
Limbless Locomotion: Learning to Crawl with a Snake Robot

Submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy in Robotics

Kevin J. Dowling


Advised by William L. Whittaker

The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

December 1997

This research was supported in part by NASA Graduate Fellowships 1994, 1995 and 1996. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of NASA or the U.S. Government.

© 1997 by Kevin Dowling.



Limbless Locomotion: Learning to Crawl

Snake robots that learn to locomote

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Robotics by Kevin Dowling

The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213

Robots can locomote using body motions; not wheels or legs. Natural analogues, such as snakes, although capable of such locomotion, are understood only in a qualitative sense and the detailed mechanics, sensing and control of snake motions are not well understood.

Historically, mobile vehicles for terrestrial use have either been wheeled, tracked or legged. Prior art reveals several serpentine locomotor efforts, but there is little in the way of practical mechanisms and flexible control for limbless locomoting devices. Those mechanisms that exist in the laboratory exhibit only the rough features of natural limbless locomotors such as snakes.

The motivation for this work stems from environments where traditional machines are precluded due to size or shape and where appendages such as wheels or legs cause entrapment or failure. Example environments include tight spaces, long narrow interior traverses, and movement over loose materials and terrains. Several applications, including industrial inspection and exploration of hazardous environments, compel serpentine robots.

This research develops a general framework for teaching a complex electromechanical robot to become mobile where sequences of body motions alone provide progression. The framework incorporates a learning technique, physical modeling, metrics for evaluation, and the transfer of results to a snake-like mobile robot. The mechanism and control of a 20 degree of freedom snake robot is described and multiple gaits are demonstrated including novel non-biological gaits. This research furthers the design and control of limbless robots.

© 1997 by Kevin Dowling

We shall not cease from exploration
And the end of all our exploring
Will be to arrive where we started
And know the place for the first time

T.S. Eliot

Acknowledgments

Research is hard but involves great joy as well. The greatest of joys has been working with the people here at CMU. An observer might have thought I was working alone - but the critical mass of people here in the Robotics Institute meant that I could always find the help and encouragement on all issues.

Red - Friend, mentor, and force of nature. Thank you Red.

Hans Moravec - Once upon a time, Hans hired yours truly, an eager but inexperienced undergrad, to help build his robots. Hans always has a fresh perspective, new insight, and a wonderful way of looking at things. I will dearly miss the discussions.

Mike Blackwell - Friend and officemate of fifteen years, Mike understands the combination of hardware and software better than anyone I know and has been an invaluable help in answering a zillion questions over the years. Thanks Mike.

John Bares - John, Your review, help and advice and friendship have been invaluable.

Dave Wettergreen - The clearest of voices, the best of reviewers and good friend.

Dave Simon - Always weighing possibilities and thinking out issues. Thanks for the advice and support over the years.

Hagen Schempf - A potent combination of enthusiasm, talent and experience. Thanks Hagen, for advising, helping and making things happen.

Ben Brown - The best designer I know, and a sounding board on many technical issues herein including metrics and design.

Chris Leger - A remarkable programmer and Quake enthusiast. Chris developed a software toolkit that I used in this work.

Anton Staaf - A remarkable undergrad who is destined to do great things. Thanks for the caterpillar and the discussions, Anton.

Sundar Vedula, David Baraff, Martial Hebert, Andrew Moore, Howie Choset and Joel Burdick all provided advice and insights into several areas of this research. I greatly appreciate their time, perspectives and assistance.

Acknowledgments

Thanks to Tony Nolla, Jesse Eusades and Dave Vehec for their assistance on some wiring and drawings.

Thanks also to Takeo and Raj who have also advised, mentored, and supported me through the years. It's been an enormous and beneficial influence.

The members of the Field Robotics Center, the Robotics Institute, and friends throughout Carnegie Mellon. This is the best place in the world for robot research.

Most of all, Mary Jo and Ashlenn, and our most recent research project: Aidan. Your love, support, advice and understanding are monumental. I love you.

Contents

Overview and Rationale	3
Introduction	3
Why Serpentine Locomotion?	5
Applications	9
Challenges of Limbless Locomotion	12
Summary	12
Background	13
Biological Systems	13
What's Missing?	20
Robotic Systems	20
Learning	32
Summary	34
Framework	37
Overview	37
Performance Metrics	43
Metrics	43
Efficiency	44
Dimensionless Metrics	49
Scale	53
Summary and Selection	55
Learning and Optimization	57
Optimization Techniques	57
Representation	61
Summary	67

Implementation	69
Actuation	69
Design	77
Electronics	87
Sensing	89
Other Subsystems	90
Experimental Setup	92
Physical Modeling	93
Summary	94
Locomotion	97
Power Calculations	97
Units	100
Velocity Calculations	101
Gaits	101
Summary	119
Summary and Conclusions	121
Contributions	121
Future Work	122
Wisdom	123
Servo Evaluation	125
Link Weight Distribution	127
Derivation of Actuator Parameters	129
References	135

Chapter 1

Overview and Rationale

Overview and Rationale profiles the content of this dissertation and examines the rationale for serpentine robots and their application. This chapter offers strong motivation for serpentine mechanisms; this includes the advantages and disadvantages of serpentine locomotion as well as application areas where such mechanisms can make a powerful impact.

Introduction

Biological snakes are pervasive across the planet; their diverse locomotion modes and physiology make them supremely adapted for the wide variety of terrains, environments and climates that they inhabit. It would be wonderful to capture these broad features of movement and capability in man-made equivalents. While wheeled and walking machines have undergone decades, even centuries, of development, they are still limited in the types of terrain they can traverse. A snake-like device that could slide, glide and slither could open up many applications in exploration, hazardous environments, inspection and medical interventions.

Serpentine robots have a number of useful features and applications, are fascinating to observe, and may answer questions about biological equivalents. One of the fundamental issues is understanding their locomotion. In a qualitative sense, propulsion with wheels or legs is more readily apparent and understandable versus the movement of limbless locomotors. A wheel turns; the vehicle moves. A leg pushes; the vehicle moves. How a snake moves is not so evident. This dissertation addresses robots that crawl and slither without the use of wheels or legs; *where body motions alone enable progression*.

A worthwhile snake robot has the ability to wriggle into confined areas and traverse terrain that would pose problems for traditional wheeled or legged robots. The useful

features of snake robots include stability, terrainability, good traction, high redundancy and completely sealed mechanisms. Robots with these properties open up several critical applications in exploration, reconnaissance, medicine and inspection.

This research creates a robot snake that can locomote in novel ways, develops a framework for teaching the snake to locomote and also develops and integrates the many technologies required to make this happen. The research culminates in the demonstration of an effective mechanism and the enabling gait generation techniques. The research is distinguished from prior work by the development of more general locomotion, a better mechanism capable of traversing more complex terrain and the demonstration of multiple locomotion modes.

Prior work in understanding real snakes has been limited to providing qualitative descriptions of snake movement and the classifications of gaits. While a few snake robot devices have been built, they have shown only limited locomotion in the form of single types of gaits.

Another hurdle is the evaluation of gaits and the determination of what is a good gait. After examining and evaluating many criteria and measures, a non-dimensional metric, specific resistance, offers the most useful measure of gait efficacy. However, gait evaluation requires that gaits be generated so they can be tested. Most prior work in gait generation provides an explicit description of a gait. However, this assumes a good model or good imagination to provide a variety of gaits. Another way, developed here, is to let a learning process generate gaits and test them using the performance metric as an evaluation function.

Rather than simply generating gaits in a random fashion however, a guided learning process is utilized. Stochastic learning processes, genetic algorithms (GA) and probabilistic-based incremental learning (PBIL) methods provide the most general technique for generating and testing of gaits. The results of this work are not only several snake-like gaits but several novel gaits, not found in real snakes.

Because real mechanisms are susceptible to breakdown and wear, especially when confronted with rapid changes and extreme ranges of inputs, it would not be possible to test all gaits on a real mechanism. Thus, physical simulation is required; using the computer to accurately simulated the physics of the body, commands and resulting motions from those commands. Physical simulation offers the potential of quickly evaluating many gaits and motion sequences.

The design and implementation of a snake robot is the confluence of several technologies: actuation, form and structure, electronics, control, sensing, etcetera. Mature technologies for emulating characteristics of biological muscle lie in the future. However, electromagnetic motors offer many options, and the robot snake is designed around these technologies.

The framework that ties all of this together is, in itself, a separate development. By incorporating learning, evaluation, simulation and the robot, a complex system is made to locomote in surprising and new ways. Another surprising development is the resulting simplicity of the gait formulation; it took a long march to come back to a simple formulation.

Outline

Overview and Rationale presents the advantages, disadvantages and applications of snake-type robots. **Background** looks at prior efforts in understanding limbless locomotion in animals such as snakes and discusses what is still not understood in these animals. Structure and locomotion modes, in particular, are examined in detail. Additionally, several prior snake-like robots, their mechanism and control contributions as well as their limitations are examined.

Framework presents a structure or architecture for learning locomotion and control of a serpentine robot. The structure or framework itself represents the outline for the remaining chapters and each component is treated at length in the following chapters.

What is a good way to evaluate locomotion? What makes one gait or sequence better than another? Learning requires a quantitative measure of performance and

Performance Metrics examines ways to evaluate performance in locomotion and selects *specific resistance* as a metric for learning limbless locomotion. This selection is preceded by an examination of a variety of units, dimensions and dimensionless metrics.

Learning and Optimization explores machine learning techniques to maximize movement. A stochastic method, probabilistic-based incremental learning, is selected from several learning methods.

Implementation describes the robot developed in the course of this research. The configuration and design includes actuation selection, mechanical design and electronics and hardware control. Additionally, the experimental set-up and other details of implementation are shown.

Locomotion codifies the outcome of the experiments in simulation and with the physical robot and reports on the successes and failures of the many experiments. A number of gaits, both snake-like and non snake-like, are revealed.

Finally, **Summary & Conclusions** codifies the results and contributions of this work, and looks towards the future where new contributions can be made. A little wisdom is also dispensed.

The appendices include a detailed look at evaluation of a class of actuators, parameter derivation and weight tables. Finally, a detailed bibliography completes the dissertation.

Why Serpentine Locomotion?

For centuries, people have created a menagerie of machines whose appearance and movement have mirrored animals to an astonishing degree. There are anthropomorphic figures that resemble man and mobile machines that resemble animals. However, the strongest reactions are not simply to outward appearance; after all, costumes, statues and puppets are extant throughout history. What attracts and holds attention are the animated behaviors and motions without apparent human connection.

However, these historical devices are mere simulcra, controlled by unseen hands; only in the last few decades have researchers and designers begun to replicate the general movements of animals in mechanisms.

The general motivation for serpentine locomotors are environments where traditional machines are precluded due to size or shape and where appendages such as wheels or legs cause entrapment or failure. Example environments include tight spaces, long narrow interior traverses, and travel over loose materials and terrains.

Serpentine mechanisms hold particular fascination due to the singular motions usually associated with animals such as snakes and tentacles. Few terrestrial mobile devices move without the use of wheels or legs; those that exist in the laboratory have exhibited only the rough features of natural limbless locomotors such as snakes. Serpentine features include serial chains of actuators capable of subtending small curvatures. However many of these prior efforts incorporated non-biological features: the use of casters for support and propulsion or the use of fixed pins for support and traction.

Other broad features of these prior robots include the use of models that explicitly describe the shape of the robot, the use of tensor mechanisms that limit curvatures and forms and mechanism designs that are impractical for application. There are significant challenges in designing, building and controlling practical limbless mechanisms that are capable of locomoting without traditional forms of propulsion and actuation. These challenges include configuration, design and geometry of the form, determining the number and arrangement of actuators, routing power and signal distribution, and robot control.

Wheels offer smooth and efficient locomotion but often require modifications to terrains for best use; even all-wheel-drive mechanisms are limited in the type and scale of terrain that can be traversed [Bekker 61]. Walking mechanisms offer extreme terrain negotiation for a given scale and provide discrete, rather than continuous, contact. This, in principle, benefits efficiency and traversability. The current state of the practice, outside of the laboratory, are walkers such as the Komatsu RECUS, CMU's Dante, the Finnish Plustech forestry machine, and Honda's anthropomorphic biped [Ishino 83][Apostolopoulos 95] [Plustech 96][Honda 96]. The number of efficient and practical walkers is small and there is much development and incentive necessary for walkers to become viable.

Advantages of Serpentine Robot Locomotion

Serpentine locomotors possess a number of potential advantages beyond the capabilities of most wheeled and legged vehicles.

Stability

Unless a serpentine robot purposefully slithers off a cliff, it can't fall over. In contrast, stability is of great concern to wheeled and legged machines in rough terrain; they can fall over. Terrain contacts in vehicles form a constellation of points on the terrain; if the center of gravity moves beyond the bounds of the convex polygon formed by these points, it tips over. In a serpentine robot, the potential energy remains low in most situations; therefore there are few concerns for stability and no need for the support polygons formed by wheel or leg contact points. Even in the case where a free-fall occurs, the serpentine device may survive better than most mobile devices because potential failure points such as the connections between the body and wheels or legs do not exist.

Terrainability

Terrainability is the ability of a vehicle to traverse rough terrain. Terrain roughness is often measured by scale of features, power spectral density, distribution of obstacles such as rocks and geographic forms [Bekker 69], or even its fractal dimension [Arakawa 93]. A serpentine mechanism holds the promise of climbing heights many times its own girth; this feature can enable passage through terrain that would encumber or defeat similarly scaled wheeled and legged machines.

Additionally, a serpentine robot can climb steps whose heights approach its longest linear dimension. This is an attribute that few, if any, wheeled or legged mechanisms possess. This assertion assumes quasi-static systems; mobile leaping systems, like the Russian Phobos vehicle, might jump many times their height or length [Klaes 90]. While there have been numerous wheeled trackless “trains,” or coupled-mobility devices, that use powered wheels, they still suffer from the limitations of wheel traction and terrain shear [Hirose 93][Odetics 88][Gowenlock 96][Haddock 94]. Many coupled-mobility devices make use of active or passive wheels to move and body joints to accommodate obstacles. However, the wheel still limits locomotion over soft and viscous materials. Additionally the serpentine mechanism has no appendages that can become stuck unlike the shank-rocking of a leg or the descent of a wheel into a hole.

Traction

Traction is the force that can be applied to propel a vehicle. Traction is usually limited to the product of the vehicle weight and the coefficient of friction. Tractive forces can be quite high for natural snakes; a moving snake can exert a force up to a third of its own weight. The distribution of the snake mass over such a large area, in comparison to mass equivalent legged or wheeled vehicles, results in forces that can be below the thresholds of the plastic deformation of the soil. In comparison, load concentration resulting from most wheels or leg designs results in soil work. Because of the large contact area, serpentine vehicles may result in little or no soil work. If restrained or locomoting in certain modes, natural snakes can sustain a pull of up to four or five times its weight [Parker 63]. As an unscaled comparison, large man-made vehicles under good conditions and slow speed may exert drawbar pull of 90% of their weight [Caterpillar 94]. Limbless locomotion may prove superior in marginal or soft terrains where plowing and shearing actions restrict wheel mobility.

Efficiency

Snakes achieve efficiencies and performance equivalent to biomechanisms of similar scale and mass [Walton 90]. Reasons include reduced costs associated with less lifting of the center of gravity as compared to legged animals, elimination of the acceleration or deceleration of limbs, and low cost for body support. This begs the question: why wouldn't natural snakes be *more* efficient than similarly sized animals? The answer is that energy losses in snakes include greater frictional losses into the ground, lateral accelerations of the body that do not contribute to forward motion, and the cost of body support for partial body elevations during movement. These advantages and disadvantages appear to balance total energy use to a comparable level as animals of similar mass [Secor 92].

Size

Depending on the mechanism design, the small frontal area of snake mechanisms allows penetration of smaller cross-sectional areas than mass-equivalent legged or wheeled vehicles. If the volume of a snake, a cylindrical form, is kept the same and the diameter is reduced by half, the length becomes four times greater. Cross-sectional area for mechanisms of similar density and mass may result in very long vehicles.

Redundancy

Candidate configurations for serpentine robots may employ many simple motion actuators in sequence. During operation, the loss of short segments would still permit mobility and maneuverability. That is, the mechanism is still able to move and maneuver even if a number of actuators failed. The penalty, of course, is reduced efficacy in mobility.

Sealing

With its continuous unperforated surface a serpentine device has no appendages to impede progress or be exposed to surroundings. This allows better sealing between the mechanism internals and the environment. This provides advantage to applications in hostile environments.

Disadvantages of Serpentine Robot Locomotion

If snake mechanisms are so good, why aren't there lots of them? One answer is they are difficult to design, build and control. Another is that there are some disadvantages to these configurations.

Payload

Much locomotion has to do with work; the transport of materials from one place to another. There is no integral platform for attaching payloads. It's difficult to envision transport of materials using snake-like robots unless an integral conduit can be used to deliver materials.

Degrees of Freedom

To subtend the various curves needed for locomotion requires a larger number of actuators than most wheeled or legged vehicles. The number of DOFs in vehicles can range from two up to eighteen and even more for some walkers. However, a relatively flexible snake mechanism may require even more. A large number of DOFs may introduce reliability problems; if one actuator has a given failure rate then robots with large numbers of units have a higher chance of having any unit fail. Fortunately, for the serpentine mechanism, sufficient redundancy can allow the robot to continue to function in a limited manner. While the control for a serpentine mechanism involves more motions to control, an advantage is that complex planning for footholds and wheel contacts is obviated; the system can simply follow its head.

Related to this issue is designing actuators and structures that are strong, efficient, and elegant; the need is for high forces in small packages. However, this need is not unique to serpentine mechanisms, and many applications await the emergence of actuators with these properties!

Thermal Control

Obviously, snakes are not even close to being spherical; the volume to surface area ratio is worse than for animals of similar mass. Even though limbed animals have protruding limbs and appendages, the surface area to volume ratio is significantly less than for snakes. The Meeh coefficient, k , in the equation $S = kM^{0.67}$, where S is surface area and M is body mass, is higher for snakes than for many other mammals and fish. [Schmidt-Nielsen 84]. The effect of this may be that thermal control is more difficult in a serpentine mechanism. On the other hand, if the application allows the use of the environment as heat-sink or heat-source, then this works in the snake's favor and is of great benefit to actuator systems.

Speed

The fastest natural snakes, under ideal conditions, can move at 3.0 m/s and appear to have a length to circumference ratio of about 10-12 [Bauchot 94]. It seems unlikely that a robot system, in the near future, will develop speeds anywhere near this. Most snake locomotion is fairly slow, but the motion is deceptively fast however; the lateral motions of the body often give the impression of higher speeds. However, the bottom line is that robot snakes are likely to be slower than their natural counterparts.

Applications

What good are snake robots? Where would they be used? Consultation with potential users, and examination of many application areas suggests a number of areas where serpentine robots can make an impact.

In the past, a recurring litany of robotics application areas included nuclear plants, medical applications and the inspection of hard-to-reach areas. The difficulty in accepting these speculations resulted from immaturity of the technology and techniques. Robots prefer strongly structured applications, and many applications do not offer structured environments.

However, maturing and evolving technology in sensing, control and machine learning has enabled the successful deployment of operational field robots in unstructured environments and this will be true of serpentine robots as well. Each of the following applications offers a compelling scenario for self-propelled serpentine devices. Each application offers pratfalls and failure for wheeled or legged robots; problems and issues where serpentine robots could succeed.

There is a separate issue of fixed-base serpentine devices and several of these applications would benefit from serpentine manipulators as well as serpentine locomotors. However, this work is concerned with locomotion and not simply manipulation.

Exploration

In unpredictable environments, there are zones of uncertainty and footing is insecure or unknown. A snake-like device can distribute its mass over a large area for support so that even if footing gives way, self-support between secure points enables continued operation. Such environments include planetary surfaces and extreme terrains with loose rubble and inclines near the slope of repose.

Inspection

Many inspection techniques in industry and medicine rely on fixed-base mechanisms such as borescopes, videoscopes and fiberscopes. These devices are primarily used to inspect cavities that cannot be seen directly by the eye. Inspection applications include airline engine maintenance, quality control in manufacturing, and process monitoring and inspection in utilities and chemical plants. Simple direct-view borescopes have proven useful, but articulated self-advancing devices forming and following complex paths could open many more applications.

To eliminate some of the difficulties with current borescope use, plant equipment is modified with portals, but this requires additional design and manufacturing resources but doesn't address needs of older or legacy plants. Such equipment would not require such alterations if a device capable of reaching those points were available. Another real need is the inspection of power station cooling tubes which can be up to 18m long and only 10mm in diameter. [Olympus 94][VIT 95][Westinghouse 97].

Utilities, chemical processors and manufacturers have large and complex pipe networks that often require inspections or determination of blockage. Guesswork, followed by trenching or cutting operations, is a very expensive technique even discounting the associated downtime costs. Mobile pipeline devices are used by pipeline service companies but these pipe pigs, as they are termed, are of limited use. A snake-like device would prove very useful. With a serpentine tool, in-situ inspections and accurate localization could lower costs and downtime significantly.

Aware of their limitations, developers of inspection equipment are keenly interested in self-propelled inspection devices. However, the industry is highly competitive, and it is difficult to get information on their efforts. Yet, their interest is evidenced by a growing number of patents on 'walking' and self-propelled videoscopes [Welch-Allyn 94][Olympus 94]. None of these devices are yet mature.

Medical

Snake-like devices have received attention as a potential medical technology. Minimally-invasive surgery reduces or eliminates the need to cut open large sections of skin and tissue. It is currently estimated that 35% of the 21,000,000 surgeries performed each year could be done with minimally invasive techniques [Grundfest 94]. There could be dramatic reductions in hospital stays, patient suffering, and costs. Laparoscopic devices, which are rigid tools inserted into the abdominal wall, and endoscopic devices are used in these types of surgical procedures. A snake-like robot could subtend the curvatures of interior tissues and enable further diagnosis and treatment.

In recent years, non-invasive surgery has met with wide acceptance and produced phenomenal results. The surgical tools of the trade, however, are often difficult to manage and have their limitations. There are many needs for dextrous and articulated tooling and surgical devices that can advance through organs and tissue. As one example, about 60% of the gastro-intestinal tract is inaccessible to conventional endoscopic tooling.

However, substantial impact to existing procedures could be made in other areas, and gastro-intestinal endoscopy is one such example. Two leading companies have several

patents in this area, but neither have self-propelled endoscopic devices on the market yet. According to one company, a market does not yet exist, but the primary reason is that the costs are high [Welch-Allyn 94][Olympus 94].

Hazardous Environments

Human activity is precluded in many areas where there are extremes of radiation, temperature, chemical toxicity, pressure or structural weakness. However, it is often necessary to explore and survey these areas to insure safety and ascertain status. A variety of small tracked or wheeled machines have been constructed for such applications, but these have limitations in their ability to traverse and maneuver through hazardous terrain [VIT 95] [Gothard 90][RSI 94][Sasaki 85][Eguchi 84]. A serpentine mechanism could fare better due to the advantages cited earlier.

Other dangerous areas include those following disasters such as earthquakes, explosions, cave-ins, hurricanes, fires etcetera. The search for survivors and removal of material is often thwarted by loose rubble that might be penetrable by a snake. Outfitted with sensors such as ammonia or pyroelectric IR detectors, a snake-like mechanism would enable sensing of humans in rubble. These are applications that would eliminate life-endangering alternatives such as using heavy construction equipment to move loose material from accident sites.

Another application is a mine accident probe. Following a cave-in or roof collapse, a small articulated device would penetrate and burrow through the loose material to effect a survey and establish communication to survivors.

Reconnaissance

Subterfuge and reconnaissance offer some novel applications of serpentine mechanisms. The ability to command small roving eyes and ears offers attractive possibilities to law enforcement agencies; the penetration of dense vegetation by serpentine robots could provide information not otherwise possible to obtain. My own work has resulted in inquiries from law enforcement agencies including the FBI and Special Forces.

Routing

Much effort in the wiring of existing structures requires routing of cables and lines through narrow passages behind existing walls and through pipes. A variety of manual tools for feeding the lines, such as fishtapes (metal bands), are useful for short runs but become more difficult to use in longer runs. While some specialized devices have been designed for wire and cable routing they are not used in practice [Hill 65].

These tasks involve long reaches, wrestling with tools, and stretching and working in awkward positions. In practice, snake-like devices would maneuver through crowded plenums and pull the initial lightweight tapes that are then used to pull the actual cables. I have shown a variety of applications where serpentine robots could provide significant advances in productivity over existing methods.

Challenges of Limbless Locomotion

While the features and advantages and the applications for serpentine robots are attractive, there remain many challenges in realizing such robots. To create a truly successful snake robot requires that all areas be addressed and solved. These must be pondered and evaluated concurrently; design affects function. Integration is complicated, even intractable, if individual areas are not thought of in the whole.

Configuration and Design

The challenge of configuration is determining the form of a robot. The challenge of actuation is determining the technology that drives the mechanism. The questions are sometimes mundane but essential to answer: How long should segments be? What angle should they subtend? Are there actuation techniques that can provide smoother curves? Determining both the result and implications of each decision is a challenge. These are addressed in later chapters.

Infrastructure and Electronics

Supplying and routing power and signals in complex robots is often underestimated as a design task. Serpentine robots must be compact and small to accrue the advantages shown in the previous section. Small size burdens the tasks of wire routing and actuation support.

Control and sensing

Finally, the greatest challenge: how to learn to control such a device? A larger issue is determining the process, method and framework to achieve this.

Summary

The advantages of snake locomotion suggest a number of applications for their use. The application areas detailed here are useful and compelling for serpentine mechanisms, but exploration and inspection are probably the initial venues for a serpentine robot. Other applications, such as medical, introduce other issues such as miniature scale and government regulation. Although the scale for devices across these applications varies from tiny medical devices to large inspection devices, the principles of the configuration remain the same. Without a doubt, the development of small self-propelled limbless devices would open up areas currently intractable to the tools and technologies available today.

There are a number of serpentine applications that could provide opportunities that are both technically tractable and economically attractive. The application areas need not be exotic to make sense, and many of these areas compel further serpentine developments.

The challenges in development of a robot that can fulfill these promises are many, and I address them in this dissertation. It is possible that a serpentine robot can be built and can learn to locomote.

Chapter 2

Background

There is prior work in snake robots and snake locomotion. However, efforts and results in these areas are relatively limited in terms of scope, understanding and results.

Background examines prior work in three areas: biological snakes, robot snakes and machine learning for physical simulation and real devices. The biological history provides few insights into design but great deal of information on the varied forms of limbless locomotion. The next section, prior work in serpentine locomotors, is surprising in its breadth, but little of the research builds on prior work. As a result the work does not have a history of continued or incremental development. For further background on serpentine robots, including manipulators, see [Dowling 97b].

Snakes are the ultimate example of limbless animals; the modes and quality of their locomotion exceeds all other biological limbless locomotors. I have avoided review of invertebrate limbless locomotors such as worms, however, because of their limited mode of locomotion and because their hydrostatic structures do not parallel or correspond to the mechanical robots I propose. I also briefly cover skeletal structure, musculature, surfaces, and forms of serpentine locomotion.

Biological Systems

Biological snakes, as existing limbless locomotors, offer lessons in design and function. The difficulty, as shown, is the codification and extrapolation from biological animals to man-made mechanisms.

There is a temptation to use the animal as a model or blueprint for robotic mechanisms. It is inherent in the nomenclature; the use of the terms such as ‘serpentine’ implies that the study of snakes can lead to ideas and forms for such mechanisms. There is some danger in this assumption. The canonical example is that of bird flight; manned flight bears little resemblance to bird flight with the exception of curved wing surfaces. In

addition, many biological forms are scale dependent, and biological selection commonly reflects compromises among multiple events or influences in biological evolution.

Commonly, biological evolution also leaves vestiges and forms that do not directly relate to a particular function such as running or slithering, so emulating these structures may be a misdirected effort [Bertram 94].

It's worth keeping this in mind as lessons and ideas are drawn from snake morphology.

Skeletal Structure

The snake is a vertebrate, an animal with a backbone, and has the largest number of vertebrae of any animal: between 100-400 vertebrae, depending on the species. Snake vertebral articulation is one of the most complex of all vertebrates. Although only a few limited motions and amplitudes are possible between adjacent vertebrae, concatenation of these articulations can produce large angular excursions. The vertebrae of the snake form ball-and-socket joints with additional projections that eliminate torsional motion to protect the spinal cord. This remarkable design uses a series of surfaces to allow the limited lateral and ventral excursions (respectively 10-20 degrees and 2-3 degrees for most snakes) but eliminate torsion which would otherwise twist the spine. The projections can be seen below in Figure 2-1.

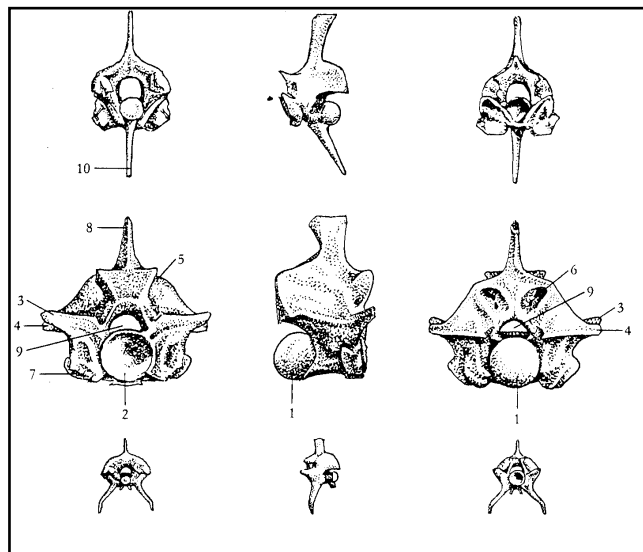


Figure 2-1: Snake vertebrae provide lateral and ventral flexing without permitting torsion.

The backbone stretches very little; snakes are not elastic like many invertebrates and they remain a constant length. Snake skeletal form and structure is quite simplified in number and type in comparison to other vertebrates. Unlike limbed vertebrates, whose skeleton has many different parts, snake skeletons have only three types of bones: skull, vertebrae and ribs.

The interesting lessons from snake skeletons are the simplicity of a repeated structure and the relatively limited motions between adjacent pieces. These aspects are worth examining in a mechanism design.

Forms of Limbless Locomotion

Snakes and other limbless animals have been objects of study for centuries. However, until recently, little research has focused on the detailed mechanics of serpentine locomotion. Yet, there is a fair amount of information on the qualitative aspects of snake locomotion. There are several broad classes of limbless locomotion; these include concertina, lateral undulation, sidewinding, rectilinear, slide-pushing and other less common forms. These classes are, in fact, gaits, a term normally associated with legged animals. Gaits are repetitive patterns of movement used to change speed, adapt to terrain, and improve stability. Gaits are often chosen because they are more economical for a particular situation [Alexander 92].

Lateral Undulation

Lateral undulation is the most frequently used form of snake locomotion for most snakes. All parts of the body move simultaneously, experiencing continuous sliding contact with the ground. It is a sliding motion with all parts moving at the same speed that occurs through the propagation of waves from the front to rear of the snake. The snake remains in contact with surface and the motion is similar to a swimming motion. As shown in [Walton 90], energy consumption is comparable to that of legged animals of similar scale. During lateral undulation, the snake pushes against features in the environment to facilitate forward movement.

Lateral undulation is the only form of biological snake locomotion that doesn't use static contacts between snake and substrate. The ideal path is a single track along which the snake slides. Lateral undulation requires a minimum of three contact points for continuous forward progress: two to generate force and the third to balance forces to move in a particular direction [Gray 68].

Lateral undulation is unsuited for smooth, low-friction surfaces and narrow corridors. Nor is it well suited for short stouter animals or for large heavy-bodied snakes because they are unable to either subtend the curves required or the body mass and environment tend to significantly reduce its efficacy [Gans 74].

Both wheels and legs use static contacts for propulsion but lateral undulation in snakes offers an interesting variant using sliding or dynamic friction. This is not as inefficient as it might first appear. However, the complexity of snake anatomy may make it difficult to realize these advantages in mechanisms.

Concertina

The concertina gait derives its name from a small accordion-like instrument because of the shape and motion of the snake body. Concertina progression provides a base in which parts of the body stop for purchase and other parts move forward. The sequence repeats, and the snake moves forward. It is usually used in confined areas, such as tunnels, where the snake cannot utilize the full amplitude of other gaits. As shown in Figure 2-3, the trunk straightens forward of each contact site and is simultaneously set down in a curved pattern at the rearward end of each site. As a result the musculature needs to be activated at or near moving portions of the trunk. The key element of concertina locomotion is the utilization of the difference between high forces with the static coefficient of friction and low forces with the dynamic coefficient of friction along different parts of the body.




Figure 2-2: Lateral undulation uses continuous sliding contacts to propel the body.

Due to momentum changes, static friction, and slower speeds, concertina is a relatively inefficient mode of locomotion [Walton 90], but forms of concertina allow traverses not otherwise possible, such as moving along wires and cables as well as through tree branches.




Figure 2-3: Concertina locomotion is usually used in enclosed areas.

Concertina movement resembles, in some ways, the motion of worms; parts of the body remain in place and other parts move forward. It would also appear to be simpler, perhaps, to implement in a mechanism, than other forms of snake locomotion.

Sidewinding

Sidewinding is probably the most enchanting gait to observe; among all serpentine gaits it evokes the most curiosity. Sidewinding is the use of continuous and alternating waves of lateral bending. A downward force is exerted for purchase on low shear surfaces like sand or loose soil; this mode establishes rolling static contacts to cross relatively smooth substrates. There are only two contact patches while the snake is in motion. The technique minimizes slippage and is even more efficient than lateral undulation [Secor 92]. Some sidewinding snakes have been observed to travel kilometer-length distances continuously [Mosauer 30]. Sidewinding is used primarily by snakes in desert regions where loose soils and sands are prevalent. The development of sidewinding may be related both to the need for traction on low shear surfaces such as sand and the need to avoid the high temperatures of desert terrain. As shown in Figure 2-4, sidewinding can




Figure 2-4: Sidewinding locomotion results from rolling or static contacts.

be thought of as the ‘peeling’ of the body from one track to the next. The tracks, or lines, show the rolling of the body contacts during locomotion.

Rectilinear

Rectilinear progression uses movements of skin with respect to the skeleton to ‘ratchet’ the body along the ground. Rectilinear motion is a slower, creeping motion using the belly to provide traction through anchoring and is typically used by larger snakes. Rectilinear motion was once conjectured to result from ‘Rib-walking,’ an active movement of the ribs. However, this was conclusively disproved in [Lissman 50] through x-ray observations of a snake in motion. Muscles connected from the ribs to the elastic skin provide the propulsive motions through reciprocating or ratcheting movements.

In rectilinear locomotion, several portions of the body are in contact with the ground at any moment, and the gait uses symmetrical rather than staggered waves of contraction. A section of the skin of the belly is drawn forward so belly scales are bunched. This part of the body is then pressed down, and ventral edges engage the surface. Then the body slides forwards within the skin until it is in normal alignment with skin, and the motion repeats. Only small vertical motions are needed for rectilinear locomotion.




