The Interplay between Mathematics and Robotics

Summary of a Workshop

National Science Foundation
Arlington, Virginia
May 15, 16, and 17, 2000

Preface

In the last two decades the field of robotics has largely been developed within the computer science and engineering communities. However, many open problems of great practical significance remain that require the application of modern mathematical and computational tools. Members of the relevant scientific communities and several National Science Foundation programs share common interests in exploring potential advances in these directions and recognize the advantages in bringing together people from diverse areas to contribute to the development of fundamental research in these disciplines. Some initial collaborative efforts in recent years demonstrate that the time is right for sustained and intensive activities between the two groups of researchers.

In order to stimulate fruitful interdisciplinary investigations, the Divisions of Mathematical Sciences, Information and Intelligent Systems, and Civil and Mechanical Structures of the National Science Foundation co-sponsored a workshop on the interplay between mathematics and robotics which was held at the National Science Foundation on May 15, 16, and 17. We the participants identified fundamental and significant areas of robotics research and the mathematical tools necessary to help solve problems therein. A summary of our findings, the workshop agenda, and the list of attendees form the basis for this report.

James S. Albus, National Institute of Standards & Technology
Mihai Anitescu, University of Pittsburgh
Anthony Bloch, University of Michigan
Roger Brockett, Harvard University
Joel Burdick, California Institute of Technology
Greg Cherikjian, The Johns Hopkins University
Mike Erdmann, Carnegie Mellon University
Robert Ghrist, Georgia Institute of Technology
Leo Guibas, Stanford University
John Harer, Duke University
Roberto Horowitz, UC Berkeley
John Horst, National Institute of Standards & Technology
Steve Kaufman, Sandia National Laboratories
P.S. Krishnaprasad, University of Maryland
Vijay Kumar, University of Pennsylvania
Ming Lin, University of North Carolina
Matt Mason,Carnegie Mellon University
Jim Milgram, Stanford University
Jean Ponce, University of Illinois, Urbana-Champaign
Florian Potra, University of Maryland Baltimore County
Daniela Rus, Dartmouth University
Ed Scheinerman, The Johns Hopkins University
Ray Sterling, Louisiana Tech University
David Stewart, The University of Iowa
Russell Taylor, The Johns Hopkins University
Sebastian Thrun, Carnegie Mellon University
Dawn Tilbury, University of Michigan
Jeffrey C. Trinkle, Sandia National Laboratories
Louis Whitcomb, The Johns Hopkins University

Table of Contents

The Vision: Objective, Difficulties, and Possible Directions

Real-world applications of intelligent machines involve scientific, mathematical, and engineering problems with enormous practical, theoretical, and economic interest. The ultimate goal of current robotics research is the creation of physical machines with near or even beyond human levels of perception, physical capability, practical intelligence, and behavior, as well as the creation of intelligent and efficient human interfaces to these complex systems. Indeed, in some areas, such as computer aided surgery, human capabilities can be surpassed.

The wider deployment of robots awaits further advances in basic unsolved problems of motion planning and contact, particularly in the area of robust solutions. Here are some of the dramatic recent successes as well as areas where the potential importance of robotics is clear:

- Industrial robots and computer vision comprise a $2.8 billion industry in the U.S., growing at more than 20% annually. In certain applications, such as automobile painting and welding and electronic assembly, robots are pervasive. Predictions as recently as 1980 foresaw even faster growth for this sector, but, technical difficulties such as a lack of standards have imposed barriers to wider adoption.
- The rapid pace with which the human genome was sequenced depended entirely upon robotics and computing technology. Automation and robotics will clearly play a role in future biological and chemical laboratory research and clinical practice.
- Underwater robots have led to spectacular discoveries (e.g., the Titanic exploration and recovery). Robots promise to extend human capability in many other hazardous domains.
- The field of computer-integrated surgery (CIS) and medical robotics is poised to undergo exponential growth. CIS will fundamentally alter medical practice, as surgeons will use CIS to carry out surgical interventions that are more precise and less invasive. Health care currently makes up 14% of US GDP. The complete adoption of minimally invasive medicine, as enabled by CIS, will save 0.75% of GDP alone.
- More generally, robotics can widely impact quality of life, serving as assistants to the elderly or incapacitated. The cost of providing care to a quadriplegic, for example, is well over $75,000 per year. We are learning how to build and program robots that can co-exist and cooperate with their human owners and this will completely change the paradigm for care giving.
- Robotics will also play a crucial role in future efforts to upgrade and maintain U.S. infrastructure:
    - Intelligent transportation initiatives will increase throughput on crowded highways.
    - Two million miles of underground piping are in need of inspection and repair. Much of this infrastructure can only be accessed by robots.
    - Thousands of crumbling bridges and nuclear and chemical pollution sites must be repaired and cleaned-up. The logical mechanisms for doing this are robots.
- In the future, robotic technology will serve as an interface between the emerging distributed computing and information networks (e.g. the World Wide Web) and the physical world, particularly household and business environments. The Web will have a physical presence largely provided by robotics.

The true potential for robotic technology is even greater than the areas outlined above. There are two emerging technologies that will have a dramatic impact on future robots --- their form, shape, function and performance --- and change the way we think about robotics. First, advances in MicroElectronicMechanical Systems (MEMS) will enable inexpensive and distributed sensing and actuation. Just as nature provides complex redundant pathways in critical processes (e.g., synthesis of biomolecules and cell cycle control) to combat the inherent noisiness in the underlying processes and the uncertainty in the environment, we will be able to design and build robots that can potentially deal with uncertainty and adapt to unstructured environments. Second, advances in biomaterials and biotechnology will enable new materials that will allow us to build softer and friendlier robots, robots that can be implanted in humans, and robots that can be used to manipulate, control, and repair biomolecules, cells, and tissue.

Realistically, we are far from realizing this tremendous potential. While some of the obstacles are technological in flavor (for example, lack of high strength to weight ratio actuators, or lack of inexpensive three-dimensional vision systems), there are several obstacles in robotics that stem from our lack of understanding of the basic underlying problems and the lack of well-developed mathematical tools to model and solve these problems.

There is a tradition of mathematicians working with roboticists. Many problems in robotics, or in the disciplines that are core to what we call robotics, have attracted mathematicians to this field. As far back as the 19^{th} century, algebraic geometers like Kemp and Tschebyshev were drawn to the beautiful mechanics of linkages (the predecessors of today's complex articulated manipulators) and geometers like Poincare were attracted by the dynamics of machines. Clearly, the development of the foundations of many classical and modern mathematical tools was spurred by technological advances in machines and mechanisms. This tradition has continued and there are several examples in the last two decades where some of the most fundamental and enduring results in robotics have come from mathematicians and their interactions with engineers. The classical general motion planning results in robotics were developed by Schwartz and Sharir, which in turn were improved on by Canny using methods from algebraic geometry. Similarly, Milgram and Trinkle have used results from modern algebraic topology to obtain an improved understanding of the configuration spaces of closed chain mechanisms, leading to improved algorithms for motion planning for such systems. Marsden, Brockett, and Sastry have used differential geometry and Lie theory to formalize the kinematics, dynamics and control of spatial linkages. This work has also sensitized the engineers to such important mathematical ideas as frame invariance and invariance with respect to parameterization. Other areas where there are strong developing ties between mathematics and robotics include, for example, Bayesian statistics to develop algorithms for perception and learning, nonsmooth optimization techniques for parameter optimization and optimal control of mechanical systems, the use of nonsmooth analysis for developing simulation methods for mechanical systems with contact or impact, and partial differential equation formulations of image segmentation and pre-processing problems.

