Mine-Sniffing Robotic Snakes and Eels: Fantasy or Reality?

Ian A. Gravagne
Clemson University
Dept. of Electrical and Computer Engineering
Clemson, SC
igravag@clemson.edu

Ronald L. Woodfin
(Retired, Sandia National Laboratories)
Wayland Baptist University
Albuquerque, NM
rwoodfin@lwol.com

Abstract

Organizations around the world interested in mine detection and elimination have come to realize there will probably be no "silver bullet" to solve the problems of proliferating anti-personnel mines. From a technogical standpoint, it seems more likely that success will require a suite of different tools including various types of sensors and the means to get the sensors where they are needed. This paper focuses on one possible avenue to accomplish this task: attaching a sensor to a robotic snake or eel. We will illustrate the advantages of this method, the state of the art on robotic snakes, the difficulties that must be overcome and the timeline to overcome them.

1 Introduction

Oftentimes deminers are faced with a situation where the area they must clear has been allowed to lie fallow for some considerable time. In some climates the native undergrowth takes over and the mines that were laid in more or less clear areas are now so covered that no equipment can safely approach them, even to survey the area for their presence. Workers at Sandia National Laboratories (and likely in other places as well) realized several years ago that the most appropriate vehicle for carrying chemical sensors, and perhaps other sensors as well, into such an area would be some sort "robotic snake." Such a device could operate in any of several ways. For example, if it carried also some sort of RF beacon, the location of suspect chemical concentrations could be located without physically penetrating the brush. Alternately, if the vehicle were expendable, it could carry a counter charge. Robotic eels, operating on the same principles as snakes, could swim in extremely shallow waters, through the grasses and other brush that grow there, and perform similar functions.

Two fundamental technological hurdles stand in the way of successfully employing snakes and eels as robotic mine detectors. First, chemical (or other) sensors must be made rugged, reliable and small enough to be carried by a snake. Second, robotic snake designs must undergo a significant reduction in complexity and cost relative to the current state of the art – the subject of this paper. We will begin by highlighting some of the specific benefits and capabilities of snakes, and then focus on current design philosophies and parameters, including a survey of some of the recent work in snake and eel robots. We will end with a summary of the remaining challenges facing robotic snake designers, and outline some conservative guidelines about the real-world capabilities of robotic snakes, such as payload capacity, power consumption, and complexity.

2 Why Snakes?

Clearly, when one observes the natural world, the vast majority of creatures at "macro" scales do not slither – they walk, run, hop or fly. The mere existence of the snake is, in many ways, a biological paradox: most snakes are less then 4 cm tall, they are typically not very fast relative to runners and hoppers, most cannot see long distances, they have no appendages for the purpose of digging out a habitat or manipulating their environment, and most are not vegetarians. Yet, snakes remain among the most prolific reptiles on the earth. They are capable of climbing trees and swimming. They make up for their lack of vision with some of the world's most sensitive noses, they can crawl through labyrinthine obstacle courses, several species are amphibious, and many are adept at strangulating prey with astonishing force. In short they are nearly the perfect mine detectors: capable of locomoting almost anywhere with an extremely high-gain chemical sensor attached, without activating trip wires or posing an obvious presence.

In a broader context, snakes and eels share many characteristics with the tentacles and trunks of creatures like elephants, octopi, and squid. For this reason, the robotics community has generally not drawn a sharp distinction between snake-like devices and trunk-like devices, except to say that the actuation mechanisms must necessarily be self-contained in a snake or eel design. We will see examples of both types of devices in this paper.

Snake and trunk robots are typically classified according to their underlying physical structure. If the device consists of very many small links connected by rotational joints (similar to an actual snake backbone), then it is called a "High Degree of Freedom" (HDOF) device. On the other hand, if the backbone consists of a continuous structure, such as a flexible elastic member or pneumatic chamber, then it is termed a "continuum" device. In the context of mine detection, the distinction is somewhat academic, though we will argue later that continuum designs may eventually offer some benefits in complexity and weight reduction.

3 Current and Previous Research

The following work represents a good cross-section of the state of the art in robotic snakes, as well as the theories that support their design and operation. This is not an exhaustive list, but it is representative of the types and variations of the eight to ten snake prototypes that have been published.