Figure 2-5: Rectilinear motion couples muscle action between ground and vertebrae.

Other Snake Locomotion Modes

Other forms of limbless locomotion include slidepushing, saltation, burrowing and climbing. Slidepushing is a gait used in times of stress where anteriorly propagating waves move more quickly backwards than the snake moves forwards. A great deal of sliding and motion occur without a corresponding forward progression. Saltation is the jumping the near-vertical walls and trunks of trees. Some saltating snakes can leap gaps of a meter or more, sometimes vertically. This requires storage and release of a lot of energy and, additionally, involves a ballistic phase of motion during which control is difficult. Other extraordinary modes are used by certain asian tree snakes that glide through the air by opening the rib cage to form a gliding surface. The amazing thing about this mode of snake travel it is not how well it flies, but that it flies at all! These modes are exotic and I do not explore these forms of locomotion in robots.

Skin and Musculature

Snakes are covered with scales whose distal sections are loose and overlap other scales. They are dry and highly polished with a coefficient of friction of between 0.3 and 0.4 [Gans 84] [Jayne 88]. The belly scales are much wider than the scales along the sides and back of the snake. The skin, to which the scales are attached, is highly elastic. This

is most evident in feeding since the snake eats everything in a single gulp; enough food for over a year in some cases!




Figure 2-6: Snake ribs, vertebrae and skin are linked by complex woven muscle bundles.

Snake posture is established by muscle groups as shown in Figure 2-6. Many such bundles interconnect vertebra to each other, to the ribs on each side and in several bands to the skin. In contrast to early observations of snake locomotion the ribs do not ‘walk’ or move while the snake moves forward [Gray 46]. Specialized musculature in some snakes allows 50% of the body length to be extended above the ground without support and constriction.

Both the skin and musculature of the snake are highly refined and specialized. It is not possible, within the scope of this research, to replicate their capability in sensing and control. There are few, if any, commercial actuators like muscles, no sensing like snake skin and no controller equivalent to the nervous system of such a complex animal. However, I show that it is possible to replicate general characteristics of serpentine locomotion.

Analysis of Limbless Locomotion

[Fokker 27] and [Jones 33] show that body curvature is a key element of the lateral undulatory form of locomotion. The snake body tends to propel best at portions of the body that are undergoing the greatest amount of curvature change [Gray 68].

To gain a better understanding of force application by limbless animals, it would be useful to externally monitor forces of snake motion. Gray did this with pendulums and force plates and Gans instrumented pegs on a low friction surface [Gray 68]. But the small number of discrete points gives only a little information about a few discrete points. No force plates or other techniques seem capable of providing this information. Full and others at Berkeley’s Polypedal lab have a number of means of monitoring forces even for small insects, but their clever photo-elastic techniques are limited to discrete contacts [Zimmer 94][Harris 78]. Force plate techniques, used in biomechanics

studies, are too coarse to provide good information but some sensing pads used in these studies may be useful tools [Novel 97]. Hirose also showed measurement techniques for measuring forces in snake locomotion using strain gauges and a support mechanism [Hirose 93].

From observation of snake motions it becomes obvious that the control mechanism utilizes local information about the terrain to quickly and effectively adapt to changing conditions to propel. Since the position of the contact sites is not known by the snake and these contact sites can move or deform, the snake must make continual selections by monitoring external forces and contact sites. In addition, a feedback mechanism exists that responds to this information so that following portions of the body adapts its curvatures and provides appropriate forces to the terrain [Gans 85]. While comparisons are often made between fish swimming and the lateral undulation of a snake, swimming is significantly different for a fish where buoyancy, fluid dynamics, and hydrodynamic mass play the critical role in propulsion.

A key difference between gait selection in snakes and most other animals is that gait selection in legged animals is a function of speed in a given environment. For snakes, gait selection appears to be more a function of environment, not of speed. This difference may strongly affect learning results in a fixed environment.

What's Missing?

The impetus for examining snakes was to discover what could be transferred to a serpentine mechanism. What has been described so far is the consensus regarding qualitative forms of limbless locomotion and when they are used. However, there is much that is unknown. For example, snakes crawl on their bellies but propulsive forces are generated along vertebrae axis. At this time it is not known clearly how this is done. It is conjectured that the complexity of muscle structure may be a reflection of the control system, not just and not simply the structural basis of force generation. This is not known although it likely that the anatomy and control co-evolved.

There are no detailed and complete analyses of the force propagation that results in motion in snakes although there are some analyses for invertebrates [Keller 83][Niebur 91]. The full sequence of muscle control is also unknown. Additionally, the feedback mechanism for contact force and slip is not well understood. More studies on efficiency are needed. Additional work is needed to make definitive comparisons of energy use in locomotion. Many other questions include: How are gait selection strategies decided? What is the distribution of control?

I do not propose to answer these questions; they are tightly related and coupled to the anatomy and nervous system of the snake. However, it may be that serpentine robots will eventually offer insights to those who study the natural organisms, although this work does not make this claim.

Robotic Systems

Nearly all mobile vehicles built by man for terrestrial use have either been wheeled or legged. Wheeled vehicles date back several thousand years; walking devices can be

traced back to the 19th century. Locomotion without the attributes of legs and wheels is represented by only a few examples, mostly within the past twenty years and almost entirely within the laboratory rather than in the field.

In my search for undulating and locomoting serpentine mechanisms, the earliest mechanisms I came across were developed by a Russian constructivist artist of the 1920's, Petr Miturich. He designed a series of utilitarian designs for undulators, termed *volnoviki*, that moved by wriggling. He made many designs for *volnoviki* that were to operate on land, in air, or water. He applied for several patents on the ideas and built models, but none included power and control [Lodder 83].




Figure 2-7: Miturich developed a wide variety of undulating designs as art.

There are also obscure references to clockwork snakes and caterpillars built by craftsmen such as Fabergé, but most had small hidden wheel drives and did not locomote through body motions alone.

Hirose

Work by Hirose and Umetami, in the early 1970's, was among the first to explore and develop limbless locomotors. Hirose has a sustaining interest in limbless locomotion and designed and built several robots over decades. He termed the devices Active Cord Mechanisms or ACMs. Hirose focused on developing robots that could perform lateral

undulation and later developed a series of wheeled coupled-mobility devices that followed from this work.




Figure 2-8: Hirose's Adaptive Cord Mechanism utilized a series of articulated links with passive wheels.

Hirose's development of modeling and control first derived expressions of force and power as functions of distance and torque along the curve described by the snake. The curve was then derived and compared with results from natural snake locomotion. The curve, termed serpenoid, has curvatures that vary sinusoidally along the length of the body axis. These equations are shown below:

$$x(s) = sJ_0(\alpha) + \frac{4l}{\pi} \sum_{m=1}^{\infty} \frac{(-1)^m}{2m} J_{2m}(\alpha) \sin\left(m\pi\frac{s}{l}\right)$$

$$y(s) = \frac{4l}{\pi} \sum_{m=1}^{\infty} (-1)^{m-1} \frac{J_{2m-1}(\alpha)}{2m-1} \sin\left(\frac{2m-1}{2}\pi\frac{s}{l}\right)$$

This curve is different from sinusoidal or even clothoid curves. Comparisons with natural snakes across constant friction surfaces showed close agreement between the serpenoid curve and the empirical data.

Hirose then went on to develop models for the distribution of muscular (actuator) forces along the body. This was done for normal and tangential forces as well as power distribution. Again, the developed models closely correlated to muscle exertion data and force measurements from natural snake movements.

The experiments to this point were primarily of a uniform nature, but Hirose recognized




Figure 2-9: A close-up of the first ACM link showing the body and drive.

that snakes quickly adapt locally to variations in terrain and environment. The next issue was to characterize this adaptation. From observation it was noticed that snake locomotion is not necessarily a two-dimensional problem; in fact during higher speed motions, snakes use ventral motions to actively distribute their weight to those areas where propulsion is maintained.

Further study developed relationships between amplitudes and wavelengths of the motion and local friction conditions, as well as morphological features of the snake such as vertebrae motion and muscles (actuators). Models for locomotion in rough terrain where obstacle contact is made were also developed and correlated with snake motions.

Hirose examined the construction of mechanisms that were able to perform lateral undulation. Several views of these machines are shown in Figure 2-8, Figure 2-9, and Figure 2-10. By calculating torques, velocities and power required, Hirose was able to provide design guidelines for the actuators and drivetrains. The next development was a distributed control scheme wherein each link could respond independently. In Hirose's work the control took the form of angle commands at each joint. The variables were simply related closely to the amplitude, wavelength and velocity of the body axis.

Steering of the robot was accomplished by biasing the control to adjust curvature in a section of the body.




Figure 2-10: Hirose's locomoting Adaptive Cord Mechanism.

A 20 link mechanism weighing 28kg was constructed. Link actuation was accomplished with DC motors coupled to a caster board and potentiometers were used for feedback. Later, after a motor change, the weight was reduced to 13kg.

To accommodate unknown environments required tactile sensing; this was the next step in Hirose's work. Small contact switches provided this information to the controller. As shown in Figure 2-10, this robot could negotiate and propel itself through winding tracks. The developments included a control technique called lateral inhibition tactile signal processing, which provided for contact and reflex motions. The shape of the body was varied according to the second derivative of the sensed contact pressure and responded appropriately to provide forward progress.

All of Hirose's locomotors used either powered wheels or passive casters and the only locomotion mode studied was lateral undulation. Hirose and his colleagues have gone on to develop an elastic elephant-like trunk, a large serpentine mechanism for interior inspection of turbines and small manipulators for surgical applications. The best overall paper on Hirose's work is [Hirose 90]. It succinctly covers many years of development in serpentine mechanisms. More recently, Hirose's book *Biologically Inspired Robots* is an excellent overview of his work [Hirose 93].

Hirose's work in serpentine robots is probably the most complete of all work in this area. He dealt with issues of mechanism, control, sensing and modeling of natural animals. However, the mechanisms used wheels, the terrains for the ACM's were 2D only, and the mechanism used only lateral undulation as the locomotion mode. The configuration, while not practical for application use, was a great advance in serpentine robots.

Burdick and Chirikjian

Joel Burdick and his students at Caltech, especially Greg Chirikjian¹, have pursued work in serpentine manipulation and locomotion for several years. Chirikjian's thesis




Figure 2-11: Burdick's Snakey, a VGT-style hyper-redundant manipulator and locomotor.

presented a framework for kinematics and motion planning of serpentine mechanisms. Curves in three dimensions, \mathbb{R}^3 , are defined to provide a general means of parameterizing curves and sets of reference frames. In addition to describing the curve shapes, they extended features to allow roll distribution along the curve and extensions and contractions along curve segments. These were then used to specify serpentine configurations [Chirikjian 92].

Since most manipulators do not describe continuous curves, there remained the problem of fitting rigid link devices to the desired curve. A general parallel algorithm was found for fitting manipulator segments to the desired curve. The modal approach was then used to resolve the 'excess' degrees of freedom (DOFs) in hyper-redundant robots to carry out specified tasks. The modal approach provides a means of characterizing the shape and motions without developing full inverse kinematics, which have an infinite number of solutions. A series of specified functions could be specified in modal form, and the problem became finding modal participation factors to satisfy, as best can be done, the task constraints. Optimal techniques for minimizing measures of bending, extension, etcetera, were then developed via the calculus of variations. One issue with these optimal techniques is the selection of cost functions to evaluate configurations. That is, how to determine the 'goodness' of a particular solution.

Chirikjian described obstacle avoidance using this set of tools and it was assumed that paths were provided through traditional motion planning techniques. An additional issue addressed is that of time, that is, velocity, for the solutions. A series of arcs and lines were used to create a path along which the manipulator sections can move. But, independent of the path formulations, the previous solutions to kinematics could fit manipulator configurations and trajectories.

Locomotion through sequences and patterns of geometries was developed next. The extensible locomotion modes were traveling wave, similar to rectilinear motion in

1. Now at Johns Hopkins University

snakes or caterpillars, and stationary wave, similar to inchworm motion where the advancing wave remains in the same position with respect to body coordinates.

The extensible modes are similar to earthworm locomotion where segments provide extension and contraction to propel the robot. To avoid the need for differential friction, portions of the body can be raised to facilitate this motion. Descriptions of techniques for non-flat floors are also developed. Intriguing ideas were also introduced using serpentine robots to provide grasping and manipulation capabilities. The mechanism could contact and wrap about an object; the propagation of a wave or extension of the links caused the object to move in a desired direction. These techniques could be used to simultaneously grasp, move and manipulate objects.

A mechanism, a variable geometry truss configuration, was designed and built and is shown in Figure 2-11. The mechanism was comprised of commercial linear actuators; a simple modular and maintainable approach to design was used. A variety of tests using the methods described above were conducted and a number of successful experiments in control and locomotion were carried out.

Key to Chirikjian and Burdick's work was the modeling of the robot as a 3D shape and sequence of shapes. This enabled a variety of techniques in trajectory planning and path generation. The mechanism also allowed exploration of non-snake-like extensible gaits. The mechanism however was primarily a fixed base device and a couple of limited gaits were demonstrated on the robot. Additionally, ratchet wheels were used in locomotion [Chirikjian 95]. Sidewinding was also formulated in piecewise continuous curves in [Burdick 95] and, although the exact shape of the body was not necessarily snake-like, the general form of the motion was identical to that of snakes.

Later work by Burdick with J. Ostrowski explored the use of geometric mechanics to formulate general notions of locomotion. Two systems were evaluated in this context: a 'snake-board' which is an actively articulated skate board, and Hirose's ACM [Ostrowski 96]. Other related work at Caltech included the work on geometric phases to describe robot locomotion [Kelly 95].

Choset also developed path planning methods for highly articulated robots such as snake robots. He developed the Generalized Voronoi Graph (GVG) and Hierarchical GVG for use in sensor-based planning motion schemes. The techniques utilize concise descriptions of the topological spaces to build paths. A key feature of the work is that it does not require a priori knowledge of the world [Choset 96a][Choset 96b].

Shan

Shan's work was primarily in obstacle accommodation: motion planning that makes use of obstacles rather than strictly avoiding them. The mechanism in this work, shown below, used a form of concertina locomotion. The device Shan built, shown in Figure 2-12, uses solenoids at the joints to drive vertical pins into the surface. This establishes

fixed contact points from which the rest of the mechanism can move [Shan 92][Shan




Figure 2-12: Shan's snake mechanism uses a concertina-like motion.

93a][Shan 93b].

The configuration and locomotion of Shan's robot were limited to flat floors and a concertina mode that required a great deal of space; far more than the cross-section of the mechanism suggests. The length of the links or, more importantly, the ratio of length to diameter, play a large role in the robot's inability to traverse tight spaces.

Ikeda and Takanashi

In 1995, the giant Japanese electronics company, NEC, announced the development of a snake robot which was dubbed 'The Quake Snake' and designed to enter the rubble-strewn aftermath of earthquakes and explosions to search for survivors. The device, called Orochi, utilized an active universal joint, a novel form of a Hooke's joint designed by Ikeda and Takanashi. The seven segment device is shown in Figure 2-13

and the joint design in Figure 2-14 [Ikeda 87] [NEC 96]. A small video camera was also




Figure 2-13: The NEC 'Quake Snake' utilized a novel universal-type joint between links for a total of 12 DOF.

deployed at the head of the mechanism and used to by the operator to assist in guiding the snake.




Figure 2-14: The rotating joint developed by Ikeda and Takanashi provides for a smooth and wide range of motion.

The mechanism is one of the best mechanical designs in serpentine robots, and some automatically generated gaits have been used on this mechanism [Burdick 97]. Control as shown in the videos is done manually and the single gait used is akin to a rectilinear or inchworm gait. Additionally, in some footage of the device, but not shown in Figure 2-13, are small brackets used to stabilize the snake as it moves; effectively they are wide 'feet.' This class of mechanism has great promise for serpentine robots in real applications. Another version of this joint was used in a snake built for JPL. The key lesson in this robot was effective packaging of the mechanism, the slim design and the modularity of the links.

Nilsson

Martin Nilsson of the Swedish Institute for Computer Science in Sweden, as part of the PIRAIA project, has developed a novel universal serpentine link that is a roll-pitch-roll joint. Multiple links give it the ability to subtend some very non-snakelike modes of locomotion that incorporate a rolling motion. In one instance, the snake might ‘hug’ a tree and, using the side rolling capability, roll directly up the tree. The joint has another nice feature in that cables passing through the joint cannot twist if the joint is controlled properly. This joint is equivalent to universal joint, but unlike a normal universal joint, the input angular velocity equals the output angular velocity for all angles. The mechanism, while relatively complex, can be realized with standard components.




Figure 2-15: The PIRAIA link provides a roll-pitch-roll capability.

Additional work by Nilsson showed learning techniques for locomotion using differences between static and dynamic coefficients of friction and sliding contact surfaces. A simple pair of paddles joined by a single degree of freedom was capable of learning to move itself through a simple physical model [Nilsson 95] [Nilsson 97a] [Nilsson 97b].

Nilsson’s mechanism is very different from other work and from natural animals. The two roll motions at each joint enable wheel-like effectiveness in locomotion, but also complicate internal mechanics. Since the intent of this dissertation is to replicate more standard snake-like locomotion modes, this design does not appear to have many parallels to this work.




Figure 2-16: The PIRAIA links incorporate power as well as drive mechanisms in a unique roll-pitch-roll joint.

Paap

Karl Paap and his group at GMD in Germany developed a snake-like device to demonstrate concepts and developments for real-time control. The device is a tensor device that uses short sections with cable winding mechanisms to effect curvatures along several segments. The device is shown in Figure 2-17 [Paap 96]. The curvatures




Figure 2-17: The GMD Snake mechanism

are continuous along those sections but the joining segments, where the drive mechanisms are located, do not bend or move. Some very limited locomotion has been shown in the mechanism and the cable drives have been a design challenge.

IS Robotics

IS robotics built a small snake-like machine, Kaa, for prehensile grasping of pipes and locomoting. Not an effective locomotor, the robot was initially designed for moving in and through networks of pipes and support structures. It is probably the first completely self-contained snake locomoting robot. Using RC-servos as actuators, the robot



Figure 2-18: The Kaa snake is a self-contained locomoting device.

propagated a ripple down the body to effect a straight-line motion on a flat surface [Desai 95]. The robot did not locomote in the position shown in Figure 2-18 but instead lay flat upon the ground so that actuator motion was in vertical plane only. The movement was limited and the large box in the middle of the robot, housing power and computing, made locomotion problematic on the ground.

Coupled-Mobility Devices

Coupled mobility devices, sometimes called overland trains, are similar to trains of vehicles linked together. Although Hirose's ACM robots resemble a coupled-mobility device, all wheels were passive and the robot 'skated' because of body movements. However, a number of machines have been built that are similar to small trains for off-road navigation and used active wheel drives. The largest of these were LeTourneau's huge Sno-trains [Gowenlock 96] and a much earlier overland steam train [Anon 95].

Others, such as the Odetics ATMS, All-Terrain Mobility System, as shown in Figure 2-




Figure 2-19: The ATMS could cross and climb obstacles.

19, were coupled-mobility devices with active couplings designed to traverse a variety of terrains. In this device, both link and wheel motions were explicitly described for movement [Odetics 88].

Coupled mobility devices, while bearing resemblance to snake-like robots, are not limbless. They use wheels to drive or support, and because of this, they do not offer some of the advantages of snake locomotion as described earlier. However, it is possible that such mechanisms could offer advantages of the wheel and some advantages of snake-like robots in future developments.

Learning

Physical modeling is a key element of a serpentine mechanism. In all prior work described for serpentine robots, control was generated in an explicit fashion. Results included limited modes of locomotion and little adaptation. There is a great deal of research in machine learning, but much of it resides within computer models and databases. Little work in learning has been applied to physical mechanisms that are more complex than, for example, a robot learning to throw a ball. Two pieces of research that are relevant to my work are the work of Karl Sims and David Barrett.

Sims

Karl Sims developed a learning tool that simultaneously evolved both morphology and control in simulation using genetic algorithms. The various creatures evolved both their geometric forms and their control structure using simple physical models to determine fitness. Metrics for fitness were mostly based on speed but also incorporated the notion of winning margins where fitness evaluation depended on the margin of victory between pairs of creatures [Sims94a][Sims94b]. However, Sims did not extend his work beyond simulation, but did illustrate the possibilities of learning using physical modeling. The physical models were not detailed and were fairly simple simulations

but the end result was both inspiring and enchanting. As can be seen from Figure 2-20, the bodies of the evolved creatures were simple constructs of intersecting rectangular volumes. No realism or accuracy was required; the work was not intended to be a true




Figure 2-20: An example of evolution from Sim's creatures.

or realistic predictive simulation. The creatures themselves were simple constructs of intersecting polygonal objects without real joints or pivot points.

Barrett

David Barrett's work at MIT used genetic algorithms to teach a swimming robot, a tunafish, to swim efficiently. A large externally powered, cable-driven mechanism was used to control the motion of the tuna. The metric for the robot was an efficiency measure formed from the ratio of the mechanical power output to the power input to undulate the tuna's body.

Locomotion was evolved in situ, not in simulation. The hydrodynamics were too complicated to model and compute so, in essence, the water acted as a large analog computer with infinite resolution. Seven parameters related to fluid flow were chosen for the optimization criteria. Small populations of 10 were used and convergence was accomplished within only a few generations [Barrett 97] [Triantafyllou 95].




Figure 2-21: The RoboTuna internals and foam latex covering.

Summary

Biological Understanding

Biological analogues to serpentine robots offer remarkable performance and many issues might be understood through the study of these natural animals. However, much is not understood in snake control and locomotion and perhaps serpentine robots will offer explanations for biologists! In the meantime, snake locomotion modes offer striking examples of the promise of limbless locomotion. Additionally, useful lessons from structure and morphology of biological snakes can be applied to the mechanisms. Skin, in particular, appears to have a strong influence on locomotion but prior robot mechanism work has not addressed this issue.

Robotics Developments

Although there have been several projects related to serpentine manipulation and locomotion there have been far fewer robots built, and little success towards practical mechanisms. In fact, only a few serpentine *manipulator* mechanisms ever made it as far as a commercial venue; the Toshiba Multijoint Inspection Robot [Asano 83][Asano 84][Nakayama 88][Toshiba 89], and the Spine robot [Drozda 84][Grunewald 84]. Neither were successful in the marketplace. I am aware of only one serpentine manipulator manufacturer, Kinetic Sciences, who is attempting to market their ‘Tentacle Arm’ [Immega 95].

The developments of Sims, Barrett and other in the machine learning community show great promise, but also show the immense amount of work and computational requirements required in modeling and evolution of locomoting vehicles. The control and planning work of Burdick, Ostrowski, Choset and others is furthering the

techniques necessary for the control of these and other highly articulated mechanisms. The most promising emerging developments are the non-linear control and geometric mechanics literature Burdick, Murray, Kelly, Ostrowski and others. They are forming general frameworks for all locomotion and have used serpentine locomotion as one of several case studies.

Technical Needs

There are many technical developments required for useful and successful serpentine mechanisms. Mechanical structures have limits to rigidity resulting from both technical and economic considerations. Additionally, long and narrow structures raise issues of dynamics and oscillations that must be addressed. Other technical areas include sensing; accurate sensing needs arise because of the indeterminacy of individual joint positions. Another need is the development of sensor ‘skins’ that allow joints to independently react to local obstacles and control. Research is underway in each of these areas but only at formative stages. Additionally, most of the technology efforts are not directly related to serpentine mechanisms but focus instead on the needs of industry and mechanism control.

The work to date shows great promise in understanding and configuring robots, but there remain many hurdles on the path to develop, control and deploy practical serpentine robots. The quest is to demonstrate that robots can learn to locomote even when they have no wheels or legs. The following sections show a new approach to the problem of teaching limbless robots to locomote.

Bridging the Gap

As shown in Figure 2-22, while several works have bridged physical simulation and learning and others have bridged learning for robots, little has been done in bridging all three areas. As systems grow in complexity and as classical modeling techniques become more intractable, then connecting these areas will be critical. This work takes a step towards bridging the areas of mechanism, learning and simulation.

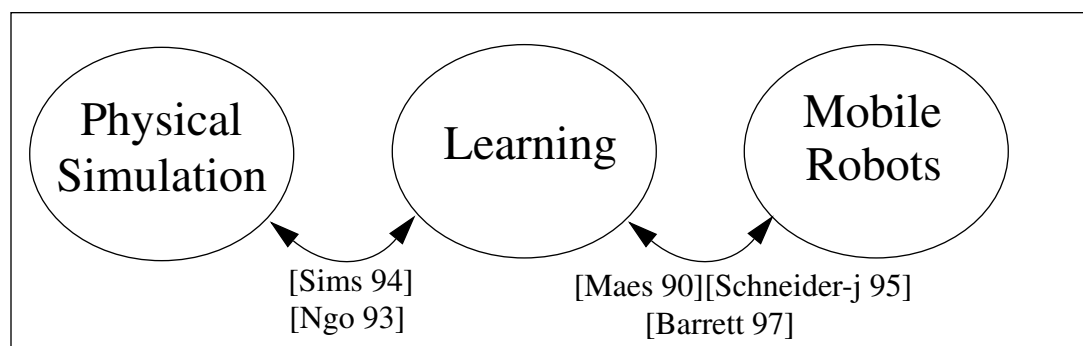


Figure 2-22: Bridging physical simulation, learning and robot mechanism is key to the control of complex mobile robots .

Each of the areas shown in Figure 2-22 are open research areas with many people working hard in the areas of physical simulation tools, learning methods and mobile robots. Rarer are the works that bridge these areas including the gap between physical simulation and learning and the gap between learning and mobile robots. A few people,

as listed in the figure, have begun the work required in these areas. What has not been bridged is the long path between simulation, learning and mobile robots. This requires a great deal of ground work in all areas. It is this bridge that I construct in this thesis.

The results of this background study provide lessons and results from prior approaches and in several cases provide inspiration that directly influenced this research. For example, simplicity in design and the avoidance. Questions arose as to why the snake robots were designed in particular ways and compelled the formulation of a methodical approach to answering questions regarding morphology and form. Other influences included the learning approaches used by Sim and Barrett for simulation and robots respectively. Snakes themselves inspire by example and the wonder they incur but also the short, repeated structure that forms their backbone. The prior work influences and inspires through example and approach.

Chapter 3

Framework

How can we teach a limbless robot to move? What is the structure, the form, or the architecture for making this happen? For a robots to learn to locomote, a structure or framework is necessary to support learning. Evaluation is also critical to making this occur; that is, how is performance evaluated? **Framework** presents such an architecture for simulating, designing and testing serpentine locomotion with a transfer of those results to a physical device. Sections of this chapter correspond to later chapters in this dissertation.

Overview

A snake robot mechanism is relatively complex; the design is a repeated structure of many identical links, all of which need to be coordinated and each of which have to be controlled. The fundamental issue is determining the sequences of motion of the elements that move the mechanism in a particular direction. While it may be possible to construct sequences by hand that move the system, this is fairly tedious, inefficient, and likely to miss a variety of interesting gaits. Schemes tried in previous works, see **Background**, appear inadequate for determining and expressing a variety of gaits for this mechanism; you need to know the gait before attempting it. Ideally, the system would move forward after learning how to move. The problem then becomes: how to teach the robot to move.

The fundamental idea then, is to develop a structure or framework to allow both representation and development of gaits. The need is then a modular framework with the following elements: learning, simulation, and evaluation. Other elements include techniques or protocol for communicating and passing information from one element to another and selection of parameters and the form of the parameters that are passed back and forth within the framework.

Physical simulation is a useful tool for the configuration and control of complex mechanisms and allows observation of these robots in a simulated environment. The rationale for physical modeling and simulation is to represent a physical device such that inputs and resultant outputs are reflected accurately and, in turn, provide an understanding of mechanism behavior. A useful simulation tool for robots provides the capabilities to model physical entities, physical laws, interactions, and incorporates monitoring tools to see input effects.

Given such a tool, it can be incorporated into a larger scheme where control and evaluation take place. This scheme would also provide a means for testing the results of various control methods. Expanding on this idea further, a machine learning technique, using an appropriately chosen metric, could interface to the modeling and control techniques for its own use.

In particular, control parameters are selected and are operated on by a learning function. These parameters, in turn, become the inputs to the physical simulation. As the simulation learns to behave in a desired fashion, using a metric to be defined later, then these results can be ported to a physical device for testing and further refinement. This testing can still be done in the context of this original framework however by simply substituting the robot for the modeled device. This assumes that the framework provides an approximation to the robot sufficient to demonstrate viability.

As shown in Figure 3-1, the framework is comprised of several areas including

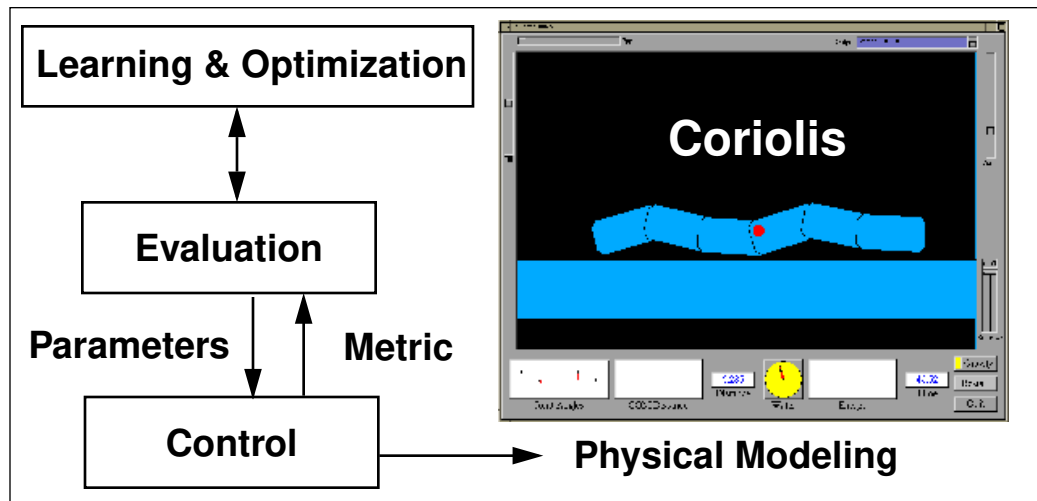


Figure 3-1: Framework for learning control of a physical device.

optimization or learning, evaluation, control and physical simulation. A loop of test and iteration is created by the results of performance in physical testing. The results in each set of tests provide input into the next series of tests. In the following subsections we will take a brief look at each component of the framework.

Evaluation and Metrics

As the physically simulated model executes, a performance measure is used to determine how well the model is doing. For example, one such measure might be the

maximum distance the device moves in a given period of time. This measurement, or metric, is a quantitative assessment of the relative ‘goodness’ of that set of parameters. This metric can take time, energy, distance and other measures into account in determining the efficacy of that set of parameters. In a sense, this is a measure of the efficiency of the particular sequence of body motions that effect forward movement.

Learning is best when the results are easily and quickly evaluated, so the metric should also be easy to evaluate. An overview and discussion of metrics as well as metric determination and evaluation is detailed in **Performance Metrics**.

Learning and Optimization

A crucial part of the framework provides the generation and selection process for new sets of parameters to determine the ‘best’ set of these values. In **Learning and Optimization**, the types and methods of machine learning techniques are detailed. These are used in the course of determining and tuning efficient gaits for the mechanism.

Control

From the optimization, a set of parameter values is chosen and these can be used to run a program or controller to implement and execute a run. Although Control is shown as a distinct box within the framework, it is really two separate components; one to interact with Physical Modeling and the other to interact with the physical robot. This work is detailed in **Implementation**.

Physical Modeling

The key element for the framework is the ability to simulate the geometry of the physical robot and its interactions with external surfaces and contacts. Physical simulation allows the modeling of complex shapes and their interaction. Physical quantities that are modeled include torques, forces, velocities, acceleration, both kinetic and potential energies. Physical simulation, until recently, was intractable for mechanisms because of algorithm complexity and the computing power required. While many of the physical principles were understood, the tools were lacking. The convergence of increased computing power and usable programs for creating physical models has enabled the use of physical simulation as a tool in research, design and testing.

Modeled quantities in physical simulation also include gravity, dynamics, masses, and material properties such as friction and density. The user is required to configure geometries, define physical relationships between geometries, set forces and any time dependent relationships. The physical modeler then creates the simulated world and allows the user to observe the subsequent behaviors and interactions. The whole process of physical modeling promises to be a wonderful extension to existing design tools; it facilitates the creation and evolution of designs.

While simulation is used for resolving inputs and outputs into the physical modeler, it is also used to learn sequences of body configurations that enable the model to move.

That is, the modeler is used in a larger framework that allows observation and reaction to gain a desired output; a closed-loop system.




Figure 3-2: An accurate 3D model of the physical snake.

The framework is also designed to be used on the robot mechanism. The physical modeler is replaced by the robot mechanism, as shown in Figure 3-3, and the same criteria and optimization techniques can be used on the real mechanism after the physical simulation produces useful results. Even with strong effort, the simulation does not model the system perfectly; the vagaries of the real world prohibit accurate and high fidelity predictions of behavior.

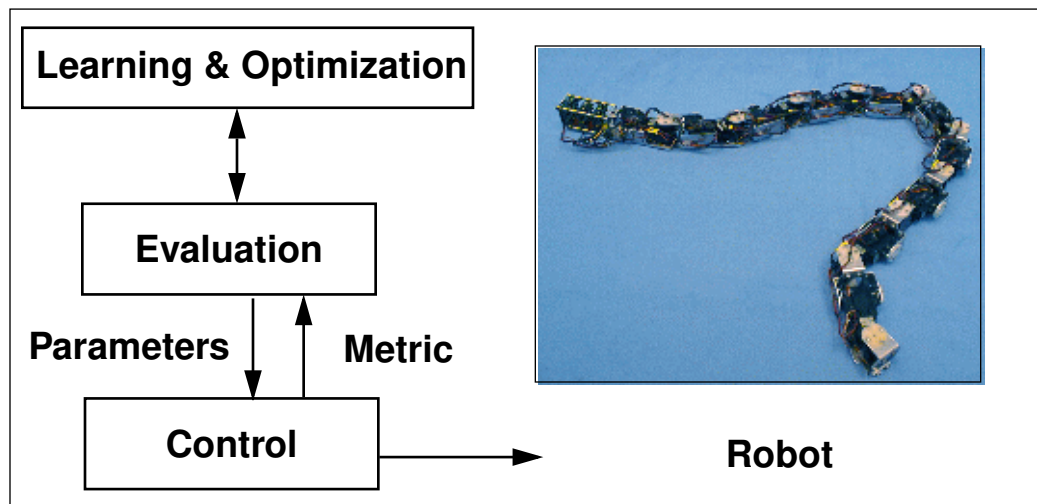


Figure 3-3: The same framework is used for the robot mechanism

However, the initial physical model is used to provide general classes of gaits that can be refined on the real system.

Summary

The confluence of a new generation design tool, physical modelers, recent advances in learning and a novel mechanism promise to bring advances in robot design. Physical modeling, a relatively new tool primarily developed for use in the graphics community, can provide the designer feedback and provide a tool in a larger context; evolutionary design. The framework shown here, comprises modeling, simulation, evaluation and control of the robot under consideration. The following chapters detail each of these aspects and their use in the design process.