While the interactions above suggest that robotics offers a fertile ground for interactions between mathematics and engineers, there are many significant challenges that require such interactions to be taken to a new level. The next section outlines and expands on the specifics of some technical areas which need immediate attention. But it is important to realize that we can ***not*** simply rely on the traditional mode of interaction in which roboticists become users of mathematical tools and results obtained from the literature. The challenges of the future will require active and equal participation on the part of the mathematics community. While robotics will clearly benefit from an effort to bring greater and closer collaboration, mathematics will also be positively affected by greater interaction. The needs of robotics will suggest new directions for mathematics research. Communication between mathematicians and roboticists will affect the questions asked, as much as the answers that are given. Mathematical tools are not necessarily a ``silver bullet" for robotics practitioners; rather the shared perspectives of these two communities are as important as the mathematical tools that are transferred.

The main thesis of this position paper is simple. If we do not create and foster an environment where mathematicians and roboticists work hand in hand to solve some of the basic problems, society will not likely reap the full benefits that the robotics technology can potentially offer.

There is a second important benefit of investment in this collaboration. Anecdotal and hard evidence both point to a world-wide brain drain from mathematics (and other theoretical disciplines) to the less technically sophisticated but more glamorous areas. This brain drain has serious adverse effects toward the long-term scientific developments for the nation. Engaging the mathematics community in robotics will provide an effective approach to attract students into the field of mathematics, thereby building and sustaining the scientific infrastructure in the country. While robotics is recognized to be an excellent vehicle for education and involving students in science and engineering, the new level of interaction envisioned here will put mathematics on an equal footing with engineering and technology and motivate promising students to pursue career tracks in mathematics and its applications.

Some Relevant Mathematical Techniques

The following mathematical disciplines are likely to have strong relevance for robotics. Along with the disciplines, we have listed in some cases the potential application areas or subtopics worthy of special consideration.

1. Algebraic and differential topology. Current work includes the application of techniques from loop space theory and low dimensional topology to understand configuration spaces of many-particle/many-body systems, as well as techniques from global differential topology to analyze configuration spaces of arbitrary length closed chains in two and three dimensions with spherical, revolute, and prismatic joints. Similar tools are also directly relevant to:

- Distributed sensing and actuation systems, such as those made possible with MEMS and nanotechnology; and
- High dimension design problems, such as intelligent vehicle-highway and air-traffic management systems.

2. Dynamical systems theory. There has already been a fruitful interaction between the dynamical systems community and robotics. Geometric mechanics has provided many insights enabling roboticists to understand robotic locomotion and manipulation and build novel systems. However, there is a need for new tools that marry the dynamical systems theory to techniques in discrete mathematics. For example, networks of cooperating robots can be viewed as dynamical systems on graphs. There is an increasing awareness of the hybrid nature of artificially engineered (and naturally occurring systems) and the fact that discrete and continuous dynamics interact in complex ways. The mathematics of such hybrid systems is not well developed. Issues like robustness to modeling uncertainty and to noise, and sensitivity analysis, are not well understood. There is no systematic approach to developing abstractions of dynamical systems and to decompose a dynamical system into a multi-scale hierarchy of subsystems. An emphasis on this set of issues will enable:

- Design and control of networks of coordinating, communicating machines for a wide range of application domains;
- Multi-scale, distributed simulation of complex systems ranging from molecular and cellular networks to networks of interacting robots sweeping mine fields or performing assembly tasks; and
- New designs for agile, dexterous robots.

3. Optimization algorithms. Robotic system design and many problems in robot task planning can be formulated as optimization problems, though they are typically ``hard" in terms of complexity and lack of readily recognizable or standard mathematical structures. Success stories include graph-theoretic and calculus of variation based approaches to determining optimal paths, randomized algorithms for finding solutions in complex spaces, optimal feedback control policies for a range of robotic tasks, and saddle-point policies for solving differential games of pursuit and evasion. However, most optimization problems are not well-posed in a number of ways. First, the domains are usually non convex and even non smooth. Second, the solutions to optimization problems may require measure-theoretic tools from real analysis, and implementations of these solutions may provide constraints that are not easily incorporated. A concentrated effort in this direction will yield new methods for developing:

- Optimal control and decision making policies; and
- Algorithms for scheduling processes, allocating resources, and decomposition of tasks and organization of teams of robots in a wide range of applications.

4. Combinatorics. Many robotic design and task planning problems involve aspects that are discrete. Problems, such as motion-planning, involve analysis of a large, continuous configuration space. Often the

configuration space is approximated by a graph, and then discrete techniques are applied. The combinatorial complexity associated with these problem formulations is often staggering. Combinatorial mathematics is directly relevant to areas such as:

- Modular robots: These are robots built of many, identical modules. Such devices can reconfigure themselves to suit their environments. The configuration space of such robots is inherently discrete.
- Discrete actuators: Smoothly actuated robotic arms (and other manipulators) can be replaced by a cheaper network of discrete actuators (devices that extend or contract into only two positions).
- Discrete sensing: Instead of determining robot position and configuration by a raster image (essentially, a continuous image of the environment), faster and cheaper robot localization can be achieved by using a modest number of discrete sensors.

5. <u>Differential algebraic inequalities</u>. Discrete and continuum mechanics play an important role in the modeling of multibody systems in contact, which in turn are central to robot manipulation. Traditionally, these systems are governed by differential algebraic equations (DAEs). In order to be able to model unilateral obstacle constraints, these DAEs have to be augmented by inequalities. The resulting differential algebraic inequalities are special instances of differential inclusions; whose computations and analysis urgently require new mathematical theories and numerical methods:

- Qualitative behavior which is useful for modeling robustness issues;
- Simulation with uncertainties; and
- Nonsmooth analysis.

6. <u>Statistical learning theory</u>. Statistical approaches to robotics explicitly represent a robot's uncertainty in perception and action selection. As a result, they are highly robust to uncertainty and dynamic environments. As robots move away from factory floors into environments populated with people, the need to cope with uncertainty is enormous. In the past few years, a range of hard robotics problems has been solved using statistical algorithms, such as the problem of mobile robot localization and mapping. However research is needed along multiple dimensions to develop more robust and more efficient algorithms. In particular, we need research on basic representations of uncertainty in robotics domains, along with fast algorithms for reasoning with uncertainty. We need better ways of integrating learning into probabilistic robotics. Basic theoretical work is needed to broaden the mathematical frameworks under which today's algorithms operate. And finally, there is a need to develop programming tools that facilitate the development of statistical algorithms for mobile robots.

