Research in the Autonomous Mobile Robotics Laboratory

J. Hendler David J. Musliner
University of Maryland

April 12, 1994

The design of robots which can autonomously function in cooperation with humans in non-constrained environments has long been a dream of both the AI scientist and the robotics engineer. From the AI perspective, for over 30 years the field of AI planning systems has been justified as "robot problem solving." The basic idea was that given a robot, a problem for it to solve, and a set of techniques it could use, a planning system would generate the ordered set of actions for the robot to take. Unfortunately, even now, when cheap robots are becoming available, very few AI planning researchers have tried to realize their ideas on realistic robots.

From the engineering direction, physical control of robots has long been pursued both on real robots and in idealized models. However, issues such as problem-solving and rapid adaptation to changing environments have been the hallmark of the field of "intelligent" control— they go beyond the relatively well-understood techniques of the college control course. Thus, the field of autonomous mobile robotic problem-solving has become the meeting place for two different sets of scientists: the AI researcher trying to ground planning research on physical devices, and the engineer interested in expanding the realm of classical control models.

The University of Maryland's Autonomous Mobile Robotics Laboratory was started this year as a workplace for research on the intelligent control of goal-based robotics and motion planning. The goal of this work is to understand how to advance control and robotics research to provide both a formal basis and a framework for practical applications. In particular, the effort is aimed at producing not only important and necessary theoretical results, but also at demonstrations reflecting the capabilities provided by these new theoretical underpinnings. Several projects along these lines have recently been launched.

Project 1: Formalization of behavior-based robotics in control terms

Two main approaches to motion planning have been proposed in the past. The first is based on finding optimal path plans using complete or near-complete information. The second is based on using reactive, sensor-based behaviors to find paths given little or no prior information. The former approach has primarily arisen from the literature on control and robotics. The latter derives mostly from the AI work of Brooks and others. The control work is advantageous in several ways: it is easily formalizable via both optimization and controltheoretic techniques, it is amenable to tests for optimality, and it can be integrated with well-understood control techniques for realization in “control laws” for simulated robots. The reactive work is advantageous in being able to deal with imperfect, limited-range sensors, with uncertain information, and with worlds in which change occurs (obstacles may be moved, etc.)

Our work focuses on providing a formal basis for behavior-based robotics using the techniques that have been successful for control-based approaches. In particular, behaviors for robots are formalized in terms of sets of kinetic state machines (KSMs) and the interactions of these state machines with information coming in from (limited range) sensors. This allows us to create a mathematics for discussing these systems, including techniques for integrating sets of behaviors. In addition, we are developing optimality criteria for comparing both individual and multiple behaviors in various environments.

This unifying formalism also allows us to develop evaluable techniques for integrating the approaches. In the absence of obstacles, smooth, well-controlled, and optimal paths are planned. When obstacles are present, particularly obstacles not provided for in an initial map, the system is able to react to their presence and to initiate avoidance behaviors. The system's behaviors can be compared with regard to the optimality criteria described above, and thus various sets of behaviors can be said to be inferior or superior to each other in a meaningful way.

Current research directions include:

- Continuing the formalization of behavior-based robotics.
- Continued development of a simulator for testing the integration of control and reactive techniques in complex (obstacle laden) problems.
- Use of this technique to solve complex motion-planning problems including “car parking,” “trailer pulling,” and other difficult non-holonomic motion planning tasks.
- Extending the technique for use in the presence of moving and/or movable obstacles.

Project 2: Use of behavior-based robotic research in realistic environments on working mobile robots

This project focuses on using the results obtained in the project above for the solution of problems in the “real” world— that is, working on a physical (as opposed to simulated) robot functioning in a non-artificial environment. There are two main parts to this project: (i) testing the ideas on an existing robot, (ii) developing a better controlled (and controllable) robot for further experimentation and demonstration.

Testing

In previous AI work, we described a simulated household cleaning domain, with a simulated robot agent which performed cleaning tasks. The agent performed tasks such as vacuuming the floor, putting dirty socks into a hamper, cleaning up spills, and even putting out a fire in the kitchen. The simulator divided the apartment into over twenty thousand spatial units, and limited sensor and effector capabilities were simulated based on these units. For example, the robot could only see what was in a projected cone of vision, could only move things it could reach, etc.

Although a "robot" was used as the domain for this work, and although the simulation was more complex than most AI micro-worlds, the simulator did not map well to real robotics. While this was an appropriate domain for AI planning research, its assumptions were unrealistic for real robotics work— sensors were assumed to return perfect information, effectors were "symbolic," etc. In addition, little effort was put into the motion planning or other control aspects (naive, non-optimal algorithms were used).

The current project focuses on mapping similar problems to a real robot, and eventually using it to directly test the ideas on comparing behaviors and hybrid control developed in Project 1. We are currently working on testing the research by developing a behavior-based[1] approach to a problem requiring a combination of goal-based and reactive behaviors, coupled with complex motion-planning needs. In particular, we are trying to get the robot to cope with the problem of vacuum-cleaning the halls of the AV Williams building (while people and obstacles may be there).

The robot used in our current research is an IS-Robotics R-2 robot. This robot is a small, wheeled, relatively cheap behavior-based autonomous robot, which is programmed using a "behavior" language (originally developed by Rod Brooks). The behavior code is (automatically) translated into a set of Augmented Finite State Automata hooked together in a subsumption-based manner. These AFSMs are then compiled into 6811 code, and downloaded to the robot. Four 6811s are used to provide control, with a master processor controlling three others, used for various sensing and effecting. The robot we use has a 2-DOF gripper, eight infrared sensors, each of which can return a value which estimates distance to an object (accuracy approximately 1 ft., with 5 foot max), and 8 bump sensors which return either a force or position reading.

To date, we have implemented a set of behaviors on the robot for vacuuming in our laboratory. It carries a small hand vacuum cleaner in its gripper and attempts to vacuum a surface in our laboratory— the robot works on a platform which has a rectangular surface and borders allowing the robot's infrared sensors to detect the end of the platform. Arbitrary obstacles on the platform can't be effectively avoided due to the robot's very limited sensing capabilities. Therefore, we limit the shapes and orientations of obstacles so that the robot can avoid them while still being able to clean the floor around them. To differentiate between obstacles and the platform border, which both appear simply as "objects" to the two front infrared sensors of the robot, we attempt to use what little dead-reckoning is available to keep track of the distance moved in the current lane. Whenever the infrared sensors pick up an object, the robot can tell if this is too close to be the opposing border or too far to be an obstacle. A border tracking behavior is used to go around obstacles, with corners counted to allow vacuuming to continue.

Currently the robot functions well in the presence of single rectilinear obstacles, or around sets of several of these placed non-contiguously. With more complex arrangements of rectangular obstacles, or with non-rectangular surfaces, the robot fares less well. To overcome this, and to allow us to take the robot out of the laboratory and into the AVW environment, we are trying to design more complex (and possibly hierarchical) behaviors,

[1] Later to become a combined control/reactive approach, see above.

and to integrate better motion-planning algorithms. In addition, we are attempting to integrate the strategic planner of the earlier AI work to allow for more complex decision-making. This work is thus expected to result in the first system which integrates intelligent control, behavior-based robotics, and AI planning in a mathematically well-motivated manner (based on the KSMs used in Project 1).

Robot Design

The robot described above was developed based on the work of MIT's Rod Brooks, with behavior-based robotics in mind. As such, it was not designed to be controlled particularly efficiently. It is particularly difficult to get the robot to perform complex navigational tasks or to be useful in more careful experimentation in the use of intelligent control techniques such as those defined in Project 1. We therefore are attempting to design and build our own robot, which will be:

- smoothly controllable,
- powerful enough (computationally) to support the behavior-based techniques,
- reliable enough for continued experimentation, and
- designed with the integration of control and behaviors in mind.

This work has a near-term and a long-term thrust. In the near term, we are using an existing 3-wheeled robot base (home made), and attempting to build a better control system for it. This will allow us to experiment with different control laws for motion planning, as well as providing a base for early work on the new control/behavior methodology. In addition, to provide for use of this base in motion-planning experimentation with limited-range sensing, the robot is being fitted with sonar range sensors.

For longer-term experimentation, we are now beginning a project to develop a better robot for this and other control/AI research. We hope to build or obtain a more precise base (possibly with an existing sensor set), and to develop on top of it an operating system which will support work in sensing, control, and artificial intelligence. This operating system is envisioned as being based on object-oriented computing techniques, allowing a number of different control and AI techniques to be easily implemented and tested. In addition, the precision of the base should make direct experimentation on motion-planning more accurate and quantifiable than it will be on the current base. This is expected to result in a better platform for exploring issues of optimality in control.

Project 3: Motion Control for Space Robotics (Joint research with the University of Maryland Space Systems Laboratory)

One of the primary areas in which robots are expected to be deployed as integral mission components is that of space exploration. Current projections call for space work situations comprised of combinations of humans, telerobotic devices, and autonomous robots. (Complex space robots, such as those expected to be used in space construction projects, are currently anticipated as being devices which combine both telerobotics and autonomous behaviors.)

Space robotics provides some novel challenges for motion planning, because of two primary differences from planetary surface motion: space motion is frictionless and 3-dimensional. In this project we are extending ideas from the prior projects (both in the underlying mathematics and in the experimental realm) to cope with these differences.

Frictionless Motion

We are beginning an exploration of frictionless motion using a unique robot called the Airborne Imaging Robotic System (AIRS). The AIRS robot uses an air-cushion (hovercraft-like) base with an air-thruster maneuvering system and supporting electrical and electronic components. Air-cushion vehicles have a very low friction coefficient, so they move and respond to outside forces more like a space robot would.

Design of the AIRS platform began in December 1991 at Brigham Young University. The basic system was completed in February 1993. (AIRS was partially financed by the BYU Honors Department with private sources providing the balance of funds. Ultimate Support Systems, Contech, and Signetics were corporate sponsors. Dr. William Barrett of the BYU Computer Science Department advised the project). AIRS now resides in the Autonomous Mobile Robotics Laboratory, and is being used for this project.

Current plans for AIRS include:

- Development of mathematical models. Models of motion planning have typically either assumed totally controlled, point-based motion (perhaps in a frictionless domain) or have had to be adopted for control parameters on acceleration, slippage, etc. Neither model is appropriate for the frictionless 2D motion of a complex vehicle such as AIRS, which is controlled by three air thrusters. We propose to examine vector-field-based motion planning where a non-holonomic robot is operating in a friction-free manner.

- Realization of mathematical models. We will be able to test these models on the AIRS vehicle, thus exploring how they transition to real control.

- Extension of the "hybrid" model (see Project 1) to frictionless motion control. The limitations of current mathematical approaches to frictionless motion include assumption of complete knowledge, perfect sensors, etc. We hope to extend the framework of Project 1 to handle AIRS-motion problems with incomplete and uncertain knowledge, with moving obstacles, etc.

Additional note on AIRS: We are preparing AIRS for a robotics competition and exhibition to be held this July at the National Conference on AI. The tasks the robot will need to complete include autonomous motion planning in an office setting given an incomplete map, and the "cleaning" of an office environment (finding and removing objects of specific types, such as paper cups, coke cans, etc.). A set of sensors and effectors for these tasks (including a low cost vision system) are being designed, and a team of undergraduates (from EE and CS), under the direction of Glen Henshaw (graduate student, EE) the designer of AIRS and Dr. David Musliner (an ISR post-doc), are working on building these devices and programming the robot for the competition.

Behavior-based 3D motion

The SCAMP robot, developed at the UMCP Space Systems Laboratory (SSL), was developed as a telerobotic camera platform for space operations. The primary use of SCAMP is to hold a camera so that an operator controling another robot can get an additional field of view. SCAMP has been developed and extensively tested in the SSL, and a space-ready version is currently under development. However, controlling SCAMP during the performance of other telerobotic operations can be difficult for a single operator, and thus we plan to develop a set of simple, motion-based behaviors that SCAMP can perform autonomously. These behaviors include:

- Following: maintaining visual contact with a possibly moving target while remaining in a set position relative to the target.
- Tracking: maintaining visual contact with a possibly moving target while maintaining a specific position relative to the work site.
- Surveying: moving around the target to provide views from different perspectives.
- Moving-to (or away) : moving towards (or away from) a possibly moving target (variables include path to take, amount of caution in approach, etc.)

These behaviors are being developed to work in conjunction with the existing SCAMP system in the SSL, and will be tested both on the actual robot and on a simulation being developed. All of these behaviors require 3-dimensional motion planning, as the target being viewed will itself be moving in 3-d. In addition, these behaviors can be seen as building blocks for more complex behaviors, thus allowing a hierarchy of behaviors for complex navigational behaviors.

Project 4: Real-Time Intelligent Control

As computational support in the running of such important applications as nuclear power plant control, aircraft operation, hospital life support systems, and military command & control has increased, the need to apply AI techniques to expand the capabilities of such systems has become clear. Unfortunately, the time-critical nature of these domains puts significant computational bounds on the running of the AI systems. In addition, the predictable behaviors required for hard real-time computations are difficult to assure given the discontinuous nature of AI decision-making systems.

To achieve the requisite combination of knowledge-based AI methods with the ability to make mission-critical real-time performance guarantees, we are developing two different real-time AI architectures. The Cooperative Intelligent Real-time Control Architecture (CIRCA) combines heterogeneous AI planning and real-time reactive subsystems in parallel, using a scheduler module and a structure interface between those two subsystems. CIRCA was originally implemented at the University of Michigan, and is currently being ported to execute on top of the MARUTI hard real-time operating system developed at the University of Maryland.

Our second real-time AI architecture is closely related to behavior-based robot programming languages, and involves using numerous concurrent, hierarchical "monitors" to detect changes in the world and react to them appropriately. We are currently developing formal methods (similar to those of Project 1) to describe the behavioral competences and resource requirements of individual monitors, and show how they can be combined into larger systems while still retaining predictable behaviors.

In order to develop and test these approaches to real-time AI, we are developing the Smart Model Autonomous Real-Time Vehicles (SMARTVs) domain as a realistic real-time application. Intelligent vehicle/highway systems are currently a major national area of research and development, exemplifying the type of mission-critical, resource-bounded environments in which real-time AI is necessary. The SMARTVs domain is designed to be a low-cost, miniature IVHS testbed which can demonstrate real-world limitations of sensors and actuators, along with true domain-imposed hard real-time deadlines. SMARTVs will consist of remote-controlled model vehicles guided over radio links by MARUTI-based workstations. On-board sensors will be used to detect information about the vehicles' environment, which will be relayed back to the "remote-brain" computer over a radio link. In one configuration, the MARUTI-based machine will run CIRCA's real-time subsystem, implementing reactive plans that are automatically constructed by CIRCA's concurrent AI subsystem. Alternatively, the MARUTI-based controller may be used to execute and test sets of scheduled monitors.

Using the robots in the SMARTVs domain, we will assess the results of our efforts to develop formal and practical mechanisms for combining aspects of intelligent behavior with real-time performance restrictions. Because the domain uses low-cost robots based on vehicles originally designed for human control, we can compare our automated systems with each other and with manual, human-controlled performance. We anticipate that the speed and accuracy of the SMARTVs under computer control will exceed human standards, demonstrating effective integration of knowledge-based planning, real-time reactive control, and dynamic attention focusing.