Chapter 4

Performance Metrics

Teaching a serpentine robot to crawl requires evaluation of locomotion strategies or the use of metrics. Metrics are measures of performance. The selection of good metrics is critical to learning methods because they provide a means to compare the results, techniques and methods of locomotion. In addition to providing assessment and comparison, they should be easily calculated and measured. Metrics are used to evaluate and drive the learning. Thus, the simpler and more straightforward the metric, the easier it is to track and guide learning.

The general question is: *how to evaluate the performance of mobile robots?* Specific questions for evaluating performance include: how fast does it move? how much energy does it use? how far can it go? how well does it carry out a task?

Performance Metrics examines criteria for evaluating the performance of locomoting vehicles including limbless robots. Performance metrics can be used for objective analyses of different locomotion techniques and be used to optimize these criteria during learning. I show several formulations of metrics used in the analysis of vehicle and animals, and I also derive useful measures and provide insight into scaling issues.

Metrics

Metrics are used to compare the performance of vehicles, tools, computers and human endeavor. As an example, a simple metric of performance might be *speed*. However, speed, by itself, is both insufficient and misleading because, without compensating for size, scale or energy use, the comparison between different systems is comparing apples and oranges; they aren't the same. In an absolute sense, if speed is the only metric of concern, then the development of vehicles favors large size and high power.

In most forms of vehicle racing, for example, the only metric that counts is speed, but a wide variety of constraints and narrow classifications result in a relatively narrow

range of parameters that can be adjusted to maximize speed. This discussion, while interesting to racers, is moot; speed or velocity is an insufficient measure of performance because scale matters, power matters and configuration matters. Additionally, the very simplest metrics, such as speed, are often insufficient because energy sources are typically power- or energy-limited for small robots [Dowling 97a]. Thus, minimizing energy use over time and distance is important for mobile vehicles. I present a number of examples of metrics and, importantly, their dimensional units and physical meaning. Units reveal metrics in a way that is intuitive and allows comparison with other metrics.

Efficiency

Efficiency is the ratio of power output to the power input of a system. It can be the theoretical mechanical output power versus measured input electrical power. It represents how effectively a system utilizes power to effect motion or how much energy it takes to move a given distance. However, to say that one system is more ‘efficient’ than another is often invalid because of environment or other constraints. Indeed while size, weight and speed are probably related to efficiency they are usually not directly taken into account in the measure.

Also, it’s not obvious how to determine efficiency. For example, the mechanical efficiency of constant speed locomotion over level ground is zero because there is no work done. Work, the product of force and distance, is zero at constant speed and with horizontal motion. Energy is expended, of course, but no physical work is done.

Does efficiency always matter? A recurring litany in walking machine literature is the potential efficiency of walking. With few exceptions, such as planetary explorers, efficiency is not likely to be important for walkers. The metric that compels walking machines is likely to be task performance and efficacy, not how little energy is used. However, planetary robots are often energy limited and efficiency does matter for that application. However, the use of snake robots may be driven by the ability to complete tasks that are intractable to other vehicles, not energy efficiency.

The major use of energy in animals is for locomotion and thus there are distinct advantages in keeping these energy costs low. Advantages include the ability to run for periods of time from predators or surviving without food for indeterminate periods. Interestingly, for most terrestrial animals, energy use for a given distance traveled is *independent* of the speed [Schmidt-Nielsen72]. However, the *amount* of energy used is related to the size of the animal; where larger animals use less energy per weight to travel a given distance. This effect is related to the amount of time that the animal’s legs are in contact with the ground during movement and is due to the storage and release of energy between tendons and muscles during that time. Strangely enough, the muscles do not directly drive most locomotion but in fact are a mechanism to store energy in tendons which transmit the forces to ground. This technique results in high efficiencies for animal locomotion [Roberts97].

There are several ways of defining efficiency including:

- Gross efficiency Ratio of work accomplished to energy expended.

- Net efficiency Ratio of work accomplished to energy expended during rest.
- Work efficiency Ratio of physical work accomplished to energy expended with no load.
- Delta efficiency Ratio of change in power output to the change in energy expended at each power output.

These forms of efficiency, when used as design principles or performance evaluation, can affect decisions of configuration or control. Efficiency is clearly important and is also used to assess energy use in animal locomotion. A clear difference between animals and man-made vehicles is that the latter do not typically exhibit the regenerative properties of the muscle and tendons of animals.

Regeneration is the recapture of expended energy through energy storage mechanisms such as the tendons in animals or springs in mechanisms. Regeneration has been explored in robots but, for the most part, regeneration has not been successful. The additional complexity, volume and mass of regenerative systems has been a break-even proposition at best [Waldron 97]. Efficiency then, is the objective of minimizing energy use, rather than regeneration. Regeneration is likely to be an important issue in electrically driven machines; results in this area will offer tremendous benefits.

Energy

There are many parameters that affect energy usage for vehicles. These include: speed, forces or torques, terrain, and infra-structure such as computing, cooling or sensing. Energy over time gives a power measure; this is useful for instantaneous power or simply integrating energy use over time. Yet it provides no relation to size, weight or mass of the vehicle nor the distance over which it exerts this power. Simply looking at power is an ineffectual exercise.

One simple measure then, is the net amount of energy used per distance traveled as shown in Equation[4-1]. At this point, steady-state power draw is not considered, but rather the *net* power draw for locomotion.

$$\frac{\text{power}}{\text{velocity}} = \frac{\text{energy}}{\text{distance}} = \frac{\text{joules}}{\text{meter}} = \text{newtons} \quad [4-1]$$

This metric, or its reciprocal, is commonly used in assessing energy or fuel consumption. Common examples include the liters-per-100km or miles-per-gallon measures used in evaluating fuel consumption for automobiles. However, this metric doesn't indicate size, mass, weight, or even payload of the vehicle.

For animals this measure, when normalized for mass, is smaller for larger animals. Surprisingly, the cost of traveling a unit distance in animals is not proportional to body mass but to $(\text{body mass})^{2/3}$. This appears counter-intuitive since mass increases as the cube of length and you would expect the energy-per-distance measure to be proportional to mass. The reason, as pointed out earlier, is that tendons in legged animals store and return energy to the animal. In general, muscles are more efficient when they develop force more slowly. This is the case with larger animals, where the strides are longer and the time in contact with the ground is greater. [Roberts97] [Kram 90].

For animals, locomotion energetics are determined by placing the animal on a treadmill and simultaneously measuring oxygen intake over a range of speeds. Of course, with animals there is a steady state rate of oxygen consumption for basic metabolic processes but any incremental intake is due to muscle exertion. In biomechanics, this energy use is measured in terms of oxygen intake. Such measures include ml O₂, m³ O₂, moles O₂, or calorie equivalents. One such measure of energy use is the *Cost of Transportation*.

Cost of Transportation

The Cost of Transportation is defined as the mass-specific aerobic power input over the speed of locomotion [Alexander 92]. In the biological literature, this is expressed in cal g⁻¹ km⁻¹. In physical units this is equivalent to energy per distance per mass. A Net Cost of Transportation is determined by subtracting the steady state energy consumption from the total metabolic expenditure.

$$\frac{\left(\frac{\text{energy}}{\text{distance}}\right)}{\text{mass}} = \frac{\left(\frac{\text{joules}}{\text{meter}}\right)}{\text{kg}} = \text{Cost of Transportation} \quad [4-2]$$

This is also dimensionally equivalent to a measure of force per mass. This measure is germane to biomechanics studies because, for animal locomotion, exerting force is more important than doing work [Kram 90]. This is because size differences in energy cost are proportional to stride frequency at equivalent speeds. That is, smaller animals have a higher stride frequency at a given speed than larger animals. The contact time with the ground, when the support force is generated, is the determining factor in the cost of transportation. Big animals take big steps and the contact time with the ground is longer. This gives additional time for the muscles to develop force and this requires less energy. So bigger animals expend less energy per unit mass to move than smaller animals [Roberts97].

Power to Mass Ratio

Another measure, shown in Equation [4-3], is energy per mass per time; this is equivalent to a power/mass ratio.

$$\frac{\left(\frac{\text{energy}}{\text{mass}}\right)}{\text{time}} = \frac{\left\langle\frac{\text{joules}}{\text{kg}}\right\rangle}{\text{second}} = \frac{\text{watts}}{\text{kg}} \quad [4-3]$$

This use of energy per mass per time is a measure of power density and is often examined with respect to speed to give trends and indications of energy use at different speeds. These types of measures are usually more interesting and useful when using weight instead of mass. This is because locomotion is typically in a gravity field and the force exerted on terrain and energy expended is primarily due to gravity. Two such ratios incorporating weight have interesting physical dimensions: energy-to-weight and power-to-weight ratios.

Energy to Weight Ratio

The dimension of energy to weight ratio is length:

$$\frac{\text{energy}}{\text{weight}} = \frac{\text{joules}}{\text{newtons}} = \text{meters} \quad [4-4]$$

This results in a dimension of distance that is equivalent to how high an energy storage system can lift its own weight in a 1g field. This gives an intuitive feel for energy capacity and the value is independent of the size of the system. For example, using specific energy values of 50Wh/kg for lead acid batteries means they can lift their own weight 18,000 meters.

Power to Weight Ratio

Similarly, the dimension of power to weight ratio is velocity:

$$\frac{\text{power}}{\text{weight}} = \frac{\text{joules/second}}{\text{newtons}} = \frac{\text{meters}}{\text{second}} \quad [4-5]$$

This has an interesting physical meaning; it represents how quickly a system could climb vertically, given a power-to-weight ratio in a 1g field. It is the terminal velocity the system could achieve if it could devote all of its power to vertical motion against gravity. In the battery example above, if the battery could provide 500W continuously, the upward velocity would be 51 m/s for about 6 minutes. For most applications, this has more pragmatic implications for climbing hills than it does for making vertical ascents, but it can be used to directly relate power limitations to velocity limitations in traversing terrain.

Computing Metrics

Another, more general, measure relates locomotion to computing resources; the cost and complexity of planning and executing movement. A proposed NASA program of the mid-1980's involved sending a mobile robot to Mars for exploration. Space-qualified computers of the time offered little computing power compared to what was needed for sensing, control and other compute intensive areas. As a result, a computing metric was proposed to provide a measure and comparison for computing power. The metric is shown in Equation [4-6] and is in units of millions of instructions per second required per meter of travel divided by the number of seconds.

$$\frac{\left(\frac{\text{mips}}{\text{meter}}\right)}{\text{second}} = \frac{\text{instructions}}{\text{meter}} \quad [4-6]$$

This is equivalent to instructions per meter and is a measure of computation required per distance traveled. It can be used when locomotion is compute bound and not limited by mechanical or energy issues. It may be a useful measure for comparisons of navigation techniques; if computing power is at a premium, this metric could justify certain methods and algorithms over others. However, metrics for these types of missions continue to be about efficiency of vehicles and not directly with computation required. Power is typically an even scarcer resource than computing!

For snake robots, at least initially, this type of metric should not be central to evaluation but if locomotion is limited by sensing and evaluation, it may be useful to revisit this metric.

Work Metrics

Metrics are also needed for evaluating working machines that move payloads from one place to another. One such metric is called normalized work, coined by [Binnard 95]. Normalized work, as shown in Equation [4-7] is the product of the payload to mass ratio and the bodylengths per time. Unfortunately, this has a strange dimension of time^{-1} and implicitly favors shorter machines. The metric is really the frequency at which the vehicle can move a load equal to its own mass, one body length. This metric unfairly penalizes long machines so that machines such as trains, which are actually quite efficient at moving enormous payloads, have low normalized work.

$$\text{normalized work} = \frac{\text{speed}}{\text{length}} \times \frac{\text{payload mass}}{\text{mass}} \quad [4-7]$$

Because the dimensions of this metric are peculiar, normalized work is not equivalent to the physical notion of work. Additionally, the incorporation of length in the metric favors small machines. Ants can lift more than elephants because strength to weight is better at smaller scales. But elephants can move a lot more material at higher speed [Eltringham 91]. Yes, an equivalent mass of ants can move more material, but overall payload velocity must be evaluated as well. If the intention was to compare working machines to determine payload ratios and velocities then a better metric would not incorporate length at all.

I suggest a better metric: the product of velocity and payload to mass ratio as shown in Equation [4-8]. This assumes horizontal movement of payloads. In this way, a small robot that weighs half as much but carries half the payload of a large robot, but moves twice as fast, can move the same amount of material over time. The dimensions of payload velocity are distance per time but these are normalized for payload.

$$\text{payload velocity} = \text{speed} \times \frac{\text{payload mass}}{\text{mass}} \quad [4-8]$$

In this way, if an elephant-mass-equivalent number of ants can carry ten times the payload per mass but only move at one-tenth the speed of the elephant then the payload velocity is equivalent. In general, although load carrying increases energy use in all animals, the cost is relatively higher for smaller animals [Schmidt-Nielsen 84].

In many cases though, payloads are irrelevant, especially in research machines or where the sensor payload is a small fraction of overall mass. Additionally, in earthmoving vehicles, such as used in construction, the economics (another pervasive metric!) favor larger machines especially where large forces are required, where maintenance is expensive, and where scheduling and traffic control for large numbers of machines is too complex.

However, the scenario of larger numbers of smaller machine has some advantages. One advantage is overall uptime; if a few machines fail, the task can be continued, albeit at reduced productivity. Additionally, the larger numbers of small machines are decoupled and perform multiple tasks and explore more opportunities. However, these features are not readily quantified as a metric and are strongly task dependent.

Another example of a comparison might be *body lengths per time* but this introduces a scale invariant measure that, again, favors very short vehicles. Clearly, a snake is at a significant disadvantage with this measure! For further overview of performance measures for ground vehicles see [Bekker 69].

What does all this mean for snake robots? For serpentine robots, it is unlikely that delivery cycles and work are defining characteristic. Serpentine robot applications will mostly involve communications; a transfer of information that is only loosely coupled to payloads.

Dimensionless Metrics

As shown, many of these metrics can be reduced to simpler dimensions of time, mass and distance. These units constitute the dimensions of the metric. However, if a metric formulation results in no units, it is a dimensionless number. For example, the ratio of height to width has no dimension.

Dimensionless relationships are important because they can reduce the number of physical variables in a problem, thus reducing the dimensionality of the design space. This is because a dimensionless construct can minimize the number of substantial variables (force, power, velocity, etcetera) necessary to determine or measure. Since the dimensionless numbers are products or ratios of the physical variables, the construction of dimensionless metrics always reduces the number of variables in a system. They also do not require reference to some external standard; e.g. the length of someone's thumb or the wavelength of Krypton. Another reason they are important is that they can provide similitude at different scales. This allows fidelity of comparisons and predictions that extrapolate from models of these systems.

Dimensionless numbers can also provide insight into a problem and reduce experimental effort. The elimination of all but the essential variables brings simplification and the result of this is the clarification or elimination of interdependencies. These features of dimensionless numbers led to a wide variety of dimensionless variables in engineering and science. Commonly used dimensionless numbers include Mach numbers, strain, angular measurement in radians, Reynolds numbers in fluid flow, and friction coefficients [Ipsen 60].

It is possible to construct a variety of dimensionless metrics from the units of interest in any problem. One such dimensionless metric that does not utilize time in its formulation is energy per distance per weight:

$$\frac{\left(\frac{\text{energy}}{\text{distance}}\right)}{\text{weight}} = \frac{\text{energy}}{\text{weight} \times \text{distance}} = \frac{\text{joules}}{\text{newton} \times \text{meter}} \quad [4-9]$$

This metric is a non-dimensional unit akin to a coefficient of friction. The reciprocal of Equation [4-9] has been termed the Net Propulsive Efficiency, NPE, for transportation vehicles whose dimensional formula is ton-miles/gal. This is equivalent to the product of weight and distance per energy [Rice72].

A similar non-dimensional metric has been used for electric vehicles; the driving distance per unit of electric energy delivered by the power source. This is equivalent to

energy used per unit of vehicle weight and distance and is sometimes shown as Wh per ton-mile. The lower the number the higher the efficiency [Kalhammer 95]. This is simply the reciprocal of the NPE metric cited above.

Many of these dimensionless metrics involve power or energy, distance or velocity and weight. All of these appear to be useful measures of performance; combining these may provide a useful and relevant metric.

Specific Resistance

An equivalent dimensionless metric can be formed from the ratio of power to velocity, P/V , a tractive force represented in Newtons, and the weight of the vehicle. This is shown in Equation [4-10]. This metric, first proposed by Gabrielli and von Kármán in [Gabrielli 50], is similar to a global friction coefficient and is called the specific tractive force or *specific resistance*. The tractive force and weight form a thrust-to-weight ratio.

Specific resistance can be thought of as the inverse of the lift-to-drag ratio; a term used in aeronautics where ‘drag’, rather than simply referring to aerodynamic losses, is a general term that refers to all energy dissipation mechanisms. Specific resistance then becomes a measure of the energetic cost of locomotion. The results shown indicate, in general, that any mode of locomotion has a relatively narrow band of efficient locomotion and vehicles that are large and move slowly, such as ships and trains, are most efficient.

$$\text{specific resistance} = \frac{\text{power}}{\text{weight} \times \text{velocity}} = \frac{\text{watts}}{\text{newton} \times \frac{\text{meters}}{\text{second}}} \quad [4-10]$$

Note that Equation [4-10] is the same as Equation [4-9] but with the addition of the time dimension. For vertical motion, the specific resistance = 1 and for horizontal motion, assuming frictionless motion with no air resistance, specific resistance = 0. Think of this in terms of the previously discussed power-to-weight ratio; the power-to-weight ratio is a velocity corresponding to vertical motion. If this is divided by the system velocity and the result is 1, then all power is efficiently devoted to vertical motion. For horizontal motion with no external forces acting against the motion then no power is used and the specific resistance is also zero. There is also a physical limit to power and speed shown by the line in the right of the chart. This represents a limit due to aerodynamic or hydrodynamic drag. That is, there is a minimum value of specific resistance that is related to the speed. It’s likely that this limit is related to the physical limits of materials. That is, to reduce weights or increase speeds requires a more efficient use of materials or stronger materials. It is beyond the scope of this work, but it would be an interesting exercise to evaluate the specific resistance of vehicles over the past fifty years to see

how they improved, or not. Some additional figures of specific resistance for robots were made in [Gregorio 94].

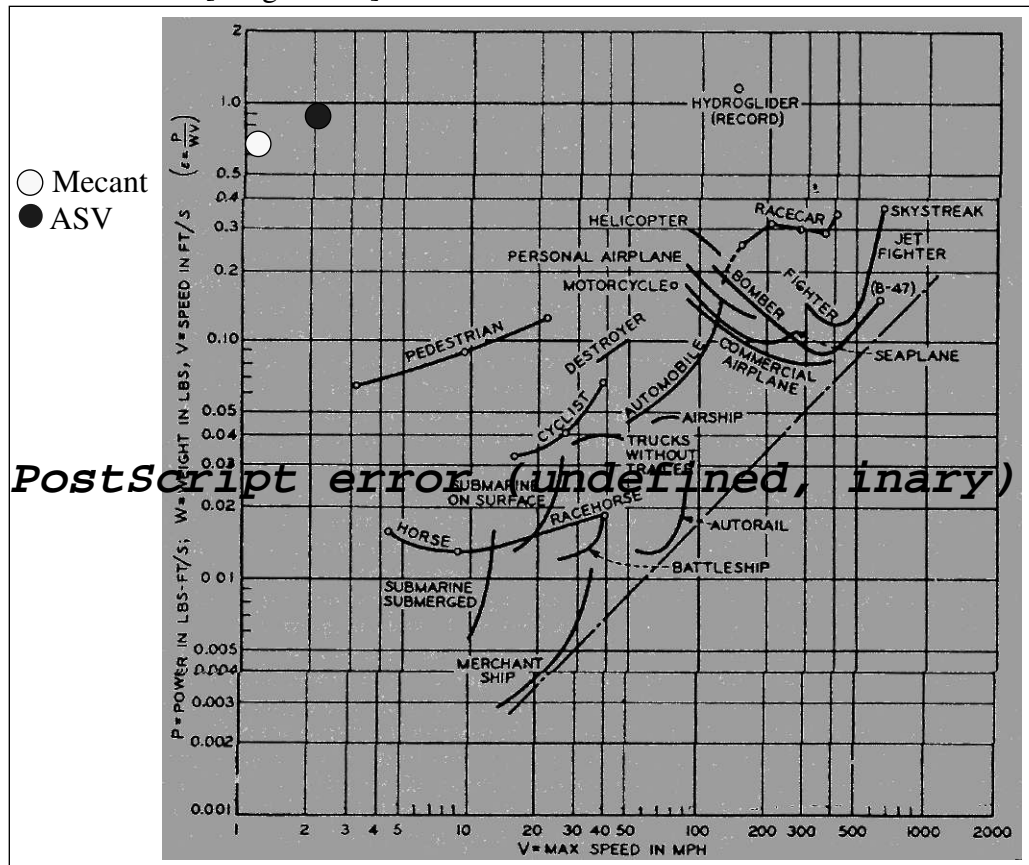


Figure 4-1: Chart of specific resistance for a wide variety of vehicles including two walking robots. From [Gabielli 50]

Although Gabrielli and von Kármán used gross weight for their analysis they also proposed using the useful load or payload in determining the specific resistance. For both animals and vehicles, the net energy and net weight can be used as an appropriate metric for specific resistance. Thus, a classification of four types of specific resistance are possible: net or total weight combined with net or total energy [Hirose 84].

Table 4-1 shows parameters from a variety of walking machines and robots and the corresponding measure of specific resistance. Two of the best performances, the Mecant and ASV walking robots, are also plotted in Figure 4-1. Even the best performing walking machines reveal that a lot of improvement is needed in the design and control of these mechanisms to approach the performance of other vehicles and animals.

The claim is often made that walking machines exhibit their mettle best in extreme terrains and areas in which other vehicles, especially wheeled vehicles, would prove inadequate at best. However, most walking machines, aside from animals, have yet to demonstrate all of the supposed benefits of walking efficiently, even in benign terrain. For smooth terrain and open areas, it is doubtful a robot snake will be as efficient as a

wheeled machine, but for the purposes of self comparison, specific resistance offers a good measure of relative efficacy of locomotion.

Specific resistance is an attractive measure to use for several reasons. The weight, obviously, is unchanging and becomes a constant in the calculations. Power is readily measured in the vehicle and can be determined in a straightforward manner from simulation as well. Velocity, determined by tracking the center of mass of the vehicle, can also be calculated. It is a metric that takes both energy and distance into account, thus providing a simple and effective measure of progress for comparison purposes.

In examining the graph in Figure 4-1, one clear trend across most vehicles is that the specific resistance goes up with velocity. This is not too surprising since energy consumption is related to the increased aerodynamic and hydrodynamic forces required to move faster. This also means that, for a given vehicle, if a low specific resistance is desirable, then this corresponds to low velocities. Thus, if the metric for a given vehicle is to minimize the specific resistance, then a low velocity is the result. For the serpentine robot, however, the effects of aerodynamics or hydrodynamics do not play a role in power consumption, so that velocity effects will be minimized.

Table 4-1: Specific resistance for a variety of walking robots. Some data from [Wettergreen 96].

Robot	Speed [m/s]	Mass [kg]	Power [W]	P/WV [1]
ARL Monopod	1.0	18	60	0.34
Mecant	0.5	1050	3500	0.68
ASV	1.0	3200	26000	0.83
Ambler	0.016	2700	1900	4.49
PV II	0.02	10	10	5.10
Aquabot	0.03	23	500	7.39
TUM	0.3	23	500	7.39
Dante I	0.02	725	1500	10.56
Dante II	0.01	770	1000	13.25
Melwalk	0.01	35	80	23.32

It's tempting, but misleading, to compare specific resistances across very different systems. The applications, tasks, missions and environments are too different for any meaningful comparisons between all the vehicles shown in Figure 4-1 or Table 4-1. Some of the robots were designed to study a particular configuration or control method and not to minimize energy use or maximize payload. In addition, some vehicles are configured for specific environments. However, for a particular vehicle in a given environment, specific resistance is a useful measure of how well the vehicle is performing.

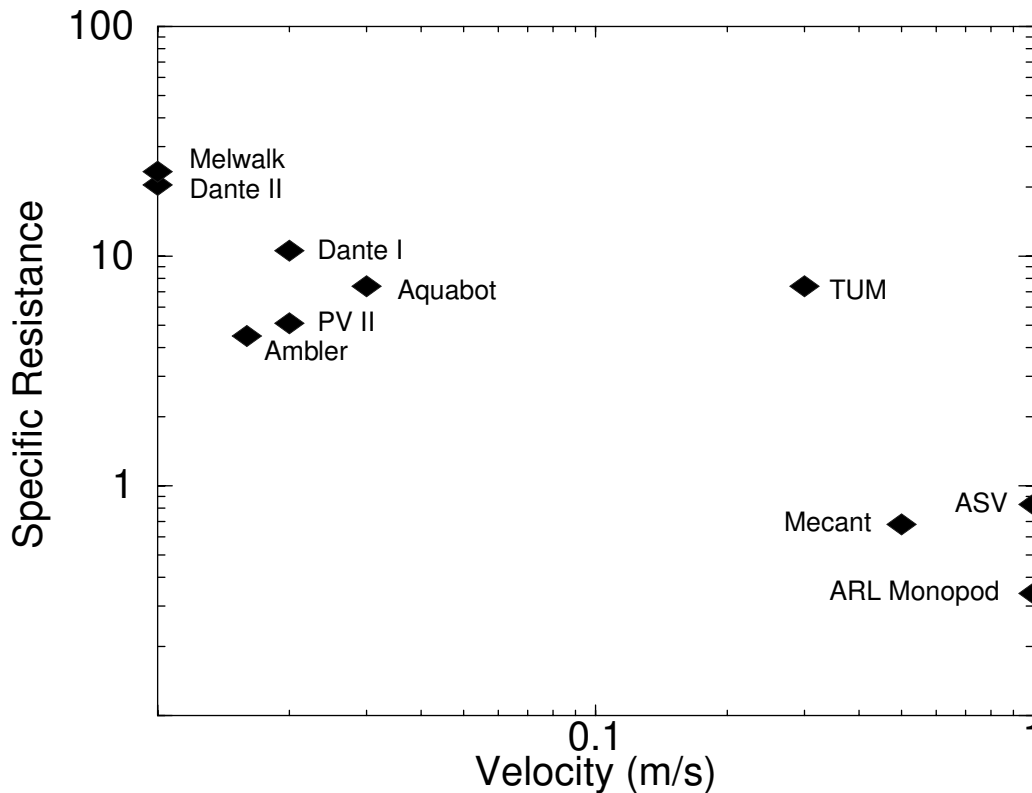


Figure 4-2: Plot of data from Table 4-1 of specific resistance versus velocity.

For a serpentine robot, specific resistance also offers a measure that is straightforward to calculate. Power consumption can be determined both in simulation and in the vehicle, weight is invariant and velocity can be tracked easily as well. The dimensionless quality is nice in that it provides a measure against other vehicles, but this is not the focus of the metric. Comparing the robot's performance against itself is the only key criteria.

Scale

Is it better to be big or small? The motions, shape and structure of animals and mechanisms depends upon size or scale. The answer for robots might depend on the task but there are advantages and disadvantages to different scales. In general, large systems use less energy per distance per mass than small systems. In animals for instance, this expression is proportional to body mass^{-1/3} and this relationship holds for insects, reptiles and mammals. Thus, as shown earlier, the energy cost of traveling a unit distance is proportional to body mass^{2/3}. Energy use does increase linearly with speed, but for smaller animals the rate is steeper than for large animals [Alexander 92] [Schmidt-Nielsen 84].

There are also issues of scale related to geometric and dynamic similarity. Shapes are geometrically similar if scaled by uniform factors of length. Two motions are dynamically similar if they can be made identical by uniform changes of the scales of

length, time and force. The metric of comparison often used for scale comparison is the Froude number, a scale-invariant measure that is often used to compare the dynamics of vehicles [Alexander 92].

The Froude number, shown in Equation [4-11], is a dimensionless number that is the ratio of inertial to gravitational forces.

$$Froude\ number = \frac{speed^2}{\sqrt{g \times length}} \quad [4-11]$$

Alexander showed that leg length in animals could be used in the Froude number formulation to compare scale. In walking and running, different animals run in dynamically similar fashions at speeds only where their Froude numbers are equivalent. For robots, the story is similar, but Froude numbers of robots tend to be very low compared to animals.

Wettergreen and Full provide an evaluation of walking robots using Froude numbers but concluded that the Froude number could not be used to compare robots because the robots were too dissimilar. The performance of the robots versus animals was quite poor. The exception, Raibert's running robots, are the only ones to approach the Froude numbers of animals. Primarily this is because most of the robots compared were statically stable and incapable of running whereas Raibert's hoppers were quite dynamic [Wettergreen 96]. Perhaps another reason is that all power and computing for the running machines was offboard the robot. The animals do better primarily due to energy recovery in muscles and tendons which allows higher speeds.

One issue with Froude numbers is that they don't reveal the suitable optimization criteria that result in different gaits. In [Alexander 84], it's postulated that two primary criteria are maximizing stability and minimizing energy consumption for different animals. Clearly, in biped and quadruped animals, gaits are a function of speed and minimize energy expenditures for that given speed, whereas serpentine gaits appear to be a function of environment.

An obvious problem with the formulation of the Froude number is that it depends on a measure or dimension of length. For walking robots, this could be the length of legs, strides, or steps. The problem for limbless robots is determining what the analogous height or length dimensions should be. Simply growing the length of a snake robot might drastically affect the dimensionless value without changing the velocity. Hence the Froude number for a snake is contrived.

Since most serpentine gaits exhibit characteristic patterns or waves, wavelength might make physical sense for limbless robots. In the example below, lambda is the characteristic wavelength of the body of the robot [McMahon 96].

$$Froude\ number = \frac{speed^2}{\sqrt{g\lambda}} \quad [4-12]$$

The problem with this, as shown in **Background**, is that snake gait selection and progression do not appear to relate directly to undulating frequency. As demonstrated

in [Secor 92], if the mean frequency and forward speeds are the same for different gaits, then the mean distance travelled per cycle must also be equal. The problem was that the energy use difference between gaits was significantly different. Thus, the Froude metric does not show the difference in its evaluation.

Another measure which may be appropriate for undulating vehicles is amplitude per wavelength. It is not clear, however, if this is to be maximized, minimized or made to approach unity!

Velocity effects from scaling can produce other effects. At higher velocities the effects of air resistance are significant. However, air resistance is much worse at *smaller* scales as a percentage of energy output if velocity is constant with size. This is because surface and cross-sectional area increase more slowly than volume (mass).

Surface area increases as the square of the size whereas volume and mass increase as the cube. Thus, wind-drag, which is dependent on surface area, is proportionally smaller for a heavier and larger object than a smaller one of similar shape and composition. An everyday example is that of falling dust and rocks. They are the same shape and composition and differ only in size, yet the dust settles much more slowly than the rocks.

For similar reasons, small things tend to do more work against friction because surface friction effects tend to be proportional to area; the ratio of viscous forces to inertial forces increases as size decreases. This has a severe damping effect on very small vehicles. In some ways, however, small scale can assist the designer. Small things have better strength-to-weight ratios and may require proportionately less structural mass. Air resistance does not affect the serpentine robots considered here, and issues with very small size, other than strength to weight ratio, are not relevant for this system.

Small animals consume more energy to carry a unit mass a given distance than larger animals. Although energy use is proportional to speed, energy use is also lower per unit mass for larger animals; smaller animals use relatively more energy to move a load over a given distance. An excellent discussion of this and other issues of scale can be found in [McMahon 83].

Summary and Selection

Metrics have been only cursorily examined in many papers, especially for robot performance evaluation. This discussion was necessary to set the stage for a final selection of a metric and provide a comprehensive background on the subject.

I selected specific resistance as the measure of performance for learning locomotion. It provides a notion both of energy, time and weight of the robot. It utilizes two measurements that can be found from both the physical model and from simulation. It is easily and quickly calculated and provides a clean and understandable metric for evaluation during the learning process.

Whatever the particular metric value, it is not a good idea to draw too many conclusions or provide close comparisons to other robots. It's too easy to contrive a metric that favors a particular robot. It's also too easy to draw conclusions about vehicles that don't take environment and task into account.

However, it is important to realize that metrics reveal only how well a vehicle did on a particular performance measure. It does not reveal why, although it can provide clues, and, finally, it does not directly reveal how to make the performance better. It can be used as a tool to ascribe trends through the use of small changes in the control techniques and hence, develop a better understanding of what makes a better gait. The metric developed is used as part of the learning process and placed into the overall framework for teach the snake robot to locomote.

Chapter 5

Learning and Optimization

Machine learning techniques evaluate past data to form insights on future performance; learning provides improved performance through experience. **Learning and Optimization** examines learning locomotion for simulated mechanisms and actual robots as well as criteria and structures for learning. This section also looks at the critical area of parameter representation. While the technique selection is important, knowing what to optimize, and knowing how to evaluate a solution are even more critical. *What to Optimize* is a metric, specific resistance. *How to Evaluate* is determining this from simulation or measuring from the robot. With the metric and the learning method combined with physical modeling a complete framework can be constructed.

Optimization Techniques

A wide variety of search methods have been developed to find solutions to problems in mathematics and computer science. A few techniques are widely used and work well in small domains but bog down to intractable levels for large high-dimensional search spaces. I examined and tested a number of these techniques in the course of this research.

Random Search

Random search is unsuited for all but the smallest problems. The technique simply chooses parameters at random and evaluates the metric by simply keeping the highest value around. If the search space is not large, then this is a simple, albeit not efficient, strategy.

Hillclimbing

Hillclimbing, a simple and sometimes very effective method, simply looks at the local gradient and moves ‘uphill’ in the direction of better solutions. Hillclimbing techniques

may take a random parameter, but use adjacent points in the space and evaluate the gradient to find the direction of higher scoring evaluations. Hill climbing is at the root of many successful techniques which use more advanced strategies including steepest ascent or conjugant gradient techniques which better optimize the direction of subsequent search. While a great deal of effort in research is focused on pther search techniques, often these more straightforward techniques still prove very successful.

A general issue with hillclimbing techniques is that the search tends to find the nearest and highest hill; the local maximum. Peaks that are higher still are not found because they lie across valley floors and other lower elevation areas. Finding a global optimum then requires new techniques or a combination of techniques.

Simulated Annealing

Based on an analogy to the slow cooling of alloys, simulated annealing is similar to hill climbing with an added ‘earthquake’ component. By adding energy to the values, the current search areas can be jumped or bounced to other search areas. It’s a stochastic technique for finding near globally minimum cost solutions to large optimization problems. Simulated annealing is sometimes much better at avoiding local minima that may trap pure hill-climbing. However, there are limits to the cost function and simulated annealing can be very computationally intensive for finding satisfactory optimal results for many problems.