Recommendations for Collaborations

Because of the very different cultures, training, and objectives of mathematicians and roboticists, collaborations between them are not always straightforward. Various obstacles exist that need to be overcome in order to foster successful collaborations. The NSF can provide a major service through both educational and research incentives to enable such interactions.

One of the largest obstacles is that collaborative research across fields requires a certain amount of broadening and interdisciplinary thinking, which is perceived as diffusing the focus and reducing the research depth of an individual research career. Although robotics can give good examples of real mathematical objects that can be studied, and will lead to new and fruitful areas of mathematics research, it does not fit well into the traditional mathematics research paradigm.

On the one hand, junior mathematics faculty, during their pre-tenure years, are reluctant to collaborate with engineers because the current criterion for achieving tenure is largely based on the quantity and quality of new research results achieved in the first three years after obtaining the Ph.D., or on exemplary teaching during the same period. Essentially no young mathematics Ph.D.'s, during this period, are willing to take the time needed to learn a new field, or attempt to apply existing mathematics in an area such as robotics.

On the other hand, it is difficult for engineering students to master the mathematics currently needed in robotics. Engineering students in the U.S. are taking less mathematics than their foreign counterparts, making research activities which rely on mathematical techniques difficult. Compounding the difficulty, many graduate mathematical curricula are either not well received or not perceived as relevant by engineering students. Thus, engineering students who do have an interest in learning modern mathematics are faced with a discouragingly steep learning curve. Notable exceptions exist, such as courses in differential geometry applied to dynamical systems (Marsden at UC-Berkeley/Caltech, Bloch at U-Michigan) which have been taken by many robotics and control students, and have fostered new research topics and new solutions to existing problems in these areas.

Active collaborations between mathematicians and roboticists, and cross-fertilization between the two cultures is critical to the nation's research and technological infrastructure. The National Science Foundation has a responsibility to create and foster a research and educational environment where such collaborations can be nucleated, nurtured and grown. We see two different ways in which NSF can help further these objectives, both by establishing new programs that will directly support collaborations in mathematics and robotics and by encouraging and fostering collaborations through existing mechanisms at NSF.

New Programs:

1. *Research funding*: One of the best top-down ways to encourage the development of a research community in a specific area or to focus an existing group in a particular direction is by creating appropriately tailored research initiatives. One good example is the Bioanalysis and Control program started approximately 10 years ago at the NSF. This program helped sensitize engineers to problems in biology requiring analytical and computer modeling and biologists to engineering methodologies and problem solving tools. We envision an initiative that encourages collaborations between mathematicians and robotics and addresses the use of novel mathematical techniques in robotics. Such initiatives are best guided by panels that consist of established researchers in both communities, which in turn will further help shape the research and educational agenda.

2. *Promoting educational and research interactions*: NSF funds workshops and conferences that help shape the national research and educational agenda. We recommend workshops that focus on mathematics and robotics that may be attached to bigger international conferences. These can be one or two day workshops. An example of a conference that can be affiliated with such workshops is the International Conference on Robotics and Automation. Such workshops will attract main-stream mathematicians to main-stream robotics conferences and promote the kinds of interactions that are essential. Obviously, a comparable effect will be achieved by hosting such a workshop at a SIAM conference or perhaps to better alert the pure mathematics community to these questions, at one of the major meetings of the AMS. Similarly, NSF can promote travel grants that will encourage researchers and educators to attend conferences in other areas.

3. *Visiting professorships*: NSF should, perhaps through existing mechanisms, encourage visiting professorships. The IGMS program pioneered by DMS inside the NSF represents one opportunity while another viable model would be to follow the NIH model that funds engineers to work in laboratories of biologists.

4. *Graduate education*: NSF has existing mechanisms that focus on education, but this focus is primarily on K-12 and undergraduate education. While this is important in every discipline, in nascent cross-disciplinary areas, it is better to focus such efforts on graduate and post graduate education. The programs that we are aware of that focus on graduate education, such as IGERT and VIGRE tend to support individual graduate students for extended periods of time, as well as creating one and two year post doctoral positions. To achieve the more broadly based exposure that we need, we also recommend that the NSF fund special short term programs. Such efforts could include:

- Short courses that train mathematicians in the methods and problem areas of roboticists; and
- Short courses that allow roboticists to overcome the steep learning curve that is encountered in subjects like topology and real analysis.

As this program matures, it may be a good idea to expand on the IGERT model and fund a number of fellowships specifically targeted at robotics students, to allow them to take a year or two off to become trained in mathematical techniques. This last item may be particularly relevant in today's research environment as many federal research grants in engineering are of sufficiently short duration as to place the investigators constantly under the pressure to generate results. Thus, these researchers do not have the luxury of encouraging graduate students to learn and become familiar with new ideas in mathematics.

Existing Programs:

There are many existing programs at NSF that could benefit from a synergy between mathematicians and roboticists. For example, the high-profile Information Technology initiative, where areas such as software engineering and human computer interfaces are publicized, does not highlight robotics and mathematics. The broader agenda of computer science is centered around computing, reasoning and interacting with the real world, which is exactly the objective in robotics. Similarly, another existing program with great potential in this regard is the Integrated Graduate Education, Research and Training (IGERT) initiative. This is a great opportunity that allows the training and education of high-caliber students without being subject to the pressures of research funding. In these and other similar programs, the main responsibility to tap into such resources lies with the community. However, NSF can be proactive in sensitizing the community to these opportunities.

Similarly, there are existing mechanisms for creating workshops. A variety of workshops are currently conducted at MSRI, UCLA, Berkeley, and Minneapolis, and it would be useful to create sessions that focus on robotics and mathematics. Once again, NSF can be proactive in advertising and encouraging the community to organize workshops and should offer targeted funding.

In conclusion, we believe NSF can adopt a two-pronged strategy in this area. One of these prongs consists of new mechanisms to promote interactions and the second prong involves the promotion of existing mechanisms. Ultimately, such initiatives can and will lead to intellectual and cultural benefits that will invigorate mathematics and robotics research and education with corresponding benefits for the nation as well.

Appendix

This appendix summarizes the discussions that occurred in a number of different working sessions. The format of these appendices varies, as the nature of the discussion varied widely by topic.

Discrete and Computational Geometry

We started our discussion by describing research problems in robotics that require discrete geometry and topology.

Motion Queries, Planning and Simulation: Some research challenges in modeling of motion for robots include proximity queries, motion and or manipulation planning, simulation of a dynamical system, motion generation and execution, acquisition of kinematic structures from motion sequences, motion capture and display of changing environments (visual, haptic and auditory rendering). These are some fundamental problems in algorithmic robots and much has been studied. However, there are still several open research issues that discrete geometry and topology that can potentially advance the state of arts significantly.

Proximity queries are ubiquitous in many robotic tasks and numerous other scientific and engineering applications. Performance of proximity queries algorithms is dependent on the combinatorial complexity of the input and the size of output. One possible approach is to simplify the geometric representations of the models, thereby reducing the intrinsic complexity of the problem. Can we design a new class of algorithms based on multi-resolution (or progressively simplified) representations of the models to accelerate the performance of proximity queries? Some real challenges include the design of the algorithms for complex geometry that is nonconvex and nonrigid. This line of research requires multi-resolution analysis and approximation theory, in addition to discrete geometry.

Analysis and understanding of configuration space for motion planning requires topology. Motion planning has already been discussed in the earlier presentations by Joel Burdick, Leo Guibas and Jeff Trinkle. We have a very fast algorithm that can exploit graphics hardware to compute the generalized Voronoi diagram of relatively complex (i.e. in the range of hundreds of thousands of primitives) environments of arbitrary geometry. We are interested in how we can use the Voronoi skeletons of the 3D workspace, the dimensions of the robot and the narrow passages, and the kinematics of the robot to design a complete and fast roadmap method for real-time planning in a dynamic or partially known environment. Other extensions include planning for non- rigid robots, assembly planning of MEMS and nano-structures, planning and coordination of multiple autonomous agents, planning with constraints, very high degrees of freedom, uncertainty, etc.

Dynamic simulation of rigid and non-rigid bodies still remains a challenge for modeling, analysis and understanding of physical interaction between the robot (in a broad sense) and the environments. Modeling of soft and flexible bodies may be of special interests to the design and planning of medical robots, as well as surgical training using haptic interface. Mathematics required includes discrete geometry for proximity queries (esp. dealing with nonconvex and deformable bodies), numerical analysis, constrained non- linear optimization techniques, etc.

Sensor Placements: The problem of sensor placements requires configuration graphs, linear programming, finding shortest paths (such as Dijkstra's algorithms). Certain situations you can use shortest path relatively easily with N-bit factor. However, there are some enormously large graphs with certain amount of structures, where you don't want to explore the entire graphs. One possible approach is getting an imbedded graph that can give you a near-optimal solution.

What characteristics in a graph structure differentiate one graph from others? If the graph is completely unknown, then the shortest path approach is the only one. Probabilistic techniques should also be considered.

Self-Configurable Robots: Counting polyominoes: Self-configurable robots can change shape based on sequences of simple translational and rotational motion. Self-configurable robots require planning of motion sequences to change from configuration i to configuration j. For lattice-based self-configurable robots, the configurations $i$and j can be thought as finite connected subgraphs of some particular lattice graph. These subgraphs are called polyominoes. We need to study complexity of such a problem. That is how many possible ways can we possibly have to reconfigure the robots so that we can change its shapes between one configuration to another. While the known algorithms requires $O(n^3)$ time for counting polyominoes where n is the number of modules, people don't know how to count the exact number for $n>30$. Scalability and parallelism are probably the keys. This may require graph theory and other discrete geometry tools.

Planning algorithm for modular self-reconfigurable robots: We need to consider the size of obstacles and robots for using self-reconfigurable robots for planning. The current system (at Darmouth) is lattice based. How about simpler prototype for basic understanding, like ameba-like or hexagonal modules. We need to study shape representations in three dimensions for better, general understanding. We are looking at a lower degree-of-freedom system to minimize the mechanical/engineering complexity.

We also have a project where we wrap a rope for a stack of cubes and move the cubes by pulling the robe. We would like to derive an analytical solution for describing the perimeter of this rope which forms the convex hull of the given collection of cubes.

Manipulation in unknown environments: Some examples include rescue missions for plane crash, bomb squad mission, and working in hazardous environments. When sensors are given incomplete information, how do we plan manipulation? We need human supervision. Sending coordinated groups of robots/vehicles for hazardous missions is another problem. Task-level programming is hard. Can some type of mathematical formalism be useful? Sensor design/cost/etc. need to be considered.

Motion Planning and Kinematics: Motion planning is exponentially hard, so there is the so-called PRM that seems to work well practice in some scenarios, but not in some cases. So, we probably need to look at lower dimension problems to find a complete solution that is output sensitive. Projection down to lower dimensions recursively and topological analysis may be the keys. Other possibilities may involve robot kinematics maps to some sort of graphs and the connectivity of the configuration space is encoded in some sort of graphs.

Topology

Overview: Interaction with mathematicians may or may not yield the precise tools needed to solve a given robotics problem. However, it is almost certain that interaction will change the perspectives that robotics researchers bring to the table. Mathematics will affect the question as much as the answer. Each of the disciplines discussed in this section possesses one or more natural features of immediate relevance to solving real-world problems: *Topology* is implicitly robust and global; *Dynamics* links topology to the concepts of stability and attraction; *Geometry* is perfectly suited to answer questions about efficiency and coordinates. Given the nature of applications desired, it is important to emphasize the computational aspects of the methods used. In particular, ties with the inchoate discipline of computational topology should be established.

Specific mathematical ideas/tools: The following ideas and tools have potential to be useful in new ways to robotics research.

- Path/loop spaces --- and related homotopy-theoretic tools --- should be helpful in C-space analysis.
- Morse theory in all its various permutations:
  - Morse and stratified Morse theory are already of use in C-space analysis to a limited extent. One goal is to increase the utility of these powerful subjects.
  - Conley-Morse theory is proving to be increasingly useful in dynamics applications. Perhaps there are problems in robotics for which these techniques would be of help.
- Geometric group theory and non-positively curved space technology can generalize hyperbolic geometry to groups and metric spaces. Can these help in C-space problems and path-planning? These ideas (NPC, CAT(0), etc.) all are tied to problems of computability (e.g., solving the word problem in group theory, etc.) --- this may be of help in path planning problems. Automatic groups should be of use in manipulation and planning.
- Gromov's perspectives on DAE/DAI type problems uses sheaf and homotopy theoretic ideas, and has its roots in prolongation methods that go back to the turn of the century. For many partial differential relations these ideas (an “h-principle”) guarantee solutions up to homotopy. These ideas for deforming an approximate solution to a true one should be relevant to other DAE/DAI techniques and problems.
- Combinatorial differential topology: There are precise notions of differential objects (vector fields, Morse functions, differential forms) for combinatorial cell-complexes. These are inherently well-suited to computation yet accurately capture the topology (homology/cohomology) of a complex. This may be one method for using smooth techniques while storing all data combinatorially.
- Computational homology: Several groups are working on fast algorithms for computing homology of maps on simplicial complexes. These have the promise of working in relatively high dimensions, or, at least, in real-time for low dimensions. Such tools should certainly be of use to vision problems (e.g., checking connectivity or changes in global structure, detecting occlusions, etc.), and hopefully will be of use in other robotics problems.

Distributed robotics

Distributed robotics is now where robotics was in the 1970s: algorithms are developed in an architecture-dependent way on a task by task basis, with no formal foundations. In the early 1980s the realization that uncertainty was the fundamental obstacle in robot control lead to the invention of configuration space, which caused a qualitative leap in the advancement of robot science.

We are at a similar crossroads in robotics today. As sensors and actuators are becoming smaller and cheaper, teams of robots working together are becoming more pervasive. Cooperative robotics has the potential of expanding greatly the application domains of robots but remains a largely unexplored field. In addition to all the challenges of single robot systems, cooperative robotics has the added difficulties of cooperation and communication, and of combining discrete and continuous systems. This has created a great opportunity for the development of a mathematical basis for distributed robotics that is grounded in engineering issues. In computer science, one often learns a lot about the structure of an algorithmic problem by parallelizing it; a similar methodology may be useful in robotics.

The three most important challenges in distributed robotics, where mathematics is likely to make significant advances are:

- Developing formal models that allow principled comparisons between distributed robot algorithms and give performance guarantees;
- Developing dynamic control for non-smooth systems; and
- Developing methods to characterize the power of modular robot systems and of matching structure to task.

Formal Models for Distributed Robotics: A key challenge in distributed robotics is the generation of distributed robot algorithms with performance guarantees, which is difficult. Structured environments, such as those found around industrial robots, contribute towards simplifying the robot's task because a great amount of information is encoded, often implicitly, into both the environment and the robot's control program. These encodings (and their effects) are difficult to measure. A possible solution is to quantify the information encoded in the assumption that (say) the mechanics are quasi-static, or that the environment is not dynamic. In addition to determining how much "information" is encoded in the assumptions, we may ask the converse: how much ``information'' must the control system or planner compute? Successful algorithms exploit mechanical computation, in which the mechanics of the task circumscribes the possible outcomes of an action by dint of physical laws. Executing such strategies may require little or no computation; in contrast, planning or simulating these strategies may be computationally expensive. Since during execution we may witness very little "computation" in the sense of "algorithm," traditional techniques from computer science have been difficult to apply in obtaining meaningful upper and lower bounds on the true task complexity. We believe that a theory is needed to measure the sensitivity of plans to particular assumptions about the world.

One possibility is to develop a notion of information invariants for characterizing sensors, tasks, and the complexity of robotics operations, similar to the Turing model developed for computation. For example, in computational geometry, a rather successful measure has been developed for characterizing input sizes and upper and lower bounds for geometric algorithms. Unfortunately, this measure seems less relevant in embedded systems, which is perhaps a reflection of change in the scientific culture. This change represents a paradigm shift from *off-line* to *online* algorithms. Increasingly, robotics researchers doubt that we may reasonably assume a strictly off-line paradigm. For example, in the off-line model, we might assume that the robot, on booting, reads a geometric model of the world from a disk and proceeds to plan. As an alternative, we would also like to consider *online* paradigms where the robot investigates the world and incrementally builds data structures that in some sense represent the external environment. Typically, online agents are not assumed to have an *a priori* world model when the task begins. Instead, as time evolves, the task effectively forces the agent to move, sense, and (perhaps) build data structures to represent the world. From the online viewpoint, off-line questions such as "what is the complexity of plan construction for a known environment, given an *a priori* world model?" often appear secondary, if not artificial.

Modular robots: The goal of modular robotics is to create more versatile robots by using reconfiguration: hundreds of small modules will autonomously organize and reorganize (automatically or manually) as geometric structures to best fit the terrain on which the robot has to move, the shape the object the robot has to manipulate, or the sensing needs for the given task.

The science-base of modular robotics is in an embryonic state, but there are lots of opportunities with huge potential impact for applications. Several mechanisms capable of reconfiguration have been proposed, and a few centralized planning algorithms: the start and goal configuration are specified and a global planner and controller synthesize motion sequences that move one module at a time. Some of the most interesting applications of this work will employ thousands of modules working together. Such systems constitute ultra-high degree of freedom systems. The off-line planning algorithms proposed above move one module at a time and may be too slow and impractical for controlling lattices made of thousands of modules. They do not take advantage of the natural distribution and redundancies in the system and presume too much global knowledge. Our goal for this part of the project is to develop distributed planning algorithms that are highly parallel, use local communication to neighbors, and minimize the required global information. In a distributed approach modules make local decisions on-line about where to move (in parallel) in each round. Complications arise from the very high degree of freedom and under-constrained nature of these systems. In addition, there are constraints on the connectivity and on the static and dynamical stability of the structure. Moreover, since these robots are not connected in a fixed topology (but rather, are allowed to reconfigure in response to tasking demands) the control, coordination, and programming of such devices remains very challenging. Finally, the control issues required by a physical implementation include synchronization at many levels: the clock used by each system, making connections, and the actual motion, which is difficult. This leads to natural problems in dynamic control.

Another important aspect of modular reconfigurable robots is the characterization of exactly what such robots could do. The modules can aggregate with other identical modules and they can move relative to the world, but what exactly can they do, and how long does it take? We would like to characterize the class of static and dynamic objects that can be assembled from Crystalline modules. For example, algebra could play an important role in characterizing the class of structures and motions that can be achieved with a given unit module. An interesting line of investigation is to treat each module that forms the basis of a reconfigurable robot as a generator for a subgroup of $SE(3)$. Motion trajectories correspond to specific paths in this subgroup according to the specific types of actuation used in the module. We would like to characterize the nature of these subgroups, and develop a formalism that makes it easy to examine different modules with different types of actuation, and compositions with modules of different types.

Dynamics and control

In mechanics, systems may be holonomic or non-holonomic. In general, a holonomic constraint is a wholly integrable sub-bundle E of the tangent bundle. The system outcome for a non-holonomic system is path-dependent. Non-holonomic systems have been studied in robotics. Examples include: car-like robots, tractor-trailers, bicycles, roller-blades, airplanes, submarines, satellites, and spherical fingertips rolling on a manipulandum. In robotics, a non-holonomic system is usually defined by a series of non-integrable constraints of the form $\Gamma_i(p,v)=0$ on the tangent bundle. For example, whereas holonomic kinematics can be expressed in terms of algebraic equations which constrain the internal, rotational coordinates of a robot to the absolute position/orientation of the body of interest, non-holonomic kinematics are expressible with differential relationships only. This distinction has important implications for the implementation of a control system.

It is well known that many under-actuated manipulation systems are naturally modeled as non-holonomic systems. An important control problem is the analysis of isotropic non-holonomic manipulation system in which multiple robots manipulate objects by using non-prehensile grasps, or by using prehensile grasps enabled by tools such as ropes.

Non-smooth analysis: Non-smooth analysis has reached a certain level of maturity in the last 10 years but still remains relatively inaccessible to roboticists, in part due to the limited communication between the areas. Consider the problem of finding the zeros of functions which are PL smooth, or maximizing an objective function which is also non-smooth. These problems have been around for 40 years without much change. But in last 10 years, new assumptions and algorithms have become available. Many people who actually pursue this within the mathematics community are somewhat marginalized but the potential interaction with robotics has the potential to speed up the development here, with consequent advantages for both areas. Example problems include: collision problems, contact problems, friction problems, grasping, manipulation, sliding, and jumping.

In a different vein, challenge coming from other areas of robotics with requirements in same area of non-smooth analysis. Examples include hybrid systems, with event driven components, e.g., motion planning in changing environments, such as assembly. Modeling of uncertainty in robotics, propagation of errors in map building for local robots in environments which are not completely known. If we use odometry in map-building, how do we update and correct for errors? How do errors accumulate to errors in linkages? This is relevant to assembly tasks. How does one acquire the bounds for errors, and figure out how to make corrections? Mathematicians in this area can well have their interests focused by the types of problems that occur in this way in robotics areas. Probability theory and statistics are clearly the dominant tools in a number of areas: Images coming from a sensor, motion sensing and pattern recognition. Activities for the mathematical community arising from these issues include probability theory on groups, and solving min-max problems.

For reconfigurable robots: just to change shape we get into combinatorial manifolds of extremely high dimension. Roboticists don't know anything about them, and mathematicians in fact do have knowledge in these areas, but not at the level of detail required for these applications.

When dealing with sensors or actuators - when we make a movement many muscle groups are brought into play, some produce fine motion, some gross effects. This type of breaking down of tasks in terms of scale is fundamental to problem solving in the biological world. In engineering world, look at signal and extract information at these same levels of scale. There is a need for multi-scale development in robotics and engineering. Robotics could be a very useful source of problems. Currently, in fluid flow, turbulence is similar in terms of these effect. Could roboticists challenge mathematicians to focus on these problems?

Contact dynamics: There have been numerous developments in contact mechanics and dynamics in the past decade, and robotics has played a major role stimulating the development of these areas in mathematics and engineering. Starting with the work of Moreau and Monteiro-Marques in the 1980's and early 1990's, new and mathematically rigorous formulations of rigid-body dynamics with impact and friction have become available. This has been closely tied to new simulation techniques based on non-smooth analysis which have been developed by Moreau, Monteiro-Marques, Jean, Paoli, and Schatzman, Stewart and Trinkle, Anitescu and Potra, and others, on the mathematical side. There has also been intense interest in this area amongst the optimization community as well. The problems of dealing with impact mechanics for *elastic* bodies has also received a great deal of attention recently by Jarusek and Eck, Cocu, Raous, Martins, and others, although most of this work is being carried out in Europe.

However, there are a number of issues that have not been properly addressed. In rigid-body dynamics the impact “law” is an “input” into the simulation, rather than a consequence of it. Recent experiments and simulations (e.g., by Hurmuzlu and Stoianovici) indicate that the standard approaches to impact laws for rigid-body dynamics are inadequate, and elastic behavior must be incorporated. The mathematics of elastic bodies in impact is also very much under-developed, and even the existence of solutions and energy conservation issues are in doubt at present. Two central issues remain in this area:

- We don't yet have rigorous and useful results for elastic bodies in impact, with *or* without friction.
- Full elastic-body models are computationally very expensive; we would like to develop low-order but physically accurate models of impact.

These issues may be resolved in the next decade, but it will require a substantial effort and some new breakthroughs in the area.

Interactions: We need to add interactions with modern topology and geometry. Here, as regards the first issue, mathematicians developed many successful techniques for doing the analogue of differential equations (flows) on piece-wise linear manifolds 35 years ago in these areas, when they were studying the global classification and structure of manifolds. In particular, a critical component of the results of that time involved a corresponding notion, that of the tangent micro-bundle. More recently, there has been an example of the potential for this area in robotics. Even the simplest kinds of problems involving motion planning for complicated (multi-link arms or two three or four link arms working in concert in an assembly environment) could not be handled since the structure of the configuration spaces was not understood. But using techniques in areas coming from Morse theory, Morse flows and handle decomposition, real progress on some of these problems has become possible, and some complete programs for motion planning with some of these mechanisms are now possible.

Mechanics and Design

Robots are electromechanical devices. Issues in the design of robotic systems, the kinematics and dynamics of their operation, and the physics of their interactions with the surrounding environment have preoccupied robotics researchers and practitioners since the earliest days of robotics research. There is a substantial accumulated body of work in these areas, and many problems are now considered solved. However, there are still important issues that will require long term attention and non-trivial mathematical expertise. Some of these issues arise because of ongoing advances in technology. The mechanics and design session identified the following broad themes and individual problems during its discussions.

Robotic System Design and Design Methodology: Electromechanical design is a formidably complex problem as the dimensionality of the ``design space'' (the space of all possible choices for system parameters) is so huge 10^6 materials, 10^2 fabrication methods) and the functions that describe design performance are highly nonsmooth and highly nonconvex. The development of formal methods for design is an active research area in the mechanical design and circuit design communities. Formal and automated design methods have been very successful in VLSI system design because the problem is highly constrained. Formal methods in mechanical design have been less successful, possibly because the design possibilities are so vast. The possibility of bringing greater mathematical rigor and fresh mathematical ideas to this subject is promising.

It was generally felt that advances in design methodology might be made by constraining the focus to a more narrowly defined set of problems. In particular, structured and automated design methods for MEMS (MicroElectroMechanical Systems) is attractive, as the design space of material (Si, GaAs, etc.) and fabrication techniques is smaller. MEMS is also likely to have a strong impact on robotics in the future in terms of cheap, distributed sensing and possibly actuation. Hence, advances in automated MEMS design capability will have a positive impact on robotics. The MEMS community is now paying some attention to this issue, but it the subject is still in its infancy.

Robustness and uncertainty is a common theme that has arisen throughout all areas of this workshop. It is impossible to manufacture components precisely. For example, even in the highly mature disk drive industry, disk drive microactuators are fabricated with 15% variation in parameters. Design methodologies that better account for these variations are clearly needed. The discussion also highlighted the following important trends and their associated needs for new research.

- It is now clear that some form of household robotics will become widespread in the future. A general issue is how to use robots to improve quality of life, particularly for disabled people. In the design context, the issue of how you design robots to work closely and safely with humans is a largely untouched subject.
- Future robotic systems will be distributed and comprised of many interacting components, or “agents.” Examples at the macroscale include “Intelligent Transportation Systems” while meso- and micro-scale examples include distributed part manipulation. Design of large-scale systems to enable their coordination is a critical point. Distinctly new design methodologies are needed for this domain. Moreover, one cannot decouple mechanical design from control and sensing in these problems; it is truly an integration problem.
- Because of advances in micro- and nano-scale fabrication, future robotic systems will be comprised of components having vastly different scales in terms of size, energy usage, and computation-communication needs. Design methodologies for integration of micro and macro systems/components are desirable. Along these same lines, multi-scale modeling is an attractive objective.
- Robots are widely used in the nuclear industry. As they are extended to other hazardous and extreme environments, such as search and rescue, new design paradigms are needed.

Finally, the goal of developing systems with minimalist actuation is attractive. In particular, exploiting nonholonomic constraints to design underactuated robots is a promising area.

Optimization and Simulation: System design is tightly linked to optimization. From this point of view, mathematics plays a role. With recent advances, optimization problems can be solved for millions of parameters, even with constraints. Coupling these recent computational advances with design is an attractive opportunity. As we expand modeling capacity and increase the complexity of system models, the associated optimization problems will only be solvable by numerical techniques. Generally, approximation methods will be used. A particularly attractive subject for research is the investigation of set-based design, simulation, and optimization methods. That is, in set-based simulation one simultaneously simulates a whole set of systems, and the output is not a trajectory, but a set of trajectories. Simulation of hybrid systems is a related and promising topic. It is likely that interval-based methods, interval arithmetic, and automatic differentiation will be useful tools.

Genetic algorithms and ``evolutionary'' design techniques have been effectively used in large-scale system design; they produce a robust solution, not necessarily an optimal one. A greater theoretical understanding of these algorithms is needed.

Self-Organizing and Reconfigurable Systems: George Whitesides has developed impressive demonstrations of self-organizing mechano-chemical systems. In these demonstrations, simply shaped objects are selectively coated with hydrophilic/hydrophobic molecules, and place in solution. Through different mechanics, such as evaporation, vibration, etc., these systems self-assemble into complex shapes and networks. Vast opportunities exist for beautiful mathematics in this area.

Within the robotics community, there has been much recent work on modular and reconfigurable systems. Realistically, our ability to design and fabricate such systems is lagging our ability to plan for and control such systems. There is clear need for new paradigms that address the issue of design for reconfigurability. In particular, the lack of tools for modeling and design of reconfigurable systems, as well as design of limited DOF systems that are reconfigurable hamper real progress in this area.

Kinematic Synthesis: Kinematic synthesis is an important enabling technology for robotic systems. There are obvious opportunities to apply methods from algebraic topology to the kinematic synthesis problem; particular in the case of closed-chain mechanisms.

Biomimetics: Borrowing concepts from biological systems as a basis for engineering system design has proven to be an appealing concept. However, ``biomimetics'' is still more of a philosophy than a rigorous discipline. Within the broad realm of biomimetic system design, there appear to be targets of opportunity for mathematical analysis that leads to useful systems. For example, fish locomotion is a mathematically complex phenomena whose greater understanding could be translated into design of new systems. As molecular biology advances, future design paradigms must also consider biological material as a potential design substrate.

Contact Modeling: Contact modeling continues to be a vexing problem. New contact models and their experimental validation is an issue.

Soft Tissue Modeling and New Modeling: It is now clear that robotics will have a significant impact on medicine. To support anticipated medical applications, modeling of soft tissues and organic structures is needed. Such models are needed for design of medical robotics and for simulating (often for purposes of training surgeons) their actions. Because of the highly anisotropic and large deformation characteristics of soft materials, their faithful simulation is difficult. Low order models are appealing, but possibly unrealistic.

As robots increasingly interact with more diverse environments, we need to move beyond rigid body and mechanical models. We need a much richer set of models of how the world behaves (mechanics), more phenomena (deformations, soft tissues) and new domains. Dynamic models are a start.

Design Metrics: How do you compare systems that solve the same task but use different sensors. (i.e. vision vs. sonar sensing) or different components? Such equivalence notions are needed for design studies.

Uncertainty

Uncertainty is pervasive throughout robotics. The vast majority of existing motion planning and control algorithms in robotics does not take uncertainty into account. A proper handling of uncertainty will almost certainly lead to significantly more robust systems, along with a better understanding on how to perceive and act in the physical world. The importance of uncertainty will increase as robots move away from factory floors into increasingly unstructured environments, such as private homes.

For example, existing robotic surgical systems face an enormous amount of uncertainty, which is currently only marginally considered during motion planning. Similarly, manufacturing, drilling, tunneling and robotic exploration are characterized by significant uncertainty, and hence would benefit from better mathematical and computational tools to make decisions under uncertainty.

Currently, a range of complimentary frameworks exist for representing uncertainty: Probabilistic methods (which include parametric and non-parametric representations), binary representations, Dempster-Shafer logic, fuzzy set theory, and others. The choice of the representation influences the difficulty of crafting models and the computational efficiency of using these models. Additionally, a range of different problems can be attacked under uncertainty, such as: prediction vs. planning vs. control; worst case vs. average case; and correctness vs. optimality.

The panel identified a range of scientific questions that warrant research. This list is meant to be as a representative sample, but should not be interpreted as complete:

- How can uncertain information be propagated through process models, and what type of bounds can be obtained?
- How can we devise systems that can reason about when and what to sense?
- How can we develop contingency plans that interleave sensing and control?
- How can we reduce the complexity of probabilistic propagation and planning?
- How can we approximate uncertainty and devise bounds for those approximations?
- What problems can be solved in closed form, and which solutions can be computed efficiently?
- How can we best represent geometric uncertainty, or shape uncertainty?
- How can we devise planners that employ feedback mechanisms for reducing uncertainty?
- How can we, on solid mathematical grounds, ascertain which uncertainties can be ignored?

The fields of statistics, operations research and computer science has long addressed some of those issues, though typically deprived of a robotic context and in relatively small worlds. Current theory often focuses on discrete problems, whereas robotics spaces are typically continuous. Examples include “Optimal Experimental Design,” “Partially Observable Markov Decision Processes,” “Sensitivity Analysis,” and “Monte Carlo Methods.” We urgently need research to develop these and similar fields in the context of robotics. To make these basic methods amenable to complex robotics problems and, thus, enable the deployment of robotic systems in uncertain, real domains.

Computer Vision

Computer vision has witnessed rapid progress in the past fifteen years, with great success in classical application fields like industrial automation, and, more recently, emerging areas such as medical robotics and the entertainment industry. State-of-the-art industrial vision systems rely on sophisticated mathematical tools from geometry and statistics to model the practical problems they address. In addition, fundamental aspects of computer vision, such as the analysis of image sequences of rigid point sets observed by a roving camera, or the characterization of the appearance of smooth surface outlines, are now fairly well understood.

Thanks to these advances and ever-increasing computer power, tackling harder and exciting problems such as building general-purpose visual inference machines has become a realistic endeavor. Yet, it is unclear how current approaches can be generalized to solve such difficult problems: indeed, computer vision has sometimes been criticized for its lack of formal foundation and empirical justification, and, in contrast with mature disciplines, like control theory and information theory, it certainly lacks a mathematical framework. Such a framework should include a common language for describing vision problems, mathematical techniques for modeling them, and algorithms that can be applied to solve them.

For example, the visual recognition of learned object models is a key problem, maybe *the* key problem, in computer vision. Solving it will require radical advances in object representation, image segmentation, and in our understanding of the recognition process. Probability theory and statistical inference form a promising foundation on which to build a mathematical framework for object recognition, but major conceptual and technical difficulties remain: for example, how can we define object models that are easy to learn and effectively support inference from pictures? How can we construct algorithms that will perform Bayesian inference in real time? How should we handle large numbers of objects and object classes?

There has recently been encouraging progress on these issues following the introduction of techniques from statistics, mathematics, and information theory in our field. Conversely, the demands of computer vision throws tough challenges and interesting problems at these disciplines. Indeed, as argued by Mumford (2000), the type of statistical inference required for computer vision may become a central theme for mathematics in the twenty first century. Likewise, we believe that a strong synergy between areas of mathematics such as geometry and topology, recent advances in computer technology, and effective engineering practice, will result in revolutionary advances in the theory and applications of computer vision to areas as diverse as visual robot navigation, medical robotics, and automated three-dimensional model acquisition for the entertainment industry.

Workshop on Mathematics and Robotics
Schedule
May 15, 16, 17, 2000
Arlington, Virginia

Monday

8:00 – 8:30	coffee and bagels (Room 1020)
8:30 – 9:00	welcome and introduction (Room 1020)
9:00 – 10:30	presentations (Room 1020)
10:30 – 11:00	break
11:00 – 12:30	presentations (Room 1020)
12:30 – 13:30	lunch
13:30 – 15:00	Dynamics and Control (Room 1020)
	(Discrete) Geometry and Topology I (Room 1060)
15:00 – 15:30	break
15:30 – 17:00	Learning (Room 1020)
	Geometry and Topology II (Room 1060)
17:00 – 17:15	break
17:15 – 18:00	summary of first-day discussion (Room 1020)
19:00	group dinner (TBA)

Tuesday

8:00 – 8:30	coffee and bagels (Room 1020)
8:30 – 10:00	presentations (Room 1020)
10:00 – 10:30	break
10:30 – 12:00	Mechanics and Design (Room 1020)
	Sensing and Vision (Room 1060)
12:00 – 13:00	lunch
13:00 – 14:30	Planning under Uncertainty (Room 1020)
	Distributed Robot Systems (Room 1060)
14:30 – 15:00	break
15:00 – 17:00	summary of second-day discussion (Room 1020)

Wednesday

8:00 – 8:30	coffee and bagels (Room 1020)
8:30 – 10:00	drafting of workshop report (Rooms 1020 and 1060)
10:00 – 10:30	break
10:30 – 12:00	concluding discussion (Room 1020)
12:00	adjourn

Workshop on Mathematics and Robotics
Participant List
May 15, 16, 17, 2000
Arlington, Virginia

James S. Albus albus@cme.nist.gov
Senior NIST Fellow, Intelligent Systems Division
National Institute of Standards & Technology

Mihai Anitescu anitescu@math.pitt.edu
Department of Mathematics
University of Pittsburgh

Radhakisan S. Baheti rbaheti@nsf.gov
Program Director ECS
National Science Foundation

Anthony Bloch abloch@math.lsa.umich.edu
Department of Mathematics
University of Michigan

Roger Brockett
brockett@yellowstone.hrl.harvard.edu
Division of Engineering and Applied Science
Harvard University

Joel Burdick jwb@robotics.caltech.edu
Department of Mechanical Engineering
California Institute of Technology

Greg Cherikjian gregc@jhu.edu
Department of Mechanical Engineering
The Johns Hopkins University

Ken P. Chong kchong@nsf.gov
Program Director CMS
National Science Foundation

Mike Erdmann me@cs.cmu.edu
Robotics Institute
Carnegie Mellon University

Alison Flatau aflatau@nsf.gov
Program Director CMS
National Science Foundation

Robert Ghrist ghrist@math.gatech.edu
School of Mathematics
Georgia Institute of Technology

Leo Guibas guibas@logos.stanford.edu
Robotics Laboratory
Department of Computer Science
Stanford University

John Harer John.Harer@Duke.edu
Professor and Vice Provost for Academic Affairs
Duke University

Roberto Horowitz horowitz@me.berkeley.edu
Mechanical Engineering,
UC Berkeley

John Horst john.horst@nist.gov
Senior NIST Fellow
Intelligent Systems Division
National Institute of Standards & Technology

Steve Kaufman sgkaufm@sandia.gov
Sandia National Laboratories
Albuquerque, NM 87185-1008

P.S. Krishnaprasad krishna@glue.umd.edu
Department of Electrical and
Computer Engineering
University of Maryland

Vijay Kumar kumar@central.cis.upenn.edu
School of Engineering and Applied Science
University of Pennsylvania

Ming Lin lin@cs.unc.edu
Department of Computer Science
University of North Carolina

Deborah F. Lockhart dlockhar@nsf.gov
Program Director DMS
National Science Foundation

Vladimir Lumelsky vlumelsk@nsf.gov
Program Director ECS
National Science Foundation

Benjamin M. Mann bmann@nsf.gov
Program Director DMS
National Science Foundation

Wen C. Masters Wen_Masters@onr.navy.mil
Program Officer
Office of Naval Research

Matt Mason Matt_Mason@cs.cmu.edu
Robotics Institute,
Carnegie Mellon University

Jim Milgram milgram@math.stanford.edu
Department of Mathematics
Stanford University

Jong-Shi Pang jpang@nsf.gov
Program Director DMS
National Science Foundation

Jean Ponce ponce@cs.uiuc.edu
Department of Computer Science
University of Illinois

Florian Potra potra@math.umbc.edu
Department of Mathematics and Statistics
University of Maryland Baltimore County

Karnlakar P. Rajurkar krajurka@nsf.gov
Program Director DMII
National Science Foundation

James L. Rosenberger jrosenbe@nsf.gov
Program Director DMS
National Science Foundation

Daniela Rus rus@cs.dartmouth.edu
Computer Science Department
Dartmouth University

Ed Scheinerman ers@jhu.edu
Department of Mathematical Sciences
The Johns Hopkins University

Christopher W. Stark cstark@nsf.gov
Program Director DMS
National Science Foundation

Ray Sterling sterling@coes.latech.edu
Trenchless Technology Center
Louisiana Tech University

Michael H. Steuerwalt msteuerw@nsf.gov
Program Director DMS
National Science Foundation

David Stewart dstewart@math.uiowa.edu
Department of Mathematics
The University of Iowa

Russell Taylor rht@cs.jhu.edu
Computer Science Department
The Johns Hopkins University

Sebastian Thrun
Sebastian_Thrun@heaven.learning.cs.cmu.edu
Center for Automated Learning & Discovery
Computer Science Department
Carnegie Mellon University

Dawn Tilbury tilbury@umich.edu
Mechanical Engineering and Applied Mechanics
University of Michigan

Jeffrey C. Trinkle jctrink@sandia.gov
Sandia National Laboratories
Albuquerque, NM

Gerard A. Venema gvenema@nsf.gov
Program Director DMS
National Science Foundation

Louis Whitcomb llw@jhu.edu
Department of Mechanical Engineering
The Johns Hopkins University

Jing Xiao jxiao@nsf.gov
Program Director IIS
National Science Foundation