Reactive Control as a Substrate for Telerobotic Systems

Ronald C. Arkin
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332

ABSTRACT

Reactive control is a recently emerged paradigm for guiding robots in unstructured and dynamic environments. It is based on the principle of task decomposition into a set of constituent behaviors that are required for the completion of a navigational task. Immediate sensory feedback (visual or otherwise) provides the stimulus necessary to evoke a set of reactive responses. This methodology has been pursued by a number of laboratories including our own.

Although this research has been intended to produce fully autonomous intelligent behavior, it is possible to use certain sets of behaviors to provide an underlying support for telerobotic applications such as are found in remote vehicle driving/piloting. In particular, the cognitive load on the teleoperator can be substantially reduced by offloading significant amounts of the local obstacle avoidance abilities to the substrate reactive control system.

In this paper we describe a schema-based reactive control system that allows teleoperators to direct a mobile platform without undue concern for local obstacle avoidance. By using an analog of the potential field methodology, repulsive forces are constructed around nearby obstacles allowing the robot to avoid them and enabling the telerobot operator to concentrate on the more global aspects of the mission. The structure for this paradigm enables real-time computation.

These same principles are extendible to manipulator operations as well. Robot simulations/experiments demonstrate the concepts described in this paper.

INTRODUCTION

In the last several years, recent advances in the area of reactive control have produced robots capable of real-time on-line interaction with a dynamic changing world. Significant inroads towards the goal of fully autonomous systems have been thus achieved by reformulating the way in which the robot interacts with the environment. Many research laboratories throughout the world now have working robot prototypes capable of coping with changing world environments.

There is much that remains to be done. In particular, the goal of fully integrating high-level reasoning into autonomous systems remains elusive. This high-level planning and capability of reasoning over failures is a desireable trait of human performance in navigational situations.

In order to bring to light the advantages of reactive control in a timely manner without waiting for subsequent development of real-time high-level reasoning systems, it is possible to allow a human operator to be embedded within the autonomous framework of a reactive control system. The term embedded is not chosen lightly; it refers to the human teleoperator working in conjunction with the existing autonomous control system, simplifying both the task of human teleoperation of a mobile platform and that of robot autonomy. In this manner, the reactive control system can simplify the teleoperator's task, by reducing the cognitive load on the operator via a shift in the low-level perceptual tasks to the autonomous reactive control system. The robot's job is facilitated by off-loading the higher level reasoning capabilities to the human. This synergistic relationship can produce practical solutions to difficult problems in hazardous environments in the relatively short-term.

The teleoperator's role can take two forms, first as an embedded agent working alongside the other existing processes in the reactive control system, or secondly as a supervisor of the reactive control system, modifying the parameters and gains of the active processes as the robot proceeds through the world.

Based on a paper presented at NTC 1991.

The reactive control system's role thus is to provide support for obstacle avoidance, goal/target selection, path following, convoying, and any other relevant low-level purposive behaviors necessary for the completion of the robot's task. In addition, smooth motor behavior transitions dependent upon internal resources of the vehicle such as fuel and temperature can automatically influence the behavior of the robotic system without the operator being concerned with this data. This extension of reactive control in this context we have previously termed homeostatic control [2], and it can assist significantly in assuring the survivability of a robot in potentially dangerous situations.

The utility of robots with these capabilities is obvious. Potential applications include; work in space, both for orbiting spacecraft as well as rover technology for planetary exploration; rescue operations in the civilian sector for situations such as fires, cave-ins, etc; service in disabled nuclear power plants, chemical facilities and other potentially dangerous environments; and both defensive and offensive military scenarios.

In the next section we provide background information on reactive control in general and related telerobotics efforts. In section 3 we describe schema-based reactive control, the methodology we have developed for producing autonomous mobile robot systems. In section 4, the teleoperator is introduced as both a schema (an embedded agent among the other reactive control processes) and as a supervisor of the active processes. This is followed by some preliminary results using our mobile robot George in situations where normal reactive control methodologies by themselves are known to have difficulty. A summary and conclusions section completes this paper.

BACKGROUND

In this section we first review what reactive control is and some of the related work in the area. This is followed by a brief survey of pertinent telerobotics research for intelligent navigation.

Reactive Control

Reactive control systems have been developed in response to the apparent lack of progress by those advocating the strictly hierarchical school of robotic control [1,21]. These new systems can be built from the bottom up, producing rapid working prototypes faster than other previous methods. Before quickly surveying some of the successes to date with these systems, we first enumerate the distinctive characteristics of reactive robotic control systems.

- *Tight coupling between sensory perception and motor action in a stimulus response manner.*

  Instead of fusing together incoming sensory information into an abstract symbolic representation, incoming sensor data is channeled directly to the appropriate motor behavior.