Neural Nets

Neural nets utilize a network of units, each of which has inputs and a simple weighting function to provide an output to other layers within the network. Neural networks work well at providing classification and regression functions but typically require a ‘master’ or training set that is used to initially provide the weights within the network. Optimization using a training set establishes the weights of a number of interacting ‘units’ within the neural net. In this case, the lack of training sets is a formidable obstacle. It is possible that the neural network can provide an architecture, however, for a local distributed control scheme for the snake wherein local joints listen and respond to adjacent joints.

Response Surface Methods

A response surface method (RSM) is a graphical representation of a relationship between some simple-to-evaluate metric such as yield and a large number of variables. Typically, you wish to find the values of the variables to maximize the metric. It differs from ordinary optimization techniques because gradients often cannot be found, it is an unknown function and the measurements of the function are typically noisy.

Response surface methods can provide good experiment design plans and good statistical analysis but the underlying structure relies on linear multivariate regression analysis. This type of analysis can behave badly when non-linearities are present. It is also typically used where there are slight variations in parameters rather than large changes. This makes it well suited for industrial process control but difficult to apply to applications where widely varying parameters are likely to be searched. Typically, RSM is used in conjunction with human judgement as well as statistical analysis [Box 87].

Genetic Algorithms

Genetic Algorithms (GAs) use a stochastic process to enable the testing and evaluation of many individuals over time. Genetic algorithms are a guided random selection process that utilize the following ideas:

- *Selection and coding of parameters*

Parameters are selected and represented with binary strings at a sufficient resolution so that rapid changes of the output do not occur with small changes in parameter values.

- *Search with population of points or ‘individuals.’* This is accomplished by generating multiple sets of parameters.
- *Objective functions to evaluate each individual.* The evaluation of each individual is carried out and a fitness value is returned.
- *Probabilistic transition rules.* These are used to cull or glean individuals that result in high evaluations and uses them to generate the next generation of individuals. This process is repeated and eventually converges to a solution, but not necessarily a global one.

parameter	magnitude	frequency	phase
genome	1111111111	1111111111	1111111111
values	0-10	0-100	0-10
resolution	~0.01	~0.1	~0.01

Table 5-1: Sample set of parameters for GA used for caterpillar robot.

Table 5-1 shows a sample genome and parameter encoding for a three parameter system. The genome values show the size of the bitstring and not the actual value used. I implemented a series of tests using GA's to control a 2D caterpillar robot. For the GA tests I instantiated a number of parameters including population size, number of generations, mutation, and crossover. For the program execution, I used tests of 100 generations and populations of 30.

The following example, in pseudo-code, shows a sample form of the objective function:

```
float Objective(GAGenome& c)
{
    write parameter file ("parameters");
    system call ("snakes -t 5 -nodisp");
    // read and return the metric result
}
```

The parameters represent a specific set of values generated by the learning technique. These values are passed through a file to the program, `snakes`, which creates and runs the simulation. The simulation runs for a fixed amount of time and the evaluation metric is then written out to a file. This occurs for each and every set of parameters that are tested in the learning algorithm. This metric value is read and then returned to the learning function.

Although writing files is not as efficient as passing information through other means, such as sockets, it makes for a method that is easy to implement, maintain and understand. The system call runs another program which, in turn, calls the physical simulation program and runs it without a display and for a specified period of time. When `snakes` exits, it writes a file that contains the value of the metric. This value is then read by the program and used in the regeneration of the parameter statistics. The overhead in writing these tiny files, less than 1K, is quite small compared to the overhead of running the simulation itself.




Figure 5-1: Example of GA convergence to maximum value of metric.

PBIL

Baluja, in [Baluja 95a], introduced Probabilistic-Based Incremental Learning (PBIL) as a means to provide the same functionality as other stochastic methods such as GAs but in a more efficient manner. Rather than implicitly maintaining statistics within a population, as GA's do, PBIL explicitly maintains the statistics for the genome.

In this pseudo-code example derived from [Baluja 95a], the program shows the sequence of generation, update, and evaluation.

```
// initialize probability vector
for i = 1 to LENGTH do P[i] = 0.5;

while (NOT termination_condition) {
  // generate samples
  for i = 1 to SAMPLES do
    sample_vectors[i] = generate_sample_vector(P);
    evaluations[i] = evaluate_solution(sample_vectors[i]);
```

```

    find_best_vector(sample_vectors, evaluations);

// update probability vector towards best solution
for i = 1 to LENGTH do
    P[i] = P[i] * (1-LR) + (best_vector[i] * LR);

// mutate probability vector
if (random(0,1) < MUT_PROB then
    if (random(0,1) > 0.5 then mutate_dir = 1;
    else mutate_dir = 0;
    P[i] = P[i] * (1-MUT_SHIFT) + (mutate_dir * MUT_SHIFT);
}

```

User defined constants (example values)

```

SAMPLE: number of vectors generated before new update (100)
LR: Learning rate, how fast to exploit search (0.1)
LENGTH: number of bits used in generated vector (30)
MUT_PROBABILITY: Probability of a mutation occurring in each
position. (0.02)
MUT_SHIFT: amount a mutation can alter a value in the bit
position.(0.05)

```

This encoding of the solutions as statistics can, in many cases, be far more efficient than traditional GA methods. The executions are more efficient, use fewer cycles, and converge more quickly. The complete coding, as shown in this example, is about all that is required for implementation. `find_best_vector`, for example, in my implementation is not written as a separate function but simply integrated into the evaluation loop. Thus, PBIL provides similar, if not better, performance than GA's with lower overhead.

Representation

Gaits or sequences of motion can be represented in a variety of ways and by several criteria, described below, can be used to evaluate different representation techniques. The desired result is a single representation that, with carefully selected or generated values for the parameters, represent all the forms of snake locomotion that are periodic in nature. This includes sidewinding, lateral undulation, rectilinear and concertina.

Although the problem is of greater dimension than a landscape, imagine a varying landscape that represents the solution space for gait generation. The peaks and valleys, by analogy, represent the good and bad metrics that result from parameter values. The task is to find peaks and ridges where high values correspond to good gaits. The task is to peer through the fog that covers the landscape and discern, not only the high peaks and ranges, but patterns and structure. Trends may make it possible to classify regions and areas into styles of gait.

However, even given a means of learning, there still remains the problem of selecting a way of representing the information presented to the technique. The general problem is to represent a wide variety of progressing waveforms that provide a systematic displacement of the robot. The waveform parameters are adjusted during the learning process to provide efficient motions of the snake joints and thus provide locomotion. There are several competing issues for any representation and these include:

- Compactness - the ability to represent the most information possible in the most concise manner.
- Calculation - how much overhead does the representation require?
- Complexity - how involved is the creation, debugging and evaluation of the representation? This is really an implementation issue.
- Comprehension - How easy is it to interpret and understand the results? This is different from complexity although it can be related.
- Correspondence - How easily is the information mapped from representation in learning to parameter values for simulation and control?

In this section we'll investigate several means of representing the parameters for learning locomotion gaits.

Trigonometric forms

In the first set of tests on the caterpillar, a simple trigonometric sine function was used for representing the traveling wave on the robot. This form has only three parameters: magnitude, frequency and phase. Magnitude and frequency are obvious and phase represents the shift of the waveform along the body. An additional parameter, offset, can be used to provide a nominal starting configuration. For example, the body may form a helix which is deformed to provide locomotion. Sinusoidal patterns are easy to represent and parameterize. The problem is they appear too restrictive in representing arbitrary time-varying waveforms.

Fourier

Ideally, a relatively simple function like a trigonometric function or a combination of such functions is both simple to represent and easy to parameterize. However, a counter example that can not easily be represented in this manner is shown in [Hirose 93] for lateral undulation; the serpenoid function that he derived represents *only* lateral undulation. The serpenoid curve is derived by assuming that curvatures vary in a uniform fashion along the length of the curve defined by the snake. The curvature of a sine function, however, does not smoothly vary *along the curve*. This derived curve, the serpenoid curve, was compared to that of natural snakes performing lateral undulation and shown to closely approximate their actual motions.

Hirose's derived formulation uses Bessel functions which are the results of an integral that cannot be expressed in a simple closed form. However, it is possible to use a Fourier

series with parameters that represent coefficients of the individual terms up to some cutoff frequency (spatial frequency).




Figure 5-2: A representation of time and joint link versus angular value.

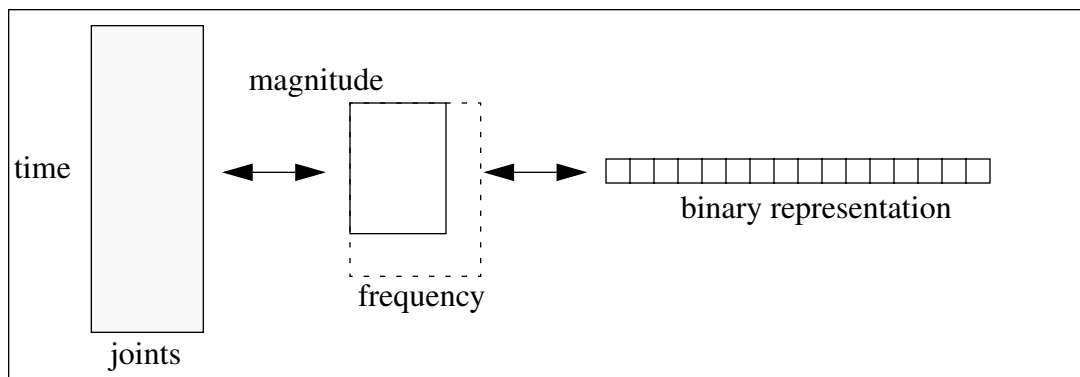


Figure 5-3: Mapping from time varying representation to fourier representation to genome.

The Fourier series which can represent any periodic function as a sum of exponential terms, usually sine curves. Figure 5-2 shows an example end result where the angular values of the angle between links are a function of time and link. Since all gaits, by definition, are periodic sequences, the Fourier series can represent them. In Figure 5-3 the mapping of the time and joints to the coefficients of the Fourier series which are magnitude and frequency values. This table of values, in turn, can be represented in a binary string which facilitates operation within the learning framework. The values in the string can be the coefficients of the values in the array.

Parametric curves

Alternative representations are the parametric forms shown in [Chirikjian 92] or the 3D splines used in graphics. Linear polynomials in \mathbb{R}^3 are another, but the problem is the number of coefficients, at least 20, that need to be represented as well as the sequence.

Higher order polynomials also require a relatively large number of parameters and representing all time varying sequences can be an issue for learning and evaluating.

Wavelets

Another means to represent forms in a concise manner are wavelets which, unlike Fourier series, can be used to represent non-periodic time-varying functions. Complicating the use of wavelets are the selection of the mother wavelet and a wide variety of choices for the representation use. Since gaits are periodic in nature, there doesn't seem to be particular advantages of the wavelet approach, although representation of gait transitions and terrain accommodation may be more easily managed in a wavelet framework.

Tables and Masks

Rather than trying to explicitly represent all manner of waveforms or trying to identify parameters to refine for different modes of locomotion, a conceptually simpler route is to directly represent the joint angles over time. By explicitly representing the angular position of the links within a one-dimensional tape, the joints can be essentially 'masked' off and the snake shifted through the tape, adjusting the positions to reflect the tape values. By adjusting the set of tape values in physical simulation and looking for effective modes, a variety of locomotion modes can be represented, including those that are not exhibited in natural snakes. This representation then forces the time-histories of the individual joints to be identical but only shifted in time.

A more general extension of this method is to represent the joint angles in a column of




Figure 5-4: A snake 'tape' defining joint angles at each time step for a periodic waveform.

a 2-dimensional array and, with each timestep, march across the columns where the column entries represent time steps for a particular waveform. This should work well for periodic forms, but also for forms such as concertina where the time history of different joints is not the same. The array becomes the representation of each joint angle or, even more concisely, the difference between the angles in adjacent time slices.

The values in the table can also be constrained to represent the likely sequences. This constrains movement between slices and even joints. This has a two-fold benefit: first

it prevents abrupt jumps and it culls unlikely gait patterns and body contortions, and finally, it reduces the gait space significantly.

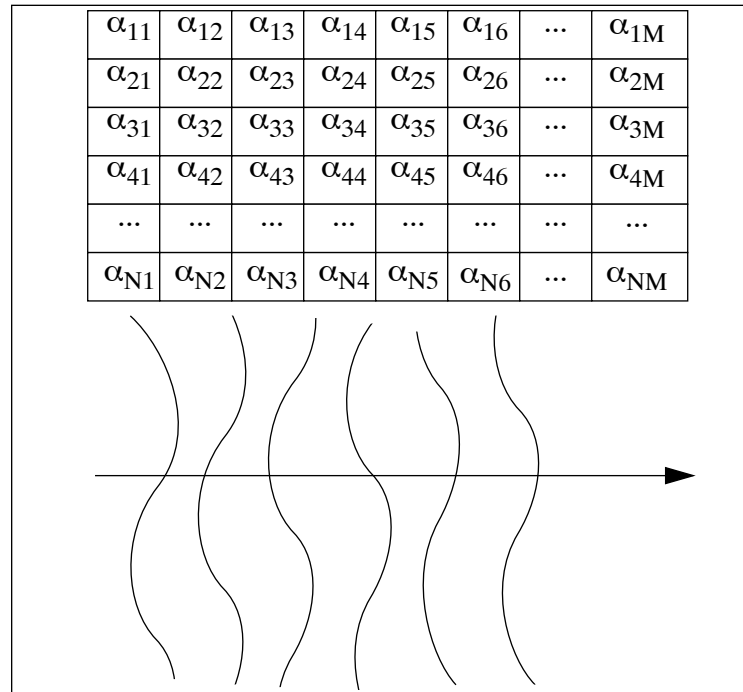


Figure 5-5: A 2D array can be used to represent all joint motions over time.

Another issue is the size of the array. Using some rough numbers, if a given gait sequence takes 2 seconds before repeating and the robot has 20 DOF, then using a 10Hz update rate gives 400 parameters. Each of the parameters might use a full 8 bits and if differences are used, 2-3 bits. This gives a total of a few thousand total bits, which creates a very large set of parameters to adjust. However, while the space of configurations described in this model is large, the representation is simple, easily understood and easily mapped to the robot. Another possibility is to operate on number values themselves rather than their binary representation.

The time to cross the table is directly related to the frequency of a gait. For legged animals the stride frequency is inversely related to the square root of the leg length, so that even for long legged animals the stride frequency is typically less than a second or so. Leg length is also proportional to the cube root of the mass, so that the stride frequency can be shown to be proportional to the sixth root of mass [Alexander 92]. If this proportion holds for snakes, then the stride (undulating) frequency appears to be less than two to three seconds even for the largest, presumably slower, snakes. Thus, the time across the table can be up to two seconds and the partitioning should be on the order of 8 or 16 numbers to capture the variations in the joints during that time.

Filling in the table

Given the tabular representation, the next issue is generating values. The most straightforward way might appear to simply fill in the table with random values and insure that the changes across the rows and columns are less than or equal to some threshold. For example, although this is straightforward to implement, the problem with this method

is that, since the table wraps around, the random walk distance is effectively cut in half. This well-known result is that the distance of a random walk is approximately the square root of the number of time steps. Thus, for 32 total time steps, the random walk distance is $\sqrt{32/2}$ or four times the change that is allowed from step to step. So unless very large steps are allowed or the array is made quite large this method does not do well for describing large motions of the joints.

The irony is that the constraint should keep the changes manageable but to also allow significant overall change in the positions of the joints. As just shown, the random walk doesn't accomplish this requirement.

To reduce the universe of search spaces it is possible to 'seed' the table with mediocre hand-tuned locomotion modes and by iterating between the simulation and physical simulation to arrive at efficient modes of locomoting. The result is likely to converge to a local maximum.

Another way is to use a fourier series, in a different manner, to represent the angle in the following manner: since a fourier series uses magnitude and phase and relates them to frequency, it is possible to define both to create a waveform with the right properties. In this case, of course, it the spatial frequency of the joint angles. The magnitude function can be defined as simply as $1/f$, the power spectrum for noise. This creates a rolling cutoff to prevent high spatial frequencies where the joint angles change abruptly. The phase is then seeded with random values and the coefficients are calculated and then the series is created.

The magnitude function can be tuned to effect a good sweep of values across the positions. Several tests resulted in only a small modification of the magnitude function to provide an earlier cutoff.

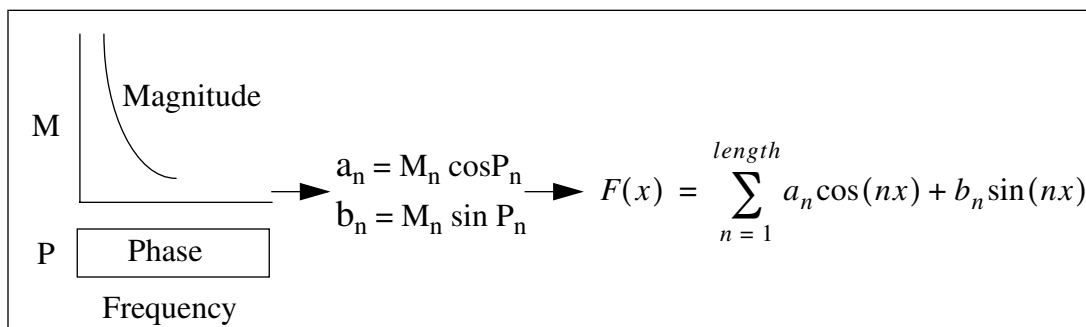


Figure 5-6: A fourier series is used to generate joint positions for tabular method.

It isn't necessary or even desirable to create an array that defines all timesteps in the process. Too many and the search space explodes and convergence, even if possible, takes an inordinate amount of time. Too few and the coarseness results in ineffective gaits.

Most natural serpentine gaits appear to take on the order of 2-3 seconds. By partitioning this time into about 10 slices/second gives 20-30 numbers to describe the gait over time. For these reasons and pragmatic coding reasons I chose 32 slices for the array. One result of the fourier seeding is shown in Figure 5-7. Again, this is one possible gait sequence generated by the fourier technique.




Figure 5-7: Table generated from the fourier technique; columns represent joints angles and rows represent time history.

The landscape of gaits in this representation is quite extensive and the process of determining good gaits in this landscape is to identify the peaks and ranges that correspond to demonstrated efficacy in simulation and in the robot.

Summary

Learning provides improved performance with frequent testing and evaluation. After examining and testing a number of the learning methods described here, I selected a stochastic technique, genetic algorithms or GAs, and, later, probabilistic-based incremental learning or PBIL. PBIL's compactness, ease-of-implementation, effectiveness and simplicity made it a good choice for evaluating metrics in physical simulation.

The representation chosen for the framework is the tabular array with entries that directly represent desired positions of body segments. This representation is the most general and does not depend on combinations of mathematical functions to attempt all possible configurations; it represents those configurations directly. The danger in the tabular approach is that it opens the search space further, but the generality appears worth the risk. The trigonometric approach to representation turns out to be surprisingly powerful one and this was used for a number of gait applications as well.

The next step in the process integrates the learning technique and learning representation into the framework.

Chapter 6

Implementation

Implementation details the physical configuration of the serpentine robot. Configuration includes actuation selection, morphology, and design of the mechanics and electronics as well as the experimental setup and physical modeling. For each area alternatives are presented that were explored and evaluated, as well as the final selection and form of the serpentine robot.

There is a cyclic design process at the core of the framework described here; the use of simulation can assist the design process. Designs that do not work in simulation are unlikely to work in the real world. Hence, an analysis of form and its effect on locomotion is productive and useful. Similarly, the choice of actuation technology can set constraints on the robot design.

As a result, actuation technologies were closely examined and a short summary of that evaluation is provided here. The final selection, small off-the shelf-servo actuators, are analyzed to provide a good model for simulation. An important geometric analysis reveals that the angular excursions of joints can be small and that robot link aspect ratios, the relationship of length to width should also be small.

The final mechanism, is a lightweight assembly of hardware, bracketry and servos connected by a very small wiring harness to provide control signals from a set of small controller boards. Skins, an aspect of snake robots that has been ignored is also examined and candidate materials were identified. The final assembly of the snake robot provides a highly articulated twenty degree of freedom machine with integral controllers and control bus.

Actuation

To create motions and sequences of body shapes requires devices that move or actuate. There are many technologies that are capable of creating motion but there are also many

other issues involved in the selection process. These include fidelity, response, power, speed, torque, required infrastructure, etcetera. The selection of actuation technology also directly affects the configuration and control and, as a result, the actuator selection process is critical and integral to the design process. A wide variety of actuation technologies were closely examined and evaluated in the course of this work.

Some of these technologies were initially examined with the intent of using scaled snake vertebrae in a robotic mechanism. I performed a high density scan of a vertebrae in a Magnetic Resonance Imaging (MRI) machine at a hospital and then constructed a 3D model of this data. The model is shown, along with the actual vertebrae, in Figure 6-1 [Carnegie 95]. The intention was to carry this into a rapidly prototyped model using stereolithography and then use muscle-like actuators for providing motions of linked vertebrae bodies. The problem, as we will see, was the immaturity of the technologies that provide the muscle-like action.

Several areas of actuation were examined in detail and the following section on actuation technologies briefly summarizes each area.




Figure 6-1: Northern Anaconda vertebrae and 3D model constructed from MRI data.

Actuation Technologies

Polymer Gels

Polymers are groups of small molecules that form long chains that are essentially repeated units of the smaller groups. They exhibit a wide range of properties and most of the synthetic materials used in our daily lives are polymer-based. Some polymers are capable of converting chemical energy to mechanical work in isothermal conditions. These polymers significantly change their length in response to chemical changes involving altered temperature, pH, or applied electric fields. Volume changes can be as high as a factor of 1000. Gel polymer networks are a balance of such properties as rubber elasticity, polymer-polymer affinity and hydrogen ion pressures and changing the balance of these determines the volume change.

For polymer gels to be useful there are many technical issues to resolve. There are issues of strength, response, stress-strain relations, fatigue life, thermal and electrical conductivity. Other issues include efficiency, power and force densities and power

limits. Finally there remain engineering considerations of supply and delivery of power, construction, manufacturing, and modeling of these actuators [Brock 91][Caldwell 89].

Shape Memory Alloys

Nickel-titanium alloys and their useful properties were discovered by the Naval Ordnance Laboratory decades ago and the material was termed NiTiNOL. These materials have the intriguing property that they provide actuation by means of current cycling through the materials. The alloy undergoes a reversible phase change exhibited as force and motion in the wire. At room temperature, nitinol wires can be easily stretched by a small force. However, when conducting an electric current, the wire heats and changes to a harder form that returns to the unstretched shape; the wire shortens in length with a usable amount of force.

Nitinol can be stretched by up to eight percent of their length and will recover fully, but only for a few cycles. However, when used at smaller strains, such as three to five percent, nitinol wires can run for millions of cycles with very consistent and reliable performance. Strain is much higher than piezos; as high as 8% with corresponding forces as high as 600N/mm².

However, the response time is contingent on heat removal and is relatively slow. As a result the efficiency is very low, as is the stiffness. Typical efficiencies of Nitinol materials are on the order of 5-6%. Hysteresis is a problem and fatigue life is relatively low. Additionally, the generated heat can be an issue in many applications.

A wide variety of grippers, manipulators, dextrous hands and even small swimming devices utilize SMA materials. For practical robot applications however, there is little beyond a few demonstration devices such as small walkers and heat engines [Dario 89][Golestaneh 84].

Piezoelectric Devices

Certain classes of crystal materials change length with applied electric potential. Conversely, they can produce electric fields when put under pressure. These materials are termed piezoelectric or PE and can produce high forces at good efficiencies.

The difficulty for some applications is that the motions produced are extremely small; The strains are on the order of 1% or less. Advantages of piezoelectrics include the ability to control small (sub-micron) displacements with applied voltages, very high stiffness and very fast response. Loads into the hundreds or thousands of Newtons are easy to achieve. They are very stiff as well; the modulus of elasticity, E, can be up to 100 GPa. As a comparison, Steel is about 200GPa and Aluminum is about 70GPa.

Disadvantages of piezoelectrics include very small displacements; 30ppm is typical. An additional concern for piezos is that high electric fields can cause breakdown and failure. Another general problem is non-linear response and high hysteresis and creep [Petrucci 94].

Electrostriction Devices

Unlike PE's electrostrictive crystals are symmetric. The electrostrictive strain is proportional to square of electric field. This property is independent of piezoelectric effect and is due to rotation of polar domains in ceramic through the field.

In general, linearity and hysteresis are better than PEs and lower voltages are used. Movement ranges to 105 microns are available in commercial products but the ratio of length to change in length is still low and is on the order of 0.1%.

The motion has low hysteresis and very small thermal expansion coefficient and the non-linearity can be overcome by operating at a bias voltage. Most electrostriction material properties are similar to piezoelectrics. A number of commercial electrostrictive devices are now offered. The differences from PE's offer advantages in some applications when there are issues with the high voltage and hysteresis associated with PEs.

The combination of high-voltages, small strain, similar to PEs make these difficult for a large actuator implementation. PEs may be the technology of choice for very small scales however.

Magnetostriction

Magnetostriction is the mechanical deformation of a ferromagnetic material when subjected to a uniform magnetic field. Unlike piezoelectrics, the displacement per unit field actually increases with length. Internal stresses in the material due to anisotropy energy are required to magnetize it in certain directions relative to the crystal axes and vice versa. The strains and displacements can be significantly more than piezoelectrics but piezoelectrics can be stacked to give nearly the same stroke per length. Terfenol-D, used in several magnetostrictive commercial products, offers high forces and good strain [Dyberg 86]

MEMS

Micro-Electrical Mechanical Systems are a relatively recent development by which fabrication techniques, normally associated with integrated circuit design, are used to build mechanical structures that can be moved and controlled. The field is rapidly developing and already micro robots on the μm scale demonstrate motion control, mobility and sensing capabilities. As of this writing however, no techniques for coupling the microscopic motions directly with macroscopic motions has been achieved.

The physical principle used in most MEMS actuation is electrostatics. This is based on the force on electrons in an electric field. Any two electrodes separated by an insulating barrier will be attracted to each other and the force is related to the square of field.

Whereas electromagnet forces depend on volume of magnet present, electrostatic forces become significant at small gaps. Traditional-sized motors using this principle result in very low forces and torques. But at very small scales, electromagnetic motors become very inefficient. Most of the power is consumed as heat and the torque is very low.

Below 1mm though, electrostatics looks very promising. The charge on a particle is large compared to particle volume if particle is small. However, there are issues of temperature and humidity suggesting use of dielectrics other than air. Electrostatic techniques can work well in vacuum too. Even though the force change is nonlinear it is a very exact relationship extending over several orders of magnitude.

Recent work by Flynn, has identified a gap in actuator technology between the MEMS technologies and traditional motor technologies. This gap is roughly between 100 μ m and 1mm. Flynn identified PE wave motors as one potential technology to fill this gap in actuation technology [Flynn 92][Stix 92][Congress 91].

Thermal Actuators

For decades, thermostats in automotive cooling systems utilize the expansion of materials as a means to actuate valves. A typical automotive thermostat uses a paraffin or wax actuator to open coolant flow to the radiator. After reaching a set temperature the wax undergoes a phase transition, i.e. it melts. The wax expands as it melts in a small confined chamber and squeezes a rubber boot that pushes out a small piston. This form of actuation has high force, long lifetime and it is very reliable. Response time depends on power input but can take many seconds. The transition temperature can be set to any value simply by changing the wax mixture.

Spacecraft use thermal actuators and many reliable designs have been built [Starsys 95]. A new development in this form of actuation is the use of thermopolymers as the phase-transition material. The key attributes of the new material are the rapid response and cycle times.

As with all thermal actuators the removal of heat is a critical issue. Based on calculations from published specs, such actuators are about 5% or less efficient. This is low, but not lower than other types of actuation such as shape memory alloys [Tcam 95][Schneider 91][Schneider 93].

Electro-magnetic Motors

The interaction of magnetic fields and current-carrying conductors produces a force which is harnessed in the form of a motor. The design of motors is relatively mature compared to the other technologies discussed here. The principles have not changed in over a century but better magnetic materials, better tolerances and improved control have resulted in the continuing evolution of small, high performance motors coupled to efficient drivetrains. Thermal considerations and magnetic field densities appear to limit motor technology at this point in time, but continued improvements in power density and the advent of superconductive motors will further performance and design.

Actuator Selection

After examining each of these technologies and the commercially available versions, I concluded that the technologies, in most cases, are either immature or unsuitable or require substantial development beyond the scope of this work. As a result, I re-examined the use of small electromagnetic DC motors and configurations of gear and lever drives. After closely examining a variety of small gear motors and drivetrains I selected a packaged actuator used in a variety of applications.

Servos

Scale models of planes, cars, boats and helicopters and other vehicles use modular actuators for steering, moving surfaces and for controlling larger actuators such as high power and high speed DC motors. Radio control, or R/C, servos are small geared DC motors that provide closed-loop position control. New generation devices are rapid,

precise, lightweight and cost effective compared to small DC motor and gearhead alternatives or packaging separate components.

R/C servos provide closed-loop position control of angular position and newer servos also provide control of linear position. As shown in Figure 6-2, the control signal is a pulse that is repeated every 10-30 milliseconds. The width of the pulse determines the position of the servo. A pulse-width change from 1 to 2 milliseconds will sweep the position of the actuator from one extreme to the other. Typical angular excursions of servos are about 60 degrees, though many servos can mechanically provide 180 degrees of motion.

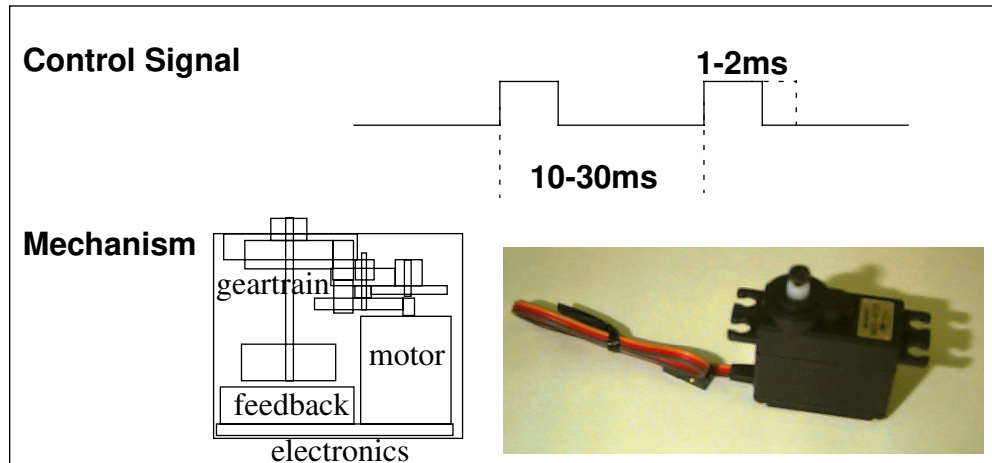


Figure 6-2: Servos use small and efficient geartrains integrated with a positioning control loop.

There have been significant improvements in R/C servo design and construction over the past several years. Recent designs use coreless motors, integrated electronics using surface mount technologies, custom integrated circuits, strong roller bearing support and O-ring-type seals for use in adverse environments. I undertook an evaluation of a variety of manufacturers products to compare the products. Appendix A details the servo specifications and resultant figures for about forty servos from six different manufacturers. Conversations with users and marketers also indicated that some manufacturers also slightly inflate specifications.

The final servo selected, the JR4721, showed significantly better performance over other servos in metrics that included power to weight and torque to volume ratios. Specifications for the servo are shown in Table 6-1. The power is calculated as one-quarter the product of speed and maximum torque. This is a typical rule-of-thumb for brushed motors; the product of the maximum no-load speed and maximum stall-torque. On a torque speed curve this is generally a linear relationship where the maximum power is halfway between the two points and, thus, is given by one fourth of the product.

Table 6-1: Specifications for the selected servo.

Model	Max Torque [Nm]	Mass [kg]	Speed [60000]	Power (W)	Power/Weight [W/N]	Torq/Weight (Nm/N)
JR4721	0.84	0.049	0.22	1.21	2.52	1.77

Servo Modeling

A model of the servo is needed for the simulated physical model of the serpentine robot. The servo is treated, appropriately enough, as a black box and the output of the actuator is examined for a specified input. The actuator is loaded and then the response observed in reaction to reference commands that move the servo to a given position. The relationship of time and angular position gives a response which is analyzed to provide parameters of the servo. Since, in the simulation, the various gains can be modeled, it remains to provide a model of the servo by testing the physical device. An experimental setup using a position tracking device fixed to a bracket and then attached to the servo is used to find the angular position versus time.

From observation, R/C servos appear to use a simple proportional control, but because of the high gear ratios, the servos do quite well in tracking the reference signal. The gearing ratio of servos is on the order of 300:1. Since the reflected inertia of the drive train is proportional to the square of the transmission ratio, this provides a fair amount of inertia but this has benefits to controlling varying loads as well.

The optical target, an active LED, traces out a section of a circle as it moves and the position is returned at about 100Hz to the tracking device. Thus, a commanded position is sent to the servo and the resultant motions are tracked with high fidelity.

The X,Y,Z positions are fitted to a plane and then these positions are fit to a circle. The angle is now be computed for each position and the relationship between time and angle is then be plotted as shown in Figure 6-4. The response of the servo, shown in the time

response curve, is used to determine the natural frequency, ω , of the servo. This is then fed into the simulation model for each servo.




Figure 6-3: Test setup for determining servo model.




Figure 6-4: The servo, as measured, exhibits a classic underdamped response.

The results of one of the tests is shown in Figure 6-4, provided the step response to an input command to move to a desired reference position. The swept angle shown is about 100 degrees. From this information, several parameters can be found, including stiffness and damping coefficients, and the inertia.

For a second order system (mass, spring, damping) the coefficients were determined by deriving these values from experimental data. See Appendix C: **Derivation of Actuator Parameters** for details. Once the parameters were found from the experimental data, they were plugged into the physical simulation for the actuator values. The derived values are for a spring damper and inertia system and the simulation values determined are the proportional and derivative gain. The behavior of the simulation had the general characteristics of the actuator. This was determined by plotting the response of the simulated snake robot to a given input and observing the angular velocity of the simulated actuator.

Design

Actuation is closely tied to the structural design that supports the robot and I examined and discarded many ideas and iterated a number of configurations to resolve this issue. Mechanisms examined included push-rods, linkages, bellcranks and clevis joints to increase leverage and provide higher torques. The additional complexity of these

mechanisms did not warrant the additional design, fabrication and maintenance that they required. By directly tying actuation to output, the mechanism was simplified and made very compact even though the torque requirements increased. A beneficial cascade effect occurred that shortened and lightened joints, thus reducing structural stresses and loads.

Geometry

In natural snakes, as we saw in **Background**, links are formed from the concatenation of many similar vertebrae. The angular motion between vertebrae is fairly small but, because the vertebrae are relatively short, they can subtend small body curvatures. Thus, the relationship between link length and the angular motion that a joint can subtend is important.

In the right hand side of Figure 6-5, a right angle corridor of equal passage width, W ,




Figure 6-5: The relationship between link aspect ratio and joint angle motion for a right angle corridor.

provides a geometry for determining the relationship between the aspect ratio of the links and the joint angle they subtend. One relationship between the joints and links is that the angle between links in a circular arc must be equal to the arc angle divided by the number of segments in that arc. However, this doesn't reveal the relationship between the joint excursions and the aspect ratio of the links.