3.1 The Tensor Arm

Perhaps the first device built using the HDOF methodology was the Tensor Arm Manipulator [1]. Dating back to the early 1960's, this arm featured 15 concentric plates ranging from about 10 to 15 cm in diameter. The center of each plate connects to its neighbors via two 2-degree of freedom joints, allowing a total of 28 degrees of freedom. Holes in the plates guide a large number of "tendons", or cables that, when pulled, can move given plates so that the entire devices has a trunk-like appearance. Not until the 1980's would any significant work in HDOF or snake robots appear again, but the concept of endowing a robot with far more degrees of freedom than strictly necessary – the HDOF concept – formed the basis for a research area termed "hyper-redundant" robotics, which includes the study of artificial snakes and eels.

3.2 The "Active Cord Mechanism"

Shigeo Hirose, of the Tokyo Institute of Technology, reinvigorated interest in snake and trunk robots by introducing the idea of the artificial articulated body in the mid

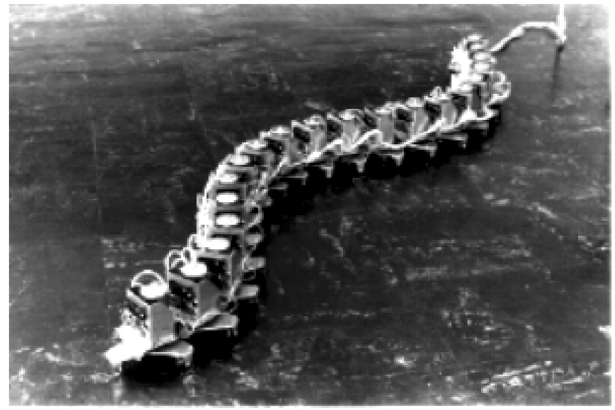

Figure 1: The ACM III snake is approximately 2 m long and weighs 28 kg.

1980's. His work rested on the premise that articulated bodies (simply structures composed of a serial chain of body segments) could perform several snake-like functions with the following technical advantages:

1. They can pass through narrow openings and over rough terrain by complying to objects on the ground.
2. The can cross soft ground because they have an extremely high weight-bearing surface area, relative to their overall weight.
3. The redundant structure increases reliability and maintainability because the device is modular, with most modules identical.
4. The robot can be easily transported by breaking it down into individual modules.

Hirose's group built several mechanisms to demonstrate the feasibility of this idea, most notably the Active Cord Mechanisms (or ACM) [8]. ACM III, perhaps the most renowned of the designs, incorporates 20 segments, each of which contains a motor capable of exerting torque against the segment ahead of it, as seen in figure 1. In order to facilitate the locomotion of the device, each segment rests on a pair of wheels and in [8] the ACM III can be seen weaving its way through several incredibly tight passageways. The ACM III was tethered for power, and because of its structure, could only maneuver across planar, flat surfaces.

3.3 The "Slim Slime Robot"

In order to address the inability of ACM III to move spatially, Hirose et. al. produced several additional gen-

Figure 2: The SSR robot is anywhere from 0.7 to 1.1 m in length, and unknown weight.

Figure 3: The Makro project's snake robot is about 1.6 m long and weighs 30 kg.

erations of snake-like mechanisms, including ACM IV-VII, ACM R1 and R2, and the Slim Slime Robot (SSR) [13], seen in figure 2. The SSR consists of at least six modules, pneumatically operated, with three internal bellows capable of generating pitch, yaw and axial extension/contraction of the module. The maximum deviation of those parameters are $\pm 30°$ for pitch and yaw, and length extension from 114mm to 177.6mm. This design represented a transition from HDOF devices to fundamentally continuous devices.

The introduction of modules capable of 3-dimensional motions opened up new possibilities for locomotive gait patterns. Where, previously, motion had been restricted to a planar surface, the primary gait generation mechanism was the propagation of a traveling wave down the spine of the snake. In fact, Hirose's team labeled a particularly useful mathematical form of this traveling wave the "serpenoid" curve. However, with the new designs, gait generation patterns could include rolling of several different varieties (due to the availability of both pitch and yaw), and "inch-worm" crawling (due to differential extension and contraction).