- *No intervening world representations.*

  World models are to be avoided at this level, as the computational cost for their construction precludes real-time response. This also enables reaction to a changing world without concern over the time stability of world models.

- *Decomposition into primitive behaviors.*

  Reaction to the world's sensations are encoded in primitive behavioral formats, describing how the robot should behave rather than following some arbitrarily determined course.

- *Arbitration strategies for behavior selection or concurrent execution.*

  Most reactive systems use arbitration mechanisms to determine which one of the behaviors is active at any given time. Our schema-based methods allow multiple behaviors to be concurrently active at any given time, allowing each to contribute to the overall behavior of the robot.

- *Parallelism inherent in design.*

  The decomposition into multiple behaviors allows for implementation on new parallel computer architectures as well as specialized distributed multi-processor systems.

- *Demonstrable robotic results.*

  Many actual robotic systems have been constructed using these methods yielding intelligent behavior in the presence of an unstructured environment. Some of these efforts are surveyed below (a review of our work in reactive navigation appears in the next section).

Most of the research in this area started in the late 1980's. Rodney Brooks at the MIT AI Lab was among the first to utilize this method. His approach termed the subsumption architecture [10] has been embedded on many robots that can accomplish various tasks such as collecting soft drink cans [15], walk on six legs [11], and move around in the interior of a lab [12]. His method utilizes finite state automata as the basic method for expressing actions, which are hard-wired to meet the specific goals of a dedicated robot. This serves as both an advantage and disadvantage: the advantage being that the real-time response of these systems is excellent, the disadvantage that they are inflexible and each machine can essentially accomplish but one task.

Work at Hughes AI Center has also utilized behavioral decomposition strategies [24]. This work has been deployed on DARPA's Autonomous Land Vehicle (ALV) for off-road navigation. This technique, which uses an arbitration strategy, similar to the subsumption architecture allows for the integration of world knowledge from maps [23].

Work by Kaelbling and Rosenschein on real-time embedded systems has led to the development of logic-based reactive systems [18]. This system deployed on a robot named Flakey has demonstrated successful navigation in an office environment.

Slack and Miller have developed a variant of reactive control that uses navigational templates [27]. This method avoids some of the difficulties found in other methods by arbitrarily making decisions regarding how to pass obstacles. Although the computational cost for this technique is considerably higher than some of the other approaches to reactive control, it has been utilized on experiments involving the Mars Rover at the Jet Propulsion Laboratory [17].

Some of these techniques have been extended to manipulator control as well. In particular, work at the University of Rochester has been concerned with the development of behavior based manipulators [13].

Despite the rapid adoption of these techniques by many research groups and the continued expansion of this technology, *pure* reactive control is not without its shortcomings, some of which are mentioned below:

- *Available world knowledge is ignored.* If reliable knowledge of the world is available beyond the scope of the sensors, it should be utilized to enhance intelligent performance.
- *Representations are potentially useful.* In situations where the robot encounters difficulty, it is useful to have the robot keep a record of where it has been (by building a map). This is useful for extrication in difficult situations.
- *Many reactive systems are inflexible.* Although they do offer real time response to a dynamic world, these systems are often not readily reconfigurable.
- *Shortsightedness.* In some cases, local minima, maxima or cyclic behavior occur which can cause failure of the robot to complete its task.

Responding to these failings, hybrid reactive-hierarchical robot navigational control systems have been developed. Our system, the Autonomous Robot Architecture (AuRA) is in this category [8]. It is capable of functioning both in the presence or absence of world knowledge, and can reconfigure its behaviors based on mission intents, environmental knowledge and success or failure of attaining the mission's goals. This hybrid approach used in our system is what enables it to be readily converted for teleoperation, with the human operator supplanting some or all of the hierarchical planner's functionality. In the next section we discuss first the role of reactive navigation in AuRA and then subsequently how teleoperation can be integrated into its reactive control framework, First however, we review the relevant work in teleoperation to provide a staging ground for our research.

Telerobotics

Considerable research has been conducted in the area of telerobotics. A recent NASA conference on the subject of space telerobotics [25] yielded a proceedings 5 volumes long. In this section we will discuss only a few of the related approaches to telerobotic control of a mobile vehicle, in particular concentrating on those efforts which are most closely related to the tasks and methods which are own research addresses. This review should by no means be construed as comprehensive.

Much of the research on telerobotics has concentrated at the interface of man and machine. In particular, many studies have been concerned with the issues of display, control and communication between operator and robot (e.g., [28,29]). These factors are indeed crucial for successful deployment of a telerobotic system, but they have not served as a focal point for our work thus far. At best, our display and communication methods are currently rudimentary. Nonetheless we claim that much of the display and communication problem can be handled by merging autonomy in with telerobotics. This can serve to ease the cognitive load posed on the operator by removing the necessity of his/her handling of low-level tasks thus reducing the demands of a display system. Also by reducing the amount of tasking required by the teleoperator, the net bandwidth of communication requirements is decreased enabling the use of slow-scan TV or other less costly methods of data transmission. Finally by permitting the robot to function in a semi-autonomous state, the importance of time delay is less critical as the operator can rely on the robot's ability to protect itself from damage through autonomous control. The ability of the teleoperator to communicate with the robot could be significantly enhanced by exploiting new communication pathways, using techniques such as speech input to free the hands of the teleoperator [26]. We leave our development of enhanced user interfaces for future work.

More recently work in the area of teleautonomous systems has emerged, systems that merge the capabilities of autonomous control with telerobotics. Teleautonomous technology has been defined as:

> *The interaction of humans with remote, intelligent, partly autonomous systems of many forms [16].*

In particular, a teleautonomous system has been said [16] to consist of the following parts: 1) a human operator communicating via a 2) remote simulator providing environmental feedback with 3) an intelligent local controller that interacts with a 4) remote intelligent robotic controller connected to the 5) robot and the controlled environment. The work described in this paper concentrates on element 4) of this chain and is intended to demonstrate the feasibility of an operator interacting with an intelligent autonomous controller.

Research at the Jet Propulsion Laboratory in Mars Rover teleoperation [14] has provided insights into the impact of extended time delays on teleautonomous operation. Here, the robot builds a local map of the environment using computer vision, and the teleoperator then traces out a path through the world. The path following commands are down-loaded to the robot which then travels autonomously through this portion of the terrain. The process is repeated until the robot reaches its goal destination. This technique is suitable for extra-terrestrial navigation where the environment can be assumed to be stable over long periods of time. The technique however is quite time-consuming, involving extensive computation for the construction of the world model and places a large burden of the low-level obstacle avoidance on the operator.

One recent paper indicates promising results for a teleautonomous mobile robot. Borenstein and Koren [9] extend their previous work using limited dynamic world model acquisition for obstacle avoidance to help an operator control a mobile robot. This work uses a certainty grid representation [22] to provide a basis for path planning by the teleoperator. The map is used for feedback to the operator to provide guidance for navigation in difficult situations. Although it is faster than most methods for driving a robot using world models, it is not a reactive methodology and does not benefit from all of the advantages of reactive control. Further they do not address the issues of extended autonomy through the use of high-level goal-oriented behaviors but rather are concerned solely with facilitating obstacle avoidance.

We now review our work in reactive control and show its applicability in the area of teleautonomous operation.

SCHEMA-BASED REACTIVE CONTROL

Schema-based reactive control is discussed at length in [4]. In this section we present only an overview of this methodology and discuss why it presents an attractive solution to the problem of teleautonomy. As mentioned earlier, schema-based control is only a component of a larger architecture (AuRA) which is described in [2]. The flexibility and robustness present in our system is made available by the integration of a hierarchical planner with the reactive control system [5].

Our version of reactive control utilizes primitive motor behaviors and perceptual strategies that are encapsulated into fundamental units called schemas. This term is derived from their relationship to cognitive psychology and neuroscience [4]. These units comprise the fundamental building blocks of intelligent behavior and perception for our robotic systems. We have developed numerous motor schemas for navigation in indoor settings [4], outdoor campus environs [2], manufacturing domains [6], rough terrain [7], and space and undersea applications [3]. These schemas include:

- *Avoid-static-obstacle*—move away from a non-threatening impediment to motion.
- *move-to-goal*—move towards an attractor.
- *move-ahead*—move in a pre-specified compass direction.
- *stay-on-path*—find a path in the environment and stay near its center.
- *noise*—move in a random direction, useful for both exploration and handling problems with local maxima.
- *docking*—move first in a ballistic then controlled motion towards a docking workstation.
- Various *maintain-altitude, move-up, and move-down* schemas useful for navigation in rough terrain.

These motor schemas are realized on the robot using a variant of the potential fields methodology [19,20]. The robot's reaction is represented as a desired velocity vector for each of the individual schemas. The outputs of each of these concurrent schemas are all summed and normalized and transmitted to the robot for execution. Three important schemas for teleautonomous control are depicted in Figure 1. It must be noted that although the entire field is represented for ease of the reader's comprehension, this field is never computed in its entirety. Only the single component where the robot is currently located needs to be determined for each schema. This makes for very rapid computation that is inherently parallel. An examples of a more complex schema field is shown in Figure 2.

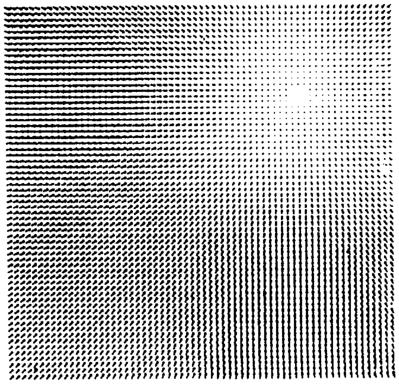

A) Move-to-Goal (Controlled) Schema

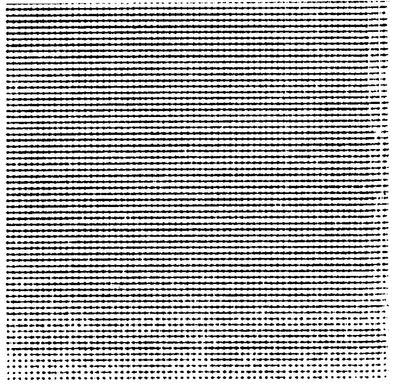

B) Move-Ahead Schema

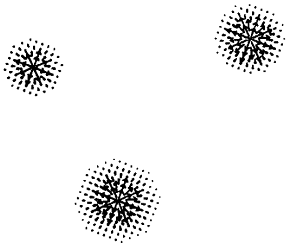

C) Three Avoid-Static-Obstacle Schemas

Fig. 1. Three Primitive Motor Schemas

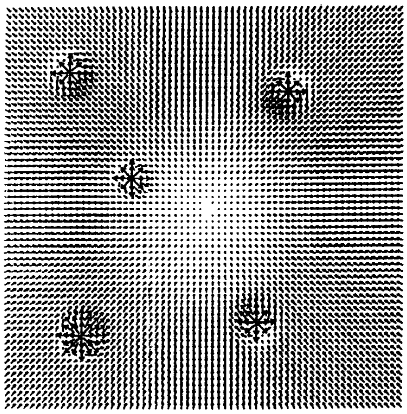

Fig. 2. An Example of Combined Schema Fields (Move-to-Goal (Controlled) + Avoid-Static-Obstable Schemas

Aside from the advantages already described for reactive control (see above), the schema-based system has many other distinguishing characteristics.

- *It is motivated by psychological, ethological, and neuroscientific studies of behavior.*
- *The schemas are concurrent and not arbitrated, allowing each of the behaviors continuously to contribute to the overall goal of the robot.*
- *Summation of the individual behaviors determines the motion of the robot.*
- *It is highly responsive to environmental conditions and copes well with dynamic and hazardous situations.*
- *Within the context of AuRA, it allows for flexible reconfiguration of behaviors based on mission goals, the robot's current conditions, and potential failures.*

It should be remembered that this form of reactive navigation by itself is limited and that a high-level hierarchical planner is used within AuRA to reconfigure the behavioral set based on

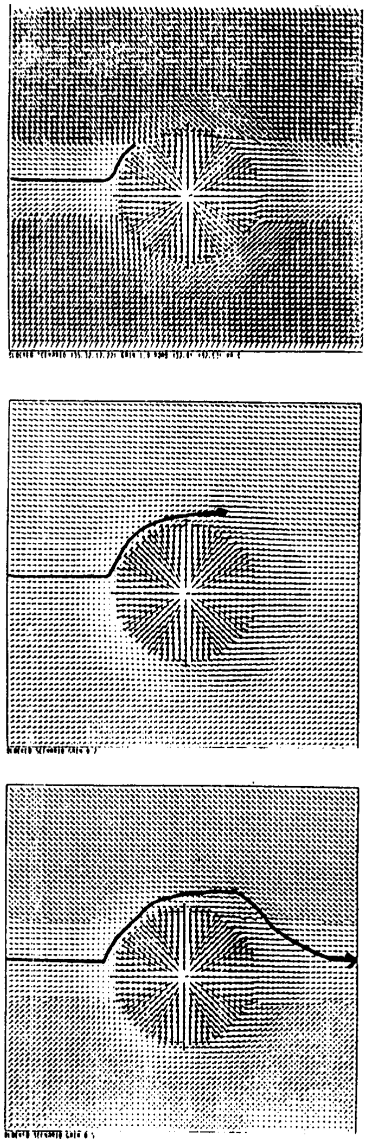

Fig. 3. Teleoperator Supervision Allowing Passage Around an Obstacle on a Path

pre-existing or acquired world knowledge. It is in this role where the teleoperator steps in, supplanting AuRA's hierarchical planner and taking responsibility for coping with potential navigational failures. In the next section we describe several forms in which this intervention can occur; from within the reactive level itself, to acting as a supervisory process for reactive control.

TELEOPERATION

The teleoperator can take several roles in interfacing with the schema-based control system. In one capacity, addressed in the immediately following subsection, the teleoperator is a peer process with the other active motor schemas. In another role, the teleoperator can take a more management-oriented approach, supervising the operation of the schema-based control system. This is discussed below. These roles are not mutually exclusive and it is possible for the teleoperator to step in and out of these capacities freely at both levels.

The Teleoperator as a Schema

We have seen in "Schema-Based Reactive Control" section above, that schema-based reactive control results in a 'sea' of forces acting upon the robot. Each schema agent provides a response to a specific environmental event and guides the robot's overall navigational reaction. It is possible to create a schema that represents the teleoperator's wishes as well, allowing him/her to exert influence within the context of the reactive control system.

This can be done readily in our reactive control regime since we exploit concurrent execution versus arbitration. Perhaps a good analogy to this operation would be *"muddying the waters."* Under normal autonomous operation the robot moves as its active schemas dictate. By allowing the teleoperator to be viewed as a schema, he/she can contribute an additional vector to be summed in with the rest of the active processes, prodding the robot in a particular direction.

The amount of influence can be controlled through the gain of this teleoperator schema. It can be slight to negligible in certain circumstances, or more overt in others, dwarfing the contributions of the other active processes. This can produce forces that range from a simple quick slight nudge (e.g., to rock the robot off of a local maxima) to a much stronger and prolonged shove (guiding the robot out of a box canyon for example).

The teleoperator schema takes as input the desired gain (magnitude), which determines the strength of interaction of the teleoperator with the autonomous control system and the desired direction that the teleoperator would like the robot to move in. The output of the schema is a vector in a frame of reference that is consistent with the other schemas and that is posted along with the other active vectors for consumption by the autonomous control system. The summation of these vectors is then normalized to fit within the constraints of the robotic system and overall mission parameters and then is forwarded to the robot for execution. The teleoperation schema is an asynchronous process and shares all of the characteristics present in the other schemas.

The Teleoperator as a Supervisor

The one thing lacking with this aforementioned method is the ability of the teleoperator to directly influence or even suspend other active behaviors. In some instances this is beneficial, as the operator need not know the machinations of the autonomous control processes but rather can try and rein in the robot much like a horseback rider trying to make his steed behave correctly. The disadvantage of this method is apparent as well: the overall behavioral composition ('personality') cannot be changed by the human operator. There are times when a gentle pony is more effective than a charging stallion. It is potentially useful for the teleoperator to be able to reach into the existing behavioral set and modify it in some manner.

This modification can occur in several ways, typically by influencing the gains and internal parameters of the active schemas. The simplest method is to turn off a troublesome be-

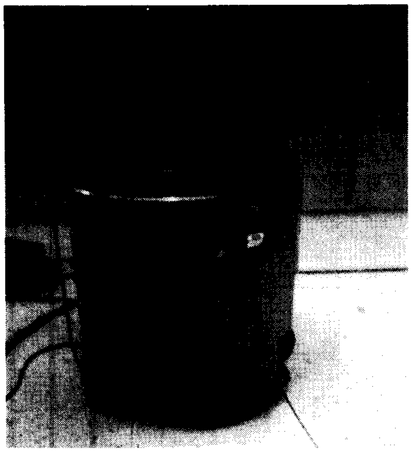

Fig. 4. George—A Denning Mobile Robot

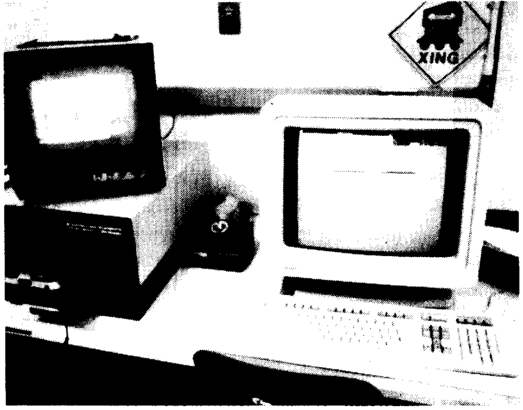

Fig. 5. Teleoperation Workstation

havior by allowing the teleoperator to set its gain to zero. The undesirable behavior no longer influences the action of the robot. One instance where this is useful includes a scenario when the robot is told to move in a particular direction, stay on a path, and avoid crashing into things. Suppose this is an outdoor sidewalk and the path is completely blocked by a parked car. By setting the gain of stay-on-path schema to zero while not influencing the other schemas the robot can go off the path and successfully circumnavigate the obstacle (Fig. 3). Notice this does not require a separate teleoperator schema but can be accomplished by simply allowing the teleoperator to change the existing behavioral configuration.

Another instance would be in allowing the robot to squeeze through tight places without explicit intervention by providing directional support. In tight quarters, if the operator reduces the sphere of influence on the avoid static-obstacle schemas and increases the directional force provided by either the move-ahead, move-to-goal, or docking schemas the robot can be drawn through tight areas that were not previously passable. From a potential fields viewpoint, this produces a stronger overall force through the obstacle field almost as if by increasing the pressure behind the robot. Reiterating, the operator does not provide any explicit information as to the direction in which the robot should move.

Fig. 6. Test Course

Summarizing, the teleoperator as a supervisor has the capability of turning on or off behaviors, modifying the strength of those behaviors (vary the gains), and reformulating the way in which the response to the stimulus of a given behavior occurs (via modification of the internal behavioral parameters). This type of control requires a deeper understanding of the schemas themselves by the teleoperator in order to be effective.

Combinations of both the teleoperator as a schema and as a supervisor are entirely possible. This way the teleoperator can influence the behavior of the autonomous system in a variety of ways, acting in a sense as perhaps a parent would in the guidance of young. This can include tugging on the sleeve of a child to initiate move from a stalled position (e.g., from in front of a toy store window), to physically holding hands through crowded areas (e.g., shopping malls), to influencing the behaviors of the children in different circumstances (i.e., telling the robot explicitly to behave itself differently in differing circumstances—Don't do that when we have company!).

RESULTS

Of course to make this all credible, we need to demonstrate this on an actual robotic system. The results described in this paper are preliminary in the sense that the user interface for the teleoperator is a far cry from being user-friendly; something we hope to address in future research. This forces us to operate the system in lurch mode where the robot moves a step at a time. Nonetheless, the results are sufficiently compelling to demonstrate the overall viability of these methods.

The experimental testbed consists of our mobile robot George (Fig. 4), a microVax II workstation and a Gould IP8432L image processor (Fig. 5). Modifications were made to the existing AuRA software to allow for the introduction of a teleoperator schema and access to the gains of several of the other motor schemas.

A course was laid out that is known to produce difficulties to fully autonomous navigation (Fig. 6). This particular box canyon configuration produces a local minima in the midst of it within which the robot would be trapped if it were not for intervention by the hierarchical planner or in this case by the teleoperator.

The three schemas used for these test scenarios were move-ahead in a direction directly past the obstacle field, avoid-static-obstacle to prevent hitting any of the obstacles (sphere of influence 3.5 feet), and teleoperation to influence either the direction of the robot or permit the adjustment of other schema gains.

Several test runs were performed, each with different degrees of teleoperation control.

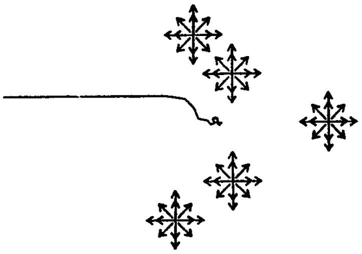

Fig. 7. Pure Autonomy. The Robot Becomes Trapped in the Local Minima and Wanders Aimlessly Without Operator Intervention

1. **Pure autonomy (Fig. 7).**
   This run demonstrates the pitfall a robot encounters without any additional intervention by a human or machine planner. The robot navigates using only the move-ahead and avoid-static-obstacle schemas and finds itself rocking around in a local potential well without satisfying its goal of navigation past the obstacle field.

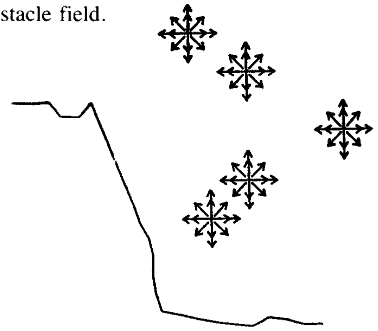

Fig. 8. Direction Intervention. The Operator Senses the Approach of a Box Canyon and Steers the Robot Away from it Before it Gets Trapped

2. **Directional Intervention (Fig. 8).**
   Here the teleoperator can see through the monitor that the robot is entering into a potential trap. By exerting a strong influence on the direction of the robot early on, the potential minima trap is avoided. No gains are changed in any of the other active schemas.

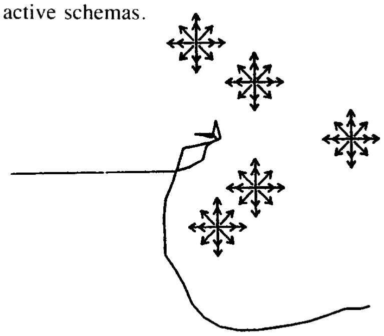

Fig. 9. Trap Extrication. The Operator is Alerted that the Robot is Trapped, at Which Time He Guides the Robot Out of the Box Canyon and Around the Obstacle Field.

3. **Trap Extrication (Fig. 9).**
   In this case the robot is already trapped in the local minima (see (1) above). The operator is notified and the robot is extricated from the problem and guided around the obstacle field at which point the teleoperator relinquished control back to the other active motor schemas. No gains are changed in any of the other active schemas.

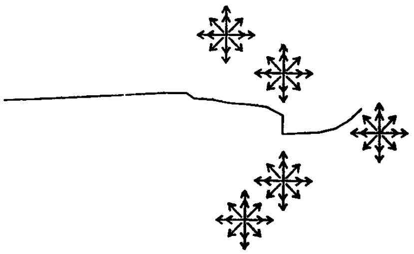

Fig. 10. Guided Narrow Passage. By Both Altering the Gains and Providing Directional Support, the Operator Guides the Robot Through the Obstacle Field.

Squeezing via gain adjustment is similar but with more rocking in the center. In this case, the operator, without providing any directional information, alters the gains on the active schemas, permitting it to safely pass through the field.

4. **Squeezing via gain adjustment only (Fig. 10).**
   In this instance the gains on the avoid-obstacle-schema (reduced) and move-ahead schema (increased) are altered to enable the robot to squeeze though the narrow opening which it would otherwise be unable to move through. There is no directional information given to the robot via the teleoperator, only instructions to change the active gains of the existing schemas.
5. **Guided narrow passage (Fig. 10).**
   Here both directional control and gain alteration of active schemas as in (4) is used. This provides the most efficient passage of these examples while still protecting the robot from collision. This also requires the most skill on the teleoperator's part.

These examples, which are by no means exhaustive, illustrate the promise of these techniques. Considerable work remains to be performed on the user interface aspects of this system such as using a joystick instead of keyboard input, better video feedback, and the use of force feedback based on the existing forces produced by the other active schemas.

SUMMARY AND CONCLUSIONS

Merging reactive control with teleautonomous control holds great promise for short-term deployment of intelligent systems in hazardous environments. In particular the use of schema-based reactive control, which exploits concurrent processing instead of arbitration mechanisms, facilitates this approach.

A methodology has been presented which shows the teleoperator in two distinct capacities: first as a behavior itself, merging and influencing the other existing behaviors; and second as a supervisor, reconfiguring the parameters and gains of the other active schemas as the robot navigates through the world.

Considerable work remains to be done on the user interface aspects of this research. This includes development of improved feedback sensory mechanisms to the teleoperator. When this is completed however, the deployment of teleautonomous-reactive-systems in hazardous environments appears highly feasible.

ACKNOWLEDGEMENTS

This research is supported in part through the Computer Integrated Manufacturing Systems Program and the Material Handling Research Center at Georgia Tech.

REFERENCES

[1] Albus, J., McCain, H., and Lumia, R., (1987). *"NASA/NBS Standard Reference Model for Telerobot Control System Architecture (NASREM),"* NBS Technical Note 1235, Robot Systems Division, National Bureau of Standards.

[2] Arkin, R.C., *"The Impact of Cybernetics on the Design of a Mobile Robot System: A Case Study,"* IEEE Transactions on Systems, Man, and Cybernetics, Vol. 20, No. 6, Nov./Dec. 1990, pp. 1245-1257.

[3] Arkin, R.C., *"Three Dimensional Motor Schema Based Navigation,"* Proceedings of the NASA Conference on Space Telerobotics, JPL Publication 89-7, Vol. 1, Pasadena, CA, pp. 291-300, 1989.

[4] Arkin, R.C., *"Motor Schema-Based Mobile Robot Navigation,"* International Journal of Robotics Research, Vol. 8, No. 4, August 1989, pp. 92-112.

[5] Arkin, R.C., *"Towards the Unification of Navigational Planning and Reactive Control,"* AAAI 1989 Spring Symposium on Robot Navigation, Stanford University, CA., March 1989.

[6] Arkin, R.C. and Murphy, R.A., *"Autonomous Navigation in a Manufacturing Environment,"* IEEE Transactions on Robotics and Automation, Vol. 6, No. 4, August 1990, pp. 445-454.

[7] Arkin, R.C. and Gardner, W., *"Reactive Inclinometer-Based Mobile Robot Navigation,"* Proc. of 1990 IEEE International Conference on Robotics and Automation, Cincinnati, OH, pp. 936-941, May 1990.

[8] Arkin, R.C., *"Integrating Behavioral, Perceptual, and World Knowledge in Reactive Navigation,"* Robotics and Autonomous Systems, 6, (1990), pp. 105-122.

[9] Borenstein, J. and Koren, Y. *"Teleautonomous Guidance for Mobile Robots,"* IEEE Transactions on Systems, Man, and Cybernetics, Vol. 6, No. 20, Nov./ Dec. 1990, pp, 1437-1443.

[10] Brooks, R., *"A Robust Layered Control System for a Mobile Robot,"* IEEE Journal of Robotics and Automation, Vol. RA-2, No. 1, pp. 14-23, 1986.

[11] Brooks, R., (1988). *"A Robot that Walks; Emergent Behaviors from a Carefully Evolved Network,"* A.I. Memo, AI Lab., Mass. Institute of Tech.

[12] Brooks, R., *"The Whole Iguana,"* in Robotics Science, ed. M. Brady, MIT Press, Cambridge, MA. 1989.

[13] Brown, C. and Nelson, R.C., *"Image Understanding at the University of Rochester,"* Proc. DARPA Image Understanding Workshop, May 1989, pp. 122-125

[14] Cameron, J., Cooper, B., Salo, R., and Wilcox, B., *"Fusing Global Navigation with Computer-Aided Remote Driving of Robotic Vehicles,"* Mobile Robots, SPIE Proc., No. 727, Oct. 1986, pp. 85-89.

[15] Connell, J., *"A Colony Architecture for an Artificial Creature,"* TR-1151, MIT-AI Lab, Cambridge, MA, 1989.

[16] Conway, L., Volz, R. and Walker, M., *"Teleautonomous Systems: Projecting and Coordinating Intelligent Action at a Distance,"* IEEE Transactions on Robotics and Automation, Vol. 6, No. 2, April 1990, pp. 146-158.

[17] Gat, E., Slack, M., Miller, D., and Firby, R., *"Path Planning and Execution Monitoring for a Planetary Rover,"* Proc. of 1990 IEEE International Conference on Robotics and Automation, Cincinnati, OH, May 1990, pp. 20-27.

[18] Kaelbling, L., *"An Architecture for Intelligent Reactive Systems,"* SRI Technical Note 400, SRI International, Oct. 1986.

[19] Khatib, O., *"Real-time Obstacle Avoidance for Manipulators and Mobile Robots,"* Proc. IEEE Int. Conf. Robotics and Automation, pp. 500-505, St. Louis, 1985.

[20] Krogh, B., *"A Generalized Potential Field Approach to Obstacle Avoidance Control,"* SME—RI Technical Paper MS84-484, 1984.

[21] Meystel, A., (1986). *"Theoretical Foundations of Decision Making Processes for Design and Control of Intelligent Mobile Autonomous Robots,"* Technical Report, Drexel University.

[22] Moravec, H., (1988). *"Sensor Fusion in Certainty Grids for Mobile Robots,"* AI magazine, Vol. 9, No. 2.

[23] *"Internalized Plans: A Representation for Action Resources,"* Robotics and Autonomous Systems, 6 (1990) pp. 89-103.

[24] Payton, D., *"An Architecture for Reflexive Autonomous Vehicle Control,"* IEEE Conf. on Robotics and Auto., pp. 1838-1845, 1986.

[25] Proceeding of the NASA Conference on Space Telerobotics, JPL Publication 89-7, Vol. 1-5, Pasadena, CA, 1989.

[26] Sato, T. and Hirai, S., *"Language-Aided Teleoperation System (LARTS) for Advanced Teleoperation,"* IEEE Journal of Robotics and Automation, RA-3, No. 5, Oct. 1987, pp. 476-481.

[27] Slack, M.G., *"Situationally Driven Local Navigation for Mobile Robots,"* Ph.D. Dissertation, Computer Science, Virginia Polytechnic Institute, April 1990.

[28] Stark, L., Kim, W., Tendick, F., Hannaford, B., Ellis, S., Denome, M., Duffy, M., Hayes, T., Jordan, T., Lawton. M., Mills, T., Peterson, R., Sanders, K., Tyler, M., and Van Dyke, S., *"Telerobotics: Display, Control, and Communications Problems,"* IEEE Journal of Robotics and Automation, RA-3, No. 1, Feb. 1987, pp. 67-75.

[29] Van de Vegte, J., Milgram, P. and Kwong, R. *"Teleoperator Control Models Effects of Time Delay and Imperfect System Knowledge,"* IEEE Transactions on Systems, Man, and Cybernetics, Vol. 6, No. 20, Nov./Dec. 1990, pp. 1258-1272.

Prof. Ronald C. Arkin is an Associate Professor in the College of Computing at the Georgia Institute of Technology. He received the B.S. Degree from the University of Michigan in 1971, the M.S. Degree from Stevens Institute of Technology in 1977 and a Ph.D. in Computer Science from the University of Massachusetts, Amherst in 1987. From 1977 until 1985 he was a member of the faculty of Hawthorne College in Antrim, N.H. His current research interests include integrative architectures for reactive control and hierarchical planning for mobile robot systems, robot survivