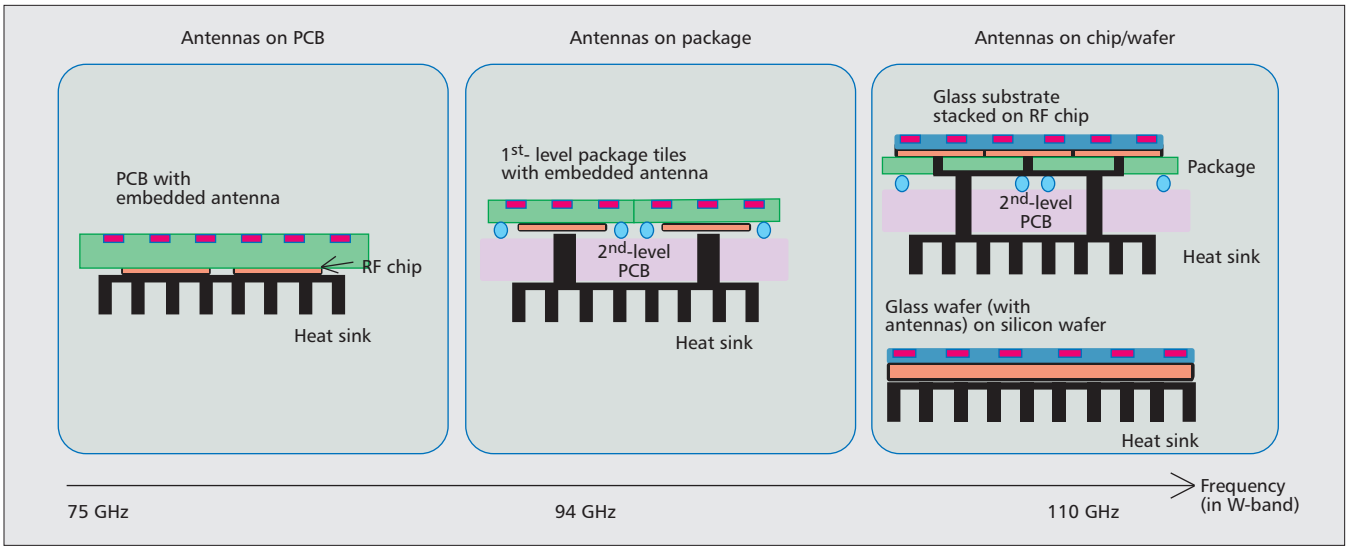

Figure 3. Antenna and package options for W-band scalable phased array.

ence [14] also proposes the concept of a wafer-scale phased array that stacks a glass wafer with antenna elements directly with a silicon wafer with phased array elements and antenna feeds. An alternative is to first have RFICs attached to the package and then a glass substrate on top of the RFICs to form a large array. In this case, through-silicon-vias (TSVs) are required for the RFICs to feed signals and deliver power and ground depending on the chip configuration (face-up or face-down). The system integration complexity of this third option is expected to be the highest among all three options.

Other important challenges in system integration from an electrical perspective include the power integrity of different voltage and ground domains, the signal integrity of the chip-to-antenna transition, as well as high-speed signal wiring (e.g., for baseband, IF, and LO signals). Characterization of on-chip and off-chip interconnect losses is required based on conductor and dielectric material properties at the frequencies of interest.

Thermal management is another key aspect of the system design. Figure 3 illustrates a conceptual placement of heat sinks. Multi-physics co-simulations with the active IC power budget are required in the design phase to accurately evaluate on-die temperature based on different cooling options and the projected ambient environment. For example, the power density for 60 GHz and 94 GHz phased-array RFICs [1, 4] operating in the receive mode can reach 87 mW/mm^2 and 77 mW/mm^2, respectively. Although these numbers are lower than what a typical server processor consumes (e.g., ~200 mW/mm^2), cooling the phased array system is expected to be more challenging: compared to a server processor, more total heat needs to be dissipated due to a large number of active elements, and many more ICs need to be thermally controlled simultaneously to support the scalability of the system. On the other hand, advances in technology such as SiGe BiCMOS processes with higher cutoff frequency (e.g., IBM SiGe 8XP or 9HP), in combination with new breakthroughs in circuit design, are expected to reduce power consumption.

The three antenna and package options in Fig. 3 are viable approaches to support the implementation of W-band scalable phased arrays. Generally speaking, due to the $\lambda/2$ pitch requirement, the on-chip antenna approach works better for higher frequencies in the band, whereas on-board and on-package antenna approaches work better for lower frequencies. In the following section, a prototype scalable 64-element phased array at 94 GHz based on a fully integrated antenna-in-package solution is presented.

Dual Polarized 94 GHz 64-Element Scalable Phased Array

Our approach to realizing a dual-polarized scalable phased array is illustrated in Fig. 4a, which follows the antenna-in-package approach described in Fig. 3. Each transceiver IC contains 16 dual-polarized RF phase-shifting front-ends. RF-path phase shifting was selected to achieve minimum hardware and power consumption at the IC level. All of the mmWave functions are integrated monolithically. Each package houses four ICs and includes 64 dual-polarized antennas. The antennas are placed at a $\sim \lambda/2$ pitch at 94 GHz in both the x and y dimensions, and the antennas on the perimeter are placed $\sim \lambda/4$ away from the package edge. By tiling the packages adjacent to one another on a PCB, phased arrays of large aperture can be created to support long-distance communication and high-resolution imaging. The first two steps of this approach to implementing scalable phased arrays (transceiver IC and package integrating ICs and antennas) have been demonstrated in hardware [4, 15] and are described below in more detail. The final step (board-level array tiling multiple packages) is the subject of future work.

The multi-function dual-polarization phased

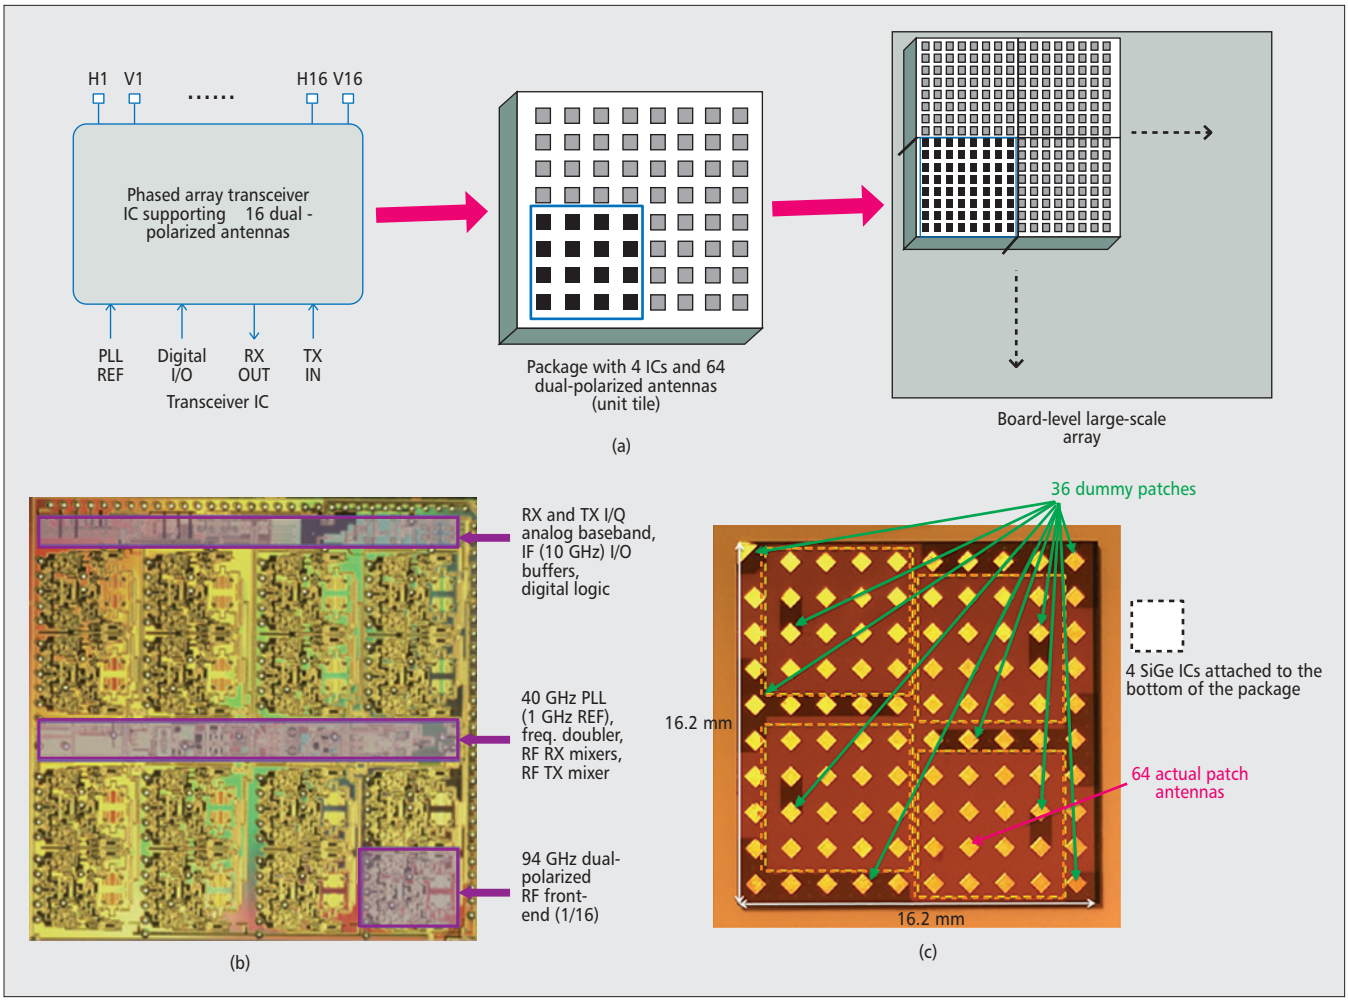

Figure 4. a) A prototype scalable 64-element phased array at 94 GHz; b) dual-polarized 16-element 94 GHz phased array transceiver IC photograph (6.6 mm × 6.7 mm); c) a close-up view of the 4-chip package with actual patch antennas and dummies.

array transceiver IC supports both radar and communication applications at W-band [4]. 32 receive elements and 16 transmit elements with dual outputs are integrated to support 16 dual polarized antennas in a package. As illustrated in Fig. 4b, the IC includes two independent 16:1 combining networks, two receiver down-conversion chains, an up-conversion chain, a 40 GHz PLL, an 80 GHz frequency doubler, extensive digital control circuitry, and on-chip IF/LO combining/distribution circuitry to enable scalability to arrays at the board level. The fully integrated transceiver is fabricated in the IBM SiGe BiCMOS 8HP 0.13 µm process, occupies an area of 6.6 mm × 6.7 mm, and operates from 2.7 V (analog/RF) and 1.5 V (digital) supplies. Multiple operating modes are supported, including the simultaneous reception of two polarizations with a 10 GHz IF output, transmission in either polarization from an IF input, or single-polarization transmission/reception from/to I&Q baseband signals.

Iterations of circuit-package-antenna co-design were performed under severe physical dimension constraints to support array scalability at the package and board levels. Figure 4c shows a close-up view of antenna patches at the top of the package. 100 (10 × 10) patch structures at 1.6 mm spacing (λ/2 at 94 GHz) cover the surface of the package. The IC area is very close to that required for 16 (4 × 4) antennas with l/2 spacing for all the functionality; as a result, a multi-chip 16.2 mm^2 × 0.75 mm package containing 4 SiGe-based RFICs and a 292-pin 0.4 mm-pitch BGA was designed to achieve as high an array fill factor as possible [15]. The multi-chip package approach also mitigates the board-level integration risks compared to the single-chip package approach. Out of the 100 patches for each package, 64 are actual dual-polarized patch antennas and 36 are dummy structures (which do not have actual antenna features other than the surface patch). Therefore, the effective array fill factor is 64 percent. The dummy structures are placed at pseudo-randomized locations to minimize the impact of the reduced fill factor on side lobes. The copper balance in terms of metal percentage per layer also increases with the dummies, which improves manufacturability of the package. Figure 5a illustrates an array simulation in MATLAB with 1024 isotropic radiators based on the antenna pattern, which is equivalent to tiling 16 (4 × 4) 94 GHz packages. Notice that by choosing the patch locations carefully, empty rows or columns of active radiation elements can be avoided. The

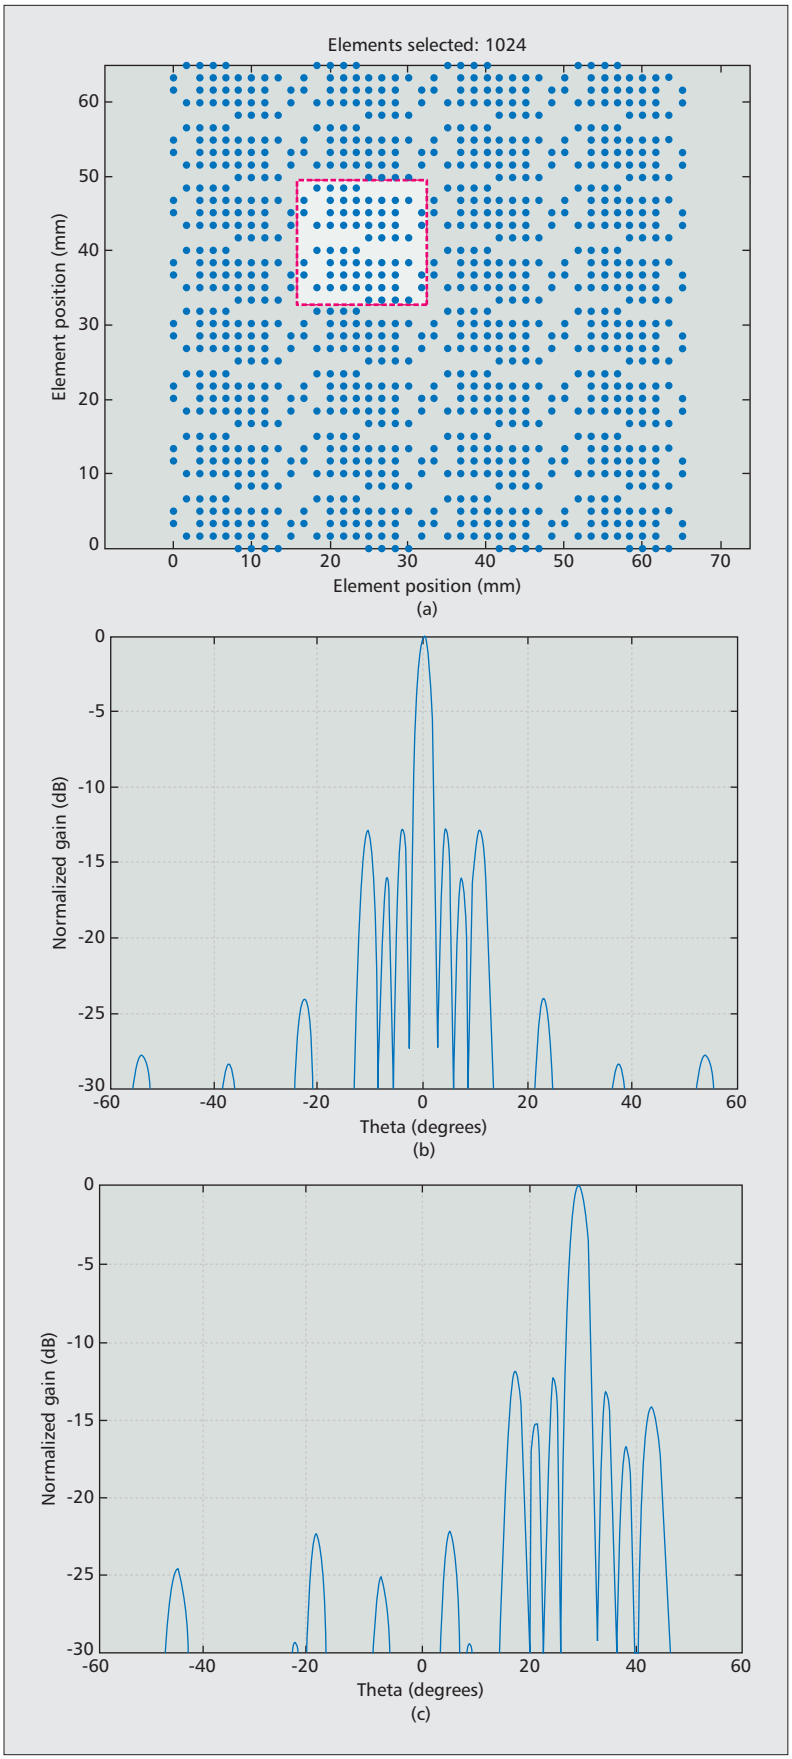

Figure 5. a) Array simulation with 1024 isotropic radiation elements based on the 94 GHz package antenna pattern (as highlighted); b) simulated 1024-element radiation pattern ($\theta = 0°$); c) simulated 1024-element radiation pattern ($\theta = 30°$).

simulated radiation patterns are plotted in Figs. 5b and 5c for $\theta = 0°$ and $\theta = 30°$, respectively.

The overlay of a SiGe die image with a quadrant of the package layout is illustrated in [15]. Two rows and columns of BGA pins provide all signal, power, and ground connections to the C4s (controlled collapse chip connections) on the north and east sides of the die. Signal and power integrity are taken into account in the IC-package co-design. For example, high-speed differential signals are routed from the inner BGA row as short microstrip pairs on the bottom surface layer to avoid via transition, whereas low-speed single-ended signals are routed from the outer BGA row as striplines. Dielectric properties (relative permittivity and loss tangent) of the buildup and core layers of the organic package are measured and characterized up to 110 GHz so that accurate full-wave electromagnetic models for the interconnects and antennas can be constructed. Furthermore, two groups of voltage supply pins, as well as ground pins, are placed evenly on the periphery to ensure good power distribution to the chip. The RF front-end C4s for the W-band antenna feed are laid out using a 225-µm-pitch GSGSG configuration. In order to minimize the RF antenna feed line length, the locations of these C4s were optimized together with the circuit layout for the front-end, core, and digital macros.

Figure 6a illustrates a conceptual view of the measurement setup for one assembled multi-chip module with 64 embedded antennas. The patch antenna array is on the top of the package. Four SiGe ICs are flip-chip attached to the bottom of the package. The module is mounted to a system board via a pogo-pin-based interposer, which allows air cooling and supports easy removal for screening. The IC package assembly was performed using standard flip-chip attach processes with lead-free solder reflow and underfill. SMP connectors are populated on the board to provide PLL reference and IF signals to the four ICs. In addition, a daisy-chain configuration is implemented as illustrated in Fig. 2c, so only one PLL reference input and one IF input are required from external sources for antenna pattern and radiated power measurement.

The test board with the phased array module is further mounted to a field programmable gate array (FPGA) board through which the digital circuits are programmed and controlled by a PC. Two motors are used to drive the rotation of the test board (azimuth angle) and the receive horn (elevation angle), respectively. For transmit-mode testing, the 16 elements of each IC are turned on sequentially while measuring the receive power in the broadside direction. The optimum phase coefficients for spatial power combining are found by sweeping the phase of each element with a 22.5° step. Next, a phase offset between groups of the four 16-element ICs can be found to achieve the module-level 64-element power combining. The measured 16-element and 64-element special power combining results in terms of normalized EIRP are plotted in [15]. To enable calibration, gain control is also applied to each element and tuned to compensate the radiated power variation (e.g., due to the intrinsic non-uniformity of

antenna gains). Finally, beam steering can be enabled by adjusting the phase coefficients, which are calculated analytically based on the target direction. Figure 6b shows the measured radiation patterns after spatial power combining of all 64 elements for both H and V polarizations. Good correlation with a simulated ideal radiation pattern is shown in [15]. In addition to the broadside radiation patterns, patterns with 15° beam steering are also demonstrated for both polarizations with side-lobe levels lower than 10 dB. A wider steering angle (e.g., 30°) is achievable at the expense of a higher side lobe level, which in turn can be overcome by using different tapering techniques.

Conclusions

Millimeter-wave phased array technologies are rapidly emerging in the areas of automotive radar, satellite and backhaul communications, security scanning, and imaging. In this article, we review the existing solutions of scalable phased array and discuss the advantages and challenges of designing multi-function scalable W-band phased arrays on silicon with special consideration to packaging and antenna integration. Our research efforts in this area have demonstrated the feasibility of using silicon technology and organic package substrates to implement scalable arrays at W-band. A fully integrated antenna-in-package solution is developed to build a compact W-band dual polarized phased array transceiver module with 64 antennas and 4 SiGe ICs. The 94 GHz 16-element transceiver IC, fabricated in a mature SiGe BiCMOS process, demonstrates noise figure (<10 dB), output power (> 0dBm per element), phase shift (~11° resolution per element), and phase noise performance (–110 dBc/Hz @ 10 MHz offset) suitable to support imaging and communication applications. The results from this phased array module have demonstrated 64-element spatial power combining and electronic beam steering for both horizontal and vertical polarizations. A next important step is to further tackle system integration challenges, that is, to implement a larger array with multiple modules on the board level (e.g., 1024 elements with 16 packages), which will allow exploration of the approach for larger-scale applications.

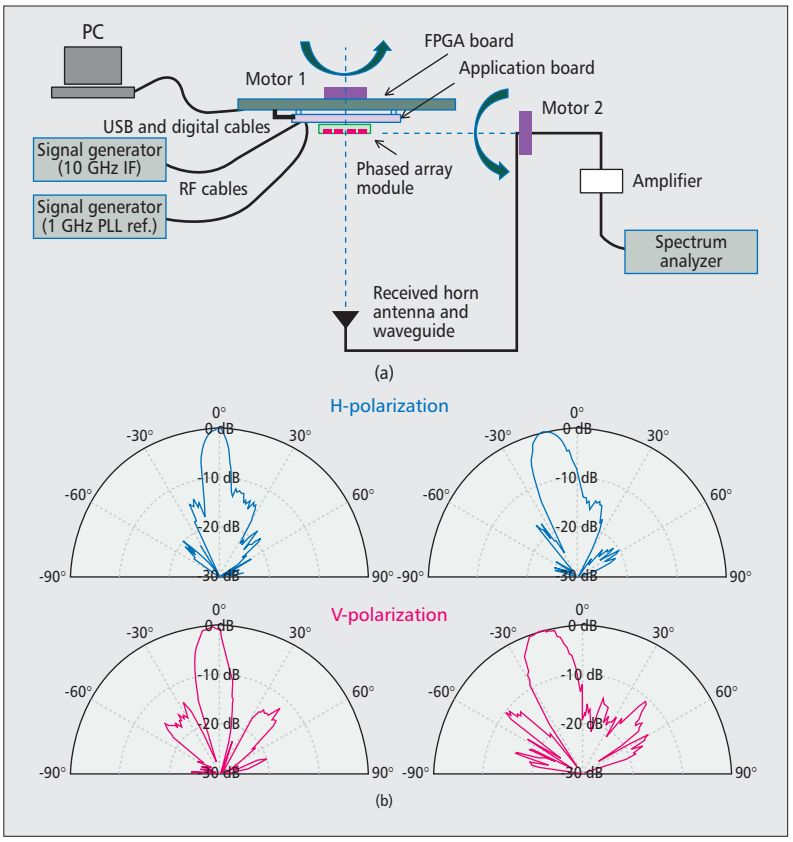

Figure 6. a) Radiated power and pattern measurement setup; b) measured 64-element array radiation patterns and beam steering.

Acknowledgments

IBM's 94 GHz work has been partially funded by DARPA Strategic Technology Office (STO) under contract # HR0011-11-C-0136 (Si-Based Phased-Array Tiles for Multifunction RF Sensors, DARPA Order no. 8320/00, Program Code 1P30). The views, opinions, and/or findings contained in this presentation are those of the authors/presenters and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense. Approved for Public Release, Distribution Unlimited.

References

[1] A. Valdes-Garcia *et al.*, "Single-Element and Phased-Array Transceiver Chipsets for 60-GHz Gb/s Communications," *IEEE Commun. Mag.*, April 2011, pp. 120–31.

[2] "International Technology Roadmap for Semiconductor 2013 Edition Radio Frequency and Analog/Mixed-Signal Technologies Summary"; http://www.itrs.net/Links/2013ITRS/2013Chapters/2013RFAMS_Summary.pdf

[3] Z. Pi and F Khan, "An Introduction to Millimeter-Wave Mobile Broadband Systems," *IEEE Commun. Mag.*, June 2011, pp. 101–07.

[4] A. Valdes-Garcia *et al.*, "A Fully-Integrated Dual-Polarization 16-Element W-Band Phased-Array Transceiver in SiGe BiCMOS," *Proc. IEEE Radio Frequency Integrated Circuits Symp.*, June 2013, pp. 375–78.

[5] M. LaManna and A. G. Huizing, "Scalable Multifunction Active Phased Array Systems: From Concept To Implementation," *Proc. 2006 IEEE Conf. Radar*, 24–27 Apr. 2006, pp. 9–14.

[6] A. G. Huizing, "Design Issues of an Open Scalable Architecture for Active Phased Array Radars," *Proc. IEEE Int'l Symposium on Phased Array Systems and Technology*, 14–17 Oct. 2003, pp. 277–82.

[7] J. Sanggeun *et al.*, "A Scalable 6-to-18 GHz Concurrent Dual-Band Quad-Beam Phased-Array Receiver in CMOS," *IEEE J. Solid-State Circuits*, vol.43, no.12, Dec. 2008, pp. 2660–73.

[8] A. Tang, et al., "A 65nm CMOS 140 GHz 27.3 dBm EIRP Transmit Array with Membrane Antenna for Highly Scalable Multi-Chip Phase Arrays," *Proc. IEEE MTT-S Int'l. Microwave Symp.*, 1–6 June 2014, pp. 1–3.

[9] J. Hacker *et al.*, "A 16-Element Tansmit/Receive Q-Band Electronically Steerable Subarray Tile," *Proc. IEEE MTT-S Int'l. Microwave Symp.*, 17–22 June 2012, pp. 1–3.

[10] C. Kim, D. Kang, and G. M. Rebeiz, "A 44–46-GHz 16-Element SiGe BiCMOS High-Linearity Transmit/Receive Phased Array," *IEEE Trans. Microwave Theory and Techniques*, vol. 60, no. 3, Mar. 2012, pp. 730–42.

[11] D. Curtis *et al.*, "32-Channel X-band Digital Beamforming Plug-and-Play Receive Array," *Proc. IEEE Int'l Symp. Phased Array Sys. and Technology*, 14–17 Oct. 2003, pp. 205–10.

[12] S. Han *et al.*, "Large-Scale Antenna Systems with Hybrid Analog and Digital Beamforming for Millimeter Wave 5G," *IEEE Commun. Mag.*, Jan. 2015, pp. 186–94.

[13] A. Hajimiri *et al.*, "Phased Array Systems in Silicon," *IEEE Commun. Mag.*, Aug. 2004, pp. 122–30.

[14] W. Shin *et al*., "A 108–114 GHz 4×4 Wafer-Scale Phased Array Transmitter with High-Efficiency On-Chip Antennas," *IEEE J. Solid-State Circuits*, vol. 48, no. 9, Sept. 2013, pp. 2041–55.
[15] X. Gu *et al*., "A Compact 4-Chip Package with 64 Embedded Dual-Polarization Antennas for W-band Phased-Array Transceivers," *Proc. IEEE Electronic Components and Technology Conf.* (ECTC), May, 2014, pp. 1272–77.

BIOGRAPHIES

XIAOXIONG GU [SM] received his Ph.D. in electrical engineering from the University of Washington in 2006. He joined IBM T. J. Watson Research Center as a research staff member in January 2007. His current research activities are focused on design, packaging, integration, and characterization of optoelectronic and mmWave communication and computation subsystems. He received SRC Mahboob Khan Outstanding Industry Liaison Awards in 2012 and 2014, IBM Invention Plateau Awards in 2012 and 2013, a Best Paper Award at IEEE EMC Symposium in 2013, the Best Conference Paper Award at IEEE EPEPS in 2011, DesignCon Paper Awards in 2008 and 2010, the Best Session Paper Award at IEEE ECTC in 2007, and the Best Interactive Session Paper Award at IEEE DATE in 2008. He is the chair of THE Professional Interest Community of Electrical Interconnect and Packaging at IBM. He currently serves on the Technical Program Committees for EPEPS, ECTC, EDAPS, and DesignCon.

ALBERTO VALDES-GARCIA is currently a research staff member and manager of the RF Circuits and Systems Group at IBM T. J. Watson Research Center. He received his Ph.D. degree in electrical engineering from Texas A&M University in 2006. His current research work is on silicon-integrated millimeter-wave systems for imaging and communications. From 2006 to 2009, he served in the IEEE 802.15.3c 60GHz standardization committee. Since 2009 he served as a Technical Advisory Board member with Semiconductor Research Corporation (SRC), where he was Chair of the Integrated Circuits and Systems Sciences Coordinating Committee in 2011 and 2012. He holds 17 U.S. patents and is a Co-Editor of the book *60GHz Technology for Gbps WLAN and WPAN: From Theory to Practice* (Wiley, 2011). He is was winner of the 2005 Best Doctoral Thesis Award presented by the IEEE Test Technology Technical Council (TTTC), the recipient of the 2007 National Youth Award for Outstanding Academic Achievements, presented by the President of Mexico, and a co-recipient of the 2010 George Smith Award presented by the IEEE Electron Devices Society. In 2013, he was selected by the National Academy of Engineering for its Frontiers of Engineering Symposium.

ARUN NATARAJAN received his B.Tech. degree in electrical engineering from the Indian Institute of Technology, Madras, in 2001, and his M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology (Caltech), Pasadena, in 2003 and 2007, respectively. From 2007 to 2012, he was a research staff member at IBM T. J. Watson Research Center, and worked on mmWave phased arrays for multi-Gb/s data links and airborne radar. In 2012 he joined Oregon State University as an assistant professor in the School of Electrical Engineering and Computer Science. His current research is focused on low-power RF transceivers for IoT as well as mmWave and sub-mmWave circuits and systems for wireless communication. He received the National Talent Search Scholarship from the Government of India (1995–2000), the Caltech Atwood Fellowship in 2001, the Analog Devices Outstanding Student IC Designer Award in 2004, the IBM Research Fellowship in 2005, and the 2011 Pat Goldberg Memorial Award for Best Paper in CS/EE/Math in IBM Research.

BODHISATWA SADHU is currently a research staff member at IBM T. J. Watson Research Center. He received his Ph.D. degree in electrical engineering from the University of Minnesota, Minneapolis, in 2012. For his Ph.D., he worked on wideband circuits and architectures for software defined radio applications. Since 2012, he has been working on frequency synthesizers and mmWave transceivers in the RF/mm-wave Communications Circuits group at IBM Research. He has authored and coauthored more than 30 papers, authored the book *Cognitive Radio Receiver Front-Ends — RF/Analog Circuit Techniques* (Springer, 2014), and holds 5 issued U.S. patents with 15+ pending. He was the recipient of the University of Minnesota Graduate School Fellowship, 2007, 3M Science and Technology Fellowship, 2009, and the University of Minnesota Doctoral Dissertation Fellowship, 2011.

DUIXIAN LIU received his Ph.D. degree in electrical engineering from the Ohio State University, Columbus, in 1990. From 1990 to 1996, he was with Valor Enterprises Inc., Piqua, Ohio, initially as an electrical engineer and then as chief engineer, during which time he designed an antenna product line ranging from 3 MHz to 2.4 GHz for the company, a very important factor for the prestigious Presidential "E" Award for Excellence in Exporting in 1994. Since April 1996, he has been with the IBM T. J. Watson Research Center as a research staff member. He has received three IBM Outstanding Technical Achievement Awards and one Corporate Award, IBM's highest technical award. He was named a Master Inventor in 2007. He has edited a book, *Advanced Millimeter-Wave Technologies — Antennas, Packaging and Circuits* (Wiley, 2009) and is a Section Editor for an upcoming Springer antenna handbook (2015). He has authored or coauthored more than 100 journal and conference papers. He received the 2012 S. A. Schelkunoff Prize Paper Award of the IEEE Antennas and Propagation Society. He has 71 patents issued or pending. His research interests are antenna design, chip packaging, and communications technologies.

SCOTT K. REYNOLDS received his Ph.D. degree in electrical engineering from Stanford University, California, in January 1988. He joined IBM in 1988, where he worked on a wide variety of IBM products, including ICs for disk drive channels, electrical and optical I/O, and RF communication. Beginning in 2003, he was engaged in development of silicon millimeter-wave ICs and packaging for high-data-rate wireless links and other applications, including imaging. He has more than 30 U.S. patents and many technical publications, including two papers on 60 GHz wireless transceiver circuits that won the best paper awards at ISSCC in 2004 and 2006. He went on to manage the RF Circuits and Systems group at IBM Research from 2010 to 2013. In 2013, he left IBM to start his own business, Tavish Design, LLC. He continues to consult for IBM on millimeter-wave IC design and packaging.

CALL FOR PAPERS
IEEE COMMUNICATIONS MAGAZINE

SOCIAL NETWORKS MEET NEXT GENERATION MOBILE MULTIMEDIA INTERNET

BACKGROUND

With ever growing popularity and widespread adoption of mobile social applications amongst users, the traffic handled by the mobile networks and the Internet has grown significantly. While researchers have been making advances in the study of social networks and independently in the area of next generation wireless networks, very little attention has been given to the interplay between the two and their impact on each other and the society. The interplay between social networks and mobile networks is compounded by the fact that the advances in smart hand-held devices and those in wireless technologies have paved the way for increasing bandwidth catering to very high data rates. It is entirely likely that such advances in turn could lead to novel social applications not yet thought of. For example, new social applications could emerge in the area of social health, social games, location-based social applications, real-time social collaboration applications, or real-time massively multi-player, multimedia, 3-D, role playing games, and m-commerce. Though there does not exist a standard definition of what constitutes a Mobile Multimedia Internet, for the purposes of this proposal, we stipulate that the Mobile Multimedia Internet is one that facilitates the convergence of various wireless networks and protocols, including 3G, 4G, and next generation 5G wireless networks and beyond, wireless LANs, Voice over IP (VoIP), cognitive radio, software-defined networking, and cloud networks and protocols.

The emergence of research in the area of 'Big Data' and its applicability to social networking can be hardly overemphasized. Together with the evolution and deployment of cloud RANs, cloud-based and software-defined networks could lead to new forms of 'Big Data,' for example, involving user locations, usage patterns, user mobility, and other user specific behaviors. Also, the ability on the part of next generation networks to facilitate peer-to-peer and ad-hoc networking paves the way for new forms of interactions with mobile social networks and applications. New application areas such as location-based social networking applications (LBSNA) and Internet of Things (IoT) based social networking (IoTBSNA) are maturing and gaining ground.

Through this feature topic we intend to create a venue to bring together researchers and practitioners from different disciplines, especially computer and information sciences and next generation mobile/wireless multimedia Internet/networks as well as other related disciplines to share, exchange, learn, and develop preliminary results, new concepts, ideas, principles, and methodologies, aiming to advance mobile social networks in the new generation of Information and Communication Technologies enabled by Web 3.0, also referred to as next generation intelligent web.

SCOPE OF CONTRIBUTIONS

Original papers are solicited that highlight new advances and directions in cross-domain convergence and blending of mobile networks and social networks to address challenges and develop opportunities in the following (but not limited to) topics:

- Mining Big Data generated by mobile social networks to understand and formulate models social theories.
- Impact of social network growth models on the mobile multimedia Internet.
- New social applications capable of exploiting the new features that next generation networks will bring about.
- Infrastructure to facilitate awareness of the capabilities of social networks within mobile networks.
- Mobile network aware social media service architecture, middleware, and framework.
- The role of IP Multimedia Subsystem in the interplay between social networks and mobile networks;
- Resilient mobile networks leveraging information diffusion, containing contagion, and disincentivization concepts.
- Network-aware search, ranking, and recommendation.
- Mobile social network data collection, analysis, trends, tools, and applications.
- Privacy/security in mobile social applications.
- Location-aware mobile social computing.
- Citizen sensing applications - Mobile sensing for community actions.
- Mobile social computing applications in crisis management.

SUBMISSIONS

This feature topic solicits original work that must not be under consideration for publication in other venues. Authors should refer to IEEE Communications Magazinesubmission guidelines at http://www.comsoc.org/commag/paper-submission-guidelines for information about content and formatting of submissions. Manuscripts must be written in English and contain substantial tutorial content and be readable by a broad general audience working in other fields. All articles must be submitted through IEEE Manuscript Central (http://mc.manuscriptcentral.com/commag-ieee) by the submission deadline. Submit manuscripts to the category "October 2015: Social Networks Meet Next Generation Internet."

SCHEDULE FOR SUBMISSIONS

- Manuscript Submission Deadline: May 15, 2015
- Author Notification: July 10, 2015
- Final Manuscript Due: August 1, 2015
- Publication Date: October 2015

GUEST EDIUTORS

Seshadri Mohan
University of Arkansas at Little Rock
sxmohan@ualr.edu

Nitin Agarwal
University of Arkansas at Little Rock
nxagarwal@ualr.edu

Ashutosh Dutta
AT&T
ad5939@att.com

THz Interconnect: The Last Centimeter Communication

Qun Jane Gu

ABSTRACT

Terahertz, sandwiched between conventional microwave and optical frequencies, has inspired increasing interest due to its uniqueness and high potential applications, such as imaging, sensing, and communications. This article, on the other hand, focuses on one emerging application of the terahertz spectrum: THz interconnect. Intra-/inter-chip communication has doubled every two years over recent decades, and the trend is projected to continue in the future. However, the bandwidth supportable by chip I/O pins cannot keep up with the requirement, which forms the increasing gap between the bandwidth requirement and support capability, or the interconnect gap. To ultimately solve the problem and close the gap, both bandwidth density and energy efficiency should be boosted. THz interconnect holds high potential to boost key performance by leveraging the advantages of both high-speed electronics devices and low-loss quasi-optical channels. This article discusses THz interconnect from different aspects: system architecture, circuit specifications, design challenges, and non-ideality effects. Particularly, this article exemplifies both active and passive circuit design techniques for THz interconnect, a 140 GHz transceiver and a terahertz generator in 65 nm CMOS technology, and a low-loss and process-compatible silicon waveguide channel. THz interconnect opens high potential new revenue to solve the long-standing interconnect issue.

INTRODUCTION AND MOTIVATION

Continuous scaling of semiconductor devices allows more processor cores and integrated functionalities into a single chip to support the growing computation demands of scientific and commercial workloads in both speed and volume [1, 2]. This trend mandates an ever increasing inter-/intra-chip communication bandwidth, which has been a big challenge over recent decades. This challenge has motivated active research to improve interconnect capacities, characterized by two key specs: bandwidth density, defined as gigabits per second per square millimeter, determining the aggregate throughput; and energy efficiency, defined as Joules per bit, indicating the overall power consumption. The required off-chip I/O bandwidth doubles about every two years, significantly exceeding the growth rate of the number of pins due to packaging/assembly limitations [3]. The gap between the interconnect requirement and the capability forms the "interconnect gap." Given state-of-art (SOA) performance of energy efficiency and bandwidth density, the power consumption and chip size to support interconnect only will be overwhelming for most computers and embedded systems [4]. In addition, cost, defined as dollars per gigabit per second, needs to scale down over the increasing interconnect bandwidth. To sustain the continuous demands for system performance unmet by the current intra- interconnect capabilities, the interconnect gap must be filled.

There are two major research areas in the interconnect: electrical interconnect (EI) [5] and optical interconnect (OI) [4, 6]. The existing SOA has demonstrated a bandwidth density of about 37 Gb/s/mm^2 for OI [4] and 8 Gb/s/pin for EI [5] with energy efficiency of about 4 pJ/bit [5, 6]. However, it is challenging for both EI and OI to completely address the interconnect issues individually. The major limitation of EI is the low bandwidth-distance product of the metallic medium. Therefore, its energy efficiency drops significantly when transmitting large throughput over a $>$ 1 mm distance due to the prohibitive channel losses. The fiber has unpreceded bandwidth-distance product, which makes it ubiquitous for long distance communication, such as wide area networks (WANs) and metropolitan area networks (MANs). However, OI faces the issues of system integration complexity and overhead, such as electronic-to-optical and optical-to-electronic (EO/OE) conversion, environment sensitivity, and high cost of short-distance communications (e.g., $<$ 10 cm). All of these render the *last centimeter* dilemma, which falls into the distance range for inter-/intra- chip communications.

THz Interconnect (TI), utilizing the frequency spectrum sandwiched between microwave and optical frequencies, has high potential to complement EI and OI by leveraging the advantages of both electronics and optics as shown in Fig. 1a. Continuous scaling of mainstream silicon technologies enables terahertz electronics in silicon, such as a terahertz oscillator [7–9] and detector [10, 11], making terahertz signal generation/detection possible and suitable in silicon

The author is with the University of California, Davis.

0163-6804/15/$25.00 © 2015 IEEE

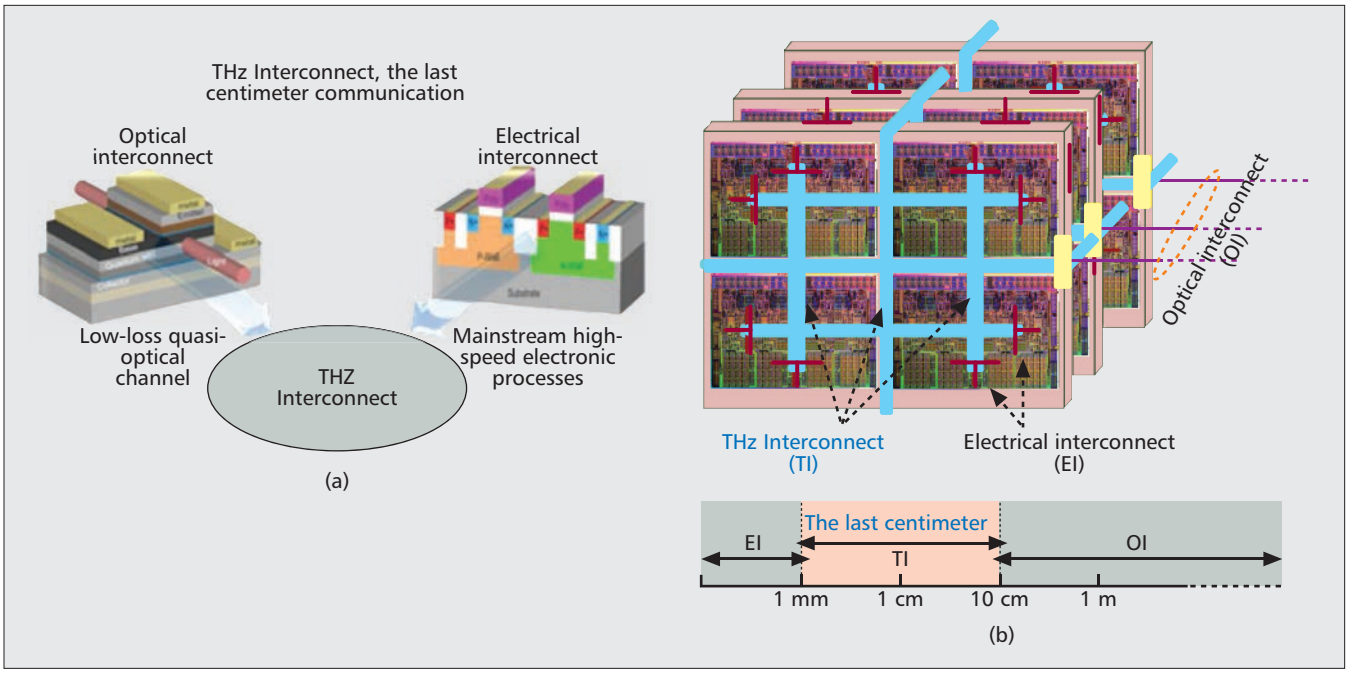

Figure 1. a) THz Interconnect leverages the advantages of both electrical interconnect and optical interconnect; b) THz Interconnect complements EI and OI to serve as the last centimeter data link.

processes. On the other hand, terahertz waveguides, similar to their optical counterparts, have small dimensions and present low loss with an attenuation factor < 0.04/cm [12, 13]. The ultra-low-loss feature alleviates the TI link budget to allow low transmission output power and improves TI energy efficiency. In addition, TI favors technology scaling because the increasing frequency supports higher communication data rates and reduces channel dimensions, thus resulting in a larger bandwidth density. These unique features of TI enable it to complement EI and OI, with one optimum architecture shown in Fig. 1b for high energy efficiency, high bandwidth density, low cost, and high resilience. As shown in Fig. 1, the TI aims to address the last centimeter dilemma between 1 mm and 10 cm, while EI and OI address the issues in their most effective operation regions to ultimately fill the interconnect gap.

THz Interconnect Link Budget Analysis

One stringent requirement of interconnect is an ultra-low bit error rate (BER), demanding high signal-to-noise ratio (SNR). To realize high-efficiency systems, simple modulation schemes, such as binary phase shift keying (BPSK) or amplitude shift keying (ASK), are preferred. Figure 2 analyzes the link budget suggested by simulation and measurement results from a 65 nm complementary metal oxide semiconductor (CMOS) technology. With the ASK scheme, a data rate of 50 Gb/s requires 50 GHz bandwidth due to the 1 b/Hz bandwidth efficiency. The assumption of 10 percent fractional bandwidth leads to 500 GHz carrier frequency. To achieve BER < 1×10^{-15}, the SNR must be > 18 dB. With a 20 dB noise figure, the receiver sensitivity is about –29 dBm. The loss from the channel, including the signal coupling from/to the transceiver and channel itself with 10 mm length, is assumed to be 10 dB. The loss is mostly dominated by the channel coupler with the loss from the channel itself < 0.08 dB/mm, based on the discussion below. Therefore, the loss does not change much with long channel length. With the link budget margin of 10 dB, the output power from the transmitter is therefore about –9 dBm. With 0.5 percent efficiency at the transmitter side, the DC power consumption is about 25 mW from the transmitter. The receiver consumes less power than the transmitter. Here, 5 mW power budget is allocated to the receiver, which leads to total power consumption of 30 mW for the entire transceiver, and the resulting energy efficiency is 0.6 mW/Gb/s or 0.6 pJ/b. With regard to the bandwidth density, the interconnect size needs to be evaluated, which is mostly constrained by the channel dimension and is about 0.25×0.25 mm^2 in Fig. 2's system assumption. Therefore, the bandwidth density can be estimated to 50 Gb/s/(0.25*0.25) = 800 Gb/s/mm^2. This link budget analysis is based on 65 nm CMOS technology speed and capabilities. The device f_T in 65 nm CMOS technologies is about 200 GHz. To support 500 GHz operation, novel circuit architecture and design techniques are required. This is exemplified below. With more advanced technologies, such as 40 nm and 28 nm processes, the increasing carrier frequency is able to support a larger signal bandwidth and higher data rate while using a smaller channel size. Therefore, the bandwidth density can increase quadratically with the carrier frequency. In addition, a higher device speed supports better DC-to-RF conversion efficiency for a better interconnect energy efficiency. Moreover, all the active circuits are based on standard mainstream processes, which provides the most cost-effective

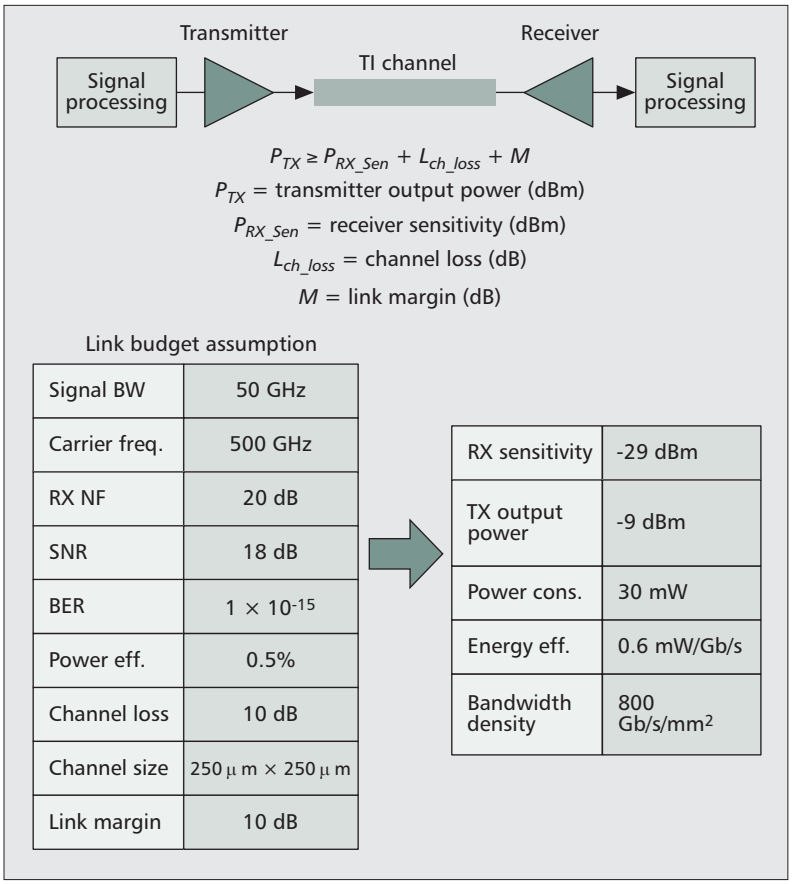

Figure 2. THz Interconnect link budget analysis, and the corresponding energy efficiency and bandwidth density based on 65 nm CMOS technology.

implementation and high resilience. Therefore, TI scales well with processes to be able to ultimately close the interconnect gap by providing scalable interconnect capabilities.

Active Circuit Design Challenges and Example

As shown in Fig. 2, active transceivers supporting high data rate with small power consumption are the key to TI. In such high frequency spectrum, spectrum usage efficiency is not as a serious concern as in the lower gigahertz range. Therefore, simple modulation schemes, such as ASK and BPSK, can be adopted to reduce hardware complexity and power consumption while still offering high communication data rates by leveraging the wide bandwidth. Among them, on-off keying (OOK) modulation, the simplest form of ASK, represents the data as the presence or absence of a high-frequency carrier through binary amplitude modulation.

To demonstrate the feasibility, we have developed a non-coherent OOK-based transmitter/receiver for short-distance communications, such as inter-/intra-chip interconnect [14]. Figure 3a sketches the system architecture. In the transmitter, a voltage-controlled oscillator (VCO) generates a high-frequency signal running at 140 GHz as the carrier, and a VCO buffer isolates the VCO core from the OOK modulated power amplifier (PA) to stabilize the carrier frequency and support a large carrier signal. The data modulates the PA input and intermediate nodes together to achieve high modulation depth, > 90 percent. The differential switching configuration by crossing the differential signals is equivalent to reduce the channel resistance to increase modulation speed. In the receiver, the low-noise amplifier (LNA) provides over 20 dB power gain with more than 20 GHz bandwidth centered at 142 GHz. An envelope detector directly demodulates the incoming OOK RF signal into a low-frequency baseband signal, which is then amplified by the following PGA before sending it off-chip.

The 140 GHz TX and RX are fabricated in a 65 nm CMOS technology. To characterize the data link, the TX and RX chips are placed in close proximity (~1 cm) and coupled with bonding wires, as shown in Fig. 3b with chip photos shown as insets. The performance is summarized in the table. The tested data rate is 2.5 Gb/s with a pattern of 2^{15}–1 pseudo-random binary sequence (PRBS). Figure 3c shows the receiver output signal eye diagram with the eye height about 32 mV and eye width about 210 ps. Figure 3d presents the comparison of the data sequence between the input data to the TX and output data from the RX, which clearly indicates a successful link. The measured BER is 4.1×10^{-6}. In this design, the 2.5 Gb/s data rate is mainly limited by the baseband VGA speed, which needs to drive off-chip 50 ohm and only provides 1.2 GHz bandwidth. In practical interconnect scenarios, driving 50 ohm may not be necessary. Therefore, the baseband speed can be designed higher to support larger throughput. In addition, due to the inefficient bonding wire coupling, which indicates > 40 dB loss from simulation, the SNR significantly drops, which degrades BER. With a better coupler design, such as patch antenna or dipolar antenna based couplers, the energy efficiency and communication distance can be improved.

One of the key issues of this transceiver is that the 140 GHz carrier frequency is not high enough to efficiently support data rate higher than 50 Gb/s. In addition, the corresponding passive component size is also large to considerably constrain the bandwidth density. For example, an on-chip patch antenna at 140 GHz occupies about 1 mm^2 chip area.

To boost the energy efficiency and bandwidth density, the carrier frequency needs to be increased to be able to support wide bandwidth and high data rates as well as reduce the passive device sizes. Traditional transit-time-based electronic devices are typically limited by their low cutoff frequencies (i.e., f_T and f_{MAX}). Despite a continuous increase in the device cutoff frequencies, deep-scaled CMOS technologies still suffer major drawbacks in realizing terahertz circuits and systems. First, external parasitics limit circuit operating speed to much lower than device-intrinsic speed due to charging/discharging parasitic capacitance. This scenario becomes even worse with technology scaling due to a larger ratio of external parasitics over intrinsic device loading. Second, a CMOS device demonstrates large losses with a combination ohmic losses

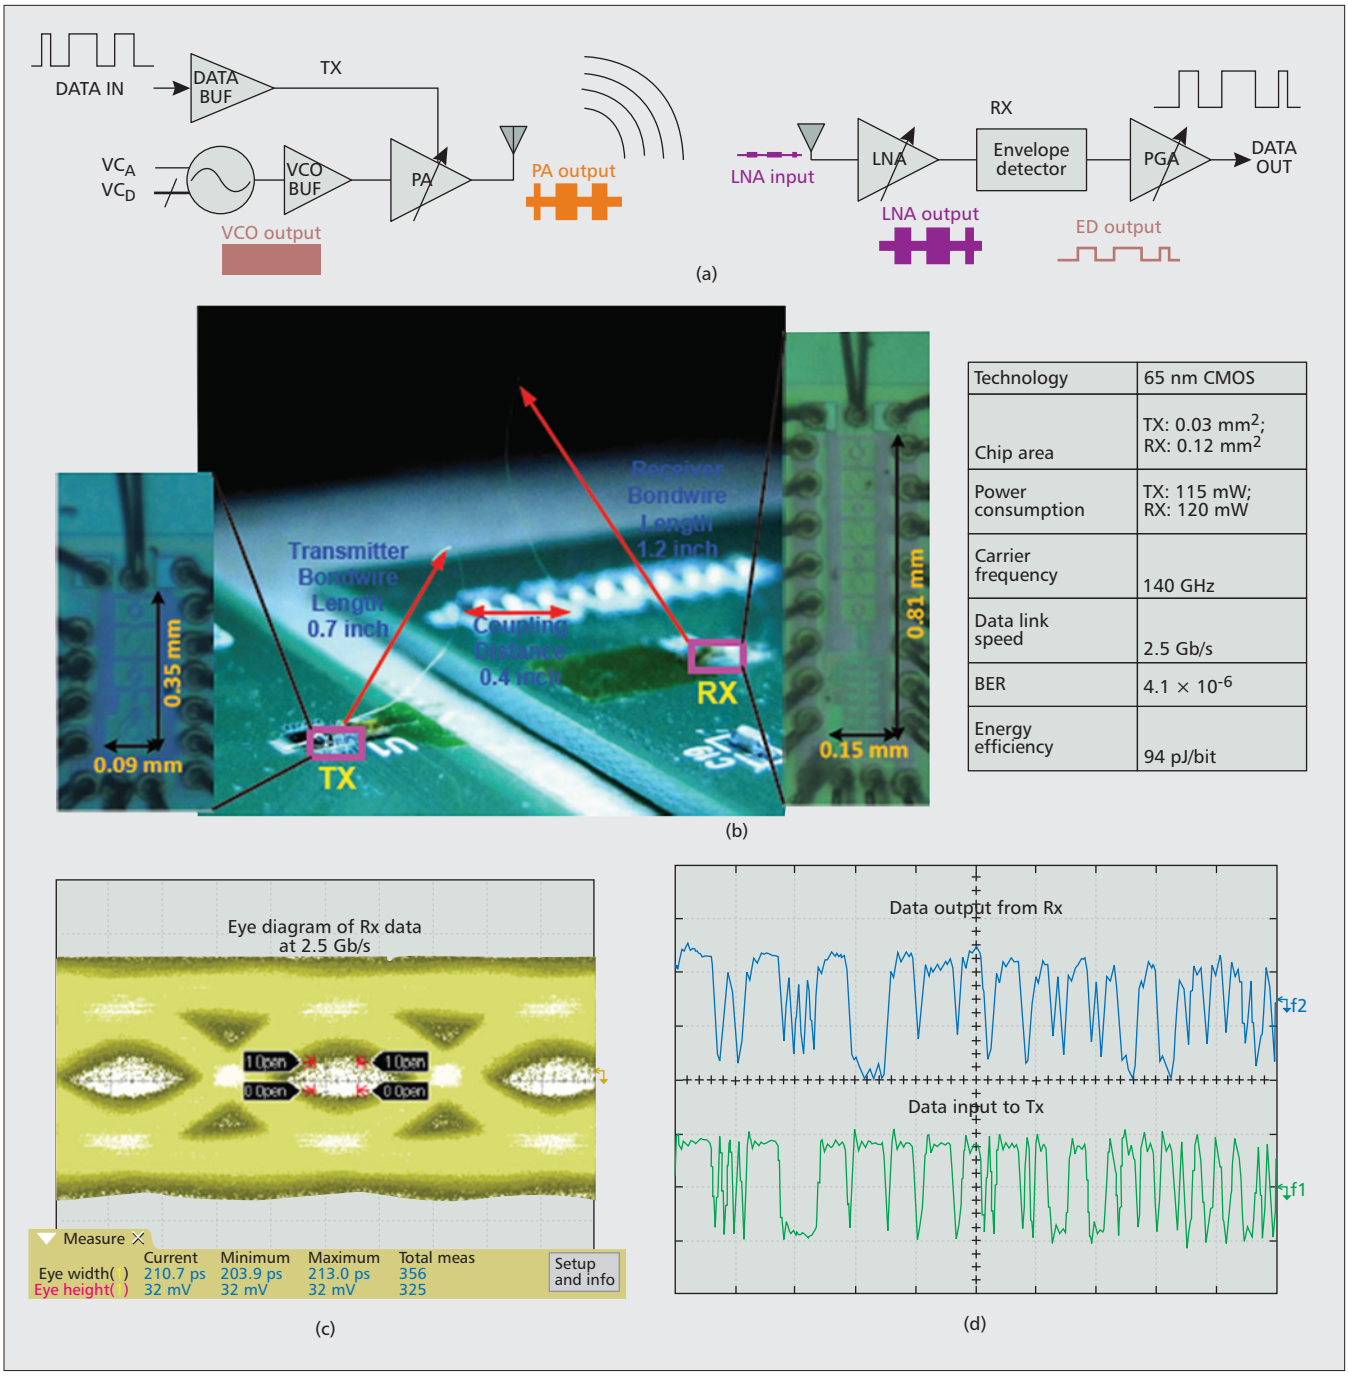

Figure 3. a) OOK-based 140 GHz transceiver for short distance communications; b) the measurement configuration coupled by bonding wires with the chip photos shown in insets and performance summary; c) measured receiver output eye diagram; d) the data sequences of transmitter input and receiver output.

from contacts and metal resistance, and dielectric losses from substrate coupling. These losses degrade circuit quality factor and reduce signal gain, which are particularly severe in terahertz operations and result in the challenge of generate terahertz signals. Third, ultra-high-frequency generation faces big challenges due to small available gain. This is because ultra-high-frequency operation requires small active and passive components, resulting in insufficient signal gain. Increasing signal gain can possibly be achieved by increasing active device size or passive device size. Increasing active device size boosts signal gain, but at the cost of larger capacitance loading, which then shifts down the operating frequency. Increasing passive device size, on the other hand, demands size reduction of the active components, which then reduces the transconductance and drops the overall signal gain. These are active terahertz circuit design dilemmas.

To overcome the above design issues, we have invented a new signal generation architecture, composed by a frequency-selective negative resistance (FSNR) tank, to generate the oscillation frequency higher than device cutoff fre-

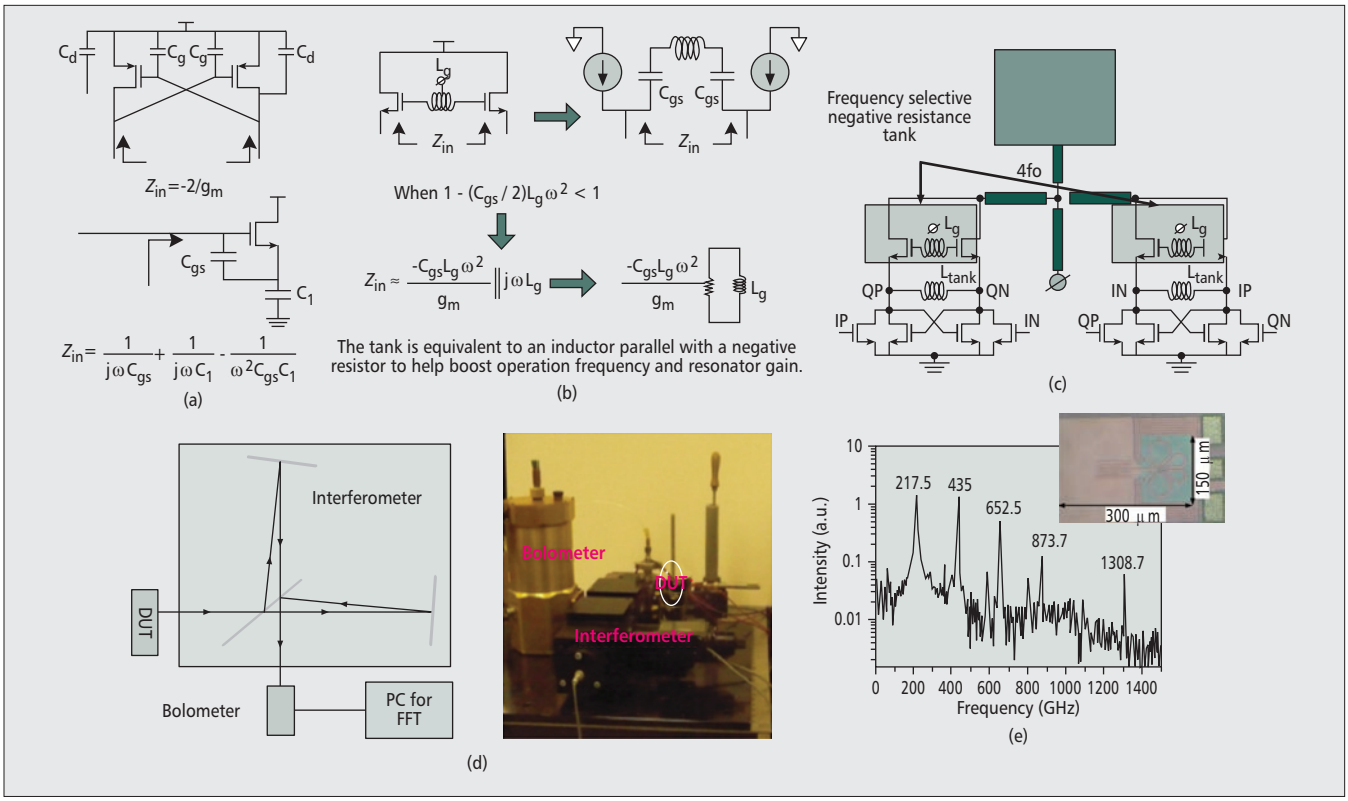

Figure 4. a) Conventional negative resistance generation mechanisms that introduce extra capacitance; b) the working mechanism of the proposed frequency-selective negative resistance tank; c) the complete terahertz generation circuit with fourth order harmonic enhancement by double push-pull scheme; d) the Michaelson interferometer measurement setup; e) the measured output spectrum with die photo shown in inset.

quency [7]. The key challenges in conventional signal generation are the small signal gain due to low quality factor Q passives and low operating frequency due to parasitic loading. To boost the signal, an extra negative resistance can be used to compensate for the tank loss. However, the well-known negative resistance generation circuits introduce extra capacitance loading. For example, Fig. 4a shows two negative resistance generation circuits: cross-coupled pairs with negative resistance of $-1/(2g_m)$, and the gate input impedance with source capacitor degeneration producing negative resistance of $-1/(\omega^2 C_{gs} C_1)$. The extra capacitance (e.g., C_g, C_d, C_{gs}, and C_1) lowers the overall operating frequency and defeats the purpose of ultra-high-frequency generation.

Therefore, the goal is to identify a circuit that can generate a negative resistance while not presenting extra capacitance. Figure 4b is the circuit configuration that satisfies this requirement. The input impedance looking into the source, Z_{in} presented in Fig. 4b, demonstrates negative resistance in a selected frequency range. When the operating frequency is high enough to satisfy $1 - (C_{gs}/2)L_g\omega^2 < 1$, the equivalent circuit can be presented as an inductor L_g parallel with a negative resistance of $-C_{gs}L_g\omega^2/g_m$, as shown in Fig. 4b. This feature exactly satisfies the requirement of generating negative resistance without adding extra loading capacitance. Moreover, this scheme provides extra inductance, which further boosts operating frequency when parallel with a tank. Because this situation only holds for a certain frequency range, it is called frequency-selective negative resistance (FSNR) tank. With the discovery of this resonant tank, a terahertz signal generator circuit can be formed by combining this tank with a conventional cross-coupled pair, shown in Fig. 4c. The fundamental resonant frequency increases to

$$1/\sqrt{C(L_g // L_{tank})},$$

which is higher than the original resonant frequency from the cross-coupled pair,

$$1/\sqrt{CL_{tank}}.$$

To further boost the operating frequency, a double push-pull structure to generate the 4th order harmonic is adopted as shown in Fig. 4c. There are a few unique features of this architecture. First, combining an FSNR tank with the primary tank boosts the resonant frequency higher than either of the individual tank resonant frequencies. This allows for larger inductance values of L_{tank} and L_g, which not only generate a larger tank impedance for higher gain, but also render on-chip inductor design more flexible and reproducible than typical super-small inductors in the terahertz frequency regime. Second, the FSNR tank provides negative resistance at the desired high resonant frequency, which ensures a high operating frequency. The additional negative resistance also relaxes a high transconductance g_m requirement of the core circuit device, which allows a

smaller device size to further boost the operation frequency and reduce the power consumption. Third, a vertical stacking structure allows a higher supply voltage without reliability concerns and increasing signal swing.

Conventional electronic apparatus is not suitable for identifying high order harmonic frequencies in the terahertz frequency range. To overcome this obstacle, a Michelson interferometer based quasi-optical measurement approach is adopted. As shown in Fig. 4d, the output signal, radiating from the vertically mounted CMOS oscillator with an on-chip patch antenna, is detected through an interferometer followed by a bolometer. The signal spectrum is then recovered through fast Fourier transform (FFT), as shown in Fig. 4e. The fundamental frequency is about 217 GHz, which is larger than 65 nm CMOS cutoff frequency, the unit current gain frequency f_T, which is about 200 GHz. The inset shows the chip photo, occupying a 0.045 mm^2 area.

The wireless communication feature of 140 GHz transceiver-based short-distance communications, shown in Fig. 3, naturally suffers from large path losses and channel interferences. In addition, the loss increases quadratically vs. the operating frequency, and is thus not favorable for frequency up-scaling. Moreover, the severe interference caused by channel crosstalk constrains simultaneous multiple channel communications and is hard to adopt in dense interconnect channel scenarios. Therefore, a wired communication channel with low loss is desired for interconnect, especially TI.

Passive Channel Design Challenges and Example

Low-loss terahertz channels have been studied extensively based on a variety of materials [15], including silicon ribbons [12], plastic ribbons and fibers [13], and so on. The silicon ribbon has demonstrated < 1 dB/m loss at frequencies up to 1 THz [12]. To reduce channel size, a subwavelength plastic fiber, with 200 mm diameter at 0.3 THz, has been demonstrated with an attenuation factor of < 0.01 cm^{-1} [13]. These low-loss channels alleviate link budget to reduce the requirements of transmitter output power and receiver noise performances to enable TI in silicon processes. However, planar silicon process compatible terahertz channels and couplers have not been investigated in previous literature. To further investigate the feasibility of terahertz channels compatible with planar silicon processes, we have designed a micro-machined dielectric waveguide-based terahertz channel.

One key specification of the TI channel is the loss. There are three major loss categories: material loss, radiation loss, and mode conversion loss. All these losses must be minimized to achieve overall low-loss performance. Material loss reduction is straightforward, which means the use of low-loss materials such as silicon [12], quartz [15], and plastic [13]. We have used high resistivity (HR) silicon as the channel material for low material loss and compatibility with mainstream silicon processes [16]. To implement the inter-/intra-chip interconnect for planar processes, a bending structure is the most intuitive and convenient approach. However, the bending structure may introduce additional losses due to radiation and mode conversion. The two losses are determined not only by the material, but also by the channel dimensions, and therefore present design trade-offs when optimizing the two loss mechanisms.

To support small form factor and package footprint, small bending structure is preferred. When the bending radius is smaller than the signal wavelength, radiation loss may exist because the portion of electromagnetic (EM) waves leaking into the air cannot preserve the phase front after bending. To minimize the radiation loss, the wave must satisfy two requirements:

- Total internal reflection (TIR) to reduce the radiation loss
- Transverse resonance condition to ensure constructive interference with itself [17]

Figure 5a illustrates the analysis of EM wave propagating in a planar slab dielectric waveguide, where n_1 and n_2 are the refractive indices, and k_1 and k_2 are the wave numbers of the dielectric waveguide material and the surrounding material, respectively. The waveguide has finite size along the x axis within the $\pm d$ region and is infinite along the y axis. The electromagnetic wave propagates along the z axis. When the incident angle θ is larger than the critical angle $\theta_c = \sin^{-1}(n_2/n_1)$ to satisfy the TIR requirement, the boundary condition determines the electrical fields on the three regions: upper, down, and inside the slab, can be represented as [17]

$$\begin{aligned} E &= E_o exp(a_x(x+d))cos(\beta z - \omega t)\ x < -d \\ E &= E_o exp(-a_x(x-d))cos(\beta z - \omega t)\ x > d \\ E &= E_o cos(\beta z - \omega t)\ -d \le x \le d \end{aligned} \quad (1)$$

where $a_x = (k_1^2 \sin^2\theta - k_2^2)^{0.5}$. Therefore, the waves outside the central waveguide area ($-d < x < d$) are evanescent waves with attenuation factor a_x. Besides the TIR requirement with the incident angle larger than the critical angle θ_C, transverse resonance needs to form constructive interference for low loss propagation. To achieve that, the phase difference needs to be an integer number of 2π after two reflections, expressed as

$$\begin{aligned} &2k_1 d cos\theta - 2\phi_r = 2\pi m \\ &tan(\phi_r/2) = (n_1^2 sin\theta^2 - n_2^2)^{0.5}/(n_1 cos\theta) \\ &\text{Therefore}: tan\left(\frac{dk_1 cos\theta}{2} - \frac{\pi}{2}m\right) \\ &\qquad = (n_1^2 sin\theta^2 - n_2^2)^{0.5}/(n_1 cos\theta) \end{aligned} \quad (2)$$

where m is the mode number starting from 0 for the fundamental mode, ϕ_r is the phase difference generated for each reflection, determined by the incident angle and dielectric constants of the two interfacing materials. Equation 2 determines the possible propagation modes supportable by the slab dielectric waveguide. FigThe graphic at the top of Fig. 5b illustrates the EM distribution along a bending dielectric waveguide channel,

The severe interference caused by channel crosstalk constrains simultaneous multiple channel communications, and it is hard to adopt in dense interconnect channel scenarios. Therefore, a wired communication channel with low loss is desired for interconnect, especially TI.

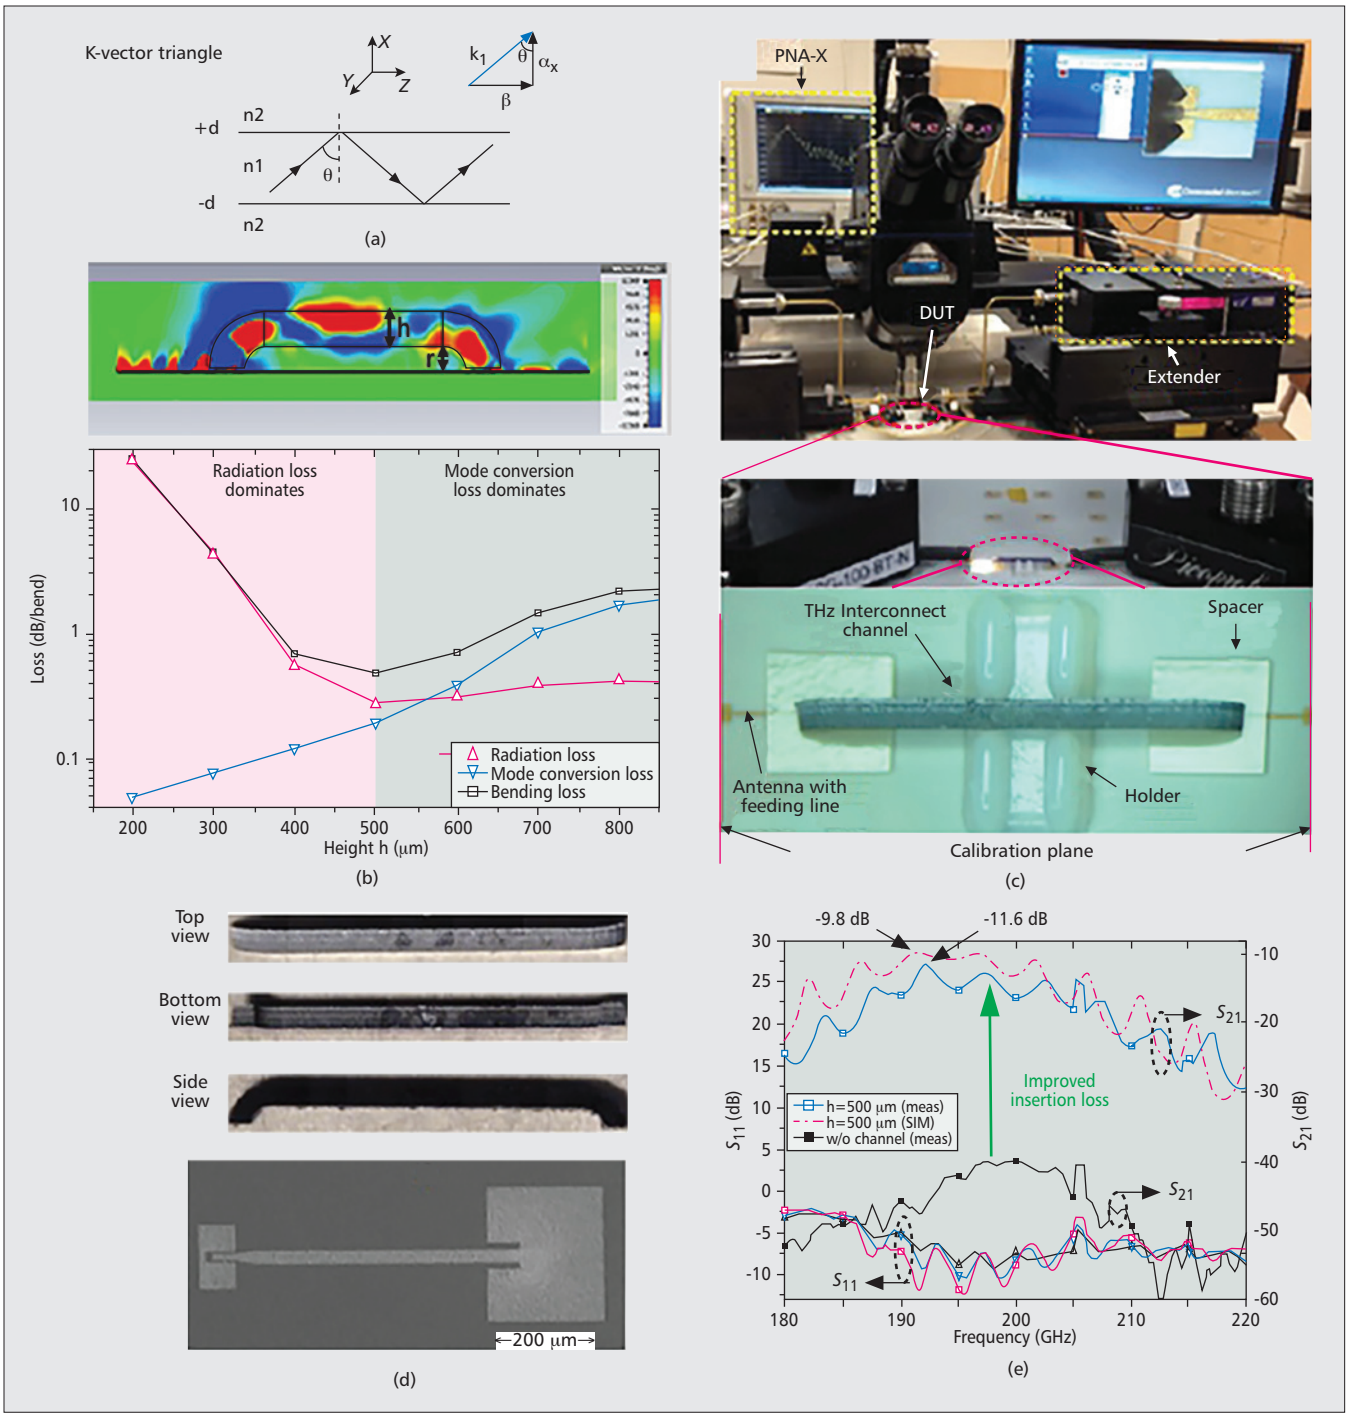

Figure 5. a) Analysis with total internal reflection and electrical field distribution in different media regions of a slab dielectric waveguide; b) simulated electrical fields and propagation waves along the channel and the simulated loss vs. different channel height; c) photograph of the test bench and a zoomed-in picture of the channel sitting in the spacer; d) photographs of the channel from different perspectives and the SEM photograph of the patch antenna based channel coupler; e) the measured and simulated S-parameters of the design silicon dielectric waveguide channel.

which indicates two major loss mechanisms: radiation loss and mode conversion loss. Increasing the channel dimension helps to confine the electrical field inside the waveguide channel to reduce the radiation loss. However, when the channel dimension becomes large enough to allow multiple modes to exist at the operating frequency, mode conversion loss will be induced. The issue is exacerbated by the bending structure. Figure 5b presents the bending loss, consisting of both radiation loss and mode conversion loss, vs. channel height h with a fixed radius r of 300 μ and a fixed width w of 300 μm at 200 GHz of a designed silicon dielectric waveguide [16]. When h is less than 500 μm, radiation loss dominates. Smaller height leads to a larger portion of the waves leaking into the air and causes larger radiation loss as shown by the curve with up triangles. When h is larger than 500 μm, higher order modes may be excited and propa-

gate, causing increasing mode conversion loss as shown by the curve with down triangles. The total loss is plotted as the curve with squares. The minimum loss of 0.3 dB occurs with h around 500 μm. To realize the signal coupling between the transceiver and the channel, this prototype adopts patch-antenna-based coupler structure to utilize its broadside radiation pattern.

Figure 5c illustrates the measurement setup and zoomed-in figure of the silicon dielectric channel [16]. Figure 5d presents the pictures of the fabricated channel and coupler. Figure 5e shows the comparison of measurement results between the cases with and without TI, which indicates dramatically improved insertion loss. The simulation results are also included for comparison, which confirms the consistency. The minimum insertion loss for h of 500 μm is about 11.6 dB. This loss can be significantly reduced due to the usage of low conductivity titanium on the feeding line and coupler to increase the fabrication stiffness, which is not needed in real interconnect scenarios. Without low-conductivity titanium, the insertion loss improves about 5 dB. Although the channel is only measured around 200 GHz due to equipment limitations, the design methodology and structure are ready to apply into the terahertz frequency range due to the benefits at higher operating frequencies. First, higher operating frequency increases signal bandwidth to support higher throughput; second, higher operating frequency reduces the device and channel size to further increase bandwidth density; third, higher frequency introduces better coupler efficiency for patch structure due to the small gaps between the top plane and the bottom ground plane in CMOS processes.

Circuit and System Non-Ideality Effects

TI performance is affected by practical circuit and system parameters and non-idealities, such as channel dispersion, circuit nonlinearity, noise, and system bandwidth. Large bandwidth facilitates high data rates. However, increasing bandwidth has realistic concerns; even the carrier frequency, channel bandwidth, and circuit bandwidth have sufficient capabilities. This is mainly because wide bandwidth results in a large integrated noise and is vulnerable to channel dispersion and circuit nonlinearity, thus degrading system SNR. When the bandwidth exceeds 100 GHz, it is extremely challenging to support the required SNR for ultra-low BER (e.g., 1×10^{-15}). Therefore, to support wide bandwidth, logically shared multi-channels through the same physical link, with relatively small bandwidth for each channel, should be adopted. The schemes to support multiple channels can be achieved through frequency-division multiplexing or other multiplexing approaches. When there are multiple channels, the interference due to channel crosstalk also degrades SNR. In addition, for all inter-/intra-chip communications, multiple parallel physical links are required to satisfy the overall extremely large aggregate data rate, such as petabyte or exabyte.

The scenario of multiple physical links and multiple logical channels sharing physical links is illustrated in Fig. 6. Due to the physically and logically adjacent channels, there is interference through channel coupling and crosstalk. A simplified scenario of channel coupling and crosstalk is illustrated in Fig. 6a. Assume a target channel, N, is coupled by eight adjacent channels. Two physically collocated channels, $N - 1$ and $N + 1$, generate interference. On the two adjacent physical links, there are also three corresponding channels close to the target channel with interference generation. Two factors — filtering, the suppression between adjacent logical channels, and crosstalk, the coupling between adjacent physical links — are critical to TI performance with multiple channel links. The coupling between adjacent channels and the target channel is illustrated in Fig. 6a. Figure 6b presents the normalized energy efficiency vs. physical link crosstalk given the filtering suppression is fixed at 40 dB. The target SNR is 18 dB to achieve a BER of 1×10^{-15}. The energy efficiency gets worse with a larger crosstalk. This is due to the fact that more crosstalk results in more interference, which therefore requires a larger signal power to maintain the same SNR, thus degrading the energy efficiency. Figure 6c illustrates the EVM and BER vs. physical link crosstalk given the same assumption of 40 dB filtering suppression. Assume initial EVM without crosstalk and channel coupling of –46 dB. The crosstalk degrades both the EVM and BER. When the crosstalk is worse than –20 dB, EVM degrades to –16.9 dB, which can no longer satisfy the BER requirement of 1×10^{-15}.

This analysis is a simplified scenario. In practical systems, there are more effects to consider, such as out-of-band intermodulation and adjacent channel spectrum regrowth. Therefore, the overall SNR and BER with practical circuit and system specifications can be presented as

$$\frac{S}{N} = \frac{P_S}{N_n \times BW + P_{nl\_self} + P_{CC}}; \quad BER = \frac{1}{2} erfc\left(\sqrt{\frac{SNR}{2}}\right) \tag{3}$$

where N_n is the noise spectrum density of a channel, P_{nl_self} is the interference from the target channel itself, such as dispersion and circuit nonlinearity, and P_{cc} is the interference from the adjacent channels, including channel crosstalk, out-of-band intermodulation, and spectrum regrowth due to nonlinearity. Good SNR and BER demand high suppression of all the non-idealities.

Conclusions

This article proposes and presents a new application for the terahertz spectrum, THz Interconnect, to potentially address inter-/intra-chip interconnect issues by leveraging the advantages from both the electronics and optics sides. TI complements electrical interconnect and optical interconnect to focus on the communication dis-

The development of THz Interconnect requires advancements from a large variety of research fields, including low-power and high-frequency active circuits and systems; low-loss, small-size, and low-dispersion channels; low-loss coupling between actives and passives; small crosstalk multiplexing techniques; and novel ideas in active and passive co-design to further boost energy efficiency.

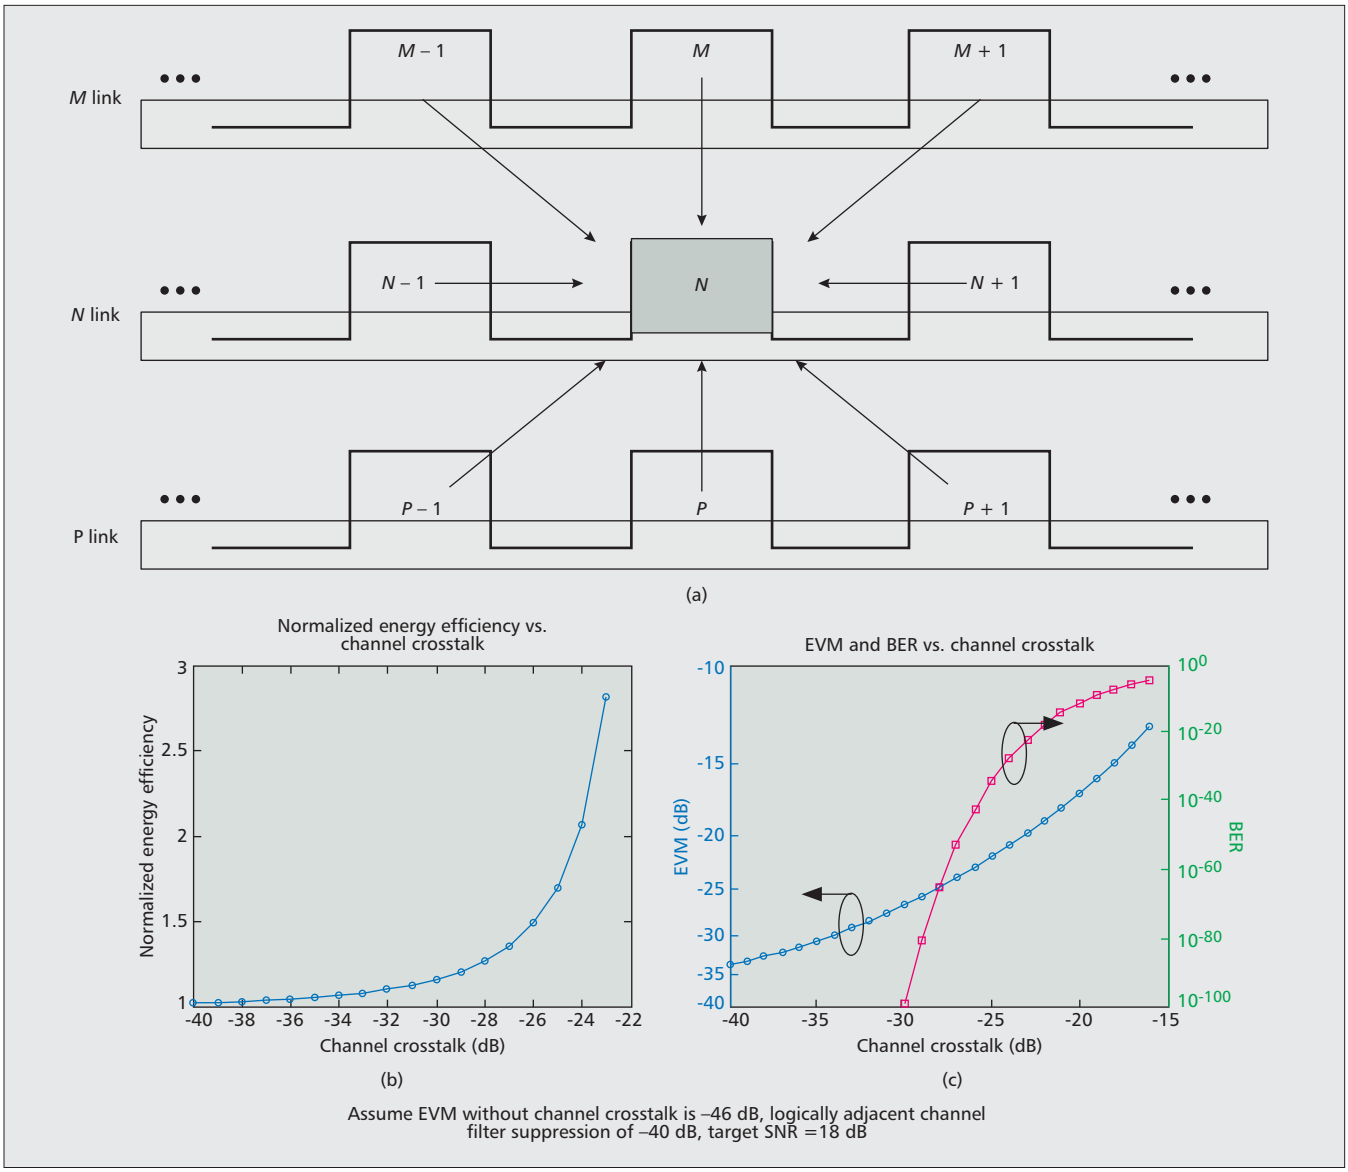

Figure 6. a) The scenario to analyze crosstalk effect among physically and logically adjacent channels; b) normalized energy efficiency; c) EVM and BER vs. channel crosstalk.

tance between 1 mm and 10 cm, which is the "last centimeter" region for inter-/intra-chip communications. The development of TI requires advancements in a large variety of research fields, including low-power and high-frequency active circuits and systems; low-loss, small-size, and low-dispersion channel; low-loss coupling between actives and passives; small crosstalk multiplexing techniques; and novel ideas in active and passive co-design to further boost energy efficiency. With the developments of the corresponding fields, we are optimistic to witness the long-standing interconnect gap to be closed.

Acknowledgments

The author is grateful to Bo Yu, Yunhao Liu, Dr. Xiaoguang Liu and Dr. Neville Luhmann Jr. for their great contributions on THz interconnect channel development and would like to thank NSF and Dr. George Haddad for funding support.

References

[1] Workshop on Emerging Technologies for Interconnects (WETI): http:// http://weti.cs.ohiou.edu/
[2] http://www.top500.org
[3] X. Zheng and A. V. Krishnamoorthy, "Si Photonics Technology for Future Optical Interconnection," *Proc. Commun. and Photonics Conf. and Expo*, 2011, ACP, Asia.
[4] A. V. Krishnamoorthy *et al.*, " Progress in Low-Power Switched Optical Interconnects," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, Mar./Apr. 2011, pp. 357–74.
[5] Y. Kim *et al.*, "Analysis of Noncoherent ASK Modulation-Based RF-Interconnect for Memory Interface," *IEEE J. Emerging Sel. Topics Circuits Sys.*, vol. 2, no. 2, June 2012.
[6] C. L. Schow, "Low-Energy VCSEL links," *Proc. IEEE Optical Interconnects Conf.*, 2012, pp. 40–41.
[7] Q. J. Gu *et al.*, "CMOS THz Generator with Frequency Selective Negative Resistance Tank," *IEEE Trans. Terahertz Science Technology*, vol. 2, no. 2, Mar. 2012.
[8] U. R. Pfeiffer *et al.*, "A 0.53THz Reconfigurable Source Array with Up to 1 mW Radiated Power for Terahertz Imaging Applications in 0.13 mm SiGe BiCMOS," *Proc. IEEE Int'l Solid-State Circuits Conf.*, 2014, pp. 256–57.
[9] Y. Tousi and E. Afshari, "A Scalable THz 2D Phased Array with +17 dBm of EIRP at 338 GHz in 65 nm Bulk CMOS," *Proc. IEEE Int'l Solid-State Circuits Conf.*, 2014, pp. 258–59.

[10] A. Tang and M.-C. F. Chang, "Inter-Modulated Regenerative CMOS Receivers Operating at 349 and 495 GHz for THz Imaging Applications," *IEEE Trans. Terahertz Science and Technology*, vol. 3, no. 2, Mar. 2013.
[11] R. Han *et al.*, "280GHz and 860GHz Image Sensors Using Schottky Barrier Diodes in 0.13 mm Digital CMOS," *Proc. IEEE Int'l Solid-State Circuits Conf.*, 2012, pp. 254–55.
[12] C. Yeh, F. Shimabukuro, and P. H. Siegel, "Low-Loss Terahertz Ribbon Waveguides," *Applied Optics*, vol. 44, no. 28, Oct. 2005, pp. 5937–46.
[13 L.-J. Chen *et al.*, "Low-Loss Sub Wavelength Plastic Fiber for Terahertz Waveguiding," *Optics Letters*, vol. 31, no. 3, 2006, pp. 308–10.
[14] Z. Xu *et al.*, "Integrated D-Band Transmitter and Receiver for Wireless Data Communication in 65 nm CMOS," *Analog Integrated Circuits and Signal Processing*, vol. 81, no. 2, Nov. 2014.
[15] R. Piesiewicz *et al.*, "Properties of Building and Plastic Materials in the Thz Range," *Int'l J. Infrared and Millimeter Waves*, vol. 28, no. 5, 2007, pp. 363–71.
[16 B. Yu *et al.*, "Micromachined Sub-THz Interconnect Channels for Planar Silicon Processes," *2014 IEEE Int'l Microwave Symp,.*
[17] J.-M. Liu, *Photonic Devices*, Cambridge Univ. Press, 2005.

BIOGRAPHY

QUN JANE GU [S'00, M'07] received B.S. and M.S. from Huazhong University of Science and Technology, Wuhan, China, in 1997 and 2000, an M.S. from the University of Iowa, Iowa City, in 2002, and a Ph.D. from the University of California, Los Angeles (UCLA(in 2007, all in electrical engineering. After graduation, she worked as a senior design engineer in the Wionics Realtek research group and a staff design engineer at AMCC on CMOS millimeter-wave and optic I/O circuits. After that, she was a postdoctoral researcher at UCLA till August 2010. From August 2010 to August 2012, she joined the University of Florida as an assistant professor. Since August 2012, she has been at the University of California, Davis as an assistant professor.

Outphasing Transmitters, Enabling Digital-Like Amplifier Operation with High Efficiency and Spectral Purity

Leo C. N. de Vreede, Mustafa Acar, David A. Calvillo-Cortes, Mark P. van der Heijden, Rosbin Wesson, Michel de Langen, and Jawad Qureshi

ABSTRACT

An overview of outphasing transmitters is given. Starting from its basic principles, we discuss its advantages and drawbacks compared to other high-efficiency amplifier configurations. Next, innovations in outphasing architecture, design, and implementation are given that overcome the difficulties related to the traditional outphasing configuration. Using the latest insights in drive schemes for input signals and quasi load-insensitive class-E output matching, very compact high-power high-performance digital-like transmitters are enabled that offer high efficiency and high spectral purity while being reconfigurable in their operating frequency.

INTRODUCTION

Modern wireless communication transmitters need to be energy-efficient when handling non-constant envelope signals. As such, they must provide improved efficiency vs. power backoff operation compared to their class-B oriented predecessors. The two most common techniques to achieve this are voltage supply and load modulation. For the first one, envelope tracking (ET) and envelope elimination and restoration (EER) are prime examples, and for the second Doherty and outphasing [1, 2]. It is the general perception that due to their lower RF complexity, ET and EER can support larger RF operating bandwidths than load modulation oriented approaches. On the other hand, since Doherty and outphasing transmitters do not need a dynamically operated DC-to-DC converter, they can offer higher instantaneous video bandwidths and higher output powers, performance parameters that are essential to base station applications. As such, currently the base station market is dominated by Doherty solutions, while ET seems to find its way into handset applications where RF configurability is more important. In contrast, outphasing as a linearity and efficiency enhancement technique, in spite of its early invention, simplicity, and inherent elegance, has not found its way yet into volume applications.

At first sight, this lack of popularity is a bit surprising because outphasing can offer some key advantages over Doherty and ET oriented transmitters:

- Higher efficiencies due to switch-mode operation of the active devices but with linear amplification
- Digital compatibility (due to phase-only control), enabling greater system integration
- Low overall system complexity

These properties make outphasing an attractive candidate for modern wireless nodes that aim for full digitalization of their TX line-up. Note that such a digital outphasing approach, where the active devices are only fully switched on/off, can offer many advantages like reduced sensitivity to drift and device degradation, which makes the amplifier operation less dependent on actual device characteristics, yielding much more straightforward pre-distortion. Furthermore, it is important to note that outphasing, in contrast to many other digital TX-like solutions, does not require any high Q-filtering to reconstruct the (analog) TX signal at its output. As such, outphasing transmitters can be made reconfigurable in their operating frequency. In spite of the benefits, outphasing transmitters are not widely used, which indicates that there are some complications in the underlying concept that need to be solved to come to practical viable implementations. To accomplish this, many modifications to the basic outphasing principle at both the system level (e.g., improved input drive profiles) and circuit level (e.g., improved output matching techniques) have been proposed to overcome its traditional drawbacks. To place these modifications in the proper context, in this article we first give the basics of outphasing transmitters, followed by a review of the drawbacks related to a straightforward implementation. Next, we discuss the proposed system/design innovations to overcome these drawbacks. We conclude this article by providing an overview on the state of the art in outphasing transmitters and the latest research directions.

Leo C. N. de Vreede and David A. Calvillo-Cortes are with Delft University of Technology.

Mustafa Acar, Mark P. van der Heijden, Robin Wesson, Michel de Langen, and Jawad Qureshi are with NXP Semiconductors.

0163-6804/15/$25.00 © 2015 IEEE

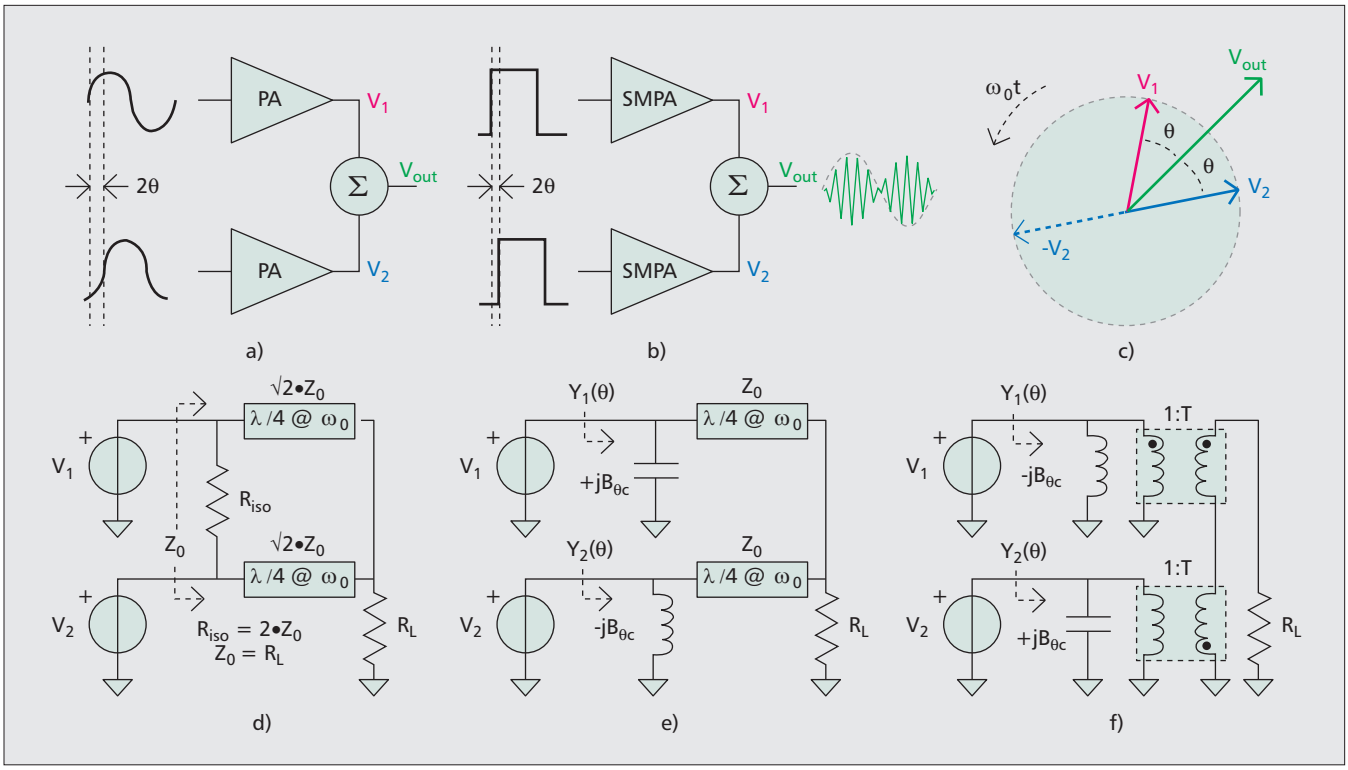

Figure 1. Outphasing amplifier topologies: a) analog; b) digital; c) signal vector representation of an outphasing amplifier. Examples of typical power combiners: d) an isolating Wilkinson power combiner; e) non-isolating transmission line-based combiner; f) transformer-based combiner.

Basics of Outphasing Amplifiers

Outphasing applied for signal amplification (Figs. 1a and 1b) was originally intended to improve the efficiency and linearity of AM broadcast transmitters [1]. In the classical two-way outphasing concept, a complex modulated signal $S_{in} = E(t) \cdot \cos(\omega_0 t + \varphi(t))$ at carrier frequency ω_0, with envelope and phase modulation $E(t)$ and $\varphi(t)$, respectively, is split into two constant-amplitude signals $S_{1,2}(t)$ having phase modulation only that are fed to the two branch amplifiers:

$$S_{1,2}(t) = V_{branch} \cdot \cos(\omega_0 t + \varphi(t) \pm \theta(t)) \qquad (1)$$

where $\theta(t) = \cos^{-1}(E(t))$ is called the *outphasing angle* (Fig. 1c) and encodes the original envelope modulation, while V_{branch} represents the constant input drive level of the branch amplifiers. The vector summation of the amplified input signals $S_{1,2}$ in Eq. 1 at the output of the branches yields the desired amplified linear replica of $S_{in}(t)$.The key advantage of this technique is that the branch amplifiers always operate at a constant power level and therefore can be of the switch-mode type. As a consequence, AM-AM and AM-PM distortion of the branch amplifiers no longer influences the overall outphasing amplification. Note that the output power combiner, which performs the vector summation of the branch amplifier signals (Fig. 1d–f), plays an essential role in reconstructing the amplified replica of the original phase and amplitude modulated input signal. Using an isolating power combiner (Fig. 1d), referred to as linear amplification using non-linear components (LINC), eliminates any interaction between the amplifier branches, and hence only experience a constant ohmic load. This makes their behavior very predictable (and linear), but unfortunately not very efficient (the curve for class-B LINC-PA in Fig. 2). To improve the efficiency in power backoff operation, Chireix [1] proposed the use of a non-isolating power combiner with compensating reactive elements (Figs. 1e and 1f). As a consequence, the effective load modulation seen by the branch amplifiers can be described by (for Fig. 1e) [4]

$$Y_{1,2}(\theta(t)) = G(\theta(t)) \pm j[B(\theta(t)) - B_{\theta_c}]$$

with

$$G(\theta(t)) = \left(\frac{2R_L}{Z_0^2}\right)\cos^2\theta(t), \qquad (2)$$

$$B(\theta(t)) = \left(\frac{2R_L}{Z_0^2}\right)\frac{\sin 2\theta(t)}{2} \text{ and } B_{\theta_c} = B(\theta_c).$$

The increase in the ohmic part, with increasing outphasing angle, yields the desired efficiency enhancement in power backoff operation if the imaginary part of this loading, Eq. 2, can be kept small. This can be accomplished by proper dimensioning of the Chireix compensating elements B_{θ_c} (Figs. 1e and 1f) when setting them equal to $B(\theta_c)$ at the desired compensation angle θ_c. Note that the load modulation for classical outphasing is infinite for both branches. In comparison, in

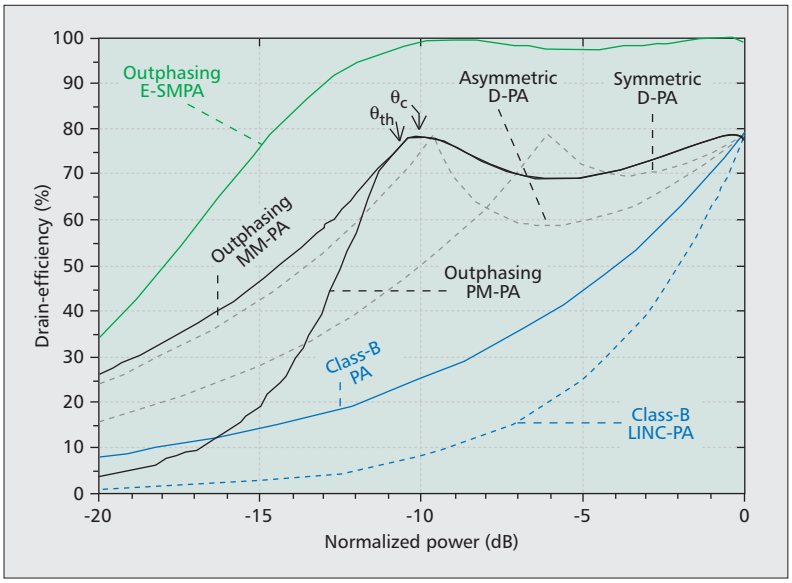

Figure 2. Ideal drain-efficiency characteristics of several high-efficiency amplifiers: (assuming class-B branch amplifier operation) two-way symmetric and asymmetric Doherty amplifiers (D-PA) [2], LINC (Fig. 1d) [2, 3], and outphasing with only phase modulation (PM-PA) [1, 2, 4] or using a combination of amplitude and phase modulation (i.e., mixed-mode outphasing) (MM-PA) [5]; and pure switch-mode outphasing with class-E branch amplifiers (E-SMPA) [6].

symmetric Doherty, the "main branch" has a load modulation factor of 2, while for the "peak branch" the load modulation is also infinite. Furthermore, in Doherty the branch amplifiers experience time-varying envelope amplitudes, while their loading (at the design frequency) is always close to ohmic as opposed to outphasing.

Assuming class-B operation for the branch amplifiers (with a maximum efficiency η_B = 78.5 percent), the resulting Chireix outphasing operation yields the classical theoretical efficiency described by [4]

$$\eta = \eta_B \cdot \cos \angle Y_{1,2} = \eta_B \frac{2\cos^2 \theta}{\sqrt{2\cos^2 \theta^2 + \sin 2\theta - \sin 2\theta_c^2}} \quad (3)$$

where $\cos \angle Y_{1,2}$ is the power factor of the combining network [7]. The expression in Eq. 3 is plotted in Fig. 2 with the curve labeled outphasing PM-PA.

When comparing this efficiency curve with that of the ideal two-way Doherty amplifier (symmetric D-PA in Fig. 2), we see that it is relatively competitive, although it does not offer any major advantage over the efficiency performance of an asymmetric Doherty (asymmetric D-PA in Fig. 2). Closer inspection reveals that the rather steep drop in the outphasing efficiency curve beyond the second high-efficiency point in power backoff (i.e., at θ_c) is mainly responsible for this. Later on, we see how mixed-mode outphasing operation or the use of quasi load-insensitive class-E switch-mode operation for the branch amplifier output stages can help us to improve this situation.

Practical Limitations of the Basic Outphasing Concept

With the basic principle of outphasing operation reviewed, it is now time to look into the practical consequences of using pure-outphasing operation with real devices and power combining networks. Note that pure outphasing in this context refers to the situation that only phase modulation is applied in the branch amplifier paths and that any kind of amplitude modulation is omitted. As such, the branch amplifiers will work under constant envelope conditions all the time. The consequence of this is that all the output power control should result from the (pure-) outphasing action only. This constraint alone has some significant consequences:

- Output power control now results from the summation/subtraction of two large vectors. This is not a problem at high output powers (when the vectors are almost in phase and constructively add), but gives inaccuracies in deep power backoff, where the vectors are almost opposite, and small errors in their magnitude directly result in dynamic range constrains [8]. Thus, to guarantee a high dynamic range some amplitude calibration of the branch amplifiers needs to be in-place [6, 9, 10].
- Active devices in general do not work optimally under highly reactive and varying loading conditions that occur in pure outphasing operation in deep power backoff (Eqs. 2 and 3). These highly reactive loading conditions are also responsible for the steep efficiency rolloff of the outphasing PM-PA curve in Fig. 2.
- Pure switch-mode operation for the branch amplifiers typically requires that the input power remains high/constant in power backoff conditions where the overall output power is reduced. As a result, the gain drops linearly with output power (Fig. 3a), directly affecting the achievable power-added efficiency and overall system efficiency.
- The use of maximum output voltage swing for all output power levels, in combination with the high load modulation used in pure outphasing operation, makes the output very sensitive to shunt losses (e.g., output conductance of the devices or shunt losses of the output matching network/power combiner). This can severely impact the achievable efficiency in power backoff operation [5].
- Bandwidth expansion of the branch amplifier signals. The fact that in outphasing systems all the information contained in the original amplitude and phase modulation needs to be represented only as phase modulation yields a considerable expansion of the bandwidth of the constant amplitude signals to be fed to the branch amplifiers. This can also be concluded from the nonlinear operations required for generating the outphasing angle, that is, $\theta(t) = \cos^{-1}(E(t))$. This bandwidth expansion puts constraints on the signal generation as well as the video bandwidth of both branches, which needs to behave identically over frequency even with the opposite varying reactive loading conditions in Eq. 2. The extended bandwidth of the branch signals is considered to be about 10–15× the original signal's bandwidth [5].

Enhancements over the Basic Outphasing Concept

Fortunately, enhancements over the basic outphasing concept can be made at various levels; for example, at the system level by using a more advanced input-signal conditioning, or at the design level by using a more robust operating class of branch amplifiers, an improved wideband power combining network, and a better device technology.

Mixed-Mode Outphasing

When considering the input-signal generation, it can be noted from Eq. 2 that in deep power backoff operation a high value for the outphasing angle translates into a highly reactive loading condition for the branch amplifiers, yielding a steep efficiency drop in deep backoff. Consequently, constraining the upper value of the outphasing angle at backoff near the second high-efficiency point (i.e., $\theta(t)|_{\max} = \theta_c$) now also yields close to ohmic loading conditions for the active devices in deep power backoff operation, which helps to improve the efficiency in this region (Fig. 2, curve MM-PA) [5]. However, doing so requires another means to control the delivered output power beyond this outphasing angle limit. A logical way to handle this situation is to introduce some amplitude control (e.g., controlling the input signal levels of the branch amplifiers) [5]. This approach requires that the branch amplifiers have an input-to-output power dependence (e.g., class-B, Class-F, or class-E with a soft switch-on characteristic). Although this approach slightly compromises the pure switch-mode outphasing concept at least beyond the second high-efficiency power backoff point, it comes, besides the avoidance of very reactive loading conditions, with some additional key advantages. Specifically, it drastically relaxes the dynamic range constraints and the related calibration requirements on the output amplitude of the branch amplifier paths, since now we no longer have to rely on the precise subtraction of two large vectors in power backoff operation. In fact, these vectors scale now in amplitude beyond the second high-efficiency point. Furthermore, reducing the input power in deep power backoff also solves the gain collapse of pure outphasing in deep power backoff operation (Fig. 3). In addition, its lower load modulation and voltage swings at the output reduce the impact of parasitic shunt losses of the active devices and matching networks [5]. Finally, but not least, the bandwidth expansion of the signals of the branch amplifiers also seems to be significantly reduced [5]. This relaxes the signal generation and also the requirement on how identical the amplifier branches need to be over frequency under varying loading conditions. Illustrations of the pure-mode and mixed-mode outphasing operation using class-B amplifiers are given in Fig. 3 and described in detail in [5, 11]. Figures 3a–3c show the outphasing PA branches input $V_{gs1,2}(t)$ and output signals $V_{1,2}(t)$, and corresponding efficiency ($\eta_{DE} = P_{out}/P_{DC}$ and $\eta_{TOT} = P_{out}/(P_{DC} + \Sigma P_{in})$) and gain ($G_{TOT} = P_{out}/\Sigma P_{in}$) performance for three different control schemes, respectively:

- Brute force phase modulation (i.e., the traditional control scheme) (Fig. 3a)
- A refined pure-mode control leading to phase-only modulation at the combining reference plane (i.e., so-called pure outphasing) that requires both amplitude and phase input control due to the transconductance nature of class-B amplifiers (Fig. 3b)
- Mixed-mode control, which leads to both amplitude and phase modulation at the combining plane (Fig. 3c)

Enhanced Pure-Mode Outphasing Operation

Although there are many benefits in mixed-mode outphasing operation, it compromises one of the key advantages of outphasing: its direct compatibility with digital-like operation (Fig. 1b). Consequently, there have also been many efforts to overcome the traditional drawbacks related to pure-mode outphasing besides the use of mixed-mode operation. The most promising ones relate to the use of quasi load-insensitive class-E operation, and the use of advanced semiconductor device technologies and duty-cycle control.

There have also been many efforts to overcome the traditional drawbacks related to pure-mode outphasing besides the use of mixed-mode operation. The most promising ones relate to the use of quasi load-insensitive class-E operation, and the use of advanced semiconductor device technologies and duty-cycle control.

Quasi Load-Insensitive Class-E — One of the first demonstrations of the excellent efficiency performance of outphasing with class-E branch amplifiers was reported in [13]. In [13], the branch amplifiers were operated in the classical RF-choke class-E mode, and the combining network was properly adapted to allow efficient outphasing operation. However, the branch amplifiers can also be directly designed and operated in the so-called quasi load-insensitive class-E mode [6, 9, 10]. Basically, quasi load-insensitive class-E enables efficient class-E amplifier operation under load-modulated conditions without any notable penalty in efficiency performance, which is a very interesting feature for outphasing systems. Since the circuit complexity of a load-insensitive class-E amplifier is quite limited, its implementation can be relatively straightforward. Its unique high-efficiency operation under changing (ohmic) loading conditions (close to 100 percent for an ideal switching device) is due to its cleverly chosen network component values, demanding that [14]

$$q_E = 1/\omega\sqrt{L_E \cdot C_E} \approx 1.3$$

for optimum quasi load-insensitive operation. In fact, this choice determines how much the resonance frequency of the device's output capacitance C_E and its DC-feed inductance L_E deviates from the operating frequency ω. In addition, the output network can have a low-Q, which facilitates wideband operation when proper filtering of the higher harmonics is accomplished [6, 9].

When quasi load-insensitive class-E operated branch amplifiers are applied in an outphasing configuration, the efficiency vs. power backoff curve is indeed very favorable (outphasing E-SMPA in Fig. 2). Closer inspection of this result indicates that reactive loading, at high outphasing angles in pure-outphasing operation and

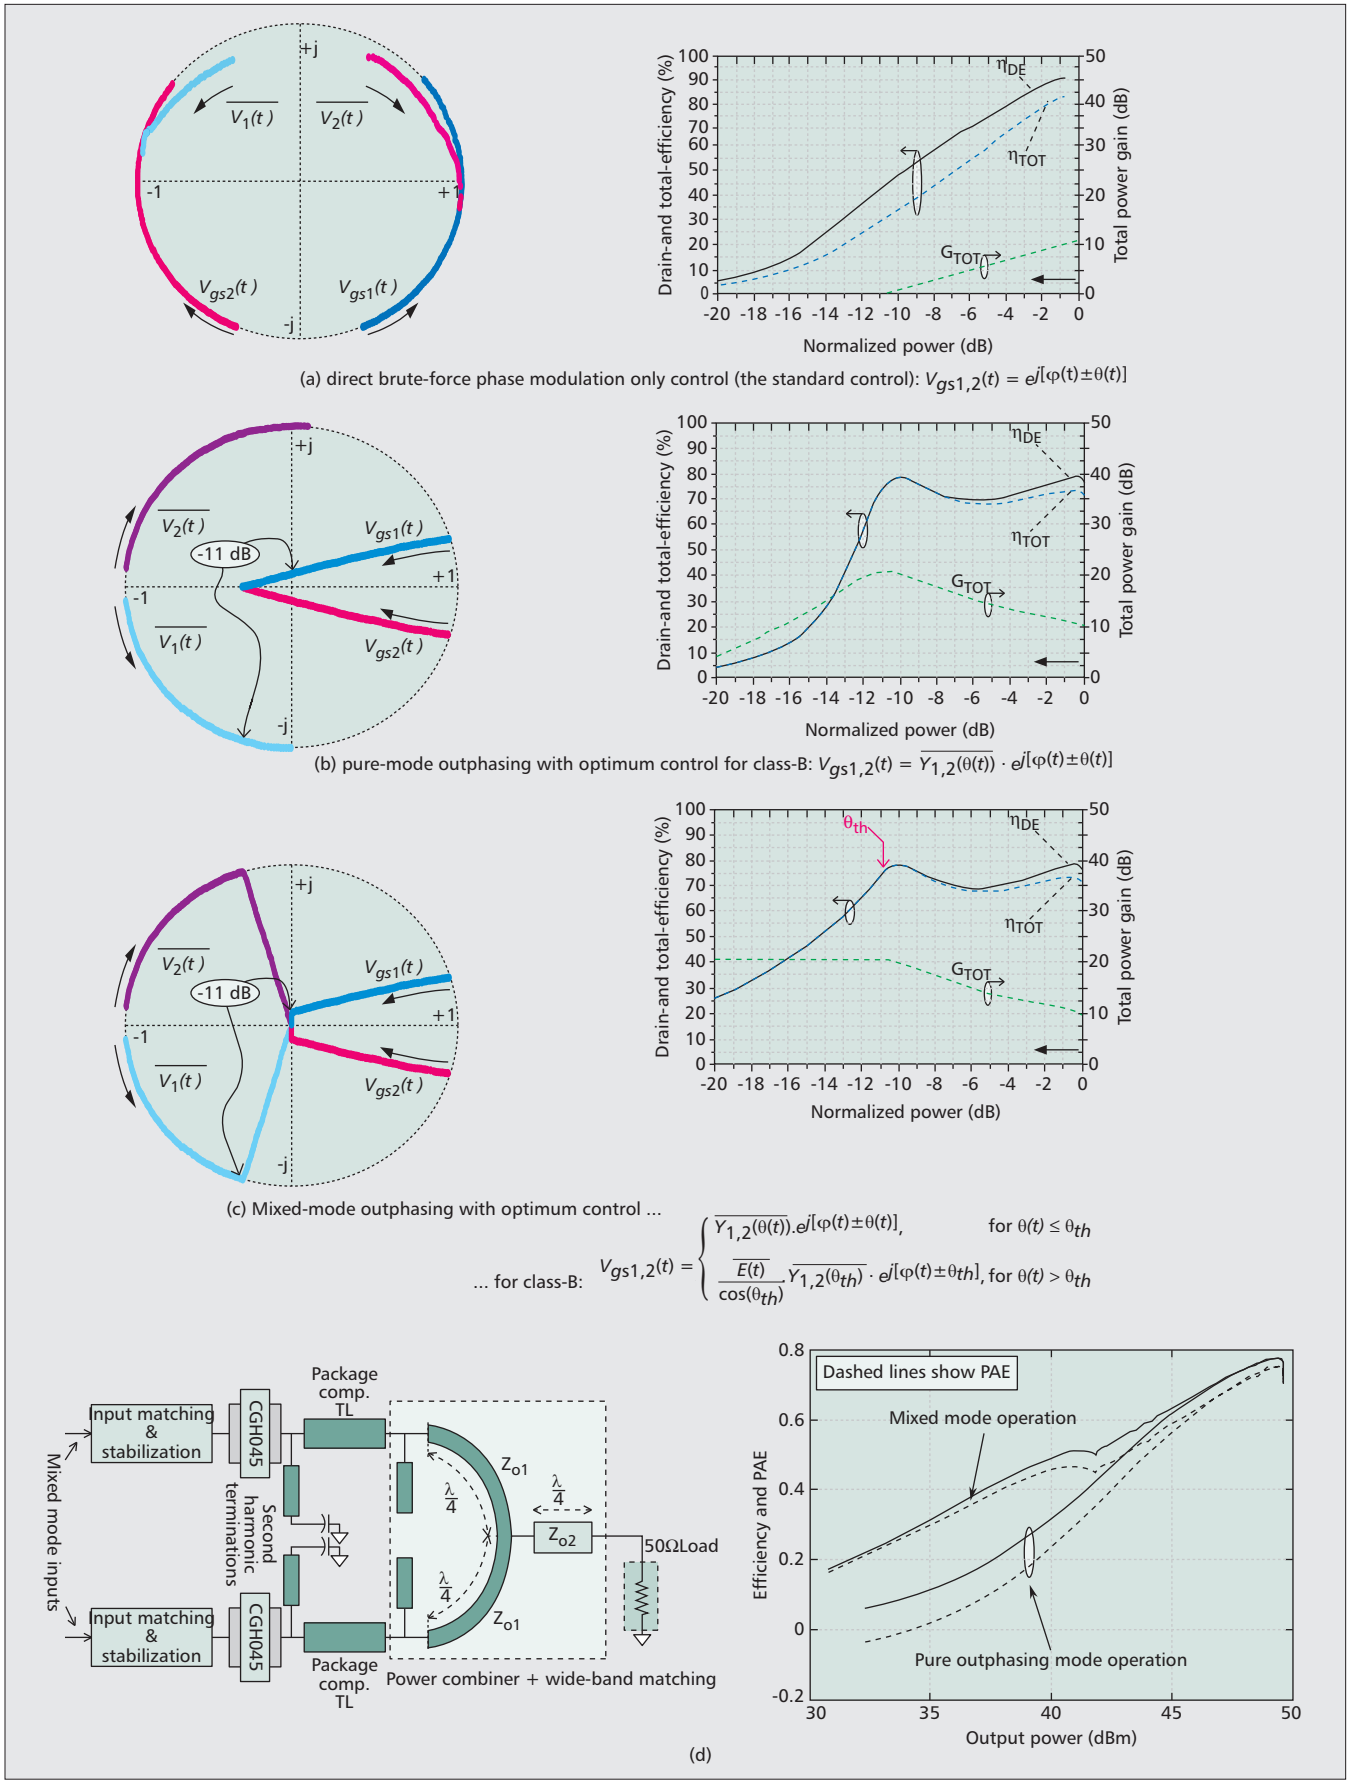

(a) direct brute-force phase modulation only control (the standard control): $V_{gs1,2}(t) = e^{j[\varphi(t)\pm\theta(t)]}$

(b) pure-mode outphasing with optimum control for class-B: $V_{gs1,2}(t) = \overline{Y_{1,2(\theta(t))}} \cdot e^{j[\varphi(t)\pm\theta(t)]}$

(c) Mixed-mode outphasing with optimum control ...

Figure 3. Outphasing using class-B branch amplifiers with different drive profiles [11, 12]: response of a) direct brute-force phase modulation only control; b) pure-mode; c) mixed-mode outphasing operation with ideal control for class-B operation. A practical class-B prototype schematic and corresponding measured results are given in d) for both pure-mode and mixed-mode operation [5].

deep power backoff, has a much less destructive impact on the active device efficiency performance [6, 9] as opposed to the use of more traditional amplifier classes like class-B. As a result, one of the most painful drawbacks of basic outphasing operation has been overcome, while the offered drain-efficiency performance vs. power backoff significantly outperforms asymmetric DPAs even in deep power backoff levels.

Use of Advanced Semiconductor Device Technologies — Although the use of load-insensitive class-E seems to solve one of the most important problems of the basic outphasing architecture, there are still other aspects that need careful attention to enable high-efficiency performance in practical situations. These include dealing with low gain at power backoff levels and sensitivity to shunt losses (output conductance). Note that the high load-modulation ratios in outphasing amplifiers make this concept much more prone to the parasitic output conductance of the active device. Hence, different than in Doherty, which typically uses rather moderate load-modulation ratios, here output losses of the active device should be kept to a minimum. A straightforward manner to overcome these drawbacks is the use of more advanced device technologies. For example, in high-power applications, GaN devices provide superior material properties and isolating substrate compared with laterally diffused metal oxide semiconductor (LDMOS) devices. Therefore, GaN is a logical technology choice to overcome the impact of output losses in pure outphasing amplifiers. Furthermore, to fully benefit from switch-mode operation, the transition from fully off to fully on, and vice versa, needs to be as short as possible. Typically, this is accomplished by overdriving the device input with a large sinewave. However, doing so significantly lowers the achievable gain, and this is highly undesirable since, as mentioned before, in pure outphasing systems the gain tends to drop linearly with the output power backoff level. One elegant way to overcome this issue is to drive the branch amplifiers with a square-like wave (e.g., by using a digital-like driver). This approach has been recently implemented in several works using tapered digital complementary metal oxide semiconductor (CMOS) buffer chains in combination with high-voltage (HV) CMOS output stages [9, 10], and it has been demonstrated that both gain and device switching are improved, maximizing the overall efficiency of the outphasing transmitter. This rather inspiring combination of GaN transistors in quasi load-insensitive class-E operation, driven by HV-CMOS technology, bridges the gap between the RF power and CMOS oriented digital world, and has already yielded impressive overall line-up efficiencies [9, 10].

Use of Duty-Cycle Control for Calibration and Operating Frequency — As earlier indicated, one of the key difficulties in using pure-outphasing operation is the dynamic range limitation which is caused by unbalances/differences in the amplitude of the branch amplifiers (recall in Eq. 2 that both amplifier branches see different imaginary loading conditions that vary opposite to the outphasing angle). To fine tune for these differences, one can slightly adjust the supply voltage, or affect the losses or matching conditions in one of the branches. However, a more elegant solution is to offset the signal input duty cycle of one branch amplifier with respect to the other. Using this method, the on/off time of the active devices in the branch amplifiers can be controlled, and hence their effective output power, without violating the purely digital-like nature of the targeted outphasing transmitter. In addition, by statically adjusting the common duty cycle of both branches, the optimum operating frequency can be shifted, so a tuned high-efficiency performance can be maintained across a large frequency range. This duty-cycle-controlled approach has been demonstrated in a CMOS-GaN class-E Chiriex outphasing amplifier operating between 1.8–2.2 GHz [10]. Highlights of this work are given in Fig. 4, including simplified schematics, an actual prototype photograph, and measurement results.

Although the use of load-insensitive class-E seems to solve one of the most important problems of the basic outphasing architecture, there are still other aspects that need careful attention to enable high-efficiency performance in practical situations.

The Power Combining Network

In addition to input signal conditioning, the power combining network in combination with the Chireix compensating elements and the actual device output matching have a big impact on the achievable bandwidth and efficiency performance of the outphasing transmitter. When aiming for a single-ended load, the power combiner also needs to fulfill a Balun function. In view of this, it can be shown that a traditional transmission-line-based power combiner (Fig. 1e) limits the achievable bandwidth of the overall outphasing amplifier [15]. Much better results can be achieved with a structure that makes use of inductively coupled elements like a transformer [6, 15] (Fig. 1f) or coupled lines like in a Marchand Balun [9, 10]. Note that although inductively coupled elements have the reputation to be lossy, this does not have to be the case when directly implemented using the bondwire connections to the active device bar. As such, the otherwise bulky power combiner can be now fully integrated inside the package itself, leading to very small-form factors at relative high power levels while exhibiting extremely low losses. This design approach has been demonstrated by the prototype highlighted in Fig. 5 and reported in detail in [6]. In such work, peak drain efficiencies for the entire outphasing prototype were measured exceeding 80 percent, while studies indicate that these numbers can still be improved by fine tuning the bondwire dimensions. From the results in Fig. 5c, the highest output power and largest backoff efficiency at 28 V and 2.3 GHz was selected for the DPD experiments with modulated signals.

Other Outphasing Related Innovations

So far, our discussion has been entirely focused on two-way outphasing transmitters that rely only on the active load modulation for their efficiency enhancement in power backoff. However,

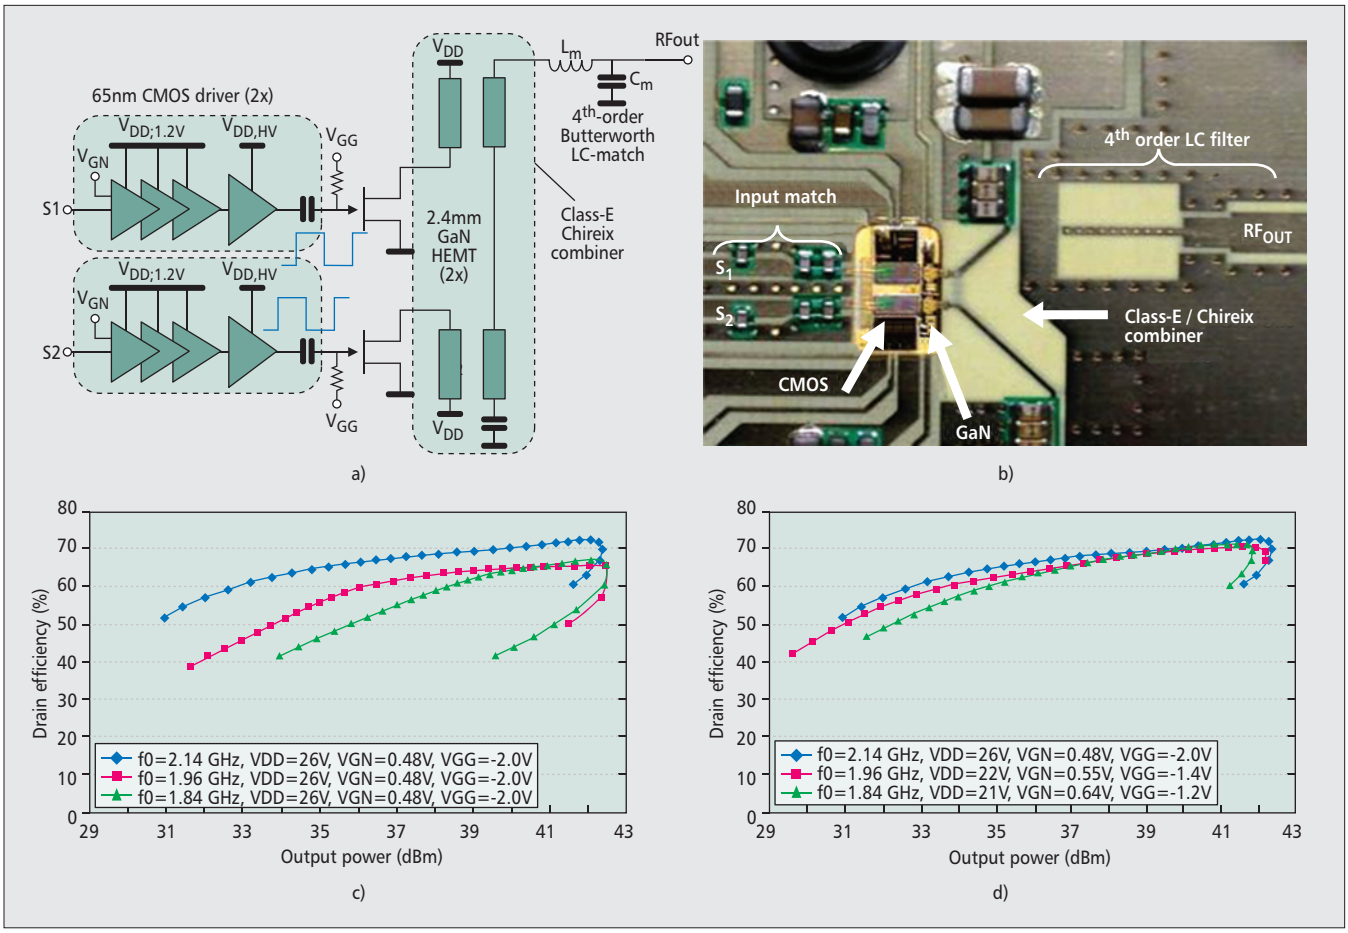

Figure 4. CMOS-driven GaN outphasing transmitter using load-insensitive class-E operation allowing duty cycle adjustment of the input signals of the branch amplifiers [10]: a) simplified schematic showing the line-up of CMOS driver and GaN output stage; b) photograph of the realized demonstrator; experimental drain efficiency performance at different frequencies c) before duty cycle adjustment; d) after duty cycle adjustment. Duty cycle is adjusted by means of the voltages V_{GN} and V_{GG} in the schematic of a).

being a generally applicable technique, outphasing is not restricted to only two-way systems. In fact, *N*-way outphasing systems are feasible [7, 16], as well as outphasing systems that utilize adaptive power combiners [17], or even outphasing systems that utilize combinations of outphasing and discrete supply modulation [3]. In all these cases, the efficiency and flexibility improvements come with a significant increase in system complexity. Consequently, it will be the actual implementation and its measured performance (not their theoretical promises) that will determine if these solutions will be successful in practice.

N-Way Outphasing Systems

One motivation to go for *N*-way outphasing systems [7, 16] is to use this extra flexibility in signal generation and power combining to limit the imaginary varying part of the load seen by the individual branch amplifiers. As such, even higher efficiencies in deep power backoff operation come into reach. However, the complexity increase is considerable, while the higher efficiencies come at even greater load-modulation ratios, making practical implementations sensitive to losses, component spread, and frequency deviations. However, its increased flexibility (and potential reconfigurability) might outweigh the increased complexity if the driving circuitry can be made fully digital.

Outphasing with Adaptive Power Combiners

In conventional two-way outphasing systems, the undesired reactance due to the outphasing action can only be compensated at two points (compensation angles). For all other angles there is a residue reactive loading that degrades the achievable efficiency performance, yielding the well-known efficiency curve of Eq. 3. However, using an adaptive power combiner or a power combiner with adaptive Chireix compensating elements (which can be varied with the envelope modulation/outphasing angle), in theory the branch amplifier loading can be kept purely ohmic, thus enabling significantly higher average efficiencies. This was demonstrated in [17] using high-Q varactors, but can also be implemented using digitally controlled solid-state switched capacitors with sufficient resolution. In addition, easy frequency adjustment seems to come into reach with this technique. Although attractive, this solution also increases

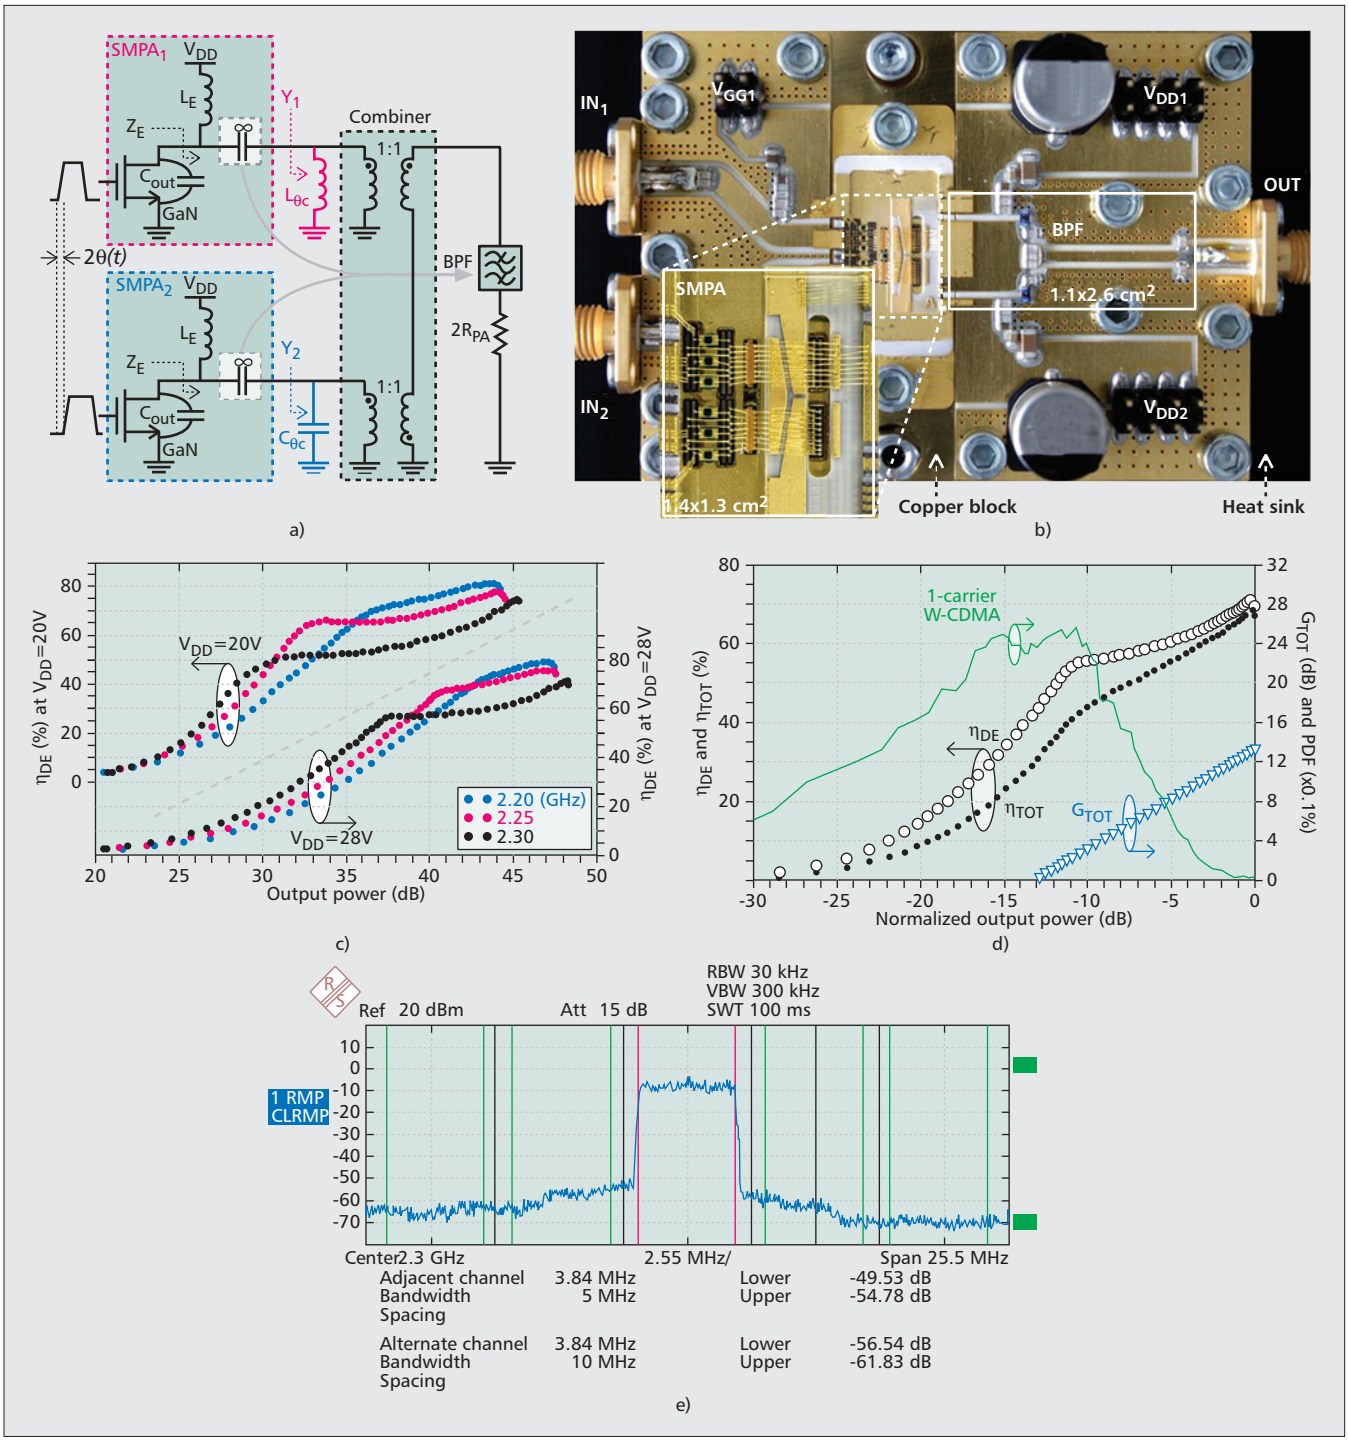

Figure 5. Package-integrated 70 W outphasing SMPA using GaN HEMT-based load-insensitive class-E branch amplifiers [6]: a) simplified schematic; b) prototype photographs; c) measured drain efficiency (η_{DE}) at different supply voltages and operating frequencies; d) measured drain efficiency (η_{DE}), total efficiency (η_{TOT}), and gain vs. normalized backoff power at V_{DD} = 28 V and 2.3 GHz, resulting in 70 W peak output power capability as used in the DPD experiments; e) resulting spectral purity meeting Third Generation Partnership Project specifications after linearization with a memoryless multi-segment polynomial DPD algorithm.

the system complexity significantly and has some practical constraints. For example, the losses in the tunable power combiner and Chireix compensating elements will have a rather dramatic impact on the achievable efficiency performance. Also, the application of this technique at high power levels might prove not to be straightforward.

Outphasing in Combination with Supply Switching

Another recently promoted technique is the combination of supply-voltage switching of the branch amplifiers with two-way outphasing using an isolating power combiner, which is referred to as multilevel outphasing [3]. By using an isolat-

Load-modulated amplifier architectures	Branch amplifier class	$\eta_{DE,avg}$	$\eta_{TOT,avg}$
		(%, simulated*)	
Pure-mode outphasing (60 %-overdrive)	Class-E	85.7	61.2
Pure-mode outphasing (onset-overdrive)		79.5	66.6
Mixed-mode outphasing	Class-B	65.9	65.8
Asymmetric (1:2) Doherty		58.8	58.7
Symmetric Doherty		50.9	50.8
Pure mode (linear) outphasing		42.7	42.7
Pure mode (saturated) outphasing		33.9	33.5
Standalone		26.4	26.4

*Assuming 20 dB power gain devices and a 10.5 dB PARW-CDMA test signal.

Table 1. Average efficiency of several PA architectures using class-B and class-E branch amplifiers [12].

ing power combiner, the problems related to the time-varying loading of the branches are eliminated. In this approach, the efficiency enhancement in power backoff operation now results from the supply-voltage adjustment of the branch amplifiers. The outphasing action itself is only used to "smoothly" connect the discrete power levels provided by the supply voltage switching. Using this approach, the difficulty of making a wideband energy-efficient DC-DC converter is omitted. By introducing a sufficient number of voltage supply levels, the achievable drain efficiency for multi-level outphasing can be made much higher than for conventional LINC transmitters [3]. Using the multi-level outphasing approach, the difficulties are now shifted from the circuit domain to the pre-distortion of this multi-level outphasing transmitter, which needs to handle the rather abrupt supply voltage changes in an appropriate and reliable manner, something that is becoming increasingly difficult with increasing modulation bandwidth. If these issues can be appropriately handled, this architecture can also offer some advantages in spite of its higher complexity.

Conclusions

As can be concluded from this overview, since its invention in 1935 and through all these years, outphasing is still a very active research topic that has gained renewed interest in the last few years thanks to technological advances in several areas. Due to the fact that outphasing allows the use of true switch-mode operation without any severe filter requirement, it provides higher efficiency potential than the currently popular Doherty amplifiers. This becomes most apparent by comparing their achievable average efficiency when assuming idealized 20 dB gain active devices with lossless power combining (Table 1 and [12]). Note that these additional gains in efficiency are very important, since they can open the road to future base station transmitters that no longer require forced air cooling, which is highly desirable from the cost, energy consumption, and maintenance points of view. To truly achieve the efficiency promise of the outphasing concept, dedicated and innovative system and circuit solutions are needed. Some of them have been discussed in this overview. Which one is your favorite depends on your attitude. Would you like to move in small steps from a Doherty-like approach to higher performance? Then take the mixed-mode approach. Are you more aggressive and aim for a fully digital approach? Then go for the pure switch-mode approach. Would you prefer something more exotic that can (at least in theory) provide you with even higher efficiency numbers, while being reconfigurable? Try duty-cycle control, adaptive combiners, or N-way outphasing. Would you like to combine outphasing with supply-voltage control? Try the multi-level outphasing approach. In the end, in spite of personal preferences, it will be the ever changing wireless market itself that will select the most cost-effective, efficient, reliable, and flexible implementation as the new favorite.

Acknowledgment

The authors would like to acknowledge Rik Jos, Fred van Rijs, and John Gajadharsing of NXP Semiconductors, Nijmegen, The Netherlands, for their many useful suggestions and guidance in related research projects. The Dutch Government, STW, NXP, and the European program Catrene are acknowledged for funding research projects on energy-efficient transmitters.

References

[1] H. Chireix, "High Power Outphasing Modulation," *Proc. Inst. Radio Engineers*, vol. 23, no. 11, 1935, p. 1370–92.

[2] S. C. Cripps, *RF Power Amplifiers for Wireless Communications*, 2nd ed., Artech House, 2006.

[3] S. Chung *et al.*, "Asymmetric Multilevel Outphasing Architecture for Multi-Standard Transmitters," *IEEE Radio Frequency Integrated Circuits Symp. Digest*, June 2009, pp. 237–40.

[4] F. H. Raab, "Efficiency of Outphasing RF Power-Amplifier Systems," *IEEE Trans. Commun.*, vol. 33, no. 10, Oct. 1985, pp. 1094–99.

[5] J. Qureshi *et al.*, "A 90-W Peak Power GaN Outphasing Amplifier with Optimum Input Signal Conditioning," *IEEE Trans. Microwave Theory and Techniques*, vol. 57, Aug. 2009, pp. 1925–35.

[6] D. A. Calvillo-Cortes *et al.*, "A Package-Integrated Chireix Outphasing RF Switch-Mode High-Power Amplifier," *IEEE Trans. Microwave Theory Techniques*, vol. 61, no. 10, Oct. 2013, pp. 3721–32.

[7] D. J. Perreault, "A New Power Combining and Outphasing Modulation System for High-Efficiency Power Amplification," *IEEE Trans. Circuits Systems I: Regular Papers*, vol. 58, no. 11, Aug. 2011.

[8] X. Zhang, L. Larson, and P. Asbeck, *Design of Linear RF Outphasing Power Amplifiers*, Artech House, 2003.

[9] M. P. van der Heijden, M. Acar, J. Vromans, and D. A. Calvillo-Cortes, "A 19W High-Efficiency Wide-Band CMOS-GaN Class-E Chireix RF Outphasing Power Amplifier," *IEEE MTT-S Int'l. Microwave Symp. Digest*, June 2011.

[10] M. P. van der Heijden and M. Acar, "A Radio-Frequency Reconfigurable CMOS-GaN Class-E Chireix Power Amplifier," *IEEE MTT-S Int'l. Microwave Symp. Digest*, June 2014.

[11] D. A. Calvillo-Cortes and L. C. N. de Vreede, "Analysis of Pure- and Mixed-Mode Class-B Outphasing Amplifiers," *IEEE Latin American Symp. Circuits and Systems*, Feb. 2014.

[12] D. A. Calvillo-Cortes, *Energy Efficient and Compact RF High-Power Amplifiers*, Ph.D. dissertation, Delft Univ. Technology, Nov. 2014.

[13] R. Beltran, F. H. Raab, and A. Velazquez, "HF Outphasing Transmitter Using Class-E Power Amplifiers," *IEEE MTT-S Int'l Microwave Symp. Digest*, June 2009, pp. 757–60.
[14] M. Acar, A. J. Annema, and B. Nauta, "Generalized Analytical Design Equations for Variable Slope Class-E Power Amplifiers," *Proc. IEEE Int'l Conf. Electronics, Circuits and Systems*, Dec. 2006, pp. 431–34.
[15] M. C. van Schie *et al*., "Analysis and Design of a Wideband High Efficiency CMOS Outphasing Amplifier," *IEEE Radio Frequency Integrated Circuits Symp. Digest*, May 2010, pp. 399–402.
[16] T. Barton, J. Dawson, and D. J. Perreault, "Experimental Validation of a Four-Way Outphasing Combiner for Microwave Power Amplification," *IEEE Microwave and Wireless Components Letters*, vol. 23, Jan. 2013, pp. 28–30.
[17] J. Qureshi *et al.*, "A Highly Efficient Chireix Amplifier Using Adaptive Power Combining," *IEEE MTT-S Int'l Microwave Symp. Digest*, June 2008, pp. 759–62.

Biographies

Leo C. N. de Vreede [M'01-SM'04] received his Ph.D. degree (with honors) from Delft University of Technology, where he became an associate professor in 1999. He is co-founder of Anteverta and DitIQ, (co)recipient of the IEEE Microwave prize 2008, the Else Kooi prize 2010, and the Dow Energy Award 2011. He has (co)authored more than 100 IEEE papers, and holds several patents and best paper awards. His interests include device optimization, characterization, and circuit concepts for wireless systems.

Mustafa Acar received his B.S. degree (with honors) from Middle East Technical University, Ankara, Turkey, in 2001, and both his M.S (high honor) and Ph.D. degrees in 2003 and 2011, respectively, from the University of Twente, Enschede, The Netherlands. Since 2007, he has been working on advanced driver and PA systems at NXP Semiconductors. Currently, he is leading the switch-mode PA activity of NXP Semiconductors.

David Angel Calvillo-Cortes received his B.Eng. degree (best student award) in communications and electronics engineering from the Universidad de Guadalajara, Mexico, in 2005, and both his M.Sc. (cum laude) and Ph.D. degrees in electrical engineering from Delft University of Technology in 2009 and 2014, respectively. Prior joining Qualcomm Inc. in late 2014, he was with NXP Semiconductors (2008–2009) and Continental Automotive (2005–2007).

Mark P. van der Heijden received his B.Sc. degree (with honors) in electrical engineering from the Haagse Hogeschool, The Hague, The Netherlands, in 1998, and his P.D.Eng. and Ph.D degrees (with honors) in electrical engineering from Delft University of Technology in 2000 and 2005, respectively. He works at NXP Semiconductors, where he is currently involved in high-efficiency RF and millimeter-wave power amplifier design and advanced transmitter architectures for wireless communication systems.

Robin Wesson received his First Class Degree in physics and electronic engineering from the Loughborough University of Technology, United Kingdom, in 1996. He is currently a system architect with NXP Nijmegen, The Netherlands. He possesses 17 years of industrial RF system, circuit, and product design experience with Motorola, Sagentia, IP Wireless, and Axis NT, among others. He is currently involved with signal processing and linearization techniques for switch-mode outphasing.

Michel T. W. de Langen received a M.Sc. degree in physics from the University of Twente in 1986. He works at NXP Semiconductors (formerly Philip Semiconductors) in The Netherlands. He was responsible for the introduction of new package concepts such as BGA, multi-chip modules, and CMOS imaging packages. Currently he is working on assembly and package research for RF power devices.

J. H. Qureshi received his B.S. degree from the University of Engineering and Technology Taxila, Pakistan, in 2000, and his Master's and Ph.D. degrees in electrical engineering from Delft University of Technology in 2006 and 2012, respectively. In 2010 he joined NXP Semiconductor, where he is currently working on advanced high-efficiency wideband power amplifiers and transmitter architectures for future base station and RF power applications.

DESIGN AND IMPLEMENTATION



Vijay K. Gurbani *Salvatore Loreto* *Ravi Subrahmanyan*

This issue presents two papers related to different aspects of application servers in communication networks.

Traditional communication system hardware and software architectures were based on resource intensive implementations that separated the control and data planes, and were built using special-purpose ASICs and hardware platforms, and proprietary operating systems. The ongoing miniaturization of communication systems, driven particularly by the demands placed by the explosion of wireless mobile data, requires new paradigms to reduce cost, size, and power while increasing capacity and scalability.

The article by Mohammed Khawer presents design strategies for application server design in next-generation operating systems. Based on an open source OS and a powerful multicore processor with hardware acceleration, the article describes design considerations for a system that combines the control and data planes within a single processor to reduce system cost of materials. As future system designs delve further into this relatively new territory for communication system implementation, the considerations explored in this article will serve as a guidepost for the implementation of these systems.

Multiuser conferencing systems with multimedia (audio, video, etc.) have become hugely popular. In particular, the number of participants associated with such events has been increasing exponentially in recent years, enabled by network capacity and cheap hardware. The Internet Engineering Task Force (IETF) has worked on frameworks for conferencing for over 10 years (e.g., the XCON framework). An important aspect of these is the Floor Control Protocol (FCP), which manages access to shared resources by multiple participants within such a conference.

The article by S. P. Romano describes an implementation of Binary Floor Control Protocol (BFCP), an IETF protocol that meets the requirements of FCP. The article gives a historical perspective of the motivation for and development of BFCP. It then describes a specific Web 2.0-based implementation, UMPIRE, of a moderation platform based on BFCP. The author also walks the reader through a sample call flow using an illustrated example. Finally, a lessons learned section gives an educational summary of the author's experiences with FCP since the initial work in the IETF.

BIOGRAPHIES

VIJAY K. GURBANI [M'98] (vkg@bell-labs.com) is a Distinguished Member of Technical Staff in wireless research at Bell Laboratories, Alcatel-Lucent. He holds a B.Sc. in computer science with a minor in mathematics and an M.Sc. in computer science, both from Bradley University, and a Ph.D. in computer science from the Illinois Institute of Technology. His current work focuses on peer-to-peer networks, Internet multimedia session protocols, and anomaly detection in such protocols. He is the author of over 50 journal papers and conference proceedings, five books, and 16 IETF RFCs. He is currently Co-Chair of the Application Layer Traffic Optimization (ALTO) Working Group in the IETF, which is designing a protocol to enable abstract network information to be provided to applications through a well defined API. He holds four patents and has 10 applications pending with the U.S. Patent Office. He is a Senior Member of the ACM, and a member of the IEEE Computer Society and Usenix.

SALVATORE LORETO [M'01, SM'09] (salvatore.loreto@ieee.org) holds an M.Sc. in computer engineering and a Ph.D. in computer networking, both from Napoli Federico II University, and an M.B.A. from Bocconi University, Milano. Currently, he works as a master researcher at Ericsson Research Finland. He has made contributions to Internet transport protocols (e.g., TCP, SCTP), signal protocols (e.g., SIP, XMPP), VoIP, Unified Communication, 3GPP IP Multimedia Subsystem (IMS), HTTP2, WebSockets, and web technologies. He is also an active contributor to the IETF, where he has coauthored several RFCs and Internet drafts. Currently he is serving within the IETF as Co-Chair of several Working Groups in the application and real-time application areas. For the IEEE Communications Society, he serves as a Design and Implementation Series Co-Editor and an Associate Technical Editor for *IEEE Communications Magazine*.

RAVI SUBRAHMANYAN [SM'97] (ravi_subrahmanyan@mac.com) is with Butterfly Network Inc., where he is involved in the architecture and design of next generation medical imaging devices. He previously worked on image sensor design at Invisage, image processing and video compression technologies at Immedia Semiconductor, and was a systems and ppplications engineering manager at National Semiconductor Corp. He was also with AMCC in Andover, Massachusetts, where he was involved in the design of communications ICs and multi-core PowerPC CPUs. He has participated in ITU-T and ATIS standards bodies on topics related to timing and synchronization in communications networks, most recently working on synchronization over packet networks. His interests are in product strategy and planning, high-speed and custom design for high-performance ASICs, image and signal processing, and communications network synchronization and architectures. He received M.S. and PhD degrees in electrical engineering from Duke University, and his B.Tech. degree (also in EE) from the Indian Institute of Technology, Bombay. He has 50 publications including conference presentations and papers in refereed journals, including invited talks at IEEE conferences such as the European Solid State Device Research Conference, International Shallow Junction Workshop, and the Workshop on Synchronization in Telecommunications Systems. He also has 20 issued or pending patents. He has served on various conference committees, including GLOBECOM and ICC, since 2008, and has been involved with ComSoc's TAOS TC since 2008, where he served as Vice-Chair, chaired the SAC-ANS track at GLOBECOM 2012, and was a presenter at the ComSoc Webinar on Next Gen Synchronization Networks held in October 2012.

Design Strategies for the Application Server Architecture/Configuration (and Its Functions) in Next-Generation Communication Systems

Mohammad R. Khawer

ABSTRACT

The market trend for next-generation communication systems has been toward miniaturization to meet the stunning ever increasing demand for wireless mobile data, leading to the need for distributed and parallel processing system configurations that are 10 times or more cost effective, flexible, high capacity, energy efficient, and scalable. Reducing cost and size while increasing capacity and scalability requires several design paradigm shifts. This article presents design strategies to meet these goals.

INTRODUCTION

The demand for mobile wireless data has been stunning in recent years, and next-generation communication systems need multiple application servers to meet the ever increasing mobile data demand for high-capacity packet processing and protocol termination.

This article presents the design strategies that are driven by requirements such as lower cost, high capacity, flexibility, scalability, and energy efficiency to reduce the capital expense (CAPEX) and recurring operational expense (OPEX) for service providers. A wireless base station is a practical real world example of next-generation communication systems that may benefit from the design strategies presented in this article.

The rest of the article is organized as follows. In the second section, we provide the rationale behind the use of an open source real-time operating system (RTOS). In the third section, we present the various multi-core processor system configurations and provide the reasoning behind the chosen configuration. In the following section, the energy efficiency consideration for the proposed design is discussed. After that, the portability and scalability of the design is presented. Following that, the critical aspect of an application server recovery mechanism is discussed. Then a discussion on testing is provided. Next, we present the state-of-the-art related work, and conclude with a summary of lessons learned.

USE OF OPEN SOURCE RTOS

Traditionally, a third party proprietary hard core RTOS such as vxWorks[1] [4, 5] made by Wind River Systems is used to serve the data plane that hosts real-time processes and threads. The control plane hosts the non-real-time processes and threads, and its performance requirement is easily met by using an open source operating system (OS) such as Linux.[2]

To reduce the overall system cost of goods sold (COGS), an open source RTOS such as symmetric multiprocessing (SMP) Linux with the PREEMPT_RT patch [5, 6] was picked to serve the data plane instead of an industry-wide popular proprietary RTOS such as vxWorks [4, 5]. The upfront development licensing cost of using open source SMP Linux with PREEMPT_RT, provided by a tier 1 OS supplier such as Wind River or MontaVista, is significantly lower than using any of their proprietary RTOSs. These companies provide extensive testing of the open source Linux code as well as patch management, and thus reduce concerns about any indemnification issues that may result from the use of open source SMP Linux with the PREEMPT_RT patch as the chosen RTOS in a commercial product. Additionally, no deployment royalty/license fee has to be paid to the OS supplier for the open source RTOS usage on each communication system or gateway that is sold and deployed commercially.

The use of the open source RTOS helps reduce the system COGS, but opens the door to serious performance challenges in meeting the stringent real-time performance requirement of the data plane application software process/threads. A number of-high performance optimization services and schemes discussed later in this article were needed to help overcome the performance issues posed by the use of a non-hard-core open source RTOS.

The author is with Alcatel-Lucent.

[1] vxWorks is a registered trademark of Wind River Systems, Inc.

[2] Linux is a trademark of Linus Torvalds.

0163-6804/15/$25.00 © 2015 IEEE

The rationale behind using a multi-core processor in our design was to consolidate, for the very first time in the industry, the control and data planes of all the configured application servers onto a single multi-core processor. Doing so significantly reduced the device count needed in the new system configuration for distributed computing and parallel processing.

Use of a Multi-Core Processor

Our architectural design required the use of a powerful multi-core processor with the processing power to support multiple high-capacity application servers. We selected Freescale Semiconductor's P4080 multi-core processor [1], mainly because it had eight e500mc processing cores, and a data path acceleration architecture (DPAA) that contained hardware acceleration engines which provide high-performance packet processing and forwarding capabilities with minimal to no impact on the processing cores. This critical packet processing offload to hardware acceleration engines means that more processing power becomes available on the cores, which may be used to increase the system capacity. The processing horsepower [1] is one of the limiting factors in determining system capacity. The combination of the more powerful e500mc processing cores and the use of hardware acceleration engines available on the multi-core processor resulted in a processing capacity increase by a factor of 1.5 [1] compared to the previous-generation single-core processor. Some of the DPAA hardware acceleration engines we employed in our design are the buffer manager (BMan), queue manager (QMan), and two frame managers (FMans). The BMan provides data buffer management in hardware for the buffers created originally by software. The QMan provides queuing and quality of service (QoS) scheduling of frames to the appropriate cores. The FMan supports multiple external Ethernet interfaces, and provides the capability to perform in-line packet parsing, general classification to enable policing, and QoS based packet distribution to central processing unit (CPU) cores for further processing.

The rationale behind using a multi-core processor in our design was to consolidate, for the very first time in the industry [7], the control and data planes of all the configured application servers onto a single multi-core processor. Doing so significantly reduced the device count needed in the new system configuration for distributed computing and parallel processing. This not only drastically reduces the COGS and size of the system, but also reduces extra hardware and software complexity, and potential failure points. The next few subsections describe the multi-core configurations that were considered and the rationale behind the chosen configuration.

Supervised Asymmetric Multi-Processing Configuration

The most obvious choice for multiple application server support using a multi-core processor is the supervised asymmetric multi-processing (S-AMP) configuration [2, 3]. A high-capacity application server requires two cores, whereas a regular-capacity application server may require just one processing core for its data plane. Each partition of the S-AMP configuration has its own dedicated OS instance, and therefore mandates the use of another system supervisory software entity called a hypervisor [2, 3] for proper system operation. The hypervisor ensures that one OS instance does not corrupt other partitions of the multi-core processor.

A system using the S-AMP configuration that supports three high-capacity and one regular-capacity application server, shown in Fig. 1, will therefore consists of five partitions: a common control plane partition and four dedicated data plane partitions, one for each configured application server. The use of a hypervisor also adds significantly to the system COGS, as a per instance licensing fee has to be paid to the hypervisor vendor for its usage. Scalability is a serious drawback of this configuration (the total number of partitions and the RTOS instance needed vary with the number of configured application servers in the system).

The cost and scalability issues associated with the S-AMP configuration forced us to look for an alternative cost-effective and scalable configuration solution for our system design.

Symmetric Multi-Processing Configuration

A less intuitive choice for multiple application server support on a multi-core processor is the symmetric multi-processing (SMP) configuration [1] with a single partition comprising all the processing cores. This configuration, shown in Fig. 2, does not need the use of a hypervisor [2, 3], and also results in a highly scalable architecture as the single partition with one OS instance serves all the control and data planes, irrespective of the number of application servers supported by the system. However, the SMP configuration loses the desired deterministic behavior of the S-AMP configuration, which has clear segregation of control and data planes. The consolidation of control and data planes under the same SMP partition using a single RTOS instance may result in mixed execution of non-real-time control plane threads/processes and the real-time data plane threads/processes on the same cores. This consolidation of control and data planes under one partition coupled with the use of an open source RTOS such as SMP Linux with PREEMPT_RT poses a serious challenge to meet the stringent real-time performance needs of the data plane real-time threads/processes. Thus, the performance issues and non-deterministic execution behavior of the SMP configuration makes it unsuitable for our system design.

Deterministic Symmetric Multi-processing Configuration

The shortcomings of the SMP configuration may be addressed by employing core reservation and core affinity constructs provided by SMP Linux with PREEMPT_RT to achieve an S-AMP-like system behavior in an SMP configuration. We call this the deterministic SMP (D-SMP) configuration. In this configuration, all non-real-time processes/threads such as operation, administration, and maintenance (OA&M) are bound using core affinity to a core that is dedicated to the control plane activities. A default affinity mask

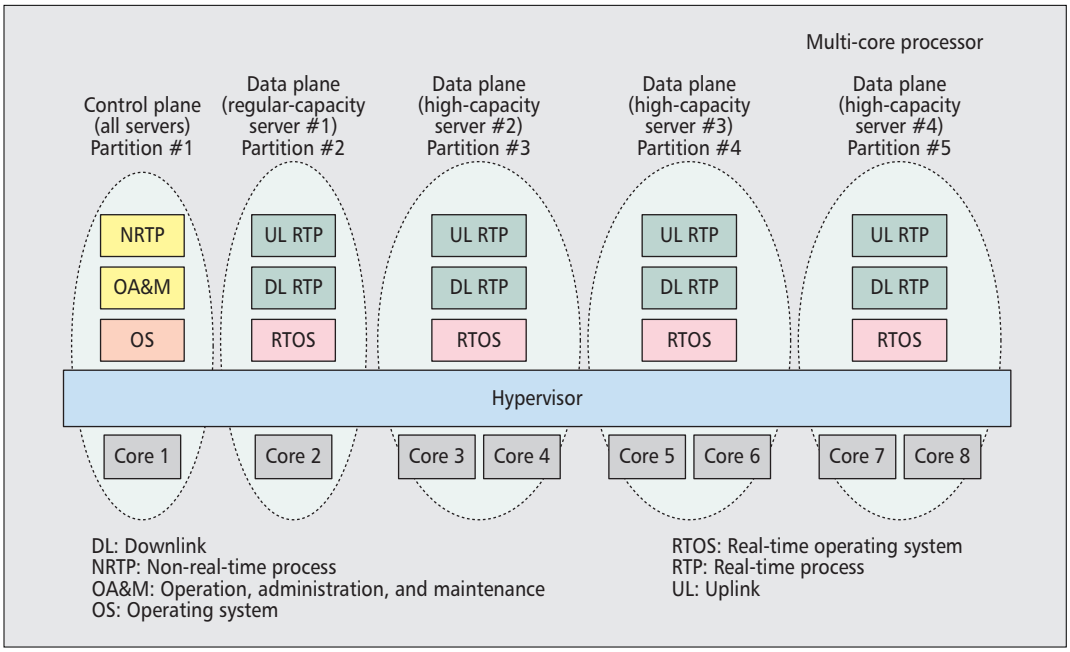

Figure 1. Supervised asymmetric multi-processing (S-AMP) configuration.

The shortcomings of the SMP configuration may be addressed by employing core reservation and core affinity constructs provided by SMP Linux with PREEMPT_RT to achieve S-AMP-like system behavior in an SMP configuration. We call this the deterministic symmetric multi-processing (D-SMP) configuration.

to control plane core is also defined to ensure that any process/thread that has no specific core bindings will default to the control plane core for execution. This ensures that the cores dedicated for the data plane activities are reserved to run only real-time processes/threads.

Figure 3 shows a D-SMP configuration that supports three high-capacity and one regular-capacity application server. The real time data plane process for each application server is bound to its own dedicated data plane core using the core affinity and core reservation construct offered by SMP Linux with PREEMPT_RT. In this configuration, the non-real-time processes/threads will not have to compete with higher-priority real-time processes/threads for processing time as they have their own dedicated control plane core for execution. Consequently, they will also not take any valuable processing time away from the real-time processes/threads executing on the data plane cores. A core abstraction layer (CAL), described in more detail in the next section, is a middleware layer defined in our architecture with an eye toward future software portability as it hides the core and hardware-specific details from the data plane application software. CAL additionally provides high-performance system services such as buffer management and messaging for the data plane processes/threads.

Without these high-performance services, it is impossible to use an open source RTOS such as SMP Linux with PREEMPT_RT in a single partition D-SMP configuration that serves all the control and data planes using the same OS instance, and still meet the stringent real-time performance needs of data plane processes/threads. The D-SMP configuration is also highly scalable as it uses just a single SMP partition and a single RTOS instance irrespective of the number of configured application servers in the system.

Energy Efficiency Considerations

Traditionally, the support of each application server requires the use of two processors to achieve physical separation of the control plane and the time-sensitive data plane. Each of these processors in turn requires additional dedicated devices such as external double data rate (DDR) memory and flash memory devices to store the generic application and the OS image, and other devices such as digital signal processors (DSPs) necessary for proper operation.

Thus, to support seven regular-capacity application servers, the previous generation communication system would require the use of 14 processors (two per application server) and the accompanying devices (DDR memory, flash memory, DSP, etc.) per processor. Traditionally, the physical hardware design for each application server is on its own dedicated board for plug and play functionality. Alternatively, a board may have support for multiple application servers with the same hardware design replicated multiple times on the same physical board for each application server. Nonetheless, such a design configuration would require either seven boards (one application server per board) or more than one board (multiple application servers per board) with backplane board connectors. This not only increases the actual size or footprint of the overall communication system, but also increases the power consumption of the communication system and thus increases the recurring OPEX for the service provider.

The chosen D-SMP configuration utilizes the P4080 multi-core processor [1] with eight cores, which is capable of supporting seven regular-capacity application servers. With this design a single multi-core device is able to replace 14 individual processors to provide support for seven application servers. This configuration also allows

To design a system of varying sizes the choice of the multi-core processor with more or less cores may vary depending upon the specific system capacity, and cost targets. Thus software portability and scalability becomes an important design consideration especially for the application software.

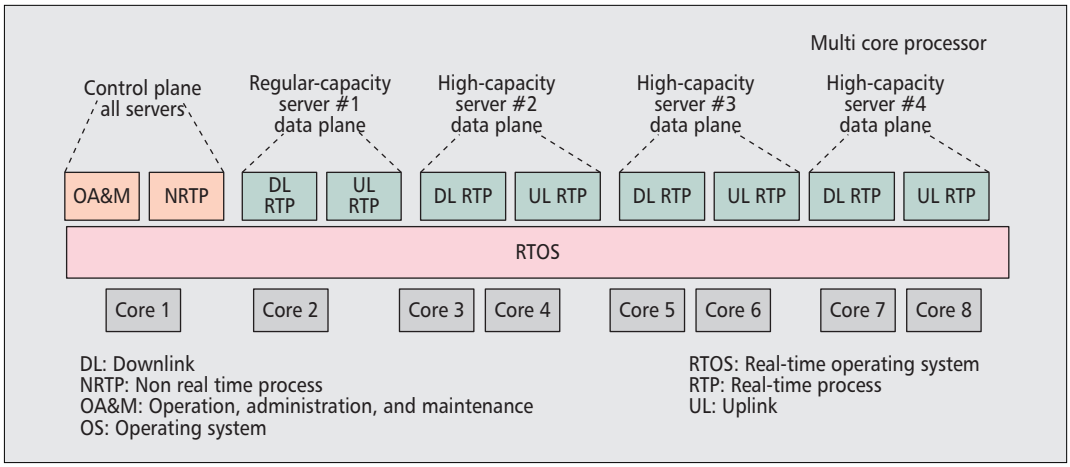

Figure 2. Symmetric multi-processing (SMP) configuration.

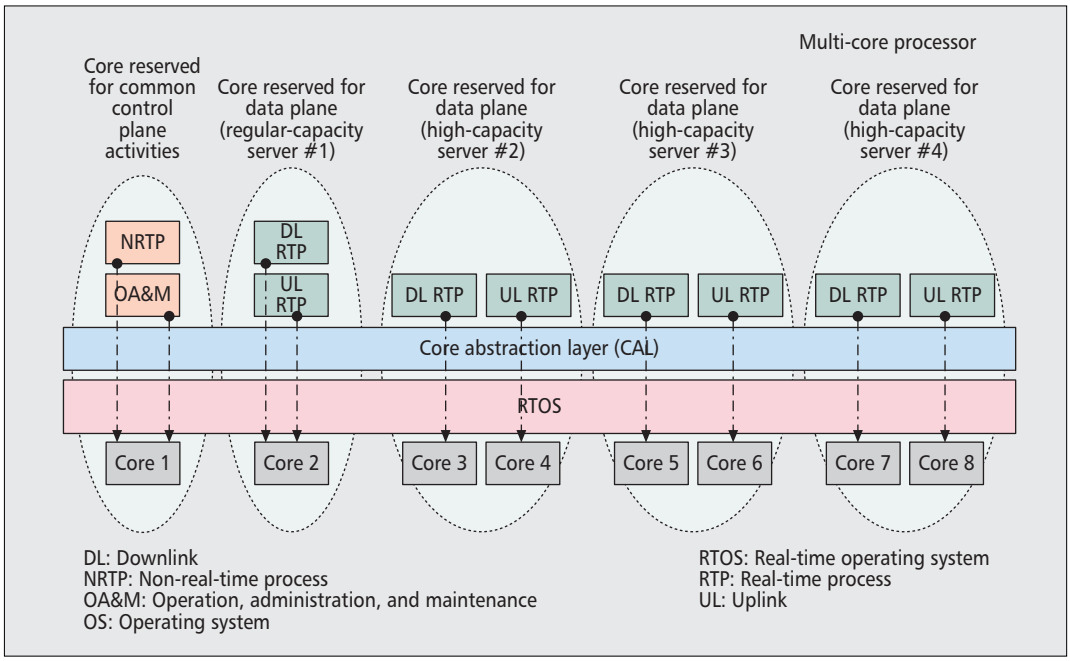

Figure 3. Deterministic symmetric multi-processing (D-SMP) configuration.

us to use the economy of scale, and use single larger-size memory devices (DDR memory, flash memory) instead of the smaller-size devices previously needed for each of the 14 processors.

Thus, the new D-SMP architecture design configuration allows us to reduce the device count by a factor of 14 compared to a similarly equipped communication system that utilizes the current prevailing design. Even higher economy of scale and device count reduction may be achieved by selecting a multi-core processor containing more than eight cores to support more than seven application servers on a single multi-core processor. This not only makes the next generation communication system 10 times or more energy efficient, but also significantly reduces the size/footprint of the communication system, leading to reduction in site leasing cost where the communication system is physically mounted/deployed, resulting in reduced recurring OPEX for the service provider.

Portability and Scalability Considerations

To design a system of varying sizes, the choice of the multi-core processor with more or fewer cores may vary depending on the specific system capacity and cost targets. Thus, software portability and scalability becomes an important design consideration, especially for the application software. It is imperative to have a middleware layer that abstracts or hides the hardware-specific details from the application software. This middleware layer may also provide performance optimized services such as buffer management and messaging service to the application software in case any performance related issues are identified during testing while using the system services provided by the native open source RTOS. The performance issue is triggered by the occurrence of unbounded latency spikes that cause the data plane processes/

threads to exceed the critical 1 ms execution time boundary, resulting in throughput degradation that eventually leads to system instabilities that lead to application server outage. These random unbounded latency spikes manifestations are mainly due to the use of system services and the protocol stack of an open source RTOS such as SMP Linux with PREEMPT_RT, as their implementation is not guaranteed to be lockless. When the same system services are used by the real-time processes/threads on the data plane cores as well as the non-real-time processes/threads on the control plane core, a software lock taken by a non-real-time process or thread on the control plane core may cause a latency spike for a real-time process or thread on the data plane core waiting for the release of the same lock. Therefore, the data plane processes/threads on their part should avoid the direct use of any system services that are offered by the native OS, and instead should rely on the alternative high-performance services offered by the custom middleware layer CAL. This is done not only to ensure that the performance criterion is met, but also to help facilitate the future portability of software to new platforms with different multi-core processors.

Core Abstraction Layer

Figure 4 illustrates the functional architecture of a core abstraction layer (CAL), which is one of the integral building blocks of our proposed system configuration architecture. This middleware layer hides the core and hardware-specific details of the multi-core processor from the user space applications, making future software portability to different multi-core processors easier.

The BMan, FMan, and QMan drivers shown in Fig. 4 are part of the standard board support package (BSP) provided by the OS vendor. The CAL framework consists of three user space modules: the CAL initialization module (CAL_INIT), CAL buffer module (CAL_BUF), and CAL messaging module (CAL_MSG), which provide a user space application programming interface (API) to be used by data plane user space applications, and a custom kernel space driver, CAL_DPA_DRIVER, which serves as an interface to the DPAA hardware acceleration engines.

CAL_INIT module is responsible for setting up the CAL infrastructure needed to support buffer management and messaging services, and initializing the CAL_DPA_DRIVER used to manage the P4080 [1] data path acceleration architecture (DPAA) resources.

The CAL_BUF module provides performance optimized lockless buffer management services to be used exclusively by the "fast path," which is the transmission path taken by the ingress and egress data plane traffic of data plane process/threads.

CAL_MSG provides performance optimized lockless zero copy messaging services for the fast path. The fast path does not use the Linux protocol stack to send or receive Transmission Control Protocol/User Datagram Protocol (TCP/UDP) Internet Protocol (IP) packets, to ensure that it does not suffer from the performance issues caused by unbounded latency spikes mentioned earlier.

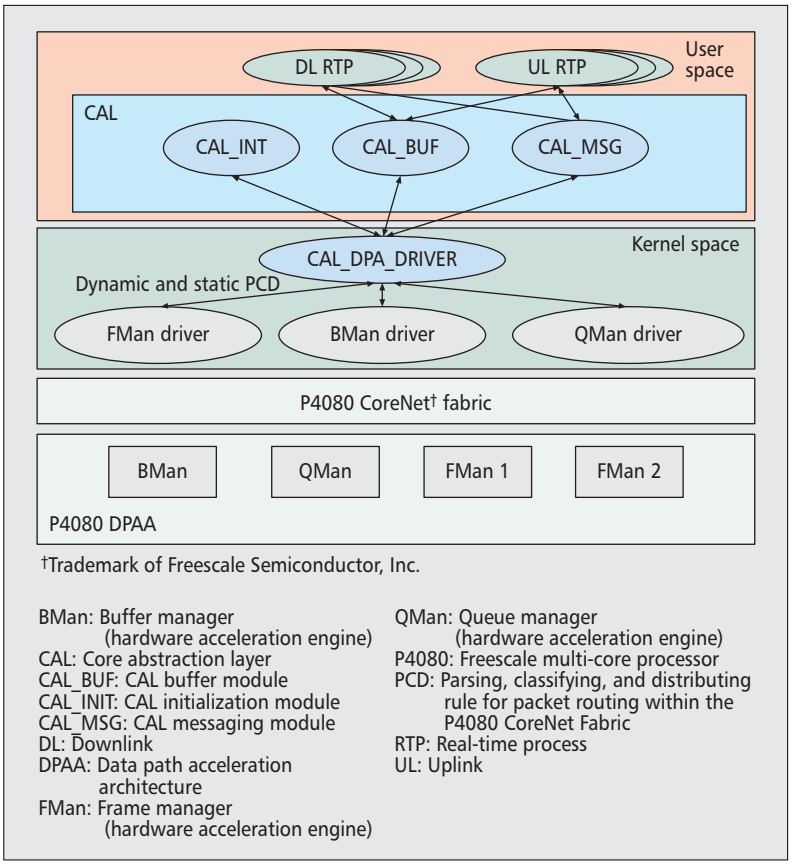

Figure 4. Core abstraction layer functional architecture.

Application Server Recovery Mechanism

Since the system supports multiple application servers on a single multi-core processor, application server recovery can become complex. In the previous implementation where each application server has its own dedicated processors, the recovery procedure following a software crash was simple. In that case, the associated processors were rebooted, and all of the hardware and software components initialized properly, ready to reconfigure the application server.

However, in our proposed architecture, rebooting the multi-core processor to recover or rescue an application server is no longer a viable option, as it will cause the common control plane and the remaining active application servers configured on the multi-core processor to reboot as well. High system availability is a critical service provider requirement, and bringing down all of the operational application servers when just one application server has to be restarted or reconfigured again after a software crash is not an acceptable option.

Thus, the general requirement for the recovery procedure is that the recovery of an application server should not impact the functionalities, behavior, and performance of the other application servers on the same board that is hosting the failing application server.

Some P4080 system resources allocated for an application server are allocated from system-wide shared pools. When such resources are allocated

A system that performs adequately only some of the time is not acceptable. End-to-end black box stress testing of the system design is essential to validate that the radical design alternatives chosen and presented here are capable of meeting the performance requirements.

(i.e., owned by an application server), they cannot be reused by other components until they are explicitly released. In the case of a failure of the software owning resources from shared pools, the platform software must clean up these resources on behalf of the owner software. Otherwise, the corresponding resources will be lost forever. Ingress buffer pools used on the ingress data plane path are an example of system-wide shared resources. They are dedicated to receive data plane packets and allocated by the FMan driver. However, releasing these buffers back to the shared pool is the responsibility of the user space process or application server. Every time an application server goes down, there is a strong likelihood that it may have owned some buffers from the shared pool. If the recovery mechanism does not recover these buffers from the shared pool, after some time and a few more occurrences of application server failure later, the shared pool may not have enough buffers left to serve the needs of the remaining operational application servers in the system.

Testing Considerations

High system availability is a key performance criterion required by service providers for commercial deployment. A system that performs adequately only some of the time is not acceptable. End-to-end black box stress testing of the system design is essential to validate that the radical design alternatives chosen and presented here are capable of meeting the performance requirements.

Off-the-shelf simulator test equipment may be employed to conduct end-to-end stress testing of the system. Such a tool provides the capability to inject into the system the maximum capacity data rate the system is configured to support, and also provides a means to validate the uplink and downlink throughput for the duration of the stress test.

The testing philosophy employed was simple: The system utilizing the architectural design concepts presented in this article should support the maximum data rate allowed by the given system configuration for an extended period of time lasting up to several days with no degradation in throughput or system instabilities leading to system outage.

As part of the new system capabilities developed, a service was also created to allow monitoring of the processing load on a per core basis for the multi-core processor. Thus, the processing loads on the data plane cores where the application servers reside could be closely monitored throughout the duration of the stress test. Such monitoring was deemed essential to identify the presence of any random unbounded latency spikes in the system so that once identified, they may be eliminated by employing appropriate optimization techniques. The design goal has been to ensure that the processing load per core should not exceed 80 percent for the full rate system configuration, thereby leaving 20 percent head room.

Extensive testing was also performed involving application server recovery scenarios to validate that an operational application server may be taken out of service and brought back up with no impact to other operational application servers configured in the system.

State of the Art/Related Work

A granted patent [7] presents the multi-cell support using the D-SMP configuration in a wireless base station modem board design. Another granted patent [8] describes a cell recovery mechanism for a wireless base-station modem board that uses the D-SMP configuration to support multi-cell configuration.

Lessons Learned

This article has described the design strategies for a low-cost high-capacity scalable and energy-efficient next-generation communication system. The architecture utilizes some radical design ideas that have never been attempted before. These include combining the control and time-sensitive data planes of multiple application servers under a single D-SMP partition on a single multi-core processor to be served by a single instance of an open source RTOS. This resulted in significant reduction (by a factor of 14) in device count and system size/footprint reduction, and led to a scalable cost-effective (10 times or more) energy-efficient system configuration solution that resulted in reduction in CAPEX and OPEX for the service provider.

The use of a multi-core processor, with powerful processing cores and hardware acceleration engines for packet processing offload, resulted in an overall capacity increase by a factor of around 1.5 for the communication system.

Several system configurations using a multi-core processor were considered. The chosen single-partition D-SMP configuration offered the desired deterministic execution behavior of the S-AMP configuration, and the scalability and cost savings of the SMP configuration, where a single open source RTOS instance is used irrespective of the number of application servers configured on the multi-core processor. This resulted in a reduction by a factor of 14 in the number of OS instances used for the seven application server example outlined in this article.

The decision to use a tier 1 OS supplier for the open source RTOS reduced concerns about any indemnification issues that may result from the use of an open source RTOS in the commercial product.

Even though the use of an open source RTOS contributed toward reducing the system COGS by eliminating the need to pay a tier 1 OS supplier a per instance RTOS deployment royalty fee, it opened the door to serious performance issues for time-sensitive data plane processes/threads. Consequently, the communication system was not able to provide high system availability for a prolonged period of time while maintaining the architected maximum system throughput. High-performance lockless zero copy buffer management and messaging services were introduced for the time-sensitive data plane processes/threads to address the performance issues to ensure that the communication system could be commercially deployed.

The choice of a multi-core processor with more or fewer cores is generally driven by the size and capacity requirements of the communication system. From a software portability per-

spective (build once, use many times) it was considered important to define a middleware layer CAL that abstracts or hides the hardware-specific details from the application software. This makes future software portability easier as the use of different multi-core processors in future products will result in minimal impact on the application software.

Since the control and data planes of all configured application servers in the communication system shared the same SMP partition and were served by a single RTOS instance, an application server recovery mechanism became quite complex. The new recovery mechanism has to be done without a multi-core processor reboot. It also required elaborate system resource cleanup procedures to ensure that when one application server goes down, other operational application servers are not affected. Thus, the new application server recovery mechanism ensures that the high system availability requirement of the next-generation communication system is not compromised.

Extensive black box testing was needed and performed to alleviate any concerns surrounding the high system availability due to the use of such radical architectural design ideas. The testing ensured that a communication system utilizing such a system configuration is not just a theoretical concept, but rather fully capable of meeting the requirements needed for commercial deployment.

Acknowledgments

The author would like to acknowledge the contributions of the following individuals to the work presented in this article: Mugur Abulius, Lina So, and Shriram K. Easwaran.

References

[1] Freescale Semiconductor, Product Brief, "QorIQ P4080 Communications Processor Product Brief," doc. no.: P4080PB, Rev. 1, 09/2008.

[2] Freescale, "The Freescale Embedded Hypervisor," Nov. 2010.

[3] Wind River, "Wind River Hypervisor," product note, http://www.windriver.com/products/product-overviews/PO_Hypervisor_0611.pdf.

[4] Wind River VxWorks Platforms 6.9, http://www.windriver.com/products/product-overviews/PO_VE_6_9_Platform_0211.pdf.

[5] G. Seiler, "Wind River Linux and VxWorks Real-Time Capabilities: A Comparison," White Paper.

[6] N. Litayem and S. Ben Saoud, "Impact of the Linux Real-Time Enhancements on the System Performances for Multi-Core Intel Architectures," *Int'l. J. Computer Applications* (0975–8887), vol. 17 no. 3, Mar. 2011.

[7] M. R. Khawer and S. K. Easwaran, "Apparatus for Multi-Cell Support in a Network," U.S. Patent # US 8,634,302 B2, Jan. 21, 2014.

[8] M. R. Khawer and M. Abulius, "Method and System for Cell Recovery in Telecommunication Networks," U.S. Patent # US 8,730, 790 B2, May 20, 2014.

Biography

Mohammad R. Khawer (Mohammad.khawer@alcatel-lucent.com) is a Distinguished Member of Technical Staff in the Wireless Product Division of Alcatel-Lucent, Murray Hill, New Jersey. Throughout his career at Alcatel-Lucent, he has worked as a lead R&D developer and platform software architect, and contributed to the design and implementation of numerous 2G/3G/4G commercial wireless base station products. He received his M.S. degree in computer science and is currently pursuing a Ph.D degree in computer and information science engineering from Syracuse University, New York. He got accepted as a member of the elite Alcatel-Lucent Technical Academy (ALTA) in 2009, and currently holds 17 granted patents.

The testing ensured that a communication system utilizing such a system configuration is not just a theoretical concept, but rather fully capable of meeting the requirements needed for commercial deployment.

UMPIRE: A Universal Moderator for the Participation in IETF Remote Events

Simon Pietro Romano

ABSTRACT

UMPIRE provides seamless meeting interaction among remote and local participants. It uses the BFCP, an IETF standard for moderation. BFCP introduces automated floor control functions to a centralized conferencing environment. This article discusses the design and implementation of the UMPIRE system and highlights the most notable solutions we devised to handle variegated requirements and constraints. We also discuss the lessons learned while experiencing in the first person how the application of research results that have eventually led to new standards still must confront a number of minor yet concrete issues that might completely undermine the overall process of wide adoption by the community.

BACKGROUND, RATIONALE AND MOTIVATION

UMPIRE provides seamless meeting interaction among remote and local participants. It uses the Binary Floor Control Protocol (BFCP), a standard for moderation. BFCP introduces automated functions for a centralized conferencing environment. The project has been motivated by the ongoing efforts within the IETF (Internet Engineering Task Force) to standardize mechanisms for enabling remote participation at meetings. At that time, work was in full swing and was being formalized in a (now expired) Internet draft[1] specifically devoted to this task. The draft in question was actually being produced at the request of the IETF Administrative Oversight Committee (IAOC), which issued an ad hoc request for proposals at the end of October 2011.[2] The draft contains discussion (section 2.3.4.1) dedicated to the task of moderation. Titled "*Floor Control for Chairs for Audio from Remote Attendees*," it contains a list of requirements, including:

> **Requirement 08-31**: Remote attendees MUST have an easy and standardized way of requesting the attention of the chair when the remote attendee wants to speak. The remote attendee MUST also be able to easily cancel an attention request.
>
> **Requirement 08-33**: The floor control portion of the Remote Participation System MUST give a remote attendee who is allowed to speak a clear signal when they should and should not speak.
> **Requirement 08-34**: The chair MUST be able to see all requests from remote attendees to speak at any time during the entire meeting (not just during presentations) in the floor control system.
> **Requirement 08-35**: The floor control system MUST allow a chair to easily mute all remote attendees.
> **Requirement 08-36**: The floor control system MUST allow a chair to easily allow all remote attendees to speak without requesting permission; that is, the chair SHOULD be able to easily turn on all remote attendees mics at once.

UMPIRE is currently capable of meeting most of the requirements in that list.

We also took inspiration from related mailing list discussions. To get an idea of such discussions, consider the following thread on the IETF's Working Group Chairs mailing list: "*Getting Taipei remote participants' input.*" It includes a message from the author of this article,[3] who was proposing to use RFID (Radio Frequency IDentification) tags to trigger BFCP requests by in-person participants, in much the same way as the floor requests generated by remote participants. In answer to such message, Dave Crocker wisely stated the following:[4]

> Suggestions like the above sound appealing. Unfortunately, they are far beyond current products and making them useful is considerably more difficult than the suggestions imply. That doesn't mean they should be ignored, but we need to be careful about slipping into the assumption that merely citing a bit of technology means that an issue is resolved. We had an experiment with RFIDs. It was awkward, at best. In the case of queue management, we have at least entering the queue, position in the queue, and the chair's control of the queue.

A further fundamental contribution to the discussion was also provided by Brian Rosen, who asserted:[5]

> People worry about RFID, but I like it because it's a faster read. All I think you want is a reader and a visible queue. The

The author is with University of Napoli Federico II.

[1] "Requirements for Remote Participation Services for the IETF," draft-ietf-genarea-rps-reqs-08.

[2] http://iaoc.ietf.org/documents/RPS-Specifications-RFP-2011-10-19.pdf

[3] http://www.ietf.org/mail-archive/web/wgchairs/current/msg10274.html

[4] http://www.ietf.org/mail-archive/web/wgchairs/current/msg10277.html

[5] http://www.ietf.org/mail-archive/web/wgchairs/current/msg10280.html

0163-6804/15/$25.00 © 2015 IEEE

queue just tells you that the reader read correctly and has you in the queue. The chair gets to change the queue, but that ought to be rare and probably just pick the next person in queue. Remote participants simply imitate the reader action.

The tool could also provide session chairs with the ability to grant "business class" requests (i.e. in the case of cut and thrust debates, or in the presence of intervention of an area director) so that individuals obtain higher priority, essentially putting such requests on top of the queue. This highlights the possibility of applying different activity templates, or paradigms, for common handling of remote and in-person participants, according to different group process modes.

We took into account these considerations and applied the usual IETF approach of having a running code prototype to identify and clarify possible issues and foster discussion. The system we developed provides a proof of concept for a moderation framework built on top of the Meetecho conferencing system [1]. Meetecho is a standards-based conferencing architecture used at IETF meetings for remote participation.

The remainder of this article is organized as follows. The next section helps the reader position this work in the context of ongoing IETF standardization efforts for multimedia conferencing. The third section provides some information about related work in the field of conference moderation. Following that we describe the overall architectural design of the UMPIRE system, whose implementation is briefly sketched in the next section. We then discuss the main issues we faced during the design and implementation phases and provide insight into the lessons learned. Finally, the conclusion summarizes the discussion stimulated by UMPIRE within the IETF community, illustrating the useful feedback we gathered and the envisioned direction of our future efforts.

Context

The IETF has devoted much effort to the specification of standard conferencing solutions. They include the Framework for Centralized Conferencing (XCON Framework) [2], which defines a signaling-agnostic architecture, naming conventions, and logical entities required for building advanced conferencing systems. An XCON-compliant framework architecture comprises several protocols, including the Binary Floor Control Protocol (BFCP) [3], which is associated with all moderation operations for a conferencing session.

As depicted in Fig. 1, BFCP models the presence of floor participants who ask for access to the conference floor (e.g. audio and/or video) by sending messages to a central entity called the floor control server. The server itself does not make decisions on its own, but rather forwards requests to the floor chair, who acts as a moderator and is in charge of making decisions.

When a conference participant asks for the floor, it sends a request message to the floor control server, which forwards it to the floor chair. When the chair makes a decision, it informs the floor control server, which in turn notifies both the requesting participant and all other participants potentially interested in receiving floor control notifications.

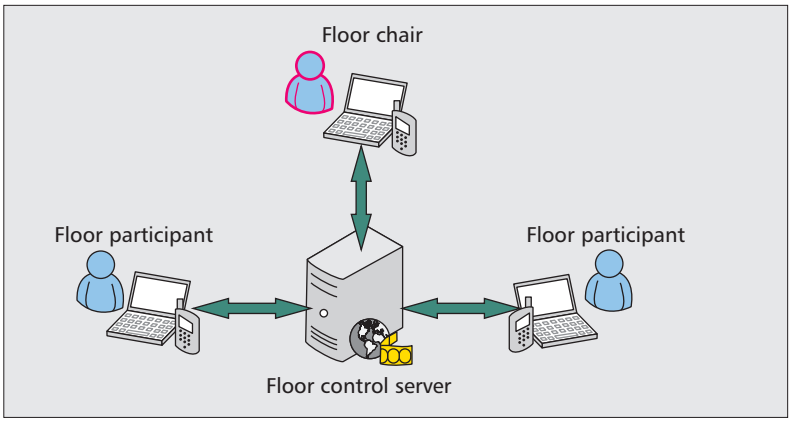

Figure 1. The Binary Floor Control Protocol architecture.

Related Work

Floor control has long been the subject of a number of research works. Before BFCP saw the light, some interesting approaches were proposed in the literature. The authors in [4] propose to use floor control in videoconferencing applications as a means to improve their scalability. With the proposed approach, moderation makes it possible to keep control over both processing and communication overhead, by allowing a maximum of two simultaneous streams at a time and hence mimicking a two-party videoconference. A different approach is embraced in [5]. Here, the authors propose an implementation of ALOHA and DQDB (Distributed Queue Dual Bus), two well-known MAC access protocols for local area networks, in a distributed overlay setting. The authors of [6] instead design an architecture for medium-sized peer-to-peer conferences. The system is equipped with a floor control mechanism to prevent too many users from speaking simultaneously (hence degrading the audio quality). To this purpose, a distributed role-based floor control protocol is introduced. The protocol leverages floor utilization statistics in order to optimize floor management activities.

More recently, the work in [7] has proposed an effective way to bring the functions made available by the BFCP protocol to the IP Multimedia Subsystem standard framework. They actually propose to use BFCP for the implementation of the IMS Fc interface envisaged by the 3GPP standard.

UMPIRE Architecture Design

UMPIRE fills the role of floor chair in a conference. More precisely, when a conference starts, the UMPIRE user will log in as floor chair. Subsequent floor requests from conference participants will be:

1) Stored (in a PENDING state) in the First-Come-First-Served queue at the centralized floor control server.

2) Forwarded to UMPIRE, which will make decisions by assigning the floor to one or more users, thereby triggering state changes at the server. As an example, if the floor control policy

Figure 2. Two local participants and a remote participant asking for the floor.

has been configured to grant the floor to one user at a time and the UMPIRE accepts, in sequence, three requests coming from three different users, the following will happen:

- The first accepted request will move from a PENDING state to a GRANTED state.
- The second and third requests will move from a PENDING state to an ACCEPTED state, indicating that they are ready to be served as soon as the currently GRANTED request has terminated (through either a floor release action from the client or a revoke action from the UMPIRE himself).
- When the GRANTED request is completed, the first ACCEPTED request in the queue is granted the floor, while the second one becomes the ready-to-be-granted request in the floor control server queue.
- If the server policy allows for a maximum of *n* requests to be granted the floor at the same time, up to *n* clients will reach the GRANTED state and will share the floor in question (e.g. in an audio conference, they will all be allowed to talk at the same time by contributing to the audio mix produced at the conference server).

3) Acknowledged to the clients through BFCP notifications (which allow participants to be kept up to date with respect to the state of their floor transactions).

The above scenario relies on the presence of a centralized floor control server with a queue managed by the chair and containing all floor requests coming from conference participants. This simple model allows for the introduction of advanced moderation functionality.

The interesting thing about the model is that the BFCP queue can be populated with requests coming from both remote participants (equipped with BFCP-enabled clients) and local participants, thanks to the utilization of an agreed-upon procedure for requesting access to the conference floors when one is physically present in the conference room. The usual way of gaining the audio floor at regular IETF meetings is for local participants to politely wait (in a First-Come-First-Served queue) at one of the conference room microphones, in order to either ask for questions or provide their own view on the topic being discussed. If the microphones themselves were equipped with some simple means for:

- Recording the presence of users.
- Sending a trigger (i.e. a floor request) to the floor control server every time a new user lines up at the microphone.

The floor control server queue could transparently (and democratically) moderate a conference envisioning the contemporary presence of both local and remote participants. Indeed, this is what we implemented. With respect to the means for recognizing the presence of users lining up at the microphone, we decided to rely on the RFID technology. We opted for the following policy:

- An RFID reader was placed close to each of the conference room microphones.
- An RFID tag was assigned to each local participant willing to actively participate in the mechanism.

With this approach, when a local participant wanted to contribute to the ongoing discussion, all they had to do was to let their RFID tag be read by the RFID reader associated with the microphone at which they lined up. As soon as the participant's tag was read, the reader would send a BFCP floor request to the conference chair and let the participant be inserted in the centralized BFCP queue.

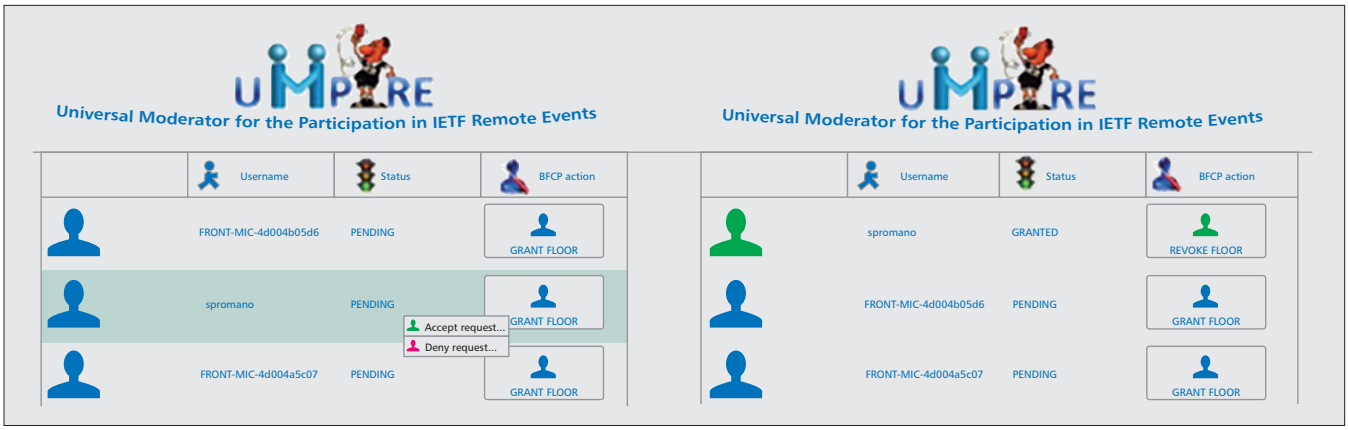

Figure 3. Accepting floor request coming from user "spromano."

An image of the BFCP queue, updated in real-time, was always projected on a screen available in the room, for inspection by meeting participants. The chair of the conference (i.e. the UMPIRE user) was allowed to manage the BFCP queue by making decisions to grant the floor to individuals. This demonstrated straightforward moderation functionality, for a conference involving both remote and local participants.

UMPIRE Implementation

UMPIRE has been implemented as a Web2.0 application, that is, a dynamic, highly interactive, web-based system. It is based on a bidirectional HTTP communication channel between the UMPIRE participant, acting as the floor chair in a conference, and the floor control server. The channel uses the COMET Server Push approach, as made available by the ZK framework.[6] Notifications from the server asynchronously arrive at the UMPIRE client and are represented on a web page providing an always up-to-date snapshot of the BFCP queue (with client requests and related BFCP states). Proactive actions undertaken by the UMPIRE (e.g. accepting or denying a PENDING request, or revoking a floor currently GRANTED to one of the participants), are immediately communicated to the floor control server, affecting the BFCP queue, as well as trigger floor notifications that are sent to clients.

Sample Call Flow

The UMPIRE functionality is demonstrated by the following simple scenario:

- Three users participate in a conference room:
  - Two local users, equipped with RFID tags:
    - User1: 4d004b05d6
    - User2: 4d004a5c07
  - A remote user (whose nickname is "spromano"), who enters the conference through a BFCP-enabled client.
- The order in which the three users ask for the audio floor is the following:
  - User1 → spromano → User2

The situation described above is illustrated in Fig. 2, which shows, respectively, the initially empty queue (top left), the first request arriving from "User1" (top right), the request from the remote participant "spromano" (bottom left), and the final request from local user "User2" (bottom right). The final state of the BFCP queue, after all these actions have been performed, shows the three users (in order of arrival) in a PENDING state, that is, waiting for the chair to take actions.

UMPIRE first decides to grant the floor to "spromano," as shown in the two snapshots in Fig. 3 ("Accept," left snapshot), which translates into the following actions:

- The BFCP queue is modified: "spromano" becomes first.
- The state of the BFCP queue is modified: the audio floor is GRANTED to "spromano" (right snapshot).
- Remote user "spromano" is unmuted.

UMPIRE now decides to grant the floor also to "User2." This is illustrated in Fig. 4.

We can observe from the pictures that the following things have happened:

- UMPIRE has accepted the request issued by "User2" (left snapshot).
- The state of the BFCP queue changes: "User2" passes from PENDING to ACCEPTED (middle snapshot) and eventually to GRANTED (right snapshot). This lets us understand that the conference in question has been configured to allow multiple users to be granted the audio floor at the same time. Were this not the case, 'User2' would have moved from PENDING to ACCEPTED and would have stayed in such a state as long as the floor was held by "spromano."

UMPIRE now decides to revoke the floor previously assigned to "spromano." This is shown in the snaphots in Fig. 5, associated, respectively, with the action undertaken by the chair (left frame) and the effect it has on both the BFCP queue at the server and the web interface (right frame), which now reports "spromano" in red, with the related REVOKED status.

Lessons Learned

UMPIRE can be regarded as an advanced chapter of the author's experiences with complex conferencing environments, such as the IETF, dating back to 2005. This chapter represents a clear example of the way research activities can be brought to the real world, if a proper engineering approach is embraced. The work done to develop the moderation platform taught us several lessons.

[6] http://www.zkoss.org/

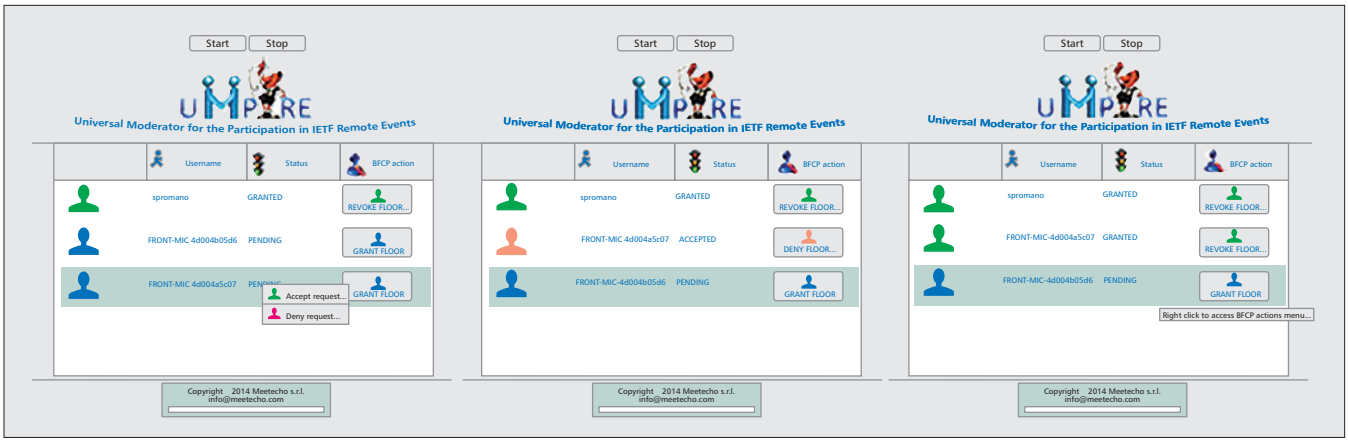

Figure 4. Accepting the request coming from the "second" local participant.

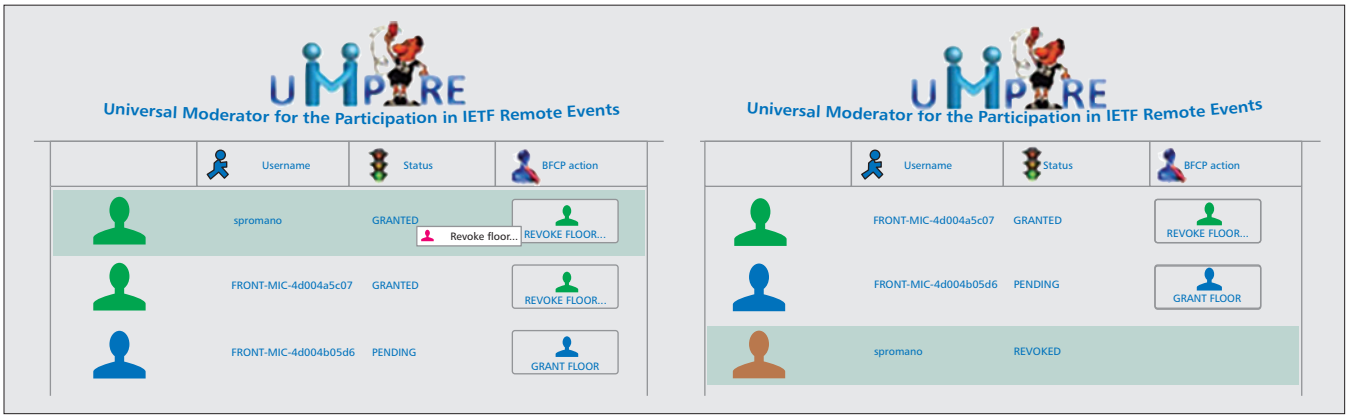

Figure 5. Revoking the floor to user "spromano."

From the design perspective, our experience demonstrates the benefits in using a "separation of concerns" pattern for re-using a specific subset of the functionality in a large and complex system. Originally we were conducting research on the scalability of conferencing frameworks, and hence seized the opportunity to contribute to the ongoing standardization work within the IETF, by actively participating in both the Centralized Conferencing (XCON) and the Media Control (mediactrl) working groups.

The eventual result was Meetecho, a spin-off of the University of Napoli offering a standards-based conferencing system used to support remote participation in IETF meetings. UMPIRE is built upon Meetecho's major component dealing with BFCP moderation, which was implemented by our research group as a joint activity with Ericsson Nomadic Lab in Helsinki. Notably, it has been conceived as an independent component, which can interoperate with any other BFCP-compliant system; Meetecho is not required. Hence, we could focus solely on the client-side of the architecture. Besides re-using the code, we were also able to take advantage of the moderation server's performance, which we have studied in detail in our previous work [1].

As the BFCP chair, UMPIRE is in charge of managing requests arriving from conference participants, while keeping an up-to-date representation of the queue, for the web-enabled GUI. The GUI itself was the most challenging part of the overall prototype, since it represents a typical example of a bidirectional, HTTP-based component and must be capable of managing input events coming both from the web interface (e.g. when the moderator clicks on a user's icon in order to perform a specific moderation action) and the server-side counterpart residing on the Meetecho conferencing server. For this part, we initially decided to use long polling HTTP requests sent by the client and responsible for the asynchronous update of the web view. This solution proved to be far from optimal, due to the unavoidable overhead of this approach. We then moved to a Comet server push approach, as described earlier. This is definitely better than polling when an application needs low latency events delivered from the server to the browser. Instead of repeatedly polling for new events, Ajax applications with Comet rely on a persistent HTTP connection between server and client.

Also worth mentioning is the communication between the RFID readers (which we install close to the room microphones) and the BFCP server. This part of the system is critical, since it raises both hardware and software issues. The reader has to interface with passive RFID tags (such as those "embedded" into conference badges worn by participants). It also has to properly communicate with the moderation server, by acting as a standard BFCP participant. Ultimately we used an effective, low-cost, programmable RFID reader made available by Phidgets Inc.,[7] a Canadian company offer-

[7] www.phidgets.com

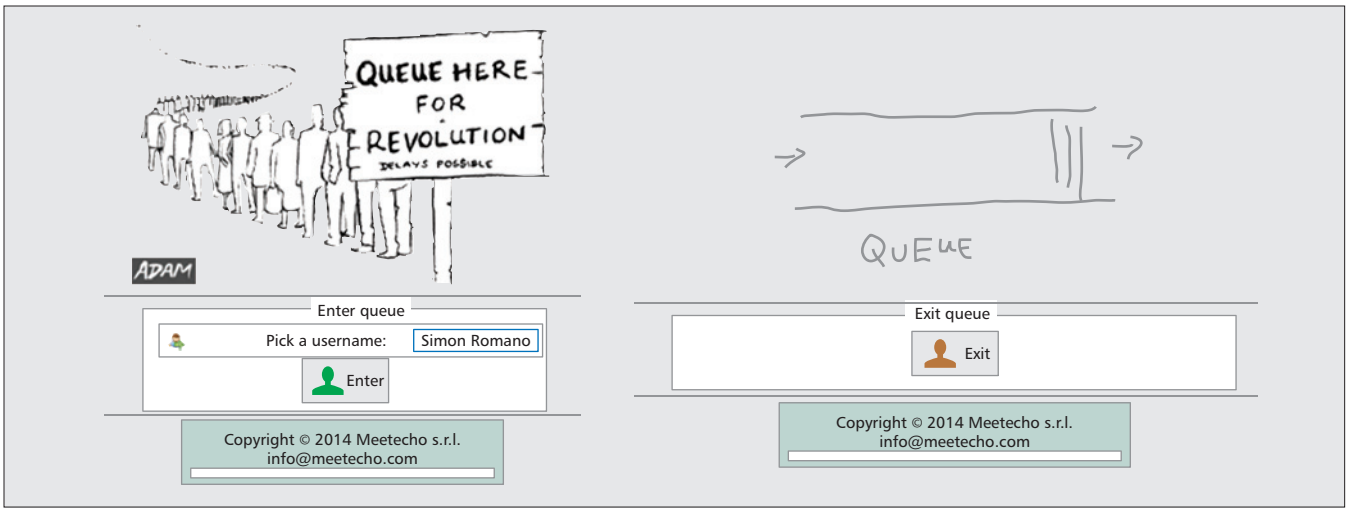

Figure 6. The simple web GUI for sending floor request and floor release messages.

ing solutions for the rapid prototyping of RFID sensing components. Fortunately they come with built-in support for an application programming interface written in a number of different languages and freeing developers from all low-level issues associated with RFID sensing capabilities, so the programmer can focus on application-level issues. In our case, we leveraged the Java API.

As a final remark, we also notice that the idea of enabling RFID-based conference moderation continues to represent a critical aspect of the UMPIRE system. As part of the feedback we received after the first experiments at recent IETF meetings, we realized that many people would actually prefer a more lightweight approach to the interaction of conference participants (both local and remote) with the floor control server. Other participant signaling mechanisms are indeed possible, such as through personal smartphones. This suggested additional experiments to consider, with the final goal to allow the possibility that an operational service would permit multiple means for participants to signal their desire for the floor. Based on the previous consideration, we recently made a further developmental step and implemented a simple cross-device (laptop, mobile, tablet, etc.) web-based floor control client that can be leveraged by conference participants to send "RFID-less" floor request and floor release messages to the floor control server. With this extension, the UMPIRE system can now moderate a unified "virtual" queue that groups together local and remote participants who make use of either the RFID mechanism, or a BFCP-enabled conferencing application like Meetecho, or the simple Web GUI showed in Fig. 6. The enhanced version of the UMPIRE system has been used at the 90th IETF meeting in Toronto (July 2014).

Discussion and Concluding Remarks

We have presented UMPIRE, a system for the automated management of floor control and moderation in a meeting room supporting the contemporary presence of local and remote participants. At the time of this writing, UMPIRE (which was first proposed at the 83rd IETF meeting in Paris, in March 2012) has not yet been used 'in the wild' as a system to moderate actual meeting sessions. However, at recent meetings it has been demo-showed during the official "Meetecho tutorial for participants and WG chairs" and has gathered consensus and appreciation, besides stimulating useful feedback from the audience.

We are currently fine-tuning UMPIRE's functionality to make it ready for official adoption within the IETF. Obviously, nothing prevents UMPIRE from being used in contexts other than the IETF. In fact, it can be employed wherever a plea for moderated access to shared resources exists.

References

[1] A. Amirante *et al.*, "On the Seamless Interaction Between WebRTC Browsers and SIP-Based Conferencing Systems," *IEEE Commun. Mag.*, vol. 51, no. 4, Apr. 2013, doi: 10.1109/MCOM.2013.6495759, pp. 42–47
[2] M. Barnes, C. Boulton, and O. Levin, "A Framework for Centralized Conferencing," RFC5239, June 2008.
[3] G. Camarillo, J. Ott, and K. Drage, "The Binary Floor Control Protocol (BFCP)," RFC4582, Nov. 2006.
[4] J. Garcia-Luna-Aceves *et al.*, "Floor Control Alternatives for Distributed Videoconferencing over IP Networks," *Proc. Int'l. Conf. Collaborative Computing: Networking, Applications and Worksharing*, 2005.
[5] S. Banik *et al.*, "Distributed Floor Control Protocols for Computer Collaborative Applications on Overlay Networks," *Proc. Int'l. Conf. Collaborative Computing: Networking, Applications and Worksharing*, 2005.
[6] Y.-J. Joung and P. H. Chien, "P2pconf: A Medium-Size P2P Internet Conference with Effective Floor Control," *Proc. Int'l. Conf. Information Networking, ICOIN 2008*, Jan. 2008.
[7] M. Al Rubaye *et al.*, "A Novel Architecture for Floor Control in the IP Multimedia Subsystem of 3G Networks," *Proc. 69th IEEE Vehic. Tech. Conf., 2009*, VTC Spring 2009, Apr. 2009.

Biographies

Simon Pietro Romano is an associate professor in the Department of Electrical Engineering and Information Technology (DIETI) at the University of Napoli. He teaches computer networks and telematics applications. He is also the cofounder of Meetecho, a startup and University spin-off dealing with WebRTC-based unified collaboration. He actively participates in IETF standardization activities, mainly in the real-time applications and infrastructure (RAI) area.

Advertisers' Index

ADVERTISING SALES OFFICES

Closing date for space reservation: 15th of the month prior to date of issue

National Sales Office
James A. Vick
Sr. Director Advertising Business, IEEE Media
Email: jv.ieeemedia@ieee.org

Marion Delaney
Sales Director, IEEE Media
Email: md.ieeemedia@ieee.org

Mindy Belfer
Advertising Sales Coordinator
Email: m.belfer@ieee.org

Northern California
George Roman
Tel: (702) 515-7247
Fax: (702) 515-7248
Cell: (702) 280-1158
Email: George@George.RomanMedia.com

Southern California
Patrick Jagendorf
Tel: (562) 795-9134
Fax: (562) 598-8242
Email: pjagen@verizon.net

Mid-Atlantic
Dawn Becker
Tel: (732) 772-0160
Fax: (732) 772-0164
Email: db.ieeemedia@ieee.org

Northeast
Merrie Lynch
Tel: (617) 357-8190
Fax: (617) 357-8194
Email: Merrie.Lynch@celassociates2.com

Jody Estabrook
Tel: (77) 283-4528
Fax: (774) 283-4527
Email: je.ieeemedia@ieee.org

Southeast
Scott Rickles
Tel: (770) 664-4567
Fax: (770) 740-1399
Email: srickles@aol.com

Midwest/Central Canada
Dave Jones
Tel: (708) 442-5633
Fax: (708) 442-7620
Email: dj.ieeemedia@ieee.org

Midwest/Ontario, Canada
Will Hamilton
Tel: (269) 381-2156
Fax: (269) 381-2556
Email: wh.ieeemedia@ieee.org

Texas
Ben Skidmore
Tel: (972) 587-9064
Fax: (972) 692-8138
Email: ben@partnerspr.com

Europe
Rachel DiSanto
Tel: +44 1932 564 999
Fax: +44 1 1932 564 998
Email: rachel.disanto@husonmedia.com

Germany
Christian Hoelscher
Tel: +49 (0) 89 95002778
Fax: +49 (0) 89 95002779
Email: Christian.Hoelscher@husonmedia.com

Currently Scheduled Topics

Topic	Issue Date	Manuscript Due Date
Social Networks Meet Next Generation	October 2015	May 15, 2015
Towards Autonomous Driving: Advances in V2X Connectivity	December 2015	June 1, 2015

www.comsoc.org/commag/call-for-papers

Instant Access to IEEE Publications

Enhance your IEEE print subscription with online access to the IEEE *Xplore*® digital library.

- Download papers the day they are published
- Discover related content in IEEE *Xplore*
- Significant savings over print with an online institutional subscription

"IEEE is the umbrella that allows us all to stay current with technology trends."

Dr. Mathukumalli Vidyasagar
Head, Bioengineering Dept.
University of Texas, Dallas

Start today to maximize your research potential.

Contact: onlinesupport@ieee.org
www.ieee.org/digitalsubscriptions

IEEE
Advancing Technology for Humanity