3.4 The "Makro" Project

The idea that snake-like robots could access tight spaces also led to a development effort in Germany, called the Makro Project [9]. The project's objective was to create a robot that can effectively examine and even repair the interior surfaces of sewage pipes, in order to mitigate the extreme expense and inconvenience associated with their failure. The Makro Project generated two similar designs of robot, culminating in Makro 1.1, in figure 3, a 15 degree-of-freedom snake with 6 modules. Each module rests on wheels, and is connected to the neighbor ahead of it via a hybrid universal joint that can exert pitch, yaw and roll motions. Because the wheels are driven on the Makro robots (as opposed to passive wheels on the ACM designs), they are able to successfully navigate pipes on the order of twice their diameter or less, while using their joints to climb over small obstacles and steps.

3.5 The NTUA Snake

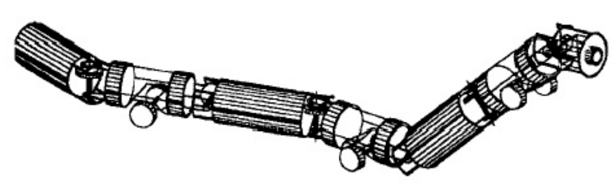

Figure 4: A drawing of the NTUA snake robot, which is 1.65 m long and weighs 16.5 kg.

Another design, produced by the National Technical University of Athens (NTUA) proposes that the wheels under each section not be directly actuated, but rather independently steered [10][14]. The NTUA prototype consists of seven segments, with six revolute joints and the ability to add more segments on as desired (as with all designs so far). While the joints in this robot are unidirectional, as with ACM III, they transpose orientations, i.e. every other joint rotates about a vertical axis, and the remaining joints rotate about horizontal axes. Only one half of the segments, or links, have wheels attached. The ability to steer the wheels permits their orientation to differ from that of the link above, and the designers argue that this allows for more efficient and effective locomotion strategies.

4 Locomotion

Clearly, the prototypes above all require locomotion schemes significantly more sophisticated than typical mobile robots, or even legged robots. To one degree or another, all of the designers of these systems have focused upon effective use of the robot kinematics for locomotion, and snake locomotion has developed as a subject unto itself, investigated even by researchers who are not participants in the construction of actual test devices. Interestingly, real snakes do not utilize all of the potential gait patterns that are available to them, and additional gait patterns may be available to devices that feature non-biomorphic designs, such as the roll capabilities of the Makro snakes. Thusfar, the primary locomotive patterns that have been identified are [3][5][8]:

1. *Lateral Undulation.* This is the typical image many people have of snakes. It is easy to mathematically quantify, and with slight modification does seem to represent the favorite locomotion mode for most snakes. The snake moves forward by the backward propagation of a single-frequency transverse (or lateral) traveling wave. Lateral undulation is also known as gliding.
2. *Side-winding.* Obviously employed by the Sidewinder snake and others, side-winding is faster and more efficient than lateral undulation. However, it requires a relatively smooth surface and a wide open area in which to “wind”.
3. *Rectilinear Propagation.* This can be imagined as lateral undulation “on edge”. Here the snake has only a few contact points with the ground, sending a vertical transverse wave propagating down the spine. This gait is very efficient, but suffers from instability because most of the snake is in the air, and the belly of the snake must be flat in order to prevent it from falling over. The snake backbone cannot have wheels for this gait.
4. *Lateral Rolling.* Rolling can occur in several ways; here one might imagine that the snake is a rope on the floor, forming a shallow U. Modulating the pitch/yaw axes of every joint causes the rope to roll in the direction of the U.
5. *Wheeled Rolling.* In wheeled rolling, the head of the snake bends up and back, while the tail bends up and forward. When the head and tail meet, the snake has formed a closed loop and can roll like a bulldozer tread. Certain caterpillars have been known to use this trick as an escape mechanism; it is very fast though suffers the same instability problems as rectilinear propagation.
6. *Axial Propagation.* As its name suggests, axial propagation requires axial compressibility, seen for instance in the SSR. Worms frequently employ axial propagation by sending an axial traveling wave down their spines in order to navigate extremely tight spaces. Several species of caterpillar use axial propagation in concert with rectilinear propagation to effect the “inch-worm” appearance.
7. *Concertina Motion.* The concertina motion is a variant of lateral undulation. However, rather than one continuous wave, the snake propagates a pulse along its spine.