In Figure 6-5, R_o is the radius of the outside arc envelope of the link configuration, R_i is the corresponding inside radius and d and L are the respective diameter (or width) and length of the individual links. Finally, θ is the angle between links. In Equation [6-1], the outside radius is shown as a function of L and d and the inside radius. The inside radius, R_i , shown in Equation [6-2], is similarly derived from the figure and finally, in Equation [6-3], the corridor width, W , is shown as a function of the two radii.

$$R_o = \sqrt{\left(\frac{L}{2}\right)^2 + (R_i + d)^2} \quad [6-1]$$

$$R_i = \frac{L}{2 \tan \frac{\theta}{2}} - \frac{d}{2} \quad [6-2]$$

$$W = R_o - \frac{\sqrt{2}}{2} R_i \quad [6-3]$$

Now substituting for R_o and R_i gives W as a function of link length and the angular excursion. If d , the link width or diameter, is set to 1, then the length of the link, L , represents the aspect ratio of the joint length to width. Substitution gives the equation below:

$$W = \frac{1}{2} \sqrt{L^2 + 4 \left(\frac{L}{2 \tan \left(\frac{\theta}{2} \right)} - \frac{1}{2} \right)^2} + \frac{4L}{\tan \left(\frac{\theta}{2} \right)} + \frac{\sqrt{2}}{2} \left(\frac{L}{2 \tan \left(\frac{\theta}{2} \right)} - \frac{1}{2} \right) \quad [6-4]$$

In Figure 6-6, this function is plotted as a function of both θ and L .

The important aspects of Figure 6-6 are that, as might be expected, the arc through which the links can move is linearly related to the length of the link. That is, the longer the link, the broader the arc. However, the angular excursion between links has an interesting property: there are rapidly diminishing returns from increasing the angular excursion of the joints. Beyond 0.3 radians, about 20 degrees, of motion for almost any link length, the corridor through which the connected links can pass does not appreciably decrease in size.

This argues for link designs to be as short as possible and that designing links whose rotation is beyond +/- 20 degrees is a misguided effort. However, there may be other reasons such as deployment and setup that require greater ranges of motion for the links. This is a general argument that shorter links are better than longer links and that large ranges of motion are probably unnecessary to subtend tight geometries.

This result is important, and although the analysis is specific to this geometry, the general lessons are: shorter links with small aspect ratios are better and the angular motion need not be large to gain benefit. More strongly, large range of motions do not appear to benefit movement through tight geometries. This argument does not reflect a particular gait, only a passage of entry and exit.

A related question is whether locomotion through open environments is served better by large angular excursions. The primary benefit of smaller angular motions is that the body of the robot can better describe arcs due to the smaller variation in distances between the links and the nominal arc. For large angular excursions the fit is worse; the maximum distance between link and arc becomes greater.

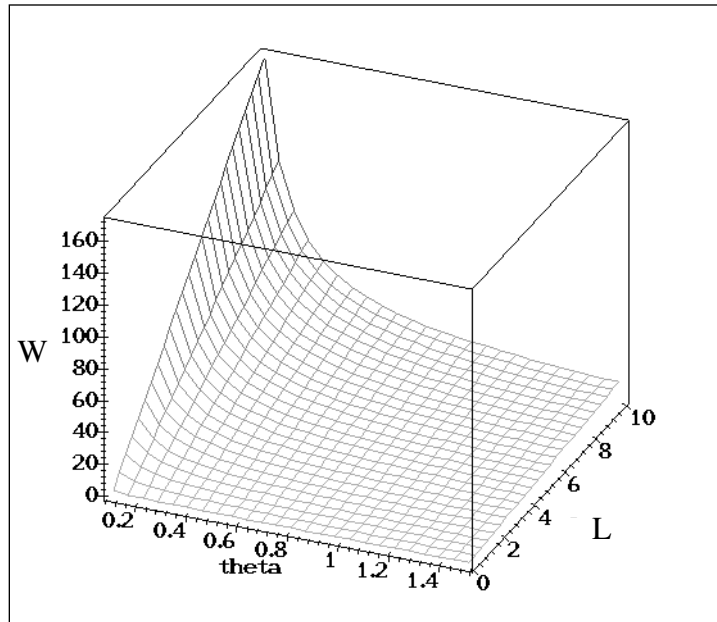


Figure 6-6: Plot of passage width as a function of link aspect ratio and angular excursion between the joints.

Mechanism

The initial link design, constructed of aluminum, used material over 3mm thick. A single link was built and constructed to test assembly, clearances and strength. This link, shown in Figure 6-9, proved the concept and provides two orthogonal motions of up to 180 degrees each. While the preceding analysis showed that the large range of motion is probably unnecessary, the motion came at little cost to the design.

The distance between parallel axes on adjacent links is 100 mm. Between adjacent links, the joint motions are 40mm apart but are 60mm apart between the two motions in a given link. This slight asymmetry doesn't appear to have any adverse effects on performance. A second generation link was designed and built using thinner material; aluminum about 1.6 mm thick with a number of modifications for reducing weight and increasing clearances in the structure. This thinner and lighter structure is the final design. Specific features included a mounting plate that is integral to the servo structure and housing. The bearing capture for the pivoting arm was also tied into this plate. Servo fasteners were used to hold the mounting plate to the servo. This eliminated modifications to the servos and makes for a strong and modular mechanism.

Following the selection of servo actuators, I completed a design of the links and mechanism for a 3D snake. An undergraduate working with me, Anton Staaf, also began the development of a 2D caterpillar system shown in Figure 6-7. The design

utilizes eight links and eight parallel degrees of freedom. The caterpillar is capable of




Figure 6-7: The Caterpillar crawling robot utilizes eight linearly-linked servos.

traveling wave gaits and enabled testing of the experimental setup while the 3D robot was developed.




Figure 6-8: Exploded view of link mechanism.

The 3D snake link design utilizes two orthogonal DOF's each with approximately 170




Figure 6-9: The first generation 3D link is comprised of two orthogonal servos.

degrees of motion limited by the mechanics of the servos. Typical servo excursions are about 90 degrees, but can be commanded to nearly 170 degrees. Figure 6-9 shows an earlier version of the link. The aluminum pieces are over 3mm thick and no weight reduction was performed on the design. The plate upon which the servos are mounted is attached to the servos using the servos own case mounting screws. This provided great simplification of attachment and a solid and direct mounting. The mounting plates on the opposing side of the servo horn has a threaded hole for mounting a shoulder screw. This attaches the rotating section to the servo very securely and takes up moment loads from the adjacent links of the mechanism.




Figure 6-10: Link mechanism design

The complete 3D snake, shown in Figure 6-11, has ten links for a total of twenty DOFs.



Figure 6-11: 3D snake utilizes 10 links with 20 servos

Improvements in the design ranged from significant reduction in the wall thickness, a corresponding reduction in the weight of the material and improved fabrication using bending instead of machining. Bearing capture and support was changed and simplified from the earlier versions and test designs.

3D Robot Specifications

Here are several general specifications for the serpentine robot:

- Mass: Mechanism is 1.32 kg. About 2/3 of the mass is the servos. The rest is metal and hardware. See **Link Weight Distribution** in Appendix B for more detail. Total mass, including wiring and controllers is 1.48kg.
- Length: 102 cm (10 links, each is 1.02cm long)
- Diameter: 6.5 cm
- Power: 24.2W max total mechanical output. ~75W max total electrical input.
Quiescent: 1.15W (9.5mA @6VDC)

The figures result in an overall robot density of about 0.39g/cm^3 , less than water, due to the spaces in the rotating sections between joints and thin sections of material. This does not include final wiring or the skin which bring this figure to about 0.5g/cm^3 .

Skin

The vehicle-terrain interface for wheeled and legged vehicles has been the topic of many research works, but is neglected for serpentine robots. Wheels, feet and ankle designs and tire tread and forms have been extensively examined in the literature. However, previous serpentine robot research has not evaluated skins and surfaces and surprisingly little attention has been paid to this area. The integument, or skin, provides terrain and environment contact and is the key component of interaction between the mechanism and the terrain.

Recall from **Biological Systems** that the skin of the biological snake is comprised of smooth, dry, highly polished overlapping scales. The underlying skin is elastic, like ours, and accommodates the varied shapes and motions of the body. To provide similar

properties for a robot, I examined, evaluated and tested a wide variety of materials for the serpentine skin.

Another effect which can be used is that of preferential friction, where the coefficient of friction varies depending on the direction. This makes it easier for a variety of gaits to demonstrate progress. The extreme case, of course, is a one-way bearing where the friction is negligible in one direction and free-rolling in the other. Materials with a nap, like velvet, show this quality and some specialty materials for material transfer exhibit this property as well.

Bellows

Bellows material is often used to protect exposed surfaces of machinery. These are usually made of heavy cloth materials or segmented plastic or metal frames. The pleating of the material provides a corrugation that gives form to the bellows as well as allowing the compression of the bellows itself. This form then, allows discrete line or point contact with the ground and a way must also be provided for the bellows to attach to the underlying mechanism.

Cable Chains

Cable chains are used for the protection of wire harnesses during the movements of machine tools. They are usually made of tough plastic materials, but are also available in metal links. Only a few cable chains allow motion out of the plane, the igus Triflex[®] series being one of them. [igus95]. This is an attractive possibility for several reasons, it provides the rotation axes, it is of the right scale for the mechanism and, through the use of sliding surfaces, provides some degree of protection from the environment. The downside is the weight of the links and a simple means to couple the actuation to the plastic frame. In addition, the angular excursion at each link is limited to about 10 degrees. Weight is also an issue; for a 1 meter length the weight of the chain alone is nearly 1.5kg for the square 50mm interior model.




Figure 6-12: igus Triflex 3D cable chain.

Flexible ducts

Ducts are often comprised of steel springs embedded within latex or vinyl cover to provide a conduit for air or wires. Dryer duct hose, commonly found in hardware stores is one example. The compound material is flexible and provides a high ratio of extension to compression. However, the rigidly rounded profile requires a means of tying the hose to the structure and does not conform well to underlying mechanism and

structure. Plastic corrugated materials such as vacuum hose were also examined but found to be too stiff.

Rubber

Natural and man-made latex materials exhibit substantial stretch and are available in very thin films. They are highly elastic and are very thin in a wide variety of forms. However, even very thin materials take a fair amount of energy to distort and stretch and prevention of wrinkling and folds is a difficult challenge.

Fabrics

A number of fabric materials were tested including spandex materials, which provide high elasticity and durability. Spandex is unusual in that it is a fiber that acts like an elastomer. Spandex materials can be stretched over 500% without breaking and completely recover. The material is a polyurea-urethane elastomer chain that has soft segments that provide elasticity and rigid segments in the chain that act like the crosslinks in natural rubber.

Spandex is stronger and more durable than rubber and also resists pilling, the build-up and consolidation of short fiber segments on the surface of the material. Compared to other fibers though, it has poor strength and so Spandex is usually blended with other fibers including polyester, nylon and cotton. Typically percentages of spandex are less than 20% in blended materials. Commercial spandex fibers include Lycra[®] by Dupont, Cleerspan[®] and Glospan[®] by Globe and Dorlastan[®] by Bayer.

Other, more exotic fabrics, tested included four-way stretch velvet materials, (92% Polyester, 8% Lycra) which have the interesting property of differential friction, commonly termed the 'nap' of the material. Another set of surfaces tested included sequined materials with 100% overlapping coverage on a Lycra-base. These are particularly gaudy materials but provided dry and polished lapped surfaces similar to scales.

Textile qualities are not only dependent on the material that they are made of but also on the fiber treatment, processing and the fabrication or weaving technique. An ideal material might also be impervious to water and, in fact, some new woven materials using microfiber (<1.0 dpf) are water resistant.

Braided materials

Another material tested, and eventually used, are polyethylene-based braids. These are typically used to bundle and protect wires, cables and hoses. They offer good abrasion protection and are very lightweight and durable. A wide variety of materials and types are available including teflon, halar, copper and steel. Additionally the weave and mesh can vary. I selected a sleeve made from a 0.25 millimeter polyethylene monofilament yarn of polyethylene terephthlate or PET. A variety of diameters were tried to evaluate fit and motion. A 3.8cm nominal diameter braid was finally selected and stretched over the 6.5 cm diameter mechanism

Final Skin

A Lycra spandex sleeve with a single dorsal seam is slid over a PET braid to form a two layer skin. The seam utilizes an elastic thread and cross-stitch to minimize loss of stretch across the seam. Figure 6-13 shows the 3D snake atop several of the skins made

for testing. All the underlying fabric material is identical with the exception of color but the surfaces are different. The surfaces range from no treatments, to small studded hard materials to sequins.




Figure 6-13: A variety of skins made using spandex fibers as the base fabric with a variety of surface treatments.

Tread

While the fiber materials provide excellent stretch characteristics and form accommodation, they also result in a smooth, low friction surface. This, in turn, can result in slipping and is an obvious problem for traction.

Surface treatments such as polymeric paints were applied in a wide variety of patterns to fabric surfaces to experiment with tread configurations. These paints or coatings are especially designed for fabric applications. In some cases, after the material was partially cured, steam is applied to the surface. The high heat and humidity combine to raise the surface of the applied coatings above the base material.

Many of these coating materials were applied by squeezing drops of the materials across the material in different patterns. Coefficients of friction were compared between the surfaces and patterns. This was done by wrapping the material around a wooden block and then laying it atop a surface and determining at what angle the material would slide at a constant velocity. The base material used was aluminum. The tangent of this angle is the coefficient of friction. For the materials shown in Figure 6-13, the coefficients of friction are shown in Table 6-2.

Spandex Material	μ_k
Patterned w/ rhinestones	0.48
Plain Spandex	0.32
Red Sequined	0.30
Larger Rhinestone patterns	0.36
Diffractive material	0.42
Polyethylene braid	0.26

Table 6-2: Skin materials and corresponding coefficients of friction.

Electronics

Electronics provide and distribute information and power to the robot. Due to the large number of actuators in this type of robot there can be a correspondingly large number of conductors carrying signals and power. However, the actuation chosen for the robot reduced some of this infrastructure because the servos provide local closed loop control of position. For the many controlled degrees of freedom in robots there are numerous feedback signals, motor winding and commutation signals etcetera. Servos, however, require only power, ground and the reference signal. Even so, this can be a difficult number; every additional conductor is multiplied by twenty.

Ideally, both power and data would be supplied over just two wires; one for ground and the other providing power with a superimposed electronic signal. The concept of a bus that multiplexes both power and data has been in use in homes and industry for a number of years: it superimposes a digital signal atop the power lines in the home and allows the control of lights, appliances, etc. from a single or multiple locations [X10 97]. The more recent IEEE 1394, Firewire, standard has a similar capability. This bus concept would be ideal for a snake-like robots. Every link and degree of freedom requires both power for motion and signal for control and the structure of serpentine robots facilities such mundane issues as routing and termination.

Digital Command Control

One such bus meant for small scale devices is Digital Command and Control, or DCC. DCC was initially designed for scale railroad control and utilizes the track for transmitting power to all devices and onboard electronics to listen to the superimposed signal. Each device is individually addressed and when a particular address is signaled the following data is directed to that particular device. The small boards provide outputs that can drive small DC brushed motors at up to an amp or so. When used with a small and efficient geared motor DCC can provide simple control of a large number of motors [DCC 94]. I built a design on a PC board with a serial interface and used a dual-H-bridge driver to provide power with a superimposed signal. Combined with control software, this provides motion control that is easy to implement and use.

The advantages include the ability to provide both power and signal information over the same two wires to each electronic device. Thus, this 'snake bus' is enormously simplified over running power and information to each actuator. The main disadvantage to DCC is that it is open loop and the commands control acceleration and velocity but not position. There are some proposed future enhancements to DCC to provide these features but no commercial versions exist at this time. I designed another bus system using RS-485, a multi-drop serial bus, in conjunction with ICs that can provide simple interfaces for A/D converters but the size, complexity and cost became prohibitive. The eventual selection of R/C servos as actuators eliminated many of these issues with feedback and control.

Servo Controller

While the input signals to the servos are logic level signals that can be provided by any digital I/O board and software routines, I chose a convenient controller board that provides serial control of up to eight servos per board [Scott 96]. The boards can be daisy-chained to provide control of up to 256 servos. The communication format is a three byte sequence for talking to any particular servo and is as follows: <255> <servo> <position> where the position is the total excursion divided into 256 possible settings. This provides servo motion and control from one extreme to the other. The servo is responsible for maintaining that position through its own feedback and control electronics. For twenty servos, the initial serpentine configuration, three boards are required. The board size is approximately 35mm by 41mm.

In wiring the robot, the power is easily parallelized and a two wire power bus was integrated into the mechanism. The signal wires also needed to be connected and there are two possibilities for integrating them. One is to put the servo controllers within the snake at intervals that minimizes the overall length of the wiring. For initial testing however, I created a 'tail' that appends the boards to the end of the robot. The tail provides a stacked arrangement of the controller boards and a tightly integrated harness to connect power and signals to all three boards.

Wiring

Since the signal wires are not multiplexed along the length of the snake it's necessary to route them along the mechanism to both minimize bulk and maximize flexibility. These constraints suggested very small and finely stranded wires. After investigating a variety of small cables including ribbon cables, small gauge wire and specially manufactured cabling, I selected a fine gauge hearing aid wire [Siemens 97]. This fine insulated wire uses seven strands, each about 0.05mm and the total cable diameter, including insulation, is less than 0.4 mm. A combined bundle of the required 20 wires is only about 2mm making the combined bundle very manageable.

A rule of thumb for small wires is that the bend radius should be no smaller than ten times the wire diameter. Thus, the bend radius in this case should be greater than or equal to about 4mm. The solution is to route the wires directly across the points of rotation at each axis minimized bending of the wire harness.

In early testing, there was significant jitter throughout the robot and a network of decoupling capacitors was devised between power and ground. However, the addition

of other return lines was sufficient to remove most of the noisy signals and crosstalk in the wires.

Sensing

A key element of biological snakes is sensing. It enables rapid adaptation to varying terrains during locomotion. Since the terrain is unknown, such sensing is necessary for traversal. If, magically, the terrain were known and both the position and configuration of the robot were also known it would be simple to provide appropriate information to guide the robot during locomotion. A minimum form of sensing is to simply provide contact sensing to determine whether or not contact has occurred. The best form, of course, is full contact sensing with knowledge of terrain and forces.

Implementation for contact and force sensing is difficult. Proximity and contact devices for industrial use are bulky and difficult to integrate into small mechanisms. The technologies exist and include optical, capacitive, force sensing resistors, and even piezo-electric pads. The difficulty is not just identifying the technology, but integrating it into the robot.

However, I identified and evaluated several candidates:

- Small tactile switches - binary, with trigger threshold
- Force sensing resistors - analog, wide range
- Capacitive arrays - flexible, good resolution, good sensitivity.




Figure 6-14: Two resistive force sensors and a small tactile switch that were evaluated.

The switches are the simplest, in terms of operation, interface and mechanism. The drawback is that they only provide contact information above a force threshold required to trip the switch. A variety of force sensing resistors were tested but for most of these devices there are significant issues with curved surfaces. The stress of forming a shape other than on a flat surface makes the response of these devices unusable for this application. One device, however, a thick film polymer with force resistive properties works well for measuring forces on a surface that is curved along one axis. This is the middle device shown in Figure 6-14. Another device uses piezoelectric-generated

acoustic signals to find deformation in a soft pad. While the resolution is high the cabling, packaging and cost make this type of sensor untenable for this application.

The final device appears to have the desired sensing properties for tactile sensing, high resolution, fast response, accurate measurements, good spatial resolution, curved surface use. The drawbacks are cost and the electronics packaging and infrastructure [Novel 97].

Additional sensing such as local range information could provide useful data for locomotion, obstacle avoidance and grasping. Small IR sensors could be used to detect local surfaces without contact. This offers other opportunities for gait selection and modification but was not explored further in this work. The sensing devices were evaluated but not emplaced on the snake robot due to cost and time considerations.

Other Subsystems

A robot is a substantial integration of several technologies. Not only mechanism, actuation and sensing but communications, power and computing. Each of these subsystems cannot be fully independent of the others and considerations for each must be taken into account during the design process.

Communications

Communications is handled by a serial-based RS-232 device. At 9600 baud, using the three byte command stream, meant that each servo could be updated about 16 times per second. This rate results from 10 bits/byte, 3 bytes/command and 20 servos. The diagram for the electronics is shown in Figure 6-15. Each link, shown as the small boxes with numbers representing the 20 servos, are connected via lines to one of the three servo controllers, C_0 - C_2 . These, in turn, are daisy-chained to a serial line connected to the computer.

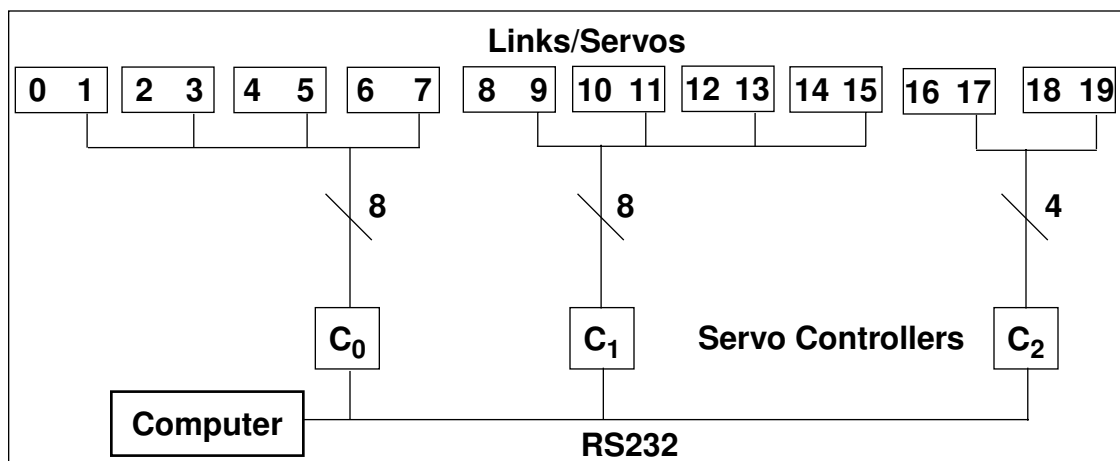


Figure 6-15: The electrical system utilizes three servo controllers with a serial connection

Computing

Since the learning experiments required high-performance platforms and the communications could be either to the simulation or the robot it made sense to use the

same platform for simulation and device control. Most of the development and testing is done on Silicon Graphics (SGI) workstations.

For these implementations, programs are written in C++ and compiled for execution on SGI platforms. However, even the high speed workstations, such as the 195MHz R10000 computers, cannot run the simulation in anything close to real-time. The framework is executed on a single processor machine for most of this research although it can be partitioned across multiple machines. In fact, optimization and learning can parcel out the tasks across multiple machines so that evaluation can be parallelized significantly. For most testing however, single machine execution is sufficient and runs take several hours or so on R5000-based computers. All display code for the 2D systems is written in C++ using OpenGL calls and an Xforms interface. For the 3D simulations, Inventor is used for display.

Power

The servos are typically powered by 4.8V batteries. However, many servos can be powered at 6V or even 7.2V with corresponding increases in power but perhaps reduced operating lifetimes due to brush arcing on the commutators. For most testing, a DC switching power supply is used but a variety of battery technologies were also investigated during the course of this work [Dowling 97a].

Key attributes of any power system are power and energy density. Energy for long term operation and power density for high demand periods. Even with the advent of many new technologies, Nickel-Cadmium (Ni-Cd) batteries offer close to the best power density of any battery technology. They are also widely available, and are available in a large number of configurations and at reasonably low cost.

Although battery testing and selection occurred, most operation was done with an offboard power supply during testing and evaluation.

Video

A small CCD camera was also added to the snake to provide video feedback for operators of the device. The device, shown in Figure 6-16, provides, a snake's eye view of the area directly in front of the robot. Because the robot is highly articulated, the

forward joints can act as positioners for the camera for pan and tilting the image. The




Figure 6-16: A camera, shown by the arrow, provides a 'snake's eye' view.

robot is tethered for power and data, as is the video. For an eventual self-contained robot, a self-contained robot could use one of several micro-miniature transmitters for wireless video transmission.

Experimental Setup

As shown in Figure 6-17, a means of measuring and calculating the metric was

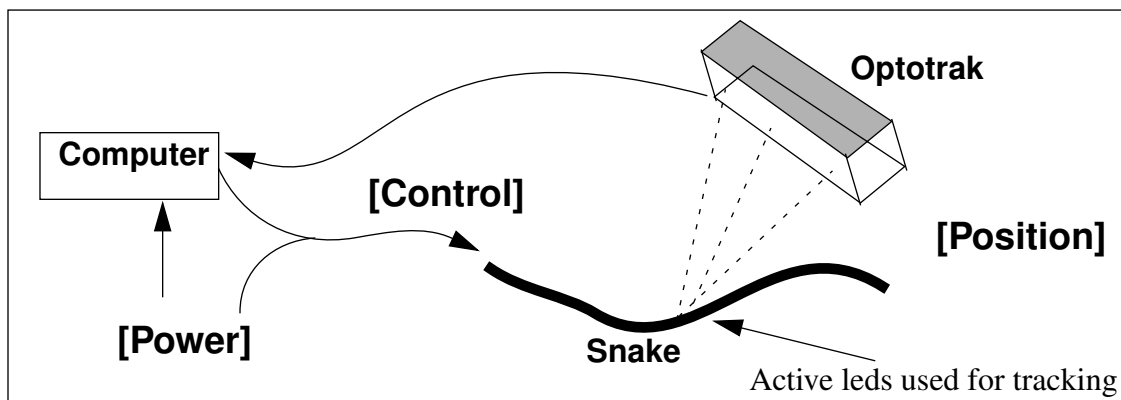


Figure 6-17: Experimental setup of the robot and tracking and power monitoring.

accomplished through power monitoring and a tracking device. The tracking device, an Optotrak device, provides data rates of up to 1000Hz and tracking accuracies to 0.1mm. It utilizes 3 very accurate PID chips to view multiple IR LED's each of whose position can be tracked. Ideally, the position data would be for the center of mass of the robot, but this would require tracking LEDs at each link and wires to connect them to the strobers. A single tracking LED was used near the middle of the robot to provide travel distance over the experimental time. Although a single LED does not provide the true center of mass, it gives a reasonable approximation to movement of the robot.

Physical Modeling

Physical modeling programs are recent developments and only one commercial package is available at the time of this writing [Knowledge 97]. Coriolis is a toolkit developed at CMU to support interactive simulation. The toolkit is implemented as a set of C++ classes that are instantiated to represent relationships and forces in the simulated world. Class types within Coriolis include `Bodies`, `Constraints` and `Influences`. `Bodies` are used to represent physical objects and their geometries. `Constraints` describe relationships between different bodies and, finally, `Influence` classes are used to describe the forces that act upon `Bodies` [Baraff 97]. Coriolis does not provide graphics, I/O formats or an interface and these must be provided by the user.

2D

My implementation for the 2D simulation uses OpenGL for the graphics and a variety of input and output methods including command-line arguments, files, and direct user interaction with keyboard and mouse. Finally, I wrote an Xforms user interface on top of the simulation which provided an OpenGL window, interaction and more importantly, a variety of monitoring tools for observation of the simulation in action. The Coriolis window in Figure 3-1 utilizes a number of these tools including power usage, distance traveled by the simulated robot, elapsed time, and real-time geometric information such as joint angles. I built in a number of options, including the ability to turn off graphics for improved performance. Other interaction includes window panning and scrolling capabilities and direct interaction with the simulation time step functions. I also enabled user interaction that allows the user to click and drag simulation pieces around to facilitate experimental setup or quick retesting without having to quit and restart the simulation.

An example of the toolkit and how it is used to create the 2D snake within the simulation is shown below:

```
void create_snake(Material *mat, UniformField *grav)
{
    for (all_links);
    AtomicShape *link_poly = makelink(geometry)
    links.set_geometry(*link_poly);
    links.set_material(*mat);
    links.compute_inertials();
    grav->add(links);
}

void SnakeMotor::update()
{
    double now = bs->current_time;
    torque = magnitude * sin(now*freq + link*phase);
    vector j(3, 0.0, 0.0, torque);
    exert_force(*body1, j);
    exert_force(*body2, -j);
}
```

`create_snake` takes as one of its arguments the material type which includes density, coefficient of friction, coefficient of restitution (bounce), color etc. Additionally the function describes the forces acting on the snake. This is set to a uniform force field in one direction, namely gravity. Then for all the links, the geometry is created, the geometry and materials are instantiated, inertial properties are calculated and the gravity property is added to all bodies.

`SnakeMotor` is a torque applied to each body segment that, in this example, describes a simple time-varying torque function for each of the joints. The torque is applied to each body and the negative to the adjacent body segment. In this example, the form of the torque function is a sinusoidal wave whose amplitude, frequency and phase are adjusted. These values can be altered for each execution of the program and the metric can be tested for efficacy. A different representation is required for more realistic and expeditious locomotion.

3D

Coriolis, augmented with supplementary classes, is used to define 3D body geometries and their connections [Leger 97]. Parts are created, connectors defined and pieces are assembled into more complex geometries. For position control of the bodies however, control is determined by setting a number of parameters, including gains, for position control. This is because, in Coriolis, all motion is effected by `Influences` and not simply by position commands. Open Inventor is used for all display functions. Figure 3-2 shows the 3D model of the physical snake structure. The model is created by individually creating the vertex and face list for each part and by specifying the relative geometries and assemblies.

The physical model of the serpentine robot is shown in Figure 3-2. The model is generated from the mechanism design and used in the physical simulation.

Summary

A robot is a complex electro-mechanical integration of many technologies and decisions on configuration affect software, planning and control. Issues as mundane as packaging and wiring can slow and arrest development unless carefully addressed. The adverse and cascading effects of improperly chosen subsystems can stop research in its tracks.

Because of the central nature of actuation (it affects mechanism, control and a host of other issues) this technology was carefully investigated. The technology chosen, small, well-packaged electro-magnetic motors and spur gear drive trains, was the most mature of the alternatives. Since the focus of this research is not the development of actuator technologies, this was a good decision for implementation. Eventually, other nascent technologies such as electrostrictive polymers, will be viable for small mechanisms but at this time the immaturity of these alternative technologies makes them unsuited for this research.

The efforts in design include a geometric analysis relating link length, diameter and angular range of motion. The result is that short links are better and that, interestingly, the range of angular motion need not be large to be effective.

The mechanism design proceeded through several iterations to simplify connection and support and reduce weight and complexity of the design. The final result, using lightweight formed plates and a simple bearing support provides good capability and reduces overhead in fabrication and assembly.

Modeling of the robot is accomplished through the use of a powerful tool kit that models a large range of physical phenomena and provides for both control and monitoring of these forces.

Electronics design proceeded from several tested concepts to a separate power bus and signal control system that is implemented in a straightforward manner. Additional work in power and sensing have provided insight into future developments. Finally I examined skins, a first for snake robots, for the purposes of providing a compliant and tractive surface.

Although design is an appreciable effort, very often testing takes the greatest amount of time and mundane issues such as connectors and cabling conspire to thwart even the best intentions in design and fabrication.

Chapter 7

Locomotion

Locomotion presents details of experimental work; calculation, evaluation and demonstration of the snake in both simulation and on the physical system. This culminates in the integration of all the elements of the framework: optimization, evaluation, simulation, and the robot. This is prefaced by a detailed look at how values for the specific resistance metric are calculated or measured.

A number of interesting gaits were developed and, surprisingly, many of these were non-snake-like. The gaits in simulation were able to provide a good measure of performance relative to specific resistance. It was also found that the ideal path from simulation coupled to learning and then learning coupled with the robot did not provide the ideal path for gait development; the ideal serial process became an iterative one.

Power Calculations

Since the metric for evaluation, specific resistance, uses a value of power in its formulation, that measure is required for both physical simulation and the robot.

For an electrically driven mechanism, the measurement of power can be determined by simply monitoring current, voltage and power factor if there are large inductive loads. The product of current and voltage is the input power to the robot, but it can be difficult to determine where the dissipative losses occur and what fraction of that power measurement goes into motion.

In simulation, the calculation of total power in the robot is the sum of the products of torques and angular velocities for all the links of the serpentine robot. In Equation [7-1] m is the number of links or motions.

$$\text{Total Power} = \sum_{n=1}^m |T_n \dot{\theta}_n| \quad [7-1]$$

Assuming a power limited source, which a robot has, this value provides both an indication of power levels and when cutoff thresholds are reached. Average power is determined by dividing the total time into the total power used. Peak power, obviously, is the maximum power reading.

However, an issue with Equation [7-1] is determining negative work where the torque applied is in the opposite direction of the motion. For a walking machine this is a critical issue and one that has adverse consequences on vehicle efficiency. An example of negative work is when an actuator is used as a brake. Power is expended, but not in a manner that enables forward motion. This issue is germane in legged robots where improperly designed configurations result in large expenditures of energy without corresponding forward progress.

Interestingly, in animals, negative work occurs when muscles are stretched while they are still developing tension [Roberts97]. This extra tension has little extra energy cost and this, in turn, has strong implications to locomotion: people can descend stairs with far less effort than it takes to ascend even though the muscles still exert considerable force. Thus, negative work requires less total energy than positive work in animals.

Geometric work in mechanisms is a form of negative work where actuator backdriving occurs opposite to the direction of motion of the vehicle. It's inefficient because other actuators have to make up backdriven energy losses in addition to the energy required to move forward [Waldron 81].

Gravity loads, while probably not significant for a snake, do contribute additional losses every time part of the robot is lifted and placed down on the ground. Gravity loads do affect contact forces and sliding friction however. For forms of sliding locomotion the additional frictional losses can be significant. It is amazing that snakes are at all efficient.

One solution is to take positive and negative work separately to assess the contributions from both types of work. The net energy input, shown in Equation [7-2], should be equal to the friction losses in the system.

Geometric work can be minimized by decoupling propulsion from support or lifting motions. In the case of serpentine robots, however, this may prove difficult to do. The opposing issues are generating sufficient traction to maximize forward progress and enabling low friction to minimize losses into the ground. One thought is to provide a skin that provides differential friction and this topic is explored in a later section on skin design.

Other significant losses can occur from isometric work where opposing forces are generated through the ground. This coupling typically results in force generation and energy losses without motion. Without good coordination this can prove to be a significant factor in energy losses. Legged machine designers must consider issues in isometric work, but designers must also address this issue for wheeled vehicles that are highly articulated, or that require multiple degree-of-freedom of control.

The most straightforward analysis of power usage is the sum of the products of torque and angular velocity for all links as shown below.

$$\text{Power input} = \sum_{n=1}^{\text{links}} T_n \dot{\theta}_n \quad [7-2]$$

This gives an instantaneous measure of power use by the robot. An average power estimate over a fixed period of time can then be calculated to give a measure of total energy used. The power for a particular joint is found over a specified time interval by assuming a fixed level during the interval. If the time intervals are short then this will generally hold true. Otherwise, the calculation is more complex and, in any case, does not reveal great difference in the calculations which are used to only determine differences between evaluations, not highly accurate values. Thus, the power calculation for a joint becomes the sum of the power during the all the intervals. Average power is determined by dividing by the number of intervals and the power for the whole robot in simulation becomes the sum of these products over all the joints.




Figure 7-1: Power is measured at discrete intervals by the product of torque applied at the joint and the angular velocity.

As shown in Equation [7-3], the total power is then summed over, n , the number of time intervals and m , the numbers of joints. The angular velocity is actually the difference between the angular velocity of the two links. As shown in Figure 7-1, this is $\omega_1 - \omega_2$. Thus, the equation for total power over the time interval becomes the formulation shown in Equation [7-3].

$$\text{Power} = \sum_{j=1}^m \left(\frac{1}{n} \sum_{i=1}^n |T_{ij} \dot{\theta}_{ij}| \right) \quad [7-3]$$

In Figure 7-2 are shown two plots of power for a robot run. The snake was a small two segment, four DOF system. The data was captured at 100Hz for thirty seconds during locomotion. Notice the power spikes, mostly negative, in the top figure. These spikes result from the surface impact when a joint section makes contact with the surface. A short rebound results in motion that is opposite that of actuation and this happens very quickly. Thus, the product of the current actuator force and the rapid change in the velocity give rise to a spike or discontinuity in the data. In the bottom graph in Figure 7-2 is shown the result of a running average (20 data points wide). The power average remains about the same in both cases, about 85mW. The power values from simulation are much smaller than that of the actual robot. This is primarily due to the friction and

inefficiencies of the robot drivetrain. Modeling all of these physical phenomena within the drivetrain is impractical, but the general results should be valid. That is, a good gait in simulation should correspond to a good gait on the physical robot.




Figure 7-2: Measured power data from run of physical simulator. Top graph is all data with impact spikes and bottom uses running average for smoothing.

Units

In physical simulation, the consistent use of units is critical for relevant modeling of the snake robot. Consistent use of physical definitions is one such requirement; units of mass, length and time can be used to define a complete physical system. For practical reasons within the simulation however, the values of these units must allow good calculations of other physical concepts such as inertias and densities. Thus, the physical units used in simulation are kg for mass, meters for length and seconds for time; an MKS system. For defining snake geometries, and for simulation purposes, I used a CGS system. This may appear to be a simple matter of scaling but at issue are the final values within the simulation that are used to determine forces, frictions, motions, moments etcetera. These values should not be small because calculation errors will creep and eventually result in erratic performance. Thus, MKS units result in very small numbers for the relatively small mechanism where links are at cm scales and masses are at gram scales.

Velocity Calculations

Another piece of the specific resistance evaluation is the velocity metric. In simulation, this is given by tracking the distance traveled by the system center of mass. This is calculated by the square root of the XY distance traveled by the center of mass. The center of mass of the whole system is determined from the average of the X and Y positions of the centers of the individual links.

For the robot, the center of mass is a little trickier. Because tracking all of the individual links is difficult, one way is an approximation using the distance traveled by one of the central links. This can be tracked using a camera-based tracking system that provides rapid and accurate feedback. The central link tracking gives only an indication of motion. However, it appears that the maximum distance that the center of mass can be from the center link is one-quarter of the length of the snake robot; this is a pathological configuration. No proof is presented but the result should be obvious. Most configurations, especially symmetric ones, should result in small differences between center of mass and center link position.

Gaits

Several types of gaits were generated including some non-biological modes of locomotion. In many cases, the gaits could also be reduced to a simpler, more general waveform which provided good insight into how the gaits work. The gaits are detailed here into snake-like and non-snake modes of locomotion.

Snake-like Locomotion

The next several pages reveal figures of various stages of locomotion modes. The first few are those snake-like modes demonstrated in simulation. These are interesting because they replicate existing modes in snakes.

Sidewinding

Sidewinding is a relatively efficient mode of locomotion with little sliding ground contact but with an odd means of moving laterally. Sidewinding is really two waves, one ventral and one lateral that are out of phase. Together they produce a motion that moves a body section to the side and the rest moves along and settles the body down to form successive parallel tracks as it moves.

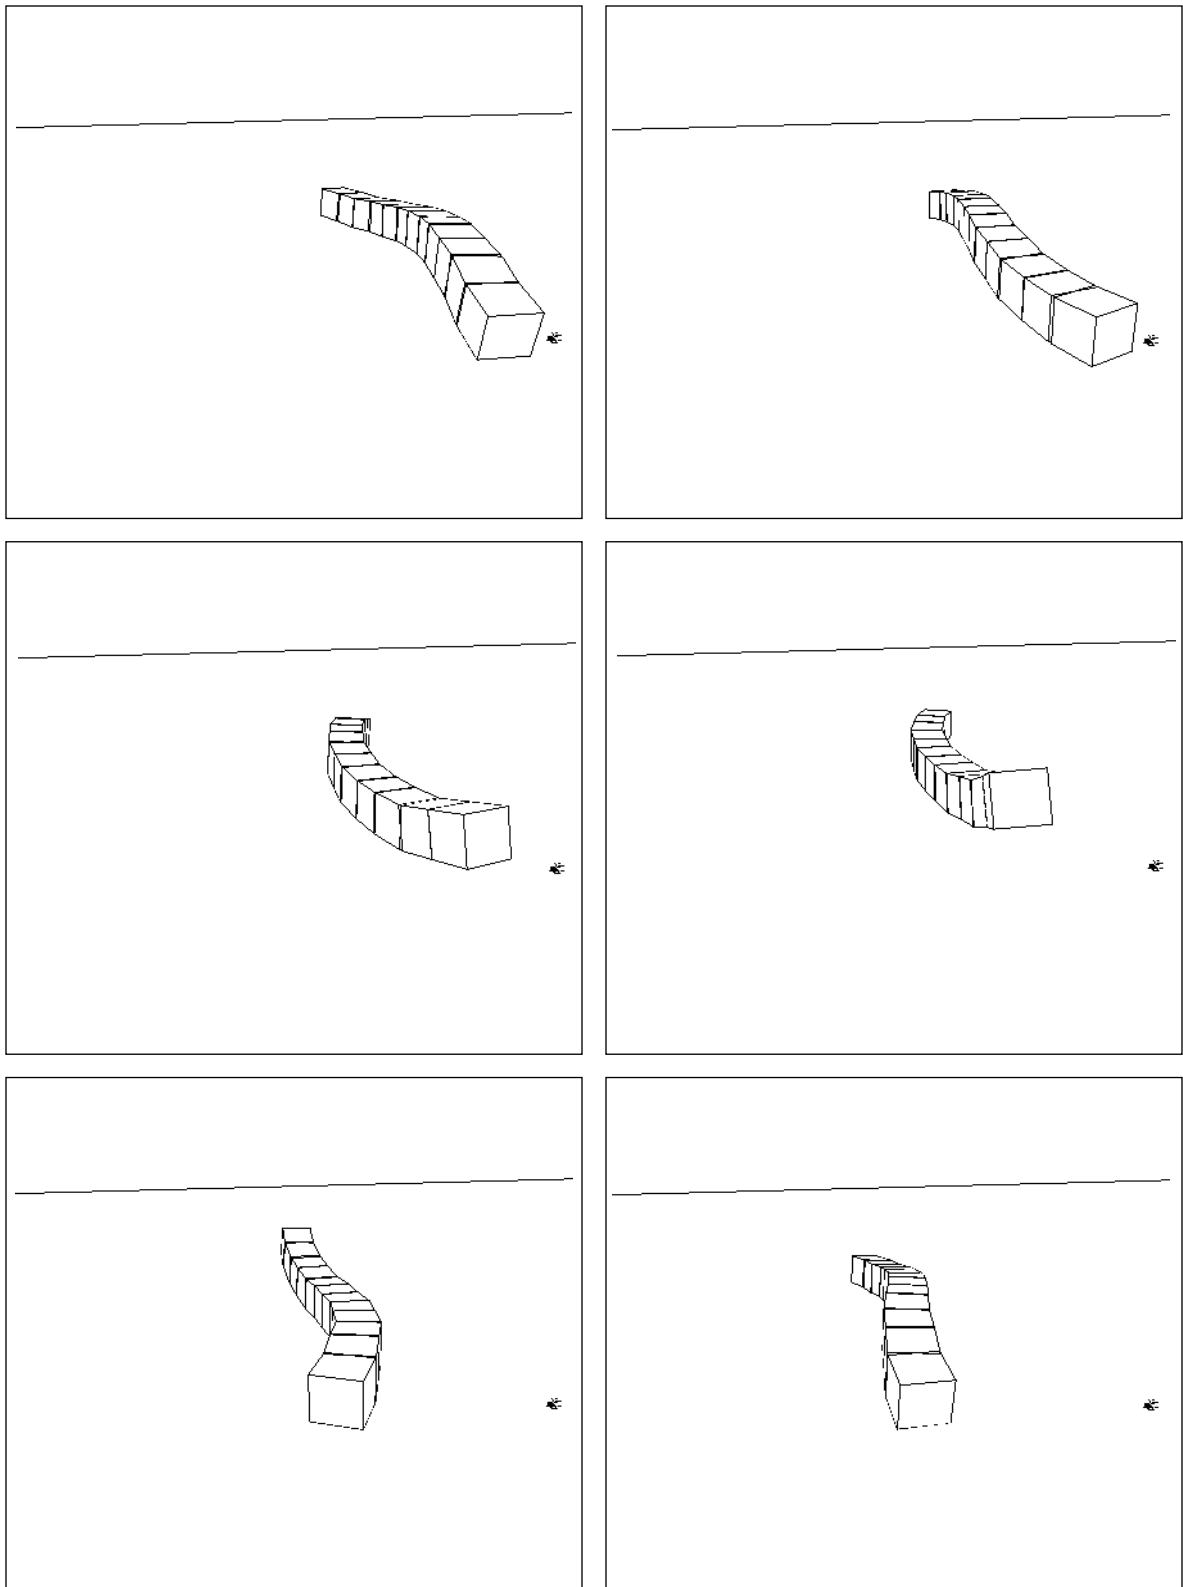


Figure 7-3: Sidewinding locomotion.

Rectilinear

Rectilinear locomotion propagates a wave along the length of the body. This reverse moving wave provides for locomotion by lifting a portion of the body and using length in the wave to move the head forward and down. An effective gait that does not slip or slide much along the ground.

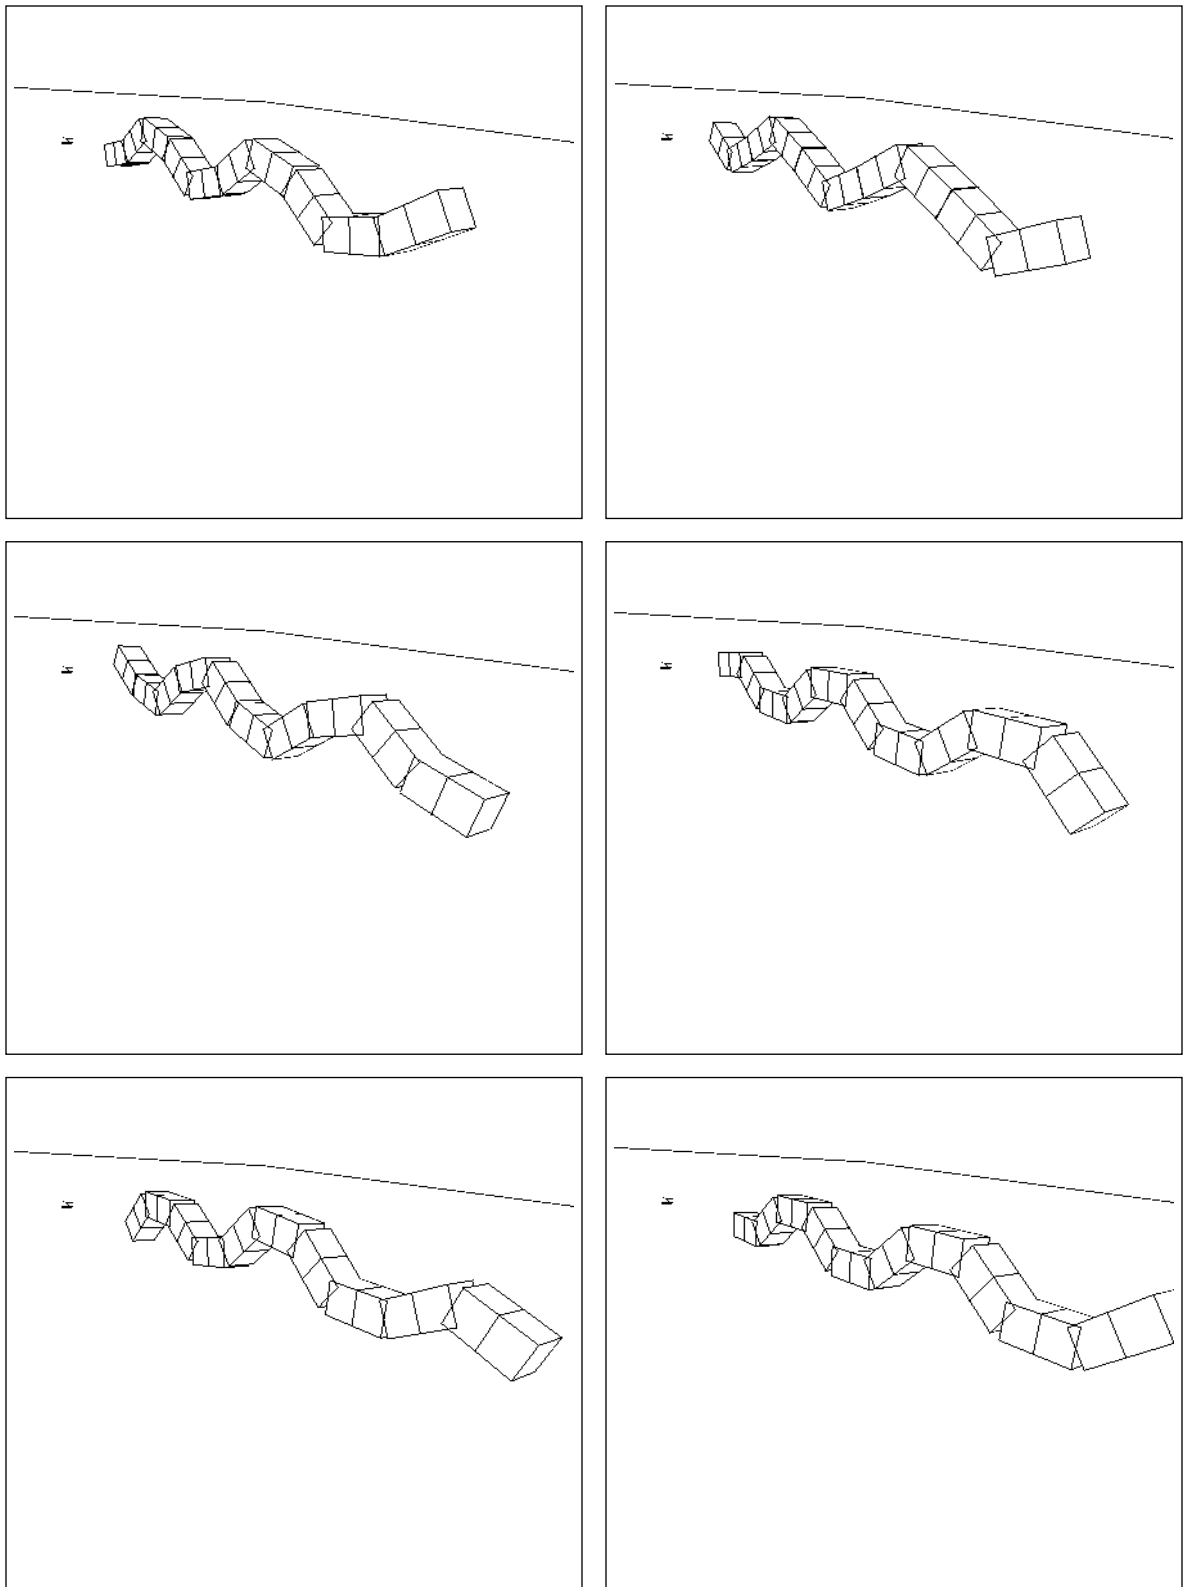


Figure 7-4: Rectilinear locomotion.

Lateral Undulation

True lateral undulation provides for a continuous sliding motion along the ground. The issue in the simulation is that the surface the robot moves along is flat; the ground plane could be populated with an array of small objects that the system could push against.

This locomotion configuration also slightly lifts the outward lateral wave. This is the the 'sinus lifting' mentioned by Hirose.

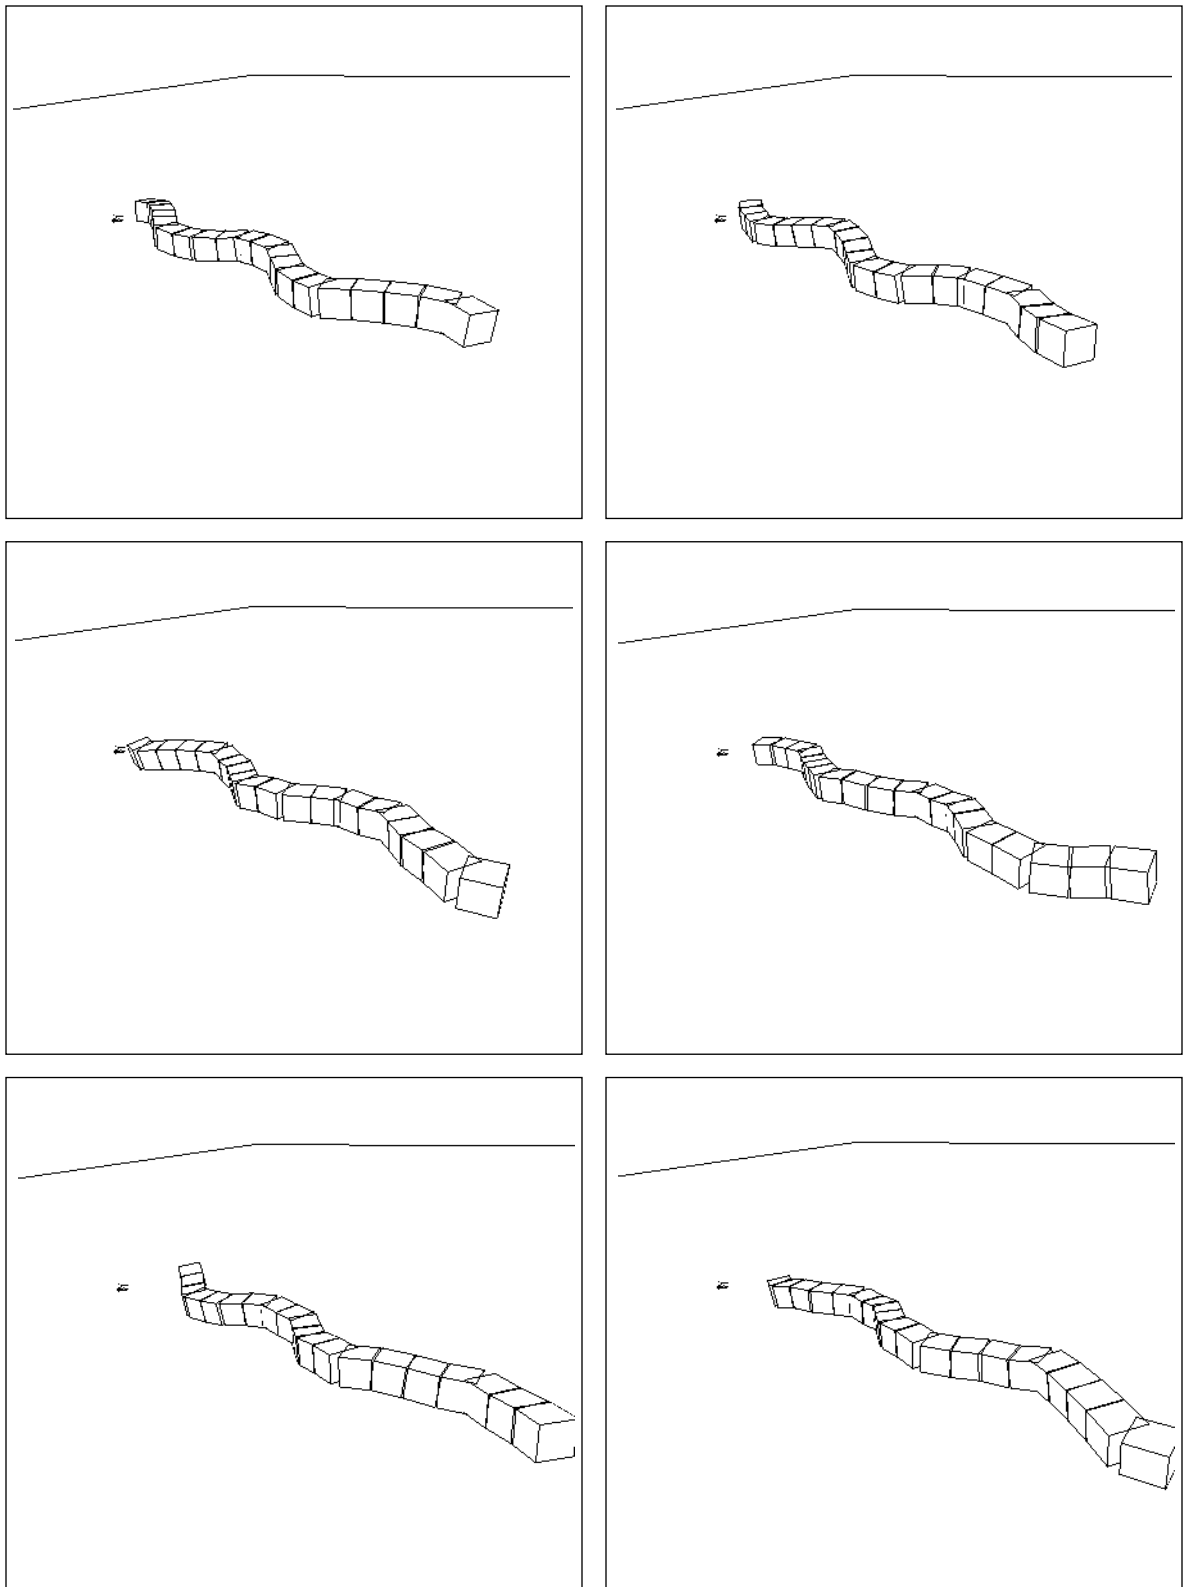


Figure 7-5: Lateral undulation.

Non-Snake Gaits

Lateral rolling

An intriguing gait is formed by a U-shaped body and providing oscillating motions about each joint. It appears, at first glance, that there is a continuous rotation of the joints but this, of course, is not possible and is unnecessary. The gait is similar to sidewinding in that it uses two waves out of phase: a lateral sine wave and a ventral cosine wave. However, in this case the phase of the waves is zero; all joints on similar axes move together.

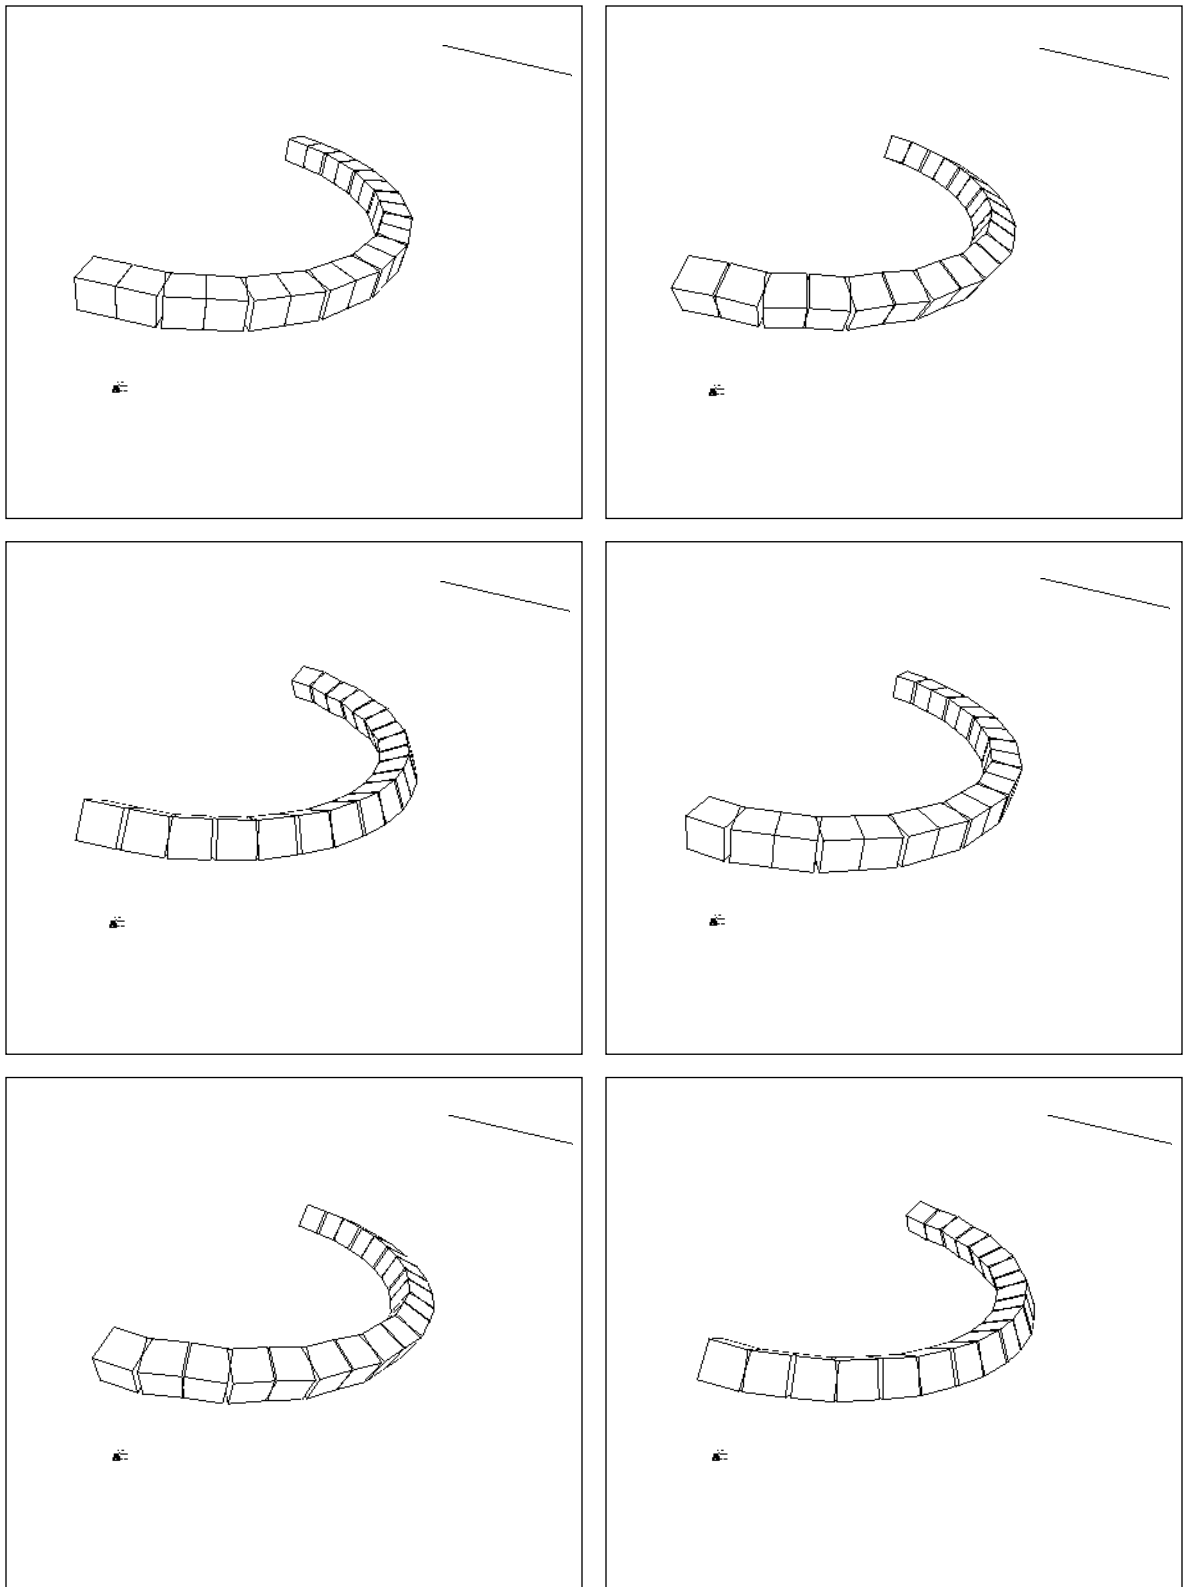


Figure 7-6: Lateral rolling.

Traveling wave rotor

This gait is similar to a spinning coin as it come slowly to rest. The ventral wave, in this case, is not a rigid body but a wave that propogates around the body of the circle formed by the snake robot body. This is also similar to the principle of ultrasonic motors that uses a traveling wave to move objects. Typically the ultrasonic motors use a low amplitude, high frequency wave to achieve motion of a plate.

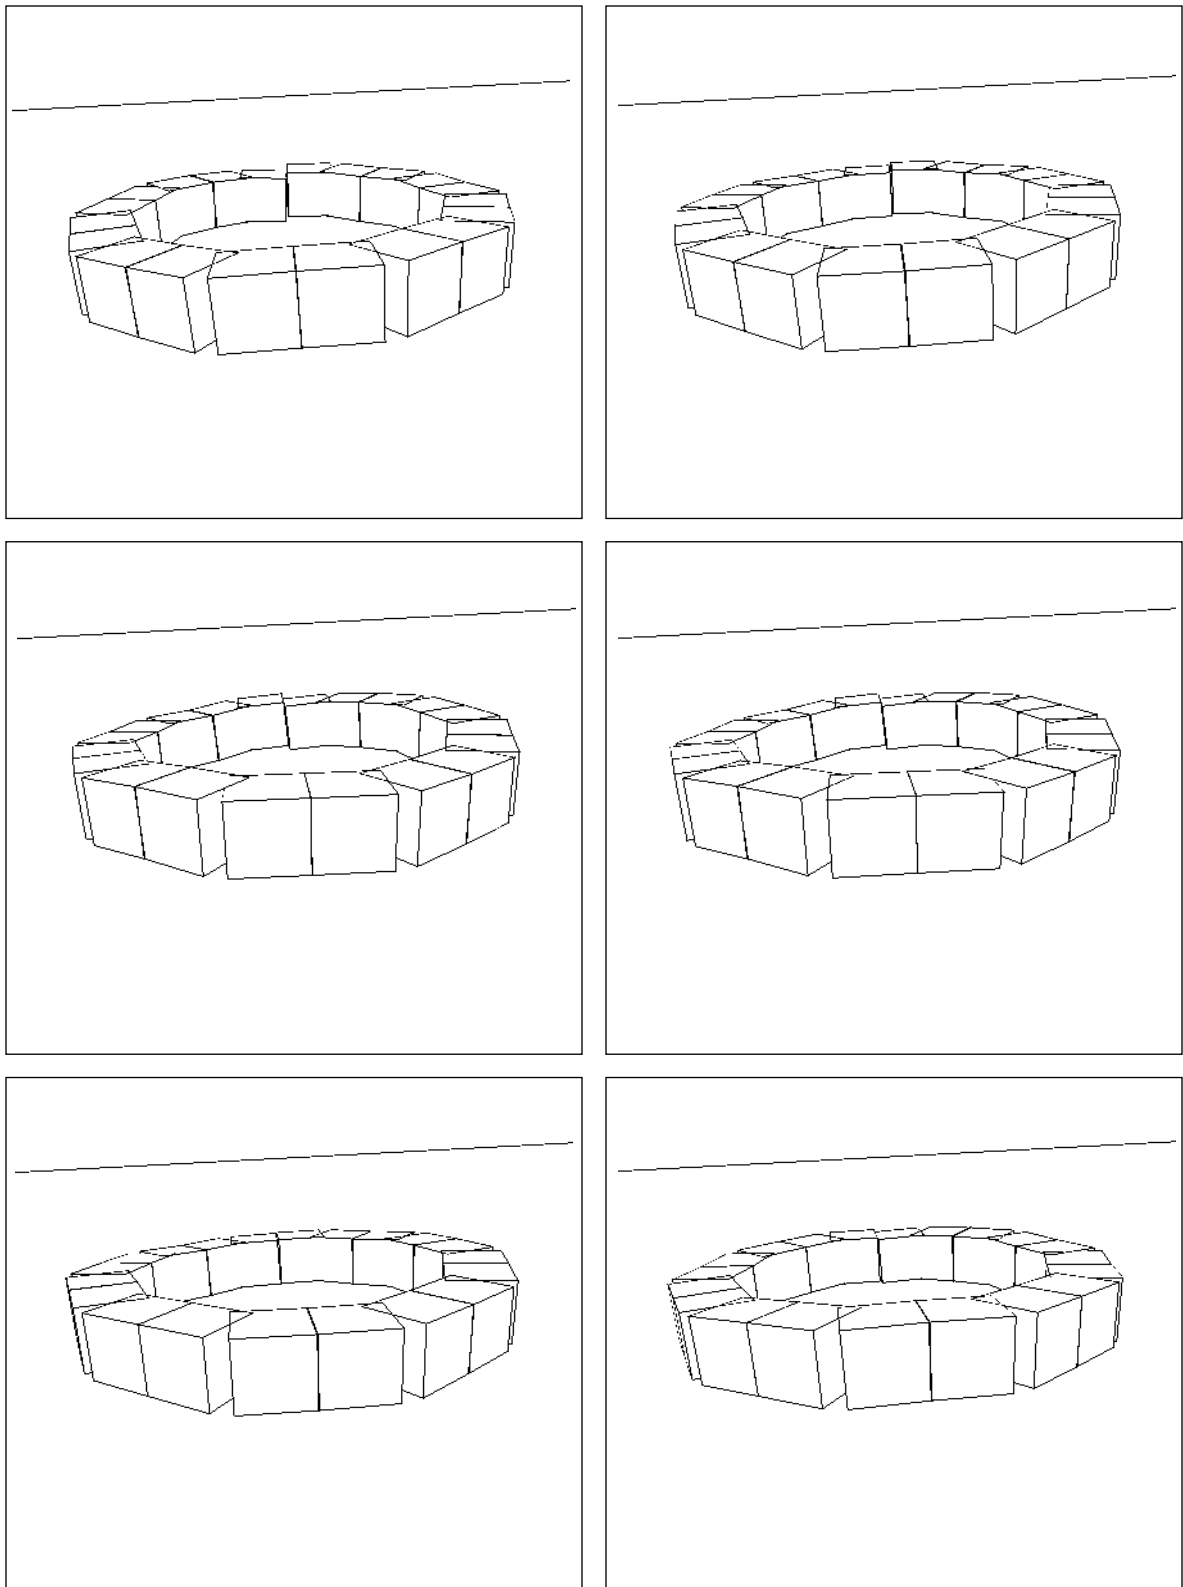


Figure 7-7: Traveling wave rotor.

Wheel

One of the most intriguing results was the reinvention of the wheel. Essentially a wrapping of the ventral joints and a coordinated motion provide a rotating section very much like a wheel or track.

[Brackenbury 97] details the odd locomotion of caterpillars that perform this feat; they can roll up and move in a protective coil. By doing so, the animal can recoil and retreat very quickly; much faster than it could locomote in a normal 'inching' manner.

This mode of locomotion was attempted on the caterpillar robot and was partially successful. The wiring complicated the attempt and the robot could quickly become unstable and fall sideways. However, in at least one test the robot rolled nearly a full revolution.

This mode could be reduced to a simple offset of the joints to form the circle and then a low amplitude wave is propagated along the periphery to roll the robot.

Several experiments revealed that interpenetration of the snake with itself occurs. In image five of the sequence, in this example, you can see the inter-penetration of the ends of the snake. To reduce computational requirements, the body segments are allowed to inter-penetrate. This reduces the amount of collision detection needed at each iteration and with most gaits this is not an issue because inter-penetration does not occur. With the wheel gait however, occasionally there is a self contact - but this does not appear to be critical to the gait.

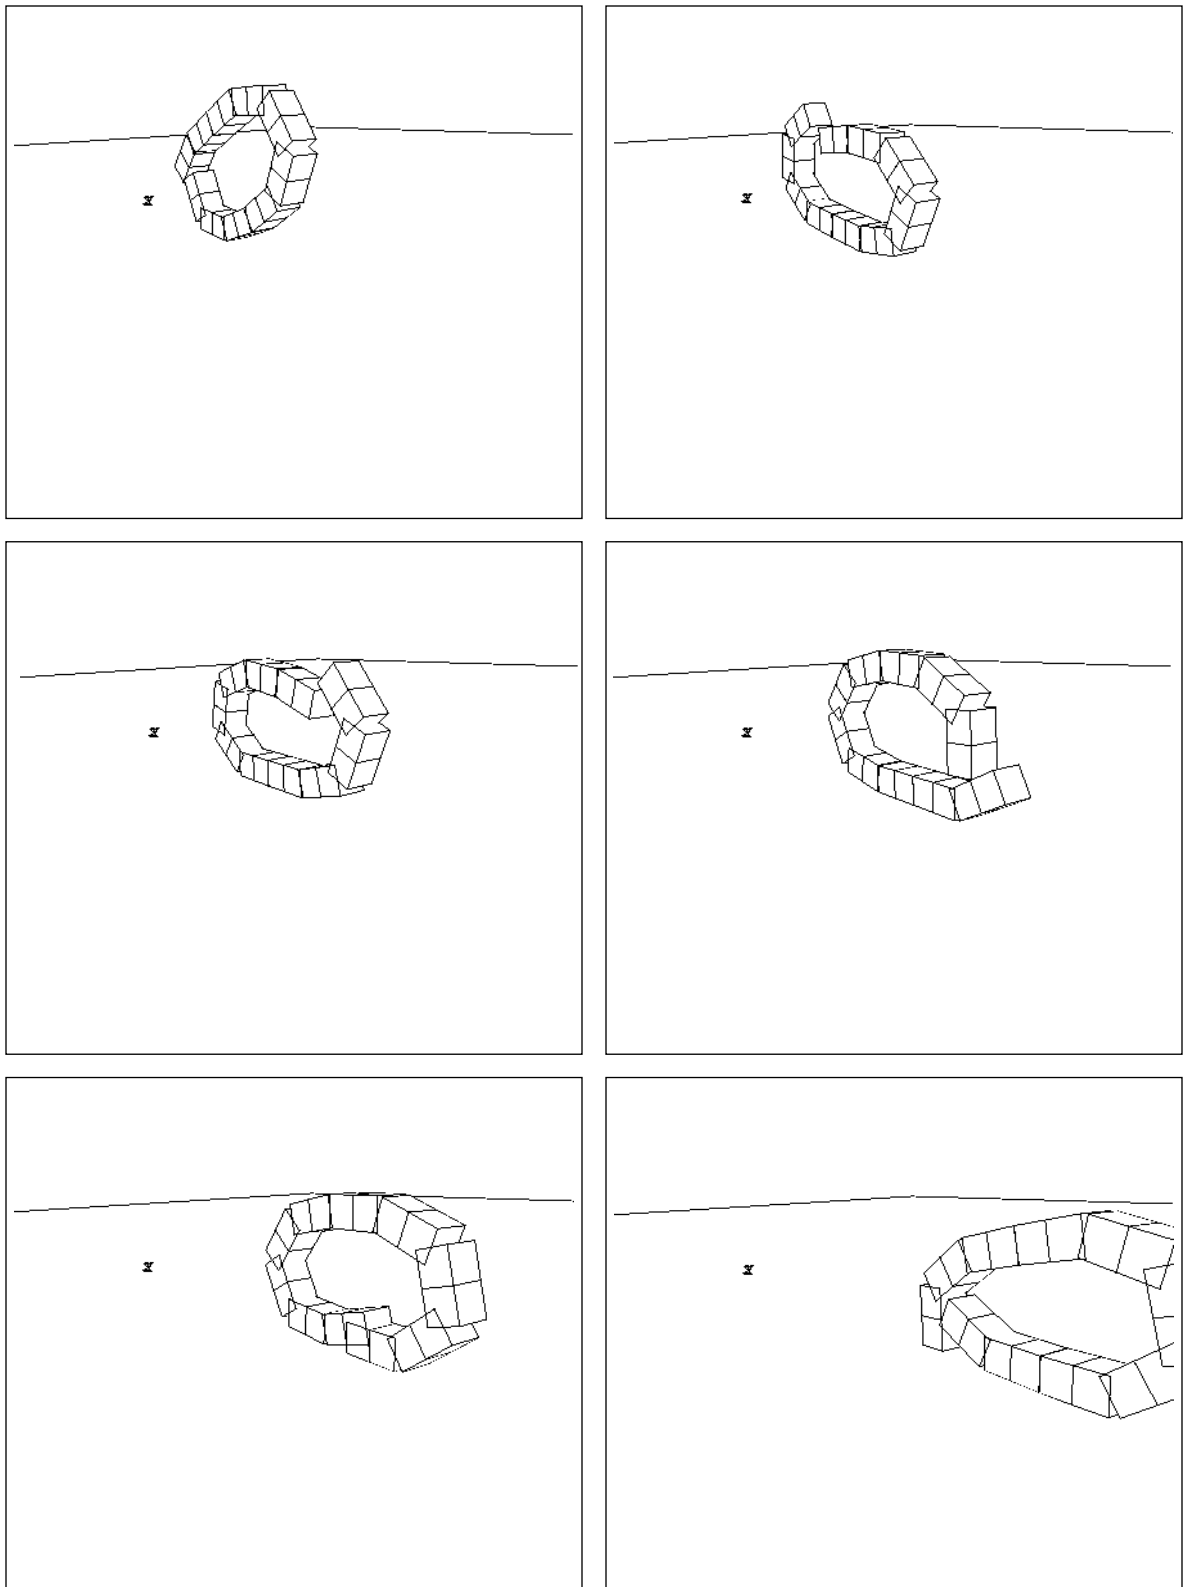


Figure 7-8: The wheel.

Flapping Locomotion

Another sideways mode of locomotion, this mode uses in-phase motions of the ends to swing forward, come down in contact and the lift or drag the center of the body forward. This is similar to the motions of a swimmer performing the butterfly stroke.

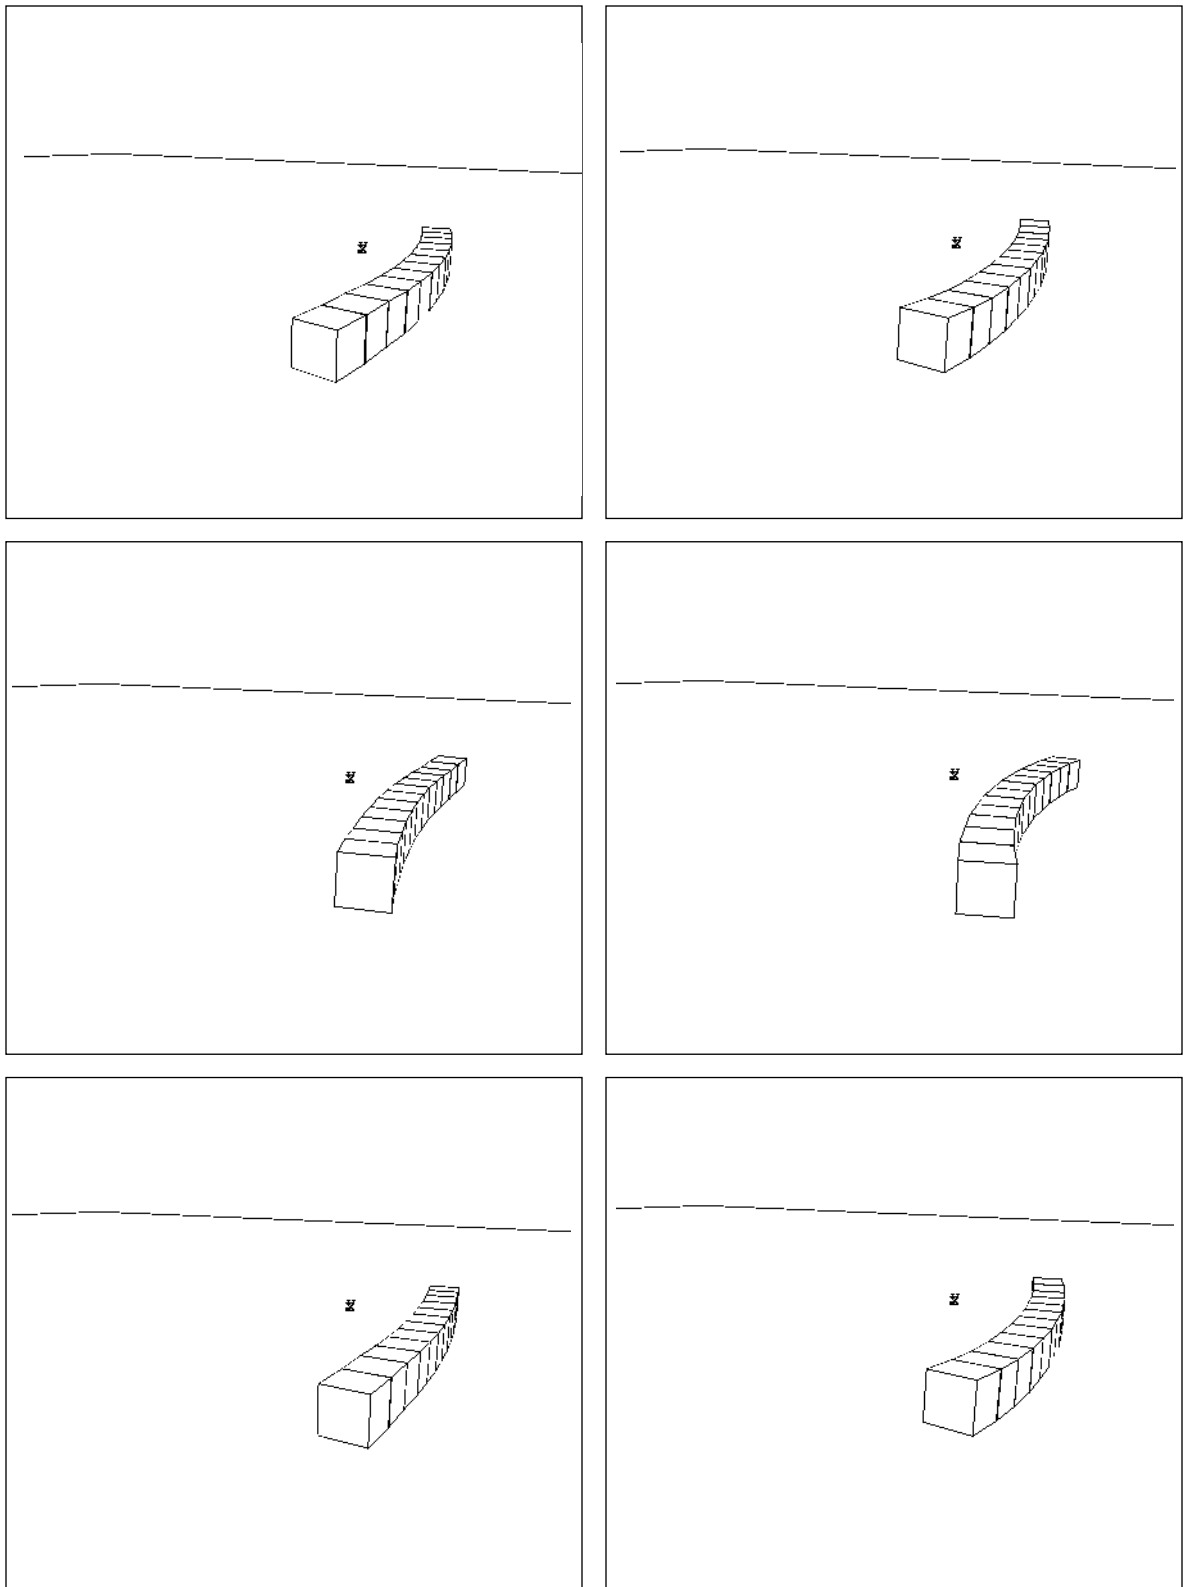


Figure 7-9: Flapping locomotion

Rolling Collar

If the configuration of lateral roller motion, shown in Figure 7-6, closes upon itself and forms a circle, it forms a rolling collar which looks similar to a smoke ring when in motion. By itself, on flat ground, this produces no locomotion. However, if this form is used to surround a pipe or other convex object or is used internally, then this motion produces locomotion. It produces a motion akin to that described by Nilsson in [Nilsson 97a], but without the use of the complex roll-pitch-roll joint. As in lateral rolling the body of the snake acts as a rolling wheel to move.

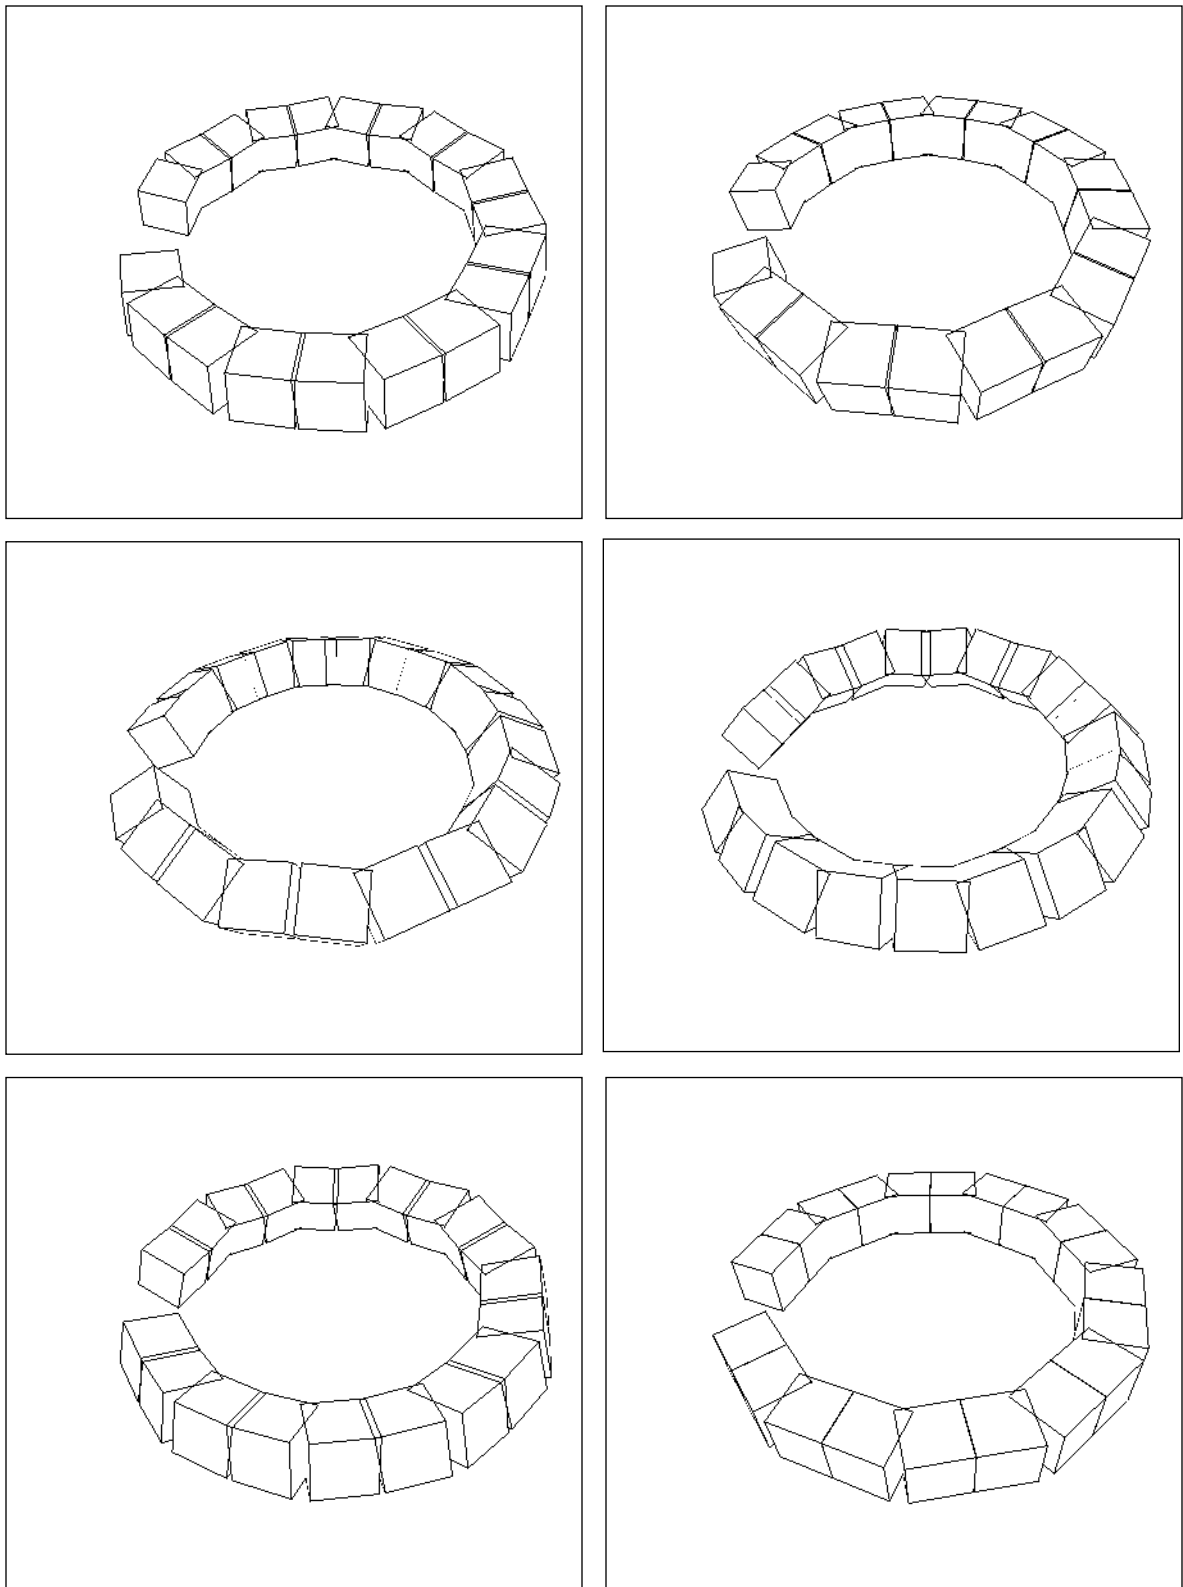


Figure 7-10: Rolling Collar locomotion

Helical rolling

Another feature is to use a constant offset for joints. For example, forming a hoop and then rippling a wave along the hoop. The formulation then becomes:

$$\text{Amplitude} * \sin(\text{time}/\text{period} + \text{phase} * i) + \text{offset}[i]$$

A slightly different class of gaits can be formed by providing fixed offsets of the joints so that a general configuration can be initiated and then operated upon during motion.

Concertina

An issue for varied gaits is that the time history of all the joints will be exactly the same. This works for many gaits, but concertina is one example where it does not. All of these gaits assume the same fixed frequency of motions along the snake, but the body frequency of an intermittent gait, such as concertina, is different. The wavelength of the body varies along the body and over time.

Discussion

The development of the gaits did not proceed in the manner originally envisioned; In many cases these gaits can be described by elementary formulations. However, it is also clear that natural snakes do not use these formulations. A number of snake gaits use a simple family of forms to describe joint and subsequent body motions.

$$\text{Amplitude} * \sin(\text{Time}/\text{Period} + \text{Phase}(\text{link number})) + \text{Offset}$$

The phase is usually the product of a phase value and the joint number. This produces a traveling or propagating wave down the body of the robot. Zero phase, obviously, produces the same motion at all joints simultaneously, whereas a phase of Pi produces alternating and opposite motions in adjacent links. The offset can also be particular to a specific joint. For example, the body could describe a nominal 3D shape, such as a helix and use that as a base from which to propagate other waveforms.

Note that this formulation describes the motion of the *links* and *not* the waveform of the body. In fact, it is the *phase* that most nearly describes the bodyform. This formulation can also be scaled to be proportional to the link number of the inverse. This gives some additional interesting motions.

180 degrees, +/-1.5 radians, of motion is the limit for individual joint excursions. In simulation, some interesting patterns emerge beyond that limit, including some intriguing lissajous-like patterns, but it makes no sense for mechanisms since hard physical constraints prevent these motions.

Additional features of these formulations include an amplitude factor that is proportional to the link number. This allows heading changes in the gait so that the system can be 'steered' in a desired direction.

Time

The time for simulation for a very small two-link snake on a powerful workstation could run twice as quickly as real time. However, longer snakes took substantially longer; run-times for longer snakes went up quadratically and, in some cases, even worse. The runtime depends mostly on the amount of contacts which, for a snake, are numerous. In addition, there is a lot of coupling between links that is closed through the ground or

surface. These closed kinematic chains plus the large number of contacts contribute to substantially longer run-times.




Figure 7-11: Performance drops significantly with increased number of segments

Figure 7-11 shows the dramatic performance drop when additional segments are used on the snake. This data was found from running a long series of dynamic simulations on snakes of varying length using a slowly propagating wave gait but not using high forces or large excursions. Each segment corresponds to two degrees of freedom. In general, the additional computation time is a function of the number of contact points; this appears to be a quadratic relationship, although it can be even worse [Baraff 97]. The numbers for actual runs became much worse, as much as 500 times longer than real-time for some configurations, taking several hours for an individual run.

Real-time playback

As a result of the runtime issue, simulation is quite slow for the larger interconnected physical systems. For this reason, a method of storing the state of the system at each timestep was developed. This stored the state vector of the entire system at each iteration of the simulation and appended it to a file. The file could later be played back, as rapidly as a purely kinematic model, to show the performance in real time.

Summary

Lessons

At this point in time the computational requirements for generalized learning of complex physical systems seems prohibitive. The 100's:1 ratio of real-time to simulated time is sufficient to explore a single model, but when thousands must be explored it becomes intractable. The obvious approach is the use of massive parallelization such as

that by Sims or the further simplification of the physical model. Simplification brings its own set of dangers for model fidelity; the results may be useless.

Another issue with the simulation system is the sensitivity to a variety of parameters. This includes the dimensioned quantities such as length, mass, density etcetera. If other values that are calculated from these values are too small then the resulting discrepancies will eventually result in poor physical modeling. The other major parameter that physical simulation is sensitive to is the size of the timestep between iterations. Although the time step is adaptive, depending on the state of the simulation, there is still a large sensitivity to timestep values. If too large, errors will almost certainly occur and if too small, the amount of time required to run will increase. The problem is determining what value of the timestep will provide the best trade-off of accuracy and speed. There is no hard and fast rule. In the graphics community, this is nearly a moot point; if it looks good, it is good.

There is no doubt that within a few years that the combination of computational power and parallelizing techniques over networks will allow rapid high fidelity simulation and testing for learning in physical robots.

Chapter 8

Summary and Conclusions

Summary and Conclusions summarizes the results and contributions of this research and looks into the future to see what research lies ahead in this area.

The old aphorism, “you have to crawl before you can learn to walk”, has not applied to mobile robotics. There have been many walking machines over the decades but very few crawling machines. In fact, crawling appears to be a harder problem. In this dissertation, research demonstrated that a snake robot can learn to crawl and it can crawl in several different ways.

The conclusion of this research is that robots can learn to locomote even when they have no wheels or legs. In this dissertation I provide a general framework to teach a complex electromechanical robot to become mobile that includes a learning method, metrics for evaluation, physical simulation and the transfer of results to a robot.

Is this thesis extendable to other mechanisms and forms? The framework and loop of learning, testing and evaluation is certainly applicable to a wide variety of domains for physical control. For locomotion, all patterns of motion, gaits, can be described in terms of cyclical or periodic forms and this architecture lends itself well to learning those modes of locomotion.

Contributions

There were a number of new and interesting developments in this work.

Contributions included:

- A novel and practical design for snake-like locomotors. With the exception of the NEC snake, prior serpentine robots have paid scant attention to practical packaging of devices such as actuators and wiring.

- Learning to locomote with a limbless locomotor. Prior works have utilized explicit models and relied on gaits contrived by humans. However, even the process described here wasn't as clean as simulate, then learn, then try on robot; it really became an iterative process.
- Varying Multiple gaits using a single mechanism. Prior locomotors have shown only one or two variants of a gait. Burdick's Snakey did show three gaits: traveling wave, stationary wave and an extensible wave; a type not possible in a snake or this robot. In other works, ground contacts used wheels or pointed metal pins to provide ground purchase. This has limited the types of locomotion that they can achieve. Snake-like gaits shown in simulation and on the robot included rectilinear, sidewinding and lateral rolling. A wide variety of novel gaits and motions were also shown including lateral undulation, variant of lateral rolling called the smoke ring and other novel gaits including the ventral wave and the butterfly gait. These extraordinary gaits can achieve locomotion without parallels in the animal kingdom.
- General learning technique and framework for representing periodic gaits. The architecture proved valuable in development. It set the stage for creating and developing gaits.
- The first development of skins for limbless locomotors. This included an evaluation of many materials for use as ground contact interface and protective sheath. This is very different from the point contacts and exposed mechanisms done by others in this area. The resulting skins made from stretchable fabrics provide protection, a low-pass filtering of the mechanics akin to a spline, and a smooth ground contact interface. This is accomplished while being faithful to the shape and movement of the underlying mechanism.
- A comprehensive examination of performance metrics and metric evaluation including the development of a new metric - payload velocity to describe capabilities of working machines that transport material.
- Novel gaits including the flapping locomotion and the lateral rolling gaits. The wheel has previously been shown in simulation by Yim.

Future Work

As with any sequence of work and discovery, you always discover how you can do it better. There are many future advances in this work including mechanism and control. I believe this mechanism works well but that there are numerous changes and refinements to the current robot to improve it. Here are a few areas for further work:

- Gaits - the testing and evolution of gaits has only begun. While a number of intriguing gaits have been shown, there is much more that can be done. Varied terrains, experimenting further with direction, vertical climbing and more. Additional work in steering and gait transitions is necessary for more general locomotion. With each test, with each gait, new ideas suggest themselves and I foresee fruitful work in this area.
- Mechanism - There are a number of potential improvements to the mechanism. These include: an easier-to-disassemble joint structure with a rapid mechanical and electric connection; perhaps similar to bayonet style connectors. The use of lighter materials,

such as polymers and composites, fabricated from molding processes will not only lighten the structure but result in a beneficial cascade effect of requiring even smaller, lighter actuators.

- Power - As with so many other applications, power is the critical technology for deploying small robots. Long term energy and short term power needs dictate limitations and capability. Recent advances in battery technologies and evolution of technologies such as fuel cells will further this and many other applications.
- Sensing - The addition of sensing takes two forms: the sensor itself and the means to use the sensor information. Sensing can utilize the small high-resolution force sensors described herein but utilizing sensing in an appropriate manner will require further research.
- Electronics- Wiring is a real issue, constant use involves wear and tear; wire exposure results in abrasion, wear and failure. One suggestion is to develop a simple bus using small PIC or ASIC devices to run motion control and feedback for each joint. This will minimize wiring and increase the control and flexibility at end joint.
- Learning - Faster computing is inevitable. This will enable exploration of even more of the search space. Stochastic learning techniques will benefit strongly from this and the use of transparent parallelization across computing platforms. Planning is also required; the ability to traverse 3D terrains will require substantial planning issues; although a case can be made for a reactive strategy for overcoming obstacles and marginal terrains.
- Physical Simulation - The first pass at real simulators is barely adequate. Coriolis and a recent commercial package, Working Model, are steps in the right direction, but computation needs and high fidelity modeling capability are sorely needed for complex systems. The selection of simulation parameters, such as the time step and iteration method can radically change the output of the simulation for the same starting conditions.

Wisdom

It is possible to go full term from design and simulation to the locomotion control of complex mobile robots. This can be done in the context of learning which can be used both in simulation and on the real machine.

From here, many possibilities suggest themselves for the design and control of complex mechanisms. Snake robots, in particular, can offer a variety of useful applications and uses ranging from exploration to inspection.

Metrics are germane to evaluation. However, if taken too seriously for all locomotion, metrics will tend to favor particular classes of vehicles. This is important to keep in mind during the design or learning process.

Appendix A

Servo Evaluation

Table 1: R/C Servo Comparison^a

Model	Torque [N-m]	Dimensions [mm]			Mass [g]	Speed [sec/60 deg] ^b	Power ^c [Watts]	Torque /Mass [Nm/kg]	Torque/ Volume [Nm/m ³ /1000]	Power/ Weight [W/N]
JR 341	0.23	12.70	28.45	29.72	17.86	0.24	1.18	12.613	20.98	6.73
JR 321	0.21	14.73	33.02	25.91	21.83	0.23	1.13	9.446	16.36	5.26
JR 3021	0.26	14.73	33.02	25.91	23.81	0.22	1.51	11.121	21.01	6.48
JR 3025	0.21	14.73	33.02	25.91	45.64	0.15	1.73	4.518	16.36	3.86
JR 901	0.30	18.03	34.80	33.53	37.71	0.27	1.42	8.072	14.47	3.83
JR 9021	0.41	18.03	34.80	33.53	42.53	0.22	2.32	9.549	19.30	5.56
JR 507	0.28	18.54	38.61	33.53	41.67	0.25	1.43	6.829	11.86	3.50
JR 517	0.28	18.54	38.61	33.53	44.79	0.25	1.43	6.354	11.86	3.26
JR 4000	0.52	18.54	38.61	33.53	49.90	0.19	3.44	10.417	21.66	7.03
JR 4131	0.64	18.54	38.61	33.53	42.53	0.23	3.49	15.012	26.60	8.36
JR 4721^d	0.84	18.54	38.61	33.53	48.76	0.22	4.82	17.321	35.19	10.09
JR 4735	0.64	18.54	38.61	33.53	48.76	0.15	5.32	13.034	26.48	11.13
JR 703	0.66	22.35	43.94	23.62	32.89	0.51	1.62	20.014	28.37	5.03
JR 7000	0.44	22.35	43.94	23.62	41.11	0.19	2.92	10.754	19.05	7.25
JR 7005	0.44	22.35	43.94	23.62	37.14	0.19	2.92	11.904	19.05	8.03
JR 3321	0.42	14.73	33.02	33.02	26.93	0.36	1.47	15.680	26.29	5.58
JR 605	0.98	32.00	63.50	58.42	134.66	0.28	4.41	7.295	8.27	3.34
Fut S125	0.91	22.35	39.62	42.93	65.21	0.62	1.85	14.004	24.02	2.89
Fut S132H	0.18	17.27	36.32	29.97	31.19	0.13	1.71	5.661	9.39	5.58

Table 1: R/C Servo Comparison^a

Model	Torque [N-m]	Dimensions [mm]			Mass [g]	Speed [sec/60 deg] ^b	Power ^c [Watts]	Torque /Mass [Nm/kg]	Torque/ Volume [Nm/m ³ /1000]	Power/ Weight [W/N]
Fut S3302	0.78	28.96	58.93	50.04	102.06	0.19	5.13	7.611	9.10	5.13
Fut S3303	1.41	28.96	58.93	50.04	107.73	0.26	6.82	13.111	16.54	6.46
FUT S9303	0.70	20.07	40.39	35.56	65.21	0.19	4.62	10.722	24.26	7.23
FUT S9304	0.49	20.07	40.39	35.56	48.20	0.22	2.80	10.184	17.03	5.93
FUT S9403	0.31	20.07	40.39	35.56	48.20	0.16	2.47	6.521	10.91	5.22
FUT S9601	0.25	15.75	30.73	29.97	31.19	0.17	1.88	8.175	17.57	6.16
Tower ts-72	0.94	58.42	27.94	50.80	99.23	0.22	5.36	9.466	11.33	5.51
Tower ts-55	0.30	40.64	20.32	38.10	45.36	0.20	1.91	6.695	9.65	4.29
Tower ts-11	0.21	27.94	13.72	27.94	17.29	0.15	1.77	12.251	19.79	10.47
Air 94831	0.27	37.08	18.03	29.97	31.19	0.17	1.98	8.605	13.39	6.49
Air 94510	0.78	47.50	22.86	39.12	65.21	0.33	2.96	11.914	18.29	4.63
Air 94501	0.14	26.92	12.45	26.92	18.43	0.33	0.54	7.665	15.65	2.98
Air 94401	0.23	30.99	14.99	30.99	26.93	0.26	1.13	8.653	16.19	4.26
Air 94403	0.18	30.99	14.99	30.99	26.93	0.20	1.11	6.555	12.27	4.20
Air 94737	0.39	39.37	20.32	35.56	53.87	0.15	3.25	7.211	13.65	6.16
Condor MS-747WB	1.18	54.61	52.07	26.67	113.40	0.26	5.70	10.400	15.55	5.13
Condor SSPS-105	35.31	130.05	55.12	111.00	779.63	0.60	73.90	45.291	44.38	9.67
Hitec 605-BB	0.54	40.89	19.81	39.88	49.05	0.16	4.27	11.087	16.83	8.88
Hitec red-apollo15	1.20	30.48	48.26	58.42	85.05	0.23	6.55	14.116	13.97	7.86
Hitec HS-615	0.76	38.10	20.32	40.64	60.10	0.21	4.52	12.573	24.02	7.67
Hitec HS-805BB	1.58	55.88	30.48	60.96	119.07	0.20	9.93	13.285	15.24	8.51
Hitec HS-205MG	0.30	33.02	17.78	33.02	31.19	0.20	1.91	9.738	15.66	6.24

a. All units converted to SI units. All figures using manufacturers specifications.

b. Servo speed is time to rotate 60 degrees. This is converted to rpm for power calculations.

c. Power is calculated using 25% of product of max stall torque and max rpm.

d. Selected for serpentine device.

Appendix B

Link Weight Distribution

Table 2-1: Weight of snake link components.

		Weight (g)	Quantity	Total weight (g)
Hardware	4mm bearing	0.60	1	0.60
	4-40 FHCS, 1/4"	0.37	1	0.37
	4-40 BHCS, 1/4"	0.35	8	2.80
	4-40 nut	0.48	8	3.84
	2-56 BHCS, 1/4"	0.20	8	1.60
	2-56 nut	0.18	8	1.44
	horn screw	0.50	2	1.00
	bearing spacer	0.50	2	1.00
	Hardware Total			12.65
Brackets	servo plate	9.40	1	9.40
	bearing plate	3.80	2	7.60
	horn plate	4.00	2	8.00
	octo plate	3.30	2	6.60
	Bracket Total			31.60
Servo	JR4721	45.00	2	90.00
	horn	1.40	2	2.80
	Servos Total			92.80
	Link Total			137.05

The weight distribution of the link components is servos 68%, hardware 9%, and brackets 23%. Since the actuators are 2/3 of the weight, there are diminishing returns in attempting to reduce the weight of the hardware or brackets.

Appendix C

Derivation of Actuator Parameters

Developing an externally-based model of a control system requires observation of input signals and a measured response of output. The actuator is treated as a second order system incorporating stiffness and damping coefficients and a mass or inertia. This model neglects non-linearities because the experimental results appeared to fit a second order system very well. To solve for these requires observations combined with a derivation based on the observable response. This is accomplished by determining the undamped natural frequency, the damping ratio, and the period from experimental data and the solving for the second-order parameters as follows:

$$w_{n_1} = \sqrt{\frac{k}{I_1}} = \text{undamped natural frequency} \quad [C-1]$$

Solving for k:

$$k = w_{n_1}^2 I_1 \quad [C-2]$$

The damping ratio, shown below, is the ratio of the actual damping value, b, over the critical damping value.

$$\xi_1 = \frac{b}{2\sqrt{kI_1}} = \text{damping ratio} \quad [C-3]$$

Solving for b:

$$b = 2\xi_1\sqrt{kI_1} \quad [C-4]$$

Now, solving for another natural frequency by adding a small mass to the system:

$$w_{n_2} = \sqrt{\frac{k}{I_1 + I_2}} \quad [C-5]$$

Solving for k gives:

$$k = w_{n_2}^2 (I_1 + I_2) \quad [C-6]$$

Now setting the two systems equal to each other:

$$w_{n_1}^2 I_1 = w_{n_2}^2 (I_1 + I_2) \quad [C-7]$$

Then solve for I_1 :

$$I_1 = \frac{I_2}{\left(\frac{w_{n_1}^2}{w_{n_2}^2} - 1\right)} \quad [C-8]$$

Now, substituting for k in Equation [C-4]:

$$b = 2\xi_1 \sqrt{I_1 w_{n_1}^2} = 2\xi_1 I_1 w_{n_1}^2 \quad [C-9]$$

Substituting for I_1 gives:

$$b = \frac{2\xi_1 I_2 w_{n_1}^2}{\left(\frac{w_{n_1}^2}{w_{n_2}^2} - 1\right)} \quad [C-10]$$

Substituting for I_1 into [C-2] then gives:

$$k = \frac{w_{n_1}^2 I_1}{\left(\frac{w_{n_1}^2}{w_{n_2}^2} - 1\right)} \quad [C-11]$$

The three parameters, k, b, and I_1 , can now be used to find the parameters for the system model. ζ is determined directly from the actuator response from Equation [C-12] where x_n is the first amplitude peak, x_n is the final value of the output and the period is the time between zero crossings. See Figure C-1.

$$\zeta = \frac{\left(\frac{1}{n-1}\right) \ln\left(\frac{x_1}{x_n}\right)}{\sqrt{4\pi^2 + \left(\left(\frac{1}{n-1}\right) \ln\left(\frac{x_1}{x_n}\right)\right)^2}} \quad [\text{C-12}]$$

To do this, however, several parameters have to be determined experimentally, The natural frequency, ω_n , and the damped natural frequency, ω_d , were determined experimentally by optically tracking the output of the actuator. A tracking LED was attached to an arm connected to the servo. Data from the position of the arc was recorded at 1ms intervals during motion. Since the motion describes an arc, a circle was fit to the data and used to determine the angle in radians as a function of time. This decaying oscillation was measured directly from the data to determine the period of oscillation (time between zero crossings). See Figure C-1. The waveform was also used to determine the damping ratio based upon the amplitude decay [Ogata 78].

$$w_n = \frac{w_d}{\sqrt{1 - \zeta^2}} = \frac{\left(\frac{1}{n-1}\right) \ln\left(\frac{x_1}{x_n}\right)}{\zeta T} \quad [\text{C-13}]$$

and finally,

$$w_d = 2\pi/T \quad [\text{C-14}]$$

Thus, the response parameters are found from experiments and the stiffness, damping




Figure C-1: Modeling the decaying oscillation of the servo actuator

and inertial coefficients are found from those values.

Now, the problem is to map these model parameters onto the gains used in a typical control system. Typical control system gains include proportional, derivative and integral terms or k_s , k_d and k_i respectively. The integral term can reduce or eliminate steady state error in a system at the expense of settling time and longer term oscillations. Even worse, it can introduce uncontrollable limit cycling. From observation of the actuator output it does not appear that integral gain is used in its internal control system; even if it is, it appears to have a negligible influence on the control system. The

system is quite stiff and the force does not appear to increase over time, only with angular error. For this reason, our model is:

$$I\ddot{x} + b_n\dot{x} + kx = k_s(x_d - x) + k_d(\dot{x}_d - \dot{x}) \quad [C-15]$$

Where b_n is the natural damping of the system as separated from the derivative gain.

Rearranging terms gives

$$I\ddot{x} + (b_n + k_d)\dot{x} + (k + k_s)x = k_s x_d + k_d \dot{x}_d \quad [C-16]$$

The Laplace transform equivalent is

$$s^2 + 2\zeta w_n s + w_n^2 \quad [C-17]$$

Dividing through by I and equating equivalent coefficients gives the following relations:

$$\frac{b_n + k_d}{I} = 2\zeta w_n \quad [C-18]$$

and

$$\frac{k_n + k_s}{I} = w_n^2 \quad [C-19]$$

These terms define the total system gains. The problem is to now distinguish the two contributions. The b , k , and I terms solved for previously are actually the numerator sums in the left hand side of [C-18] and [C-19]. For the spring stiffness, k_n , and the proportional gain, k_s , however, we will assume the contribution is entirely from the system and that the actuator stiffness is infinite. This does not neglect stiffness, it merely transfers all the effects into the overall system. Now the problem becomes the separation of system damping from actuator damping; the $b_n + k_d$ term. To eliminate the closed loop term from the natural system term requires another experiment. The actuator will be allowed to move under a gravity load, as a pendulum, and its response observed. The settling time and oscillation provides a response of the natural system and not closed loop behavior.

$$I\dot{w} + b_n w = -mgL \sin\theta \quad [C-20]$$

For small angles, $\sin\theta = \theta$ and shifting from angular velocity terms, w , to angular terms, θ , results in:

$$I\ddot{\theta} + b_n \dot{\theta} + mgL\theta = 0 \quad [C-21]$$

Again, similar to above, diving through by I, the Laplace version is as follows:

$$\theta(s) \left[s^2 + \frac{b_n}{I}s + \frac{mgL}{I} \right] = 0 \quad [C-22]$$

Then, equating like terms and also, for a pendulum, the inertia is equal to the product of the mass and the square of the length.

$$2\xi w_n = \frac{b_n}{I} = \frac{b_n}{mL^2} \quad [C-23]$$

Also

$$w_n^2 = \frac{mgL}{I} = \frac{mgL}{mL^2} = \frac{g}{L} \quad [C-24]$$

Solving for b gives

$$b_n = mL^2 2\xi \sqrt{\frac{g}{L}} \quad [C-25]$$

Thus, given the response of the system we can solve for the damping coefficient of the actuator independent of the closed loop value shown earlier. The difference of the two gives the derivative gain.

What is the end result of these derivations and underlying meaning? Can we establish that the values are sufficient for simulation and modeling purposes? The purpose, when we began, was to develop a model of sufficient fidelity that physical simulation results are valid for transfer to the real system. Vagaries and idiosyncrasies of the physical simulation tool make it difficult to ascribe figures of almost any accuracy, but the proof is in the results - they appear to approximate the motions and dynamics of the real system. With nearly all control systems, the control parameter gains are adjusted in an iterative process because the pervasive effects of friction and contact are intractable to model perfectly. It is no different in this process; values were determined and later adjusted to reflect the actual actuator performance. However, the magnitude of these adjustments was relatively small. The end result is a better understanding of actuator performance and a model sufficient for simulation purposes.

References

Following are complete references to material cited in the text.

-
- [Anon 95] Anonymous, "Steam Power 1899," *Automotive Industries*, July 1995, p36.
- [Alexander 84] Alexander, R., "The Gaits of Bipedal and Quadrupedal Animals," *International Journal of Robotics Research*, Vol. 3, No. 2, Summer 1984, pages 49-59.
- [Alexander 92] Alexander, R., *Exploring Biomechanics*, Scientific American Library, W.H. Freeman, New York, NY, 1992.
- [Apostolopoulos 95] Apostolopoulos, D., Bares, J., "Configuration of a Robust Rappelling Robot", IROS '95. Pittsburgh, PA.
- [Arakawa 93] Arakawa, K., Krotkov, E., "Fractal Surface Reconstruction for Modeling Natural Terrain," 1993 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, New York, NY June 15-18, 1993.
- [Asano 83] Asano, K. et al., "Multijoint Inspection Robot," *IEEE Transaction on Industrial Electronics*, Vol. IE-30, No. 3, 277-281, 1983. Toshiba.
- [Asano 84] Asano, K., "Control System for a Multijoint Inspection Robot," *Proceedings of the 1984 Topical Meeting on Robotics and Remote Handling in Hostile Environments*, Gatlinburg, Tenn. on April 23-27, 1984, pp 375-382. Toshiba.
- [Bauchot 94] Buchot, R., Ed, *Snakes: A Natural History*, Sterling Publishing, New York, ISBN 0-8069-0654-5, 224 pp, 1994.
- [Baluja 95a] Baluja, S., "An Empirical Comparison of Seven Iterative and Evolutionary Function Optimization Heuristics," *Carnegie Mellon University Computer Science Technical Report CMU-CS-95-193*. 1995.
- [Baluja 95b] Baluja, S., Caruana, R., "Removing the Genetics from the Standard Genetic Algorithm," *Proceedings of the Twelfth International Conference on Machine Learning*, Lake Tahoe, CA. pp. 38-46, July, 1995.

References

- [Baraff 97] Baraff, David, "Coriolis Documentation," Carnegie Mellon University and Physical Effects, Inc., 1997.
- [Bares 91] Bares, J., *Configuration of Autonomous Walkers for Extreme Terrain*, Doctoral Dissertation in Civil Engineering, Carnegie Mellon University, May 1991.
- [Barrett 97] Barrett, D., "Optimization of Swimming Locomotion by Genetic Algorithm," Recent Trends in Robot Locomotion Workshop, ICRA, 1997, Albuquerque, NM. April 21, 1997.
- [Bekker 56] Bekker, M., *Theory of Land Locomotion: The Mechanics of Vehicle Mobility*, University of Michigan Press, Ann Arbor, MI, 1956.
- [Bekker 61] Bekker, M., "Is the wheel the last word in land locomotion?," *New Scientist*, No. 248, August 1961, pages 406-410.
- [Bekker 69] Bekker, M., *Introduction to Terrain-Vehicle Systems*, University of Michigan Press, 1969.
- [Bertram 94] Bertram, John, Personal Conversation, Cornell, Department of Biology, 1994.
- [Binnard 92] Binnard, M., *Leg Design for a Small Walking Robot*, Bachelor's Thesis in Mechanical Engineering, MIT, May 1992.
- [Binnard 95] Binnard, M., *Design of a Small Pneumatic Walking Robot*, MS Thesis, Department of Mechanical Engineering, MIT, 1995.
- [Box 87] Box, G.E.P, Draper, N., *Empirical Model-building and Response Surfaces*, Wiley, New York, 1987, 669pp.
- [Brackenbury 97] Brackenbury, J., "Caterpillar Kinematics," *Nature*, Vol. 390, December 4, 1997. p453.
- [Brock 91] Brock, D., "Dynamic Model and Control of an Artificial Muscle based on Contractile Polymers," MIT AI Memo 1331. November 1991, 22 pp.
- [Brooks 89] Brooks, R., "A Robot that Walks; Emergent Behaviors from a Carefully Evolved Network," Proceedings of the IEEE International Conference on Robotics and Automation, May 1989, pages 692-696, also in *Biological Neural Networks in Invertebrate Neuroethology and Robotics*, Academic Press, Inc., Boston MA, 1993, pages 355-363.
- [Brooks 91] Brooks, R., "Intelligence Without Reason," MIT AI Memo No. 1293, for Proceedings of the International Joint Conference on Artificial Intelligence, April 1991.
- [Burdick 95] Burdick, J., Radford, J., Chirikjian, G., "A 'sidewinding' locomotion gait for hyper-redundant robots," *Advanced Robotics*, Vol. 9, No. 3, pp.195-216, 1995.
- [Burdick 97] Burdick, J., Personal conversation.
- [Caldwell 89] Caldwell, D.G., Taylor, P.M., "An Artificial Muscle Actuator for Robots," Proceedings Advanced Robotics 1989, Ed. C.W. Burckhardt., pp. 244-258.
- [Carnegie 95] Carnegie Museum of Natural History, Section of Amphibians and Reptiles, Dr. Ellen J. Censky, Assistant Curator and Section Head. *Eunectes Murinus Gigas*, Northern Anaconda, Specimen collected in Columbia, May, 1967. Borrowed 1995.
- [Caterpillar 94] Caterpillar, *Caterpillar Performance Handbook*, 1994. Peoria, IL

- [Chirikjian 92] Chirikjian, G.S., *Theory and Applications of Hyper-Redundant Manipulators*. PhD Thesis, California Institute of Technology, Pasadena, CA 1992.
- [Chirikjian 95] Chirikjian, G.S., Burdick, J.W., "The Kinematics of Hyper-Redundant Locomotion." *IEEE Transactions on Robotics and Automation*, Vol. 11, No. 6, December 1995, pp. 781-793.
- [Choset 96a] Choset, H., Burdick J., "Sensor Based Motion Planning: The Hierarchical Generalized Voronoi Graph," Workshop on Algorithmic Foundations of Robotics, 1996
- [Choset 96b] Choset, H., *Sensor Based Motion Planning: The Hierarchical Generalized Voronoi Graph*, Ph.D. Thesis, Caltech, 1996.
- [Congress 91] U.S. Congress, Office of Technology Assessment, *Miniaturization Technologies*, OTA-TCT-514 (Washington, DC; U.S. Government Printing Office, November 1991).
- [Dario 89] Dario, P. Bergamasco M., Salsedo, F., "Shape Memory Alloy Actuators for Robotic End-Effectors," Proceedings from the 4th Intl Conference on Advanced Robotics, Columbus, Ohio, June 1989, pp. 234-243.
- [DCC 94] Digital Command Control Standard, NMRA, Version S.91, August 1994. <http://www.tttrains.com/nmrdoc/>.
- [Desai 95] Desai, R., Rosenberg, C.J., Jones, J.L., "Kaa: An Autonomous Serpentine Robot Utilizes Behavior Control," 1995 International Conference on Intelligent Robots and Systems, IROS '95, Aug 5-9, Pittsburgh, PA, V.3 pp. 250-255.
- [Donner 84] Donner, M., "Control of Walking: Local Control and Real Time Systems," Computer Science Technical Report CMU-CS-84-121, Carnegie Mellon University, May 1984.
- [Dowling 97a] Dowling, K., "Power Sources for Small Robots," CMU Technical Report, RI-TR-97-02, January, 1997.
- [Dowling 97b] Dowling, K., "Survey and Evaluation of Serpentine Robots," CMU Robotics Institute Technical Report, To be published. CMU-TR-97-XX. 1997.
- [Drozda 84] Drozda, T.J., "The Spine Robot...The Verdict's Yet to Come," *Manufacturing Engineering*, pp. 110-112, Sept., 1984.
- [Dyberg 86] Dyberg, J., "Magnetostrictive Rods in Mechanical Applications," First Intl Conf on Giant Magnetostrictive Alloy and Their Impact on Actuator and Sensor Technology, Marbella, Spain, March 1986.
- [Eguchi 84] Eguchi, S., et al, "Self Propelled Monitoring System in Pipe MOGRER," Robot No. 47, 1984, (in Japanese)
- [Eltringham 91] Eltringham, S.K., *The Illustrated Encyclopedia of Elephants*, Crescent Books, New York, 1991, ISBN 0-517-06136-8.
- [Flynn 92] Flynn, A., Tavrow, L., Bart, S., Brooks, R., Ehrlich, D., Udayakumar, K., Cross, L. E. "Piezoelectric Micromotors for Microrobots," *IEEE Journal of Microelectromechanical Systems*, Vol. 1, No. 1, pp. 44-51, March, 1992.
- [Fokker 27] Fokker, A.D., "De Voortbeweging Der Slangen," *Physica*, 7, 65, 120. 1927. (In Dutch) Thanks to Chris Paredis for assistance in translation.

References

- [Gabrielli 50] Gabrielli, G., von Kármán, T.H., "What Price Speed?," *Mechanical Engineering*, 72: 775-781, 1950.
- [Gans 74] Gans, C., *Biomechanics: An Approach to Vertebrate Biology*, Lippencott, 1974.
- [Gans 84] Gans, C., "Slide-pushing: a Transitional Locomotor Method of Elongate Squamates," *The Structure, Development and Evolution of Reptiles* (M.W. J. Ferguson, ed), pp. 13-26, Symp. Zool. Soc. Lond., no 52, Academic Press, London, 1984.
- [Gans 85] Gans, C., "Motor Coordination Factors in the Transition from Tetrapody to Limblessness in Lower Vertebrates," *Coordination of Motor Behavior*, Bush, B.M.H., and Clarac, F., eds., Society for Experimental Biology Seminar 24. Cambridge University Press, 1985.
- [Golestaneh 84] Golestaneh, A., "Shape-memory Phenomena," *Physics Today*, April 1984, pp. 62-70.
- [Gothard 90] Gothard, Michael, "Tennessee Valley Authority Pipe Crawler," 13th Annual Utility/Manufacturers Robot Users Group Meeting, Chattanooga, TN, Jan 18, 1990.
- [Gowenlock 96] Gowenlock, P.G., *The LeTourneau Legend*, Paddington Publications, Pty. Ltd., Australia 4064, 1996, 294pp, ISBN 0 646 276921.
- [Gray 46] Gray, J., "The Mechanism of Locomotion in Snakes," *J. Exp. Biol.* 23:101-120. 1946.
- [Gray 50] Gray J., Lissman, H.D., "The Kinetics of Locomotion of the Grass Snake," *J. Experimental Biology*, 26:354-367.
- [Gray 68] Gray, J., *Animal Locomotion*, W.W. Norton Co., 1968, pp. 166-193.
- [Gregorio 94] Gregorio, P., "Design, Control and Energy Minimization Strategies for the ARL Monopod," MS Thesis Centre for Intelligent Machines, McGill University, Montreal, Canada. Also released as Technical Report TR-CIM-94-09.
- [Grillner 85] Grillner, S., "Neurobiological Bases of Rhythmic Motor Acts in Vertebrates," *Science*, Vol. 228, April 1985, pages 143-149.
- [Grundfest 94] Grundfest, W. S., Burdick, J.W., Slatkin, A.B., *Robotic Endoscopy*, US Patent 5337732, Aug 16, 1994. Assignee: Cedars-Sinai Medical Center, Los Angeles, CA.
- [Grunewald 84] Grunewald, P., "Car Body Painting with the Spine Spray System," 14th International Symposium on Industrial Robots, 7th International Conference on Industrial Robot Technology, Oct 2-4, 1984, Gothenburg, Sweden, pp633-639. (Spine)
- [Haddock 94] Haddock, Keith, "The Machines of R.G. LeTourneau," *Equipment Echoes*, Historical Construction Equipment Association magazine, Issue #35, December 1994, pp. 17-22.
- [Harris 78] Harris, J., "A Photoelastic Substrate Technique for Dynamic Measurements of Forces Exerted by Moving Organisms," *J. Microscopy*, vol. 114, pp. 219-228. 1978.
- [Heckrote 67] Heckrote, C., "Relations of body temperature, size and crawling speed of the common garter snake, *Thamnophis s. sirtalis*," *Copeia*, 1967(4):759-763.

- [Hill 65] Hill, *Pneumatic Self-Propelled Apparatus*, US Patent 3224734, Assignee: Her Majesty's Postmaster General, Dec. 21, 1965.
- [Hirose 83] Hirose, S., Nose, M., Kikuchi H., Umetami, Y., "Adaptive Gait Control of a Quadruped Walking Vehicle," *Robotics Research: The First International Symposium*. Edited by M. Brady, R.P. Paul, MIT Press, 1984, pp. 253-277. ISBN 0-262-02207-9.
- [Hirose 84] Hirose, S., "A Study of Design and Control of a Quadruped Walking Vehicle," *International Journal of Robotics Research*, V 3, No. 2, Summer 1984, pp. 113-133.
- [Hirose 90] Hirose, S., Morishima, A., "Design and Control of a Mobile Robot with an Articulated Body," *International Journal of Robotics Research*, vol.9, no.2, April 1990, pp. 99-114.
- [Hirose 93] Hirose, S., *Biologically Inspired Robots: Snake-like Locomotors and Manipulators*, Oxford University Press, 1993, ISBN 0 19 856261 6. (translation of 1987 Japanese edition)
- [Honda 96] Honda Motor Corporation, "Honda Biped Robot," News Release. Several news articles including Reuters from December 20, 1996.
- [Ikeda 87] Ikeda, H., Takanashi, N., *Joint Assembly Moveable Like a Human Arm*, US Patent 4683406, July 28, 1987. Assignee: NEC Corporation.
- [Immega 95] Immega, G., Antonelli, K., "The KSI Tentacle Manipulator," *IEEE International Conference on Robotics and Automation*, Nagoya, Japan, pp. 3149-3154, 1995.
- [Ishino 83] Ishino, Y., Naruse, T., Sawano, T., Honma, N., "Walking Robot for Underwater Construction," *Proceedings of 1983 International Conference on Advanced Robotics*, pp. 107-114. 1983.
- [igus95] igus Energy Chains, Catalog., igus inc., E. Providence, RI 02914.
- [Ipsen 60] Ipsen, D.C. *Units, Dimensions and Dimensionless Numbers*, McGraw-Hill, 1960.
- [Jayne 88] Jayne, B.C., "Mechanical Behavior of Snake Skin," *J. Zool., London.*, Vol 214, pp. 125-140, 1988.
- [Jones 33] Jones, A.T., "The Gliding of a Snake," *Physics*, 4:164-165, 1933. (Later *Journal of Applied Physics*.)
- [Kalhammer 95] Kalhammer, F.R., Kozawa, A., Moyer, C.B, Owens, B.B, "Performance and Availability of Batteries for Electric Vehicles: A Report of the Battery Technical Advisory Panel," Prepared for California Air Resources Board, Dec. 11, 1995
- [Keller 83] Keller, J.B., Falkovitz, M.S., "Crawling of Worms," *Journal of Theoretical Biology*, Vol. 104, No. 3, pp. 417-442, 1983.
- [Kelly 95] Kelly, S.D., Murrarby R.M., "Geometric Phases and Robotic Locomotion," *Journal of Robotic Systems*, Vol. 12, No. 6, pp. 417-431.
- [Klaes 90] Klaes, L., "The Rocky Soviet Road to Mars," *Electronic Journal of the Astronomical Society of the Atlantic*, (EJASA). Also found at <http://www.seds.org/pub/software/news/phobos.txt>.
- [Knowledge 97] Knowledge Revolution, *Working Model 3D V.2* , San Mateo, CA 1997

References

- [Kram 90] Kram, R., C.R. Taylor, "Energetics of Running: A New Perspective," *Nature*, Vol 346, 19 July 1990, pp 265-267.
- [Krotkov 92] Krotkov, E., Simmons, R., "Performance of a Six-Legged Planetary Rover: Power, Positioning, and Autonomous Walking," *Proceedings of the IEEE International Conference on Robotics and Automation*, May 1992, pages 169-174.
- [Leger 97] Leger, P.C., "Construction Toolkit," Thesis research CMU, 1997.
- [Lewis 92] Lewis, M., Fagg, A., Solidum, A., "Genetic Programming Approach to the Construction of a Neural Network for Control of a Walking Robot," *Proceedings of the IEEE International Conference on Robotics and Automation*, May 1992, pages 2618-2623.
- [Lissman 50] Lissman, H.W., "Rectilinear Motion in a Snake (*Boa Occidentalis*)," *J. Exp. Biol.* 26:368-379. 1950.
- [Lodder 83] Lodder, Christina, *Russian Constructivism*, Yale University Press 1983.
- [Maes 90] Maes, P., Brooks, R., "Learning to Coordinate Behaviors," *Machine Learning*, 1990, pages 796-802.
- [McMahon 83] McMahon, T.A., Bonner, J.T., *On Size and Life*, Scientific American Library, W.H. Freeman, New York. 1983. ISBN O-7167-5000-7.
- [McMahon 96] Personal conversation. 1996.
- [Mosauer 30] Mosauer, W., "A Note on the Sidewinding Locomotion of Snakes," *The American Naturalist*, Vol LXIV, March-April, 1930. pp.179-183.
- [Nakayama 88] Nakayama, R. et al., "Development of Mobile Maintenance Robot System 'AIMARS,'" *Japan Symposium on Flexible Automation*, 1988, p 645. (Toshiba)
- [NEC 96] NEC Corporation, "Orochi 12DOF Snake Like Robot," Press Release, NEC Corp. Melville, NY. Jan, 1996, 6 pp. Orochi was also featured in advertisements and several short mentions in popular magazines.
- [Niebur 91] Niebur, E., Erdos, P., "Theory of the Locomotion of Nematodes," *Biophysical Journal*, Vol. 60, November 1991, pp. 1132-1146.
- [Ngo 93] Ngo, J.T., Marks, J., "Spacetime Constraints Revealed," *SIGGRAPH 93*, Anaheim, CA August 1993. pp. 343-350.
- [Nilsson 95] Nilsson, M., Ojala, J., "'Self-awareness' in Reinforcement Learning of Snake-like Robot Locomotion," *Proceedings of the Third IASTED International Conference on Robotics and Manufacturing*, June 14-16, Cacun, Mexico, pp. 244-247.
- [Nilsson 97a] Nilsson, M., "Ripple and Roll: Slip-free Snake Robot Locomotion," *Proceedings Mechatronic Computing for Perception and Action, MCPA '97*, Pisa, Italy, Feb. 1997.
- [Nilsson 97b] Nilsson, M., "Snake Robot Free Climbing," *International Conference on Robotics and Automation, ICRA '97*, April 1997, Albuquerque, NM.
- [Novel 97] Novel Electronics, St. Paul, MN. "Emed and Pliance measurement system for dynamic pressure distribution," Brochures, 1997.
- [Odetics 88] Odetics, *Articulated Transporter/Manipulator System*, Video, Anaheim, CA. 1988.

- [Ogata 78] Ogata, K., *System Dynamics*, Prentice-Hall, Englewood Cliffs, NJ. ISBN 0-13-880385-4. 1978.
- [Olympus 94] Olympus, Brochures and conversations with sales representatives. 1994.
- [Ostrowski 96] Ostrowski, J.P., *The Mechanics and Control of Undulatory Locomotion*, PhD Thesis, California Institute of Technology, 1996.
- [Paap 96] Paap, K.L., Dehlwisch, M., Klaassen, B., "GMD-Snake: A Semi-Autonomous Snake-like Robot," 3rd International Symposium on Distributed Autonomous Robot Systems (DARS 96), October, 29-31, 1996, RIKEN, Saitama, Japan.
- [Parisi 97] Parisi, P., "Snake Charmers," Cinefex, No. 70, June 1997, pp. 61ff. [Anaconda movie special effects]
- [Parker 63] Parker, H.W., *Snakes*, Robert Hale Ltd. London, 1963.
- [Petrucci 94] Petrucci, R., Simmons, K., "An Introduction to Piezoelectric Crystals," *Sensors*, May 1994, pp. 26-31.
- [Plustech 96] Plustech Oy, Finland, Brochure, 1996.
- [Rice72] Rice, Richard, "System Energy and Future Transportation," *Technology Review*, January, 1972, pp. 31-37.
- [Roberts97] Roberts, T.J., Marsh, R.L., Weyand, P.G., Taylor, C.R., "Muscular Force in Running Turkeys: The Economy of Minimizing Work," *Science*, V 275, Feb. 21, 1997, pp. 1113-1115.
- [RSI 94] RSI, *Miniature Pipe Crawler*, RSI Research Inc, Sidney, B.C. Canada, Brochure, 1994.
- [Sasaki 85] Sasaki, T, Et al, "Application of Self-Propelled Devices to In-Piping Inspection Technology," *Robot* No. 51, 1985 (in Japanese).
- [Schmidt-Nielsen72] Schmidt-Nielsen, K., "Locomotion: Energy cost of swimming, flying, and running," *Science*, N.Y., 177, pp. 222-228.
- [Schmidt-Nielsen 84] Schmidt-Nielsen, K., *Scaling: Why is Animal Size so Important?*, Cambridge Press, New York, 1984.
- [Schneider 91] Schneider, E.T., *Remote Controlled High Force Actuator*, US Patent 5025627, Jun. 25, 1991.
- [Schneider 93] Schneider, E.T., *Thermochemical Actuation Method and Apparatus*, US Patent 5177969, Jan. 12, 1993.
- [Schneider-j 95] Schneider, J.G., *Robot Skill Learning Through Intelligent Experimentation*, PhD Thesis, University of Rochester, 1995
- [Schneider-j 94] Schneider, J., Gans, R.F., "Efficient Search for Robot Skill Learning: Simulation and Reality," *International Conference on Intelligent Robots and Systems*, 1994
- [Scott 96] Scott Edwards Electronics, *MiniSSC Board*, Sierra Vista, AZ, 1996.
- [Secor 92] Secor, S.M., Jayne, B.C., Bennett, A.C., "Locomotor Performance and Energetic Cost of Sidewinding by the Snake *Crotalus Cerastes*," *J. Exp. Biology*, 163, pp. 1-14, 1992.

References

- [Shan 92] Shan, Y., et al., "Obstacle Accomodated Motion Control of a Planar Manipulator," Proceedings of the International Conference Contr. Robotics, Vancouver, BC, Canada, Aug 4-7, 1992, pp. 62-65.
- [Shan 93a] Shan, Y., et al., "Design and Motion Planning of a Mechanical Snake," IEEE Transactions on Systems, Man, and Cybernetics, Vol 23. No. 4, July/August 1993, pp. 1091-1100.
- [Shan 93b] Shan, Y., *Robot Obstacle Accomodation: Mechanics, Control and Applications.*, PhD Thesis, The University of Michigan, September, 1993.
- [Siemans 97] Siemans Hearing Instruments, Piscataway, NJ 08855. Users of finely stranded wire. Product specialist sent me wire.
- [Sims94a] Sims, K., "Evolving Virtual Creatures," Computer Graphics, Annual Conference Series, SIGGRAPH '94 Proceedings, July 1994, pp. 15-22.
- [Sims94b] Sim, K., "Evolving 3D Morphology and Behavior by Competition," Artificial Life IV Proceedings, ed. by R. Brooks and P. Maes, MIT Press, 1994, pp. 28-39.
- [Starsys 95] Starsys Research, *Paraffin Actuators*, Boulder, CO, 1995.
- [Stix 92] Stix, G., "Micron Machinations," Scientific American, November, 1992, pp106-117.
- [Sturges 91] Sturges R.H., et al., "A Flexible Tendon-controlled Device for Endoscopy," Proceedings of the 1991 IEEE International Conference on Robotics and Automation, Sacramento, CA, April 1991.
- [Sturges 93] Sturges, R.H., Laowattana, S., "A Flexible Tendon-Controlled Device for Endoscopy," International Journal of Robotics Research, Vol. 12, No. 2, April 1993, pp. 121-131.
- [Sukhatme 95] Sukhatme, G.S., Lewis, M.A., Bekey, G.A., "Mission Reachability for Extraterrestrial Rovers," IEEE International Conference on Robotics and Automation, 0-7803-1965-6/95, 1995. pp.1964-1969.
- [Tcam 95] Tcam Technologies, *TCAM Solid State Actuators*, Eastlake, OH, Brochure, 1995.
- [Toshiba 89] *Robotics for Nuclear Facilities*, Toshiba Brochure, 1989.
- [Triantafyllou 95] Triantafyllou, M.S., Triantafyllou, G.S., "An Efficient Swimming Machine," Scientific American, March 1995, pp. 64-70.
- [Tucker 69] Tucker, V.A., "Energetic Cost of Locomotion in Animals," Comp.. Biochem. Physiol., 1970, Vol. 34, pp.841-846.
- [van de Panne 93] van de Panne, M., Fiume, E., "Sensor-Actuator Networks," SIGGRAPH 93, Anaheim, CA August 1993. pp. 335-342.
- [VIT 95] VIT, *Rovver and PipeCat*, Visual Inspection Technologies, Flanders NJ, Brochures, 1995
- [Waldron 81] Waldron, K., Kinzel, G., "The Relationship Between Actuator Geometry and Mechanical Efficiency in Robots," 4th CISM-IFTOMM Symposium on the Theory and Practice of Robots and Manipulators, PWN, Warsaw, 1983, pp. 366-374.

- [Waldron 84] Waldron, K.J., Pery, A., McGhee, R.B., "Configuration Design of the Adaptive Suspension Vehicle," *International Journal of Robotics Research*, V 3, No. 2, Summer 1984, pp. 37-48.
- [Waldron 97] Waldron, K.J., Personal conversation, 1997.
- [Walton 90] Walton, M., Jayne, Bruce C., "The Energetic Cost of Limbless Locomotion," *Science*, Aug 3, 1990., pp. 524-527.
- [Welch-Allyn 94] Welch-Allyn, Brochures and conversations with engineers. 1994.
- [Westinghouse 97] Westinghouse Electric Corporation. Personal conversation with managers in the Science and Technology Center in Pittsburgh. 1997.
- [Wettergreen 96] Wettergreen, D., *Robotic Walking in Natural Terrain*. Ph.D. Dissertation in Robotics, Carnegie Mellon University, May 1996.
- [Wiederick 89] Wiederick, H.D., Sherrit, S., Mukherjee, B.K., "Piezoelectric and Electrostrictive Actuators," Second Workshop on Military Robotic Applications. August 1989, pp. 214-217.
- [X10 97] X10 (USA) Inc, Seattle WA, 98188. See also <http://www.x10.com>.
- [Yim 94] Yim, M., "New Locomotion Gaits," *IEEE Robotics and Automation Conference*, 1995, pp. 2508-2514.
- [Zimmer 94] Zimmer, C., "See How They Run, Discover," September 1994, pp. 65-73.