Methods that would allow robotic snakes to climb over obstacles have also been investigated [12], as well as the gait patterns of articulated bodies in water [2][11]. At some point, all of these methods (except for climbing) have been demonstrated by machines in the laboratory. They have all met with varying degrees of success, and real snakes, caterpillars and worms make use of at least two (often more) gait modes depending on the situation. It is fair to say, however, that building a device capable of all these gait patterns probably represents an over-design, and likely a significant increase in complexity above the minimum required to simply move effectively.

5 Practical Concerns

As we have hinted throughout the text so far, a certain amount of complexity seems to be unavoidable with robotic snakes. The Engineer’s mantra applies here more than ever, "Make it as simple as possible, but no simpler." In order to successfully design a functional robotic snake, one must be able to evaluate the performance of the snake on the basis of certain metrics and then select a (not necessarily unique) minimum design that satisfies the metrics. With machines like snakes and trunks, it is tempting to succumb to the desire to build in extra functionality, but experience shows that this is generally not a wise choice.

Various researchers have investigated performance metrics, but perhaps the simplest and most useful come from Dowling [5]. As pointed out in [5], good metrics combine multiple physical measurements in meaningful ways. On their own, measurements such as power output, maximum speed, payload capacity, etc., are somewhat useful – and may dictate system design constraints – but also are often in conflict with one another. Several more sophisticated metrics include

1. *Cost of Transportation.*

   $$COT = \frac{energy}{distance * mass}$$

   Used on animals, this metric reveals the surprising result that large animals use energy more efficiently than small ones for the purpose of locomoting.

2. *Net Propulsive Efficiency.*

   $$NPE = \frac{weight * distance}{energy}$$

   NPE for vehicles is similar to the common "kilometers per liter" metric, but does not penalize heavier vehicles because they are able to carry heavier payloads. NPE is also a unitless metric, often desirable in the evaluation of machines of widely varying design.

3. *Specific Resistance.*

   $$SR = \frac{power}{weight * velocity}$$

   Another dimensionless metric, specific resistance captures the notion of the "energetic cost of locomotion".

The metrics above were selected because they all focus to one extent or another on the ability of a snake to carry a payload. However, they ignore other important realities like economics, reliability and maintainability, and environment.

Among other practical concerns in designing snakes are factors like power supply, electrical wiring, computational power, and sensing. As with many robotic applications, questions surrounding the power source can complicate robotic snake design considerably. If the power source is onboard, it must compete with the payload for space and locomotive energy. If it is offboard, a tether must accompany the snake wherever it goes. At first, the tether may sound like a drawback, but the snake's primary task may actually encourage the use of tethers. For instance, a snake designed to carry a chemical sensor into a mined area might get stuck due to a malfunction. Since the snake clearly traveled into the area because humans could not, they obviously cannot safely go retrieve it either. A strong tether may increase the chances of retrieving the snake and trying again. Also, snakes carrying onboard power sources to avoid tethers must also carry onboard computational facilities for command and control. While this has certainly been demonstrated, it naturally further restricts the payload capacity of the snake.

Computing power is another major factor in the design of a snake robots. By their very nature, snakes have a high number of actuators. At a minimum, each one requires at least one feedback mechanism, as well as a power source. Assume for the sake of argument that we design a snake with 20 degrees of freedom. Each actuator is given a position sensor with 8 bits of analog-to-digital precision, and is actuated with 8 bits of digital-to-analog precision. Thus, we need 20*8*2 = 320 bits, or 40 bytes of communication "width". At a minimum, a robust control scheme will require sampling intervals around 50 to 100Hz. Thus the bandwidth requirements are from 16 to 32 kbs. If we attempt to give the robot greater precision, say 12 bits, and a faster controller, say 500Hz updates, the bandwidth climbs to 480kbs, and that doesn't account for inputs from other sensors, of which a variety may be needed (e.g. contact sensors, force/torque sensors, the payload sensor itself), and communications overhead. Certainly these bandwidth requirements are well within the capabilities of modern microcontrollers, and some distribution of computational resources along the backbone is possible, but clearly the necessary circuitry and wiring will significantly complicate things – especially if it is all placed onboard.

It goes without saying that payload capacity is of the utmost importance in mine detection activities. Minimum payload requirements will likely be the primary constraint dictating the use of larger snake designs. To date, payload capacity has not been extensively studied, but several factors govern its upper limit. For one, snakes are generally not amenable to lumped payloads. Heavier payloads can be transported if they are distributed over the length of the snake; however, this is often not practical. Most of the gait mechanisms that snakes use suggest that payloads should ride up front, at the head, and be "pushed". On land, where gait patterns must be three-dimensional, the snake must be able to lift the payload slightly off the ground during transit (called "sinus lifting" on real snakes), and push the payload through undergrowth and over obstacles. Things are simpler in water: if the payload is at least neutrally buoyant, its primary negative contributions are viscous drag and additional inertia; size and weight are lesser concerns.

Sizing the payload, both in weight and dimensions, so as to permit effective locomotion presents a challenge. Clearly, the snake must have some amount of power to spare, over and above locomotive requirements. To get a rough estimate of the maximum payload capacity, first note Hirose's estimate that the maximum power availability for a single actuator must be greater than the product of the maximum required angular velocity times the maximum required torque, $P_{max} \geq \omega_{max}\tau_{max}$. Using his serpenoid curve, and assuming gliding motion, he then derived that

$$P_{max} \geq \left(\frac{\pi}{2}\right)^2 \left(\frac{1}{\sigma\,\mathrm{serp}(\sigma)}\right) \mu_t w v$$

where μ_t is the tangential friction coefficient (in the direction of snake motion), σ is a parameter on the order of 1 to 2, tending toward 1 in most snakes (depending on the exact geometry of the snake), $\mathrm{serp}(\sigma)$ is the serpenoid function, and w and v are the weight and desired linear speed of a given link, respectively. Given that uneven and compliant surfaces will sometimes transmit weight from the payload link to the surrounding links, and that the forces overcoming tangential friction are generated by contributions from all of the links, it is conservative but reasonable to say that each of the actuators must be "beefed up" to a level consistent with the constraint above, *as it applies to the heaviest link.*

As an example, Hirose's early ACM designs used around 20 links, each weighing around 1.4kg with a 10W DC servomotor. Given $\sigma = 1$, $\mu_t = 0.034$ (by the use of wheels), and $v = .4$ m/s, the motors were only required to output less than 1W [8]. However, if the snake was to carry a payload of twice its weight ($w = 2.8$ kg) across rough terrain ($\mu_t = 0.15$), the power requirement increases to almost the full 10W – which, we previously argued, would likely apply to all of the actuators (or at least several in the vicinity of the payload). Thus, Hirose's ACM III, at a length of 2m, weighing 28kg, would be able to safely carry slightly over 17kg over smooth surfaces on level ground. Up slight inclines, on rougher ground or through brush, the payload capacity would quickly drop to near zero for the given backbone velocity, and in fact the authors of [14] say that the NTUA snake can carry only 0.5 kg of instrumentation.

The good news is that, for all the issues surrounding complexity and capabilities, several snake robots have been built using commercially available parts and a very simple design [5][15]. In part to illustrate that it could be done, Dowling chose simple hobby servos, linking 20 of them together in an alternating pitch/yaw configuration (similar to the NTUA snake). The snake was able to demonstrate the common gaits, and even carried a small camera at its head. Very little additional onboard electronics were necessary, as Dowling chose to operate the servos at their native control rate (50Hz), using their integrated controllers, through a tether. (He did not perform experiments regarding payload capacity.) NASA Ames Research Center later demonstrated similar prototypes, called "Snakebots", also built using hobby servos.

6 The Future

It has been the assessment of the first author, as well as those of [11], that practical snake technologies will not develop further until the principles at the heart of highly-articulated and continuous bodies are better understood. For instance, currently popular models for open-chain dynamics and simulation do not lend themselves to HDOF devices with dozens of links. Better would be to develop comprehensive theories that treat the whole snake like a truly continuous device [4]. Also, most snake designers have glossed over the complexities of modeling the environment, which is why all of the snake prototypes to date have operated only in very well-defined and predictable locations (usually flat, smooth surfaces).

To address the environment problem, the authors of [11] limited their work to eels, i.e. snakes in water, building a simple prototype with 3 links and 2 joints. An aqueous environment is significantly simpler to grasp, both from a mathematical and experimental point of view. For one, there is no need for prototypes to operate spatially if they are simply swimming on the surface. For another, straightforward fluid characteristics dominate the analysis and operation, as opposed to the unpredictable and often intractable characteristics of surfaces and the different types of friction associated with sliding (not to mention objects and uneven terrain). Finally, as mentioned before, robotic eels can accommodate significantly larger payloads because, while viscous drag and inertia may be problematic, excessive weight can be offset by additional buoyancy.

Gravagne et. al. have stepped back and noted that, if we are trying to build and analyze structures that closely approximate the curves of snakes and eels, why not forgo the individual links and use one continuous elastic element as the backbone? The development of analysis and simulation techniques for continuous backbones under large displacements has taken time, but is maturing into a unified theory of "continuum robot" kinematics, dynamics and control [6][7]. For instance, the "Tentacle" manipulator in figure 5 is of very simple design, with a continuously flexible backbone. Based on experience with large-deflection flexible elements, these researchers theorize that, with a good understanding of the dynamics of large-displacement flexible elements, the undulatory locomotion of snakes and eels need not be actuated in a quasistatic manner. In other words, it is probable that snake-like undulation could be effectively reproduced without actuating every point along the backbone, but rather by using limited actuation at certain locations and relying on dynamic effects to transmit and distribute the resulting locomotive energy.

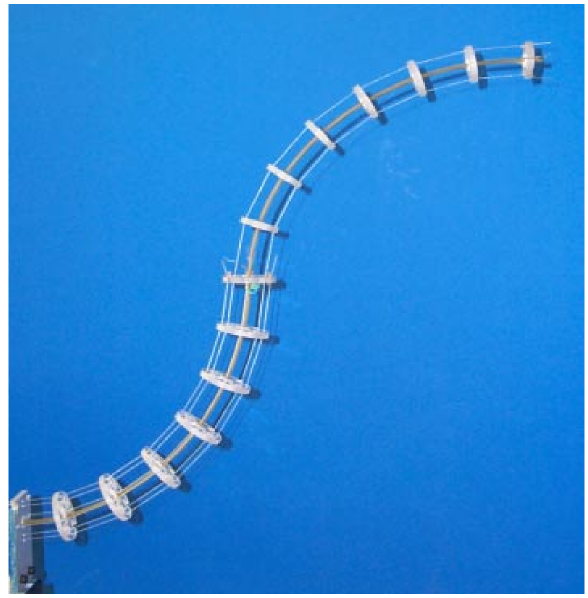

Figure 5: The Clemson tentacle robot. Unlike HDOF designs, the backbone is one long elastic rod.

The practical realization of this idea would be eels and snakes that are significantly less complex than the devices constructed to date. The next obvious step is to attempt to integrate actuation technologies that are better suited to continuous structures than are motors, e.g. artificial muscles, active fiber composites or shape memory alloys. Such hybrid actuator-backbones would be thin and flexible, have significantly less overall mass, no moving parts (in the sense of spinning shafts, bearings, gears, propellers, etc.), would be easier to water-proof, and would be suited for high efficiency propulsion in littoral environments. This type of undulatory locomotion, known as "anguilliform" locomotion, has also been studied from a biological perspective by Ayers, et. al. [2], in relation to mine detection.

The timeline that dictates successful deployment of robotic snakes for mine detection depends greatly on their exact task requirements and, of course, the environment. Obviously, functional prototypes already exist; clearly, however, they are not suited for "real-world" environments yet. Assuming that the requisite sensors can be made small enough (or the snake large enough), functional eels with payloads could be only 3 to 5 years away. For crawling snakes, it is the first author's opinion that further investigation is needed in the direction of Dowling, i.e. attempts to build robust snakes that can operate in a variety of places from easily available and replaceable parts, because complexity will be a quickly limiting factor. Especially for mine detection, the snake simply cannot cost too much because it might not return from its mission. While Dowling and others have illustrated that the mechanical components can be readily found, whether they can be appropriately scaled to fit the sensor requirements remains unknown. Further electrical and communication upgrades would be necessary as well. Such improvements could probably also result in a functional, payload-carrying prototype in a few years, but reliability field tests would be more difficult than for eels, and it is almost inevitable that field testing will spawn several iterations of redesign and further testing as the researchers gather more experience and data regarding various realistic environmental surface conditions.

7 Conclusions

Much remains to be done before robotic snakes and eels become effective "mine-sniffers". Among the most pressing needs, current snake designs (especially those that use freely available parts) need to be comprehensively evaluated and tested in realistic conditions. It is doubtful that any of the existing designs will be up to the challenge of sniffing out mines yet. Reductions in complexity would be highly desirable, perhaps by the "dynamic locomotion" strategy proffered by the first author, or simply by better design. And, of course, snake robot designers need realistic data about the dimensions and capabilities of the sensor payloads – the whole topic is mute if sensors cannot be made relatively small and portable.

Whether spurred on by applications like mine detection or not, robotic snake, eel and trunk technologies and research have gained significant momentum over the last decade. By almost any standards, monumental progress has been made – from almost no knowledge and no experimental verification in the 1980s, to something like 10 published prototypes, around 20 to 30 active researchers worldwide and dozens of research papers in the span of about a decade. Still, on the scale of other activities in robotics and its associated fields of research, snakes, eels and trunks garner relatively small overall efforts. Their benefits are numerous – the ability to crawl (or swim) on one of many different modes, operate upside down and sideways, in cramped spaces and very shallow waters, and underneath danger zones posed by trip wires, among others. The advancement of snake technologies also has direct application to other areas, like biomedical engineering (endoscope and laparoscope design), large-deflection vibration control and microelectromechanical devices (where it is extremely difficult to build motors, links, joints and bearings but relatively easy to build flexible elements). The authors hope that snakes will one day prove useful in humanitarian de-mining activities as well.

References

[1] V.C. Anderson and R.C. Horn, "Tensor Arm Manipulator Design," American Society of Mechanical Engineers, paper #67-DE-57, 1967

[2] Ayers, J, Zavracky, P., McGruer, N., Massa. D., Vorus, W., Mukherjee, R. and Currie, S., "A Modular Behavioral-Based Architecture for Biomimetic Autonomous Underwater Robots," Autonomous Vehicles in Mine Countermeasures Symposium, Naval Postgraduate School, 1998

[3] G.S. Chirikjian, "Theory and Applications of Hyper-Redundant Robotic Manipulators," Ph.D. dissertation, Dept. of Applied Mechanics, California Institute of Technology, June, 1992

[4] G.S. Chirikjian, "Hyper-Redundant Manipulator Dynamics: A Continuum Approximation," Advanced Robotics, vol. 9, no. 3, 1995, pp. 217-243

[5] K. Dowling, "Limbless Locomotion: Learning to Crawl with a Robot Snake," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, December 1997

[6] I. Gravagne, I.D. Walker, "On the Kinematics of Remotely-Actuated Continuum Robots," IEEE Int. Conf. on Robotics and Automation (ICRA), San Francisco, May 2000, pp. 2544-2550

[7] I. Gravagne, I.D. Walker, "Manipulability, Force and Compliance Analysis for Planar Continuum Robots," to appear, Transactions on Robotics and Automation

[8] S. Hirose, *Biologically Inspired Robots*. Oxford University Press, 1993

[9] B. Klaasen, K. Paap, "GMD-SNAKE2: A Snake-Like Robot Driven by Wheels and a Method for Motion Control," IEEE Int'l Conf. on Robotics and Automation (ICRA), Detroit, Michigan, May 1999, pp. 3014-3019

[10] K.J. Kyriakopoulos, G. Migadis, K. Sarrigeorgidis, "The NTUA Snake: Design, Planar Kinematics and Motion Planning," J. Robotic Systems, vol. 16, no. 1, 1999, pp. 17-72

[11] K. McIsaac, J.P. Ostrowski, "A Geometric Approach to Anguilliform Locomotion: Modeling of an Underwater Eel Robot," IEEE Int'l Conf. on Robotics and Automation (ICRA), Detroit, Michigan, May 1999, pp. 2843-2848

[12] M. Nillson, "Snake Robot Free Climbing," IEEE Control Systems Mag., Feb 1998, pp. 21-26

[13] H. Ohno, S. Hirose, "Study on Slim Slime Robot," RSJ/IEEE Int'l Conf. on Intelligent Robots and Systems (IROS), Takamatsu, Japan, October 2000, pp. 2218-2223

[14] K. Sarrigeorgidis, K.J. Kyriakopoulos, "Motion Control of the N.T.U.A. Robotic Snake on a Planar Surface," IEEE Int'l Conf. on Robotics and Automation (ICRA), Leuven, Belgium, May 1998, pp. 2977-2982

[15] "NASA Developing 'Snakebot' to Explore and Build in Space," press release 00-66AR, October 4, 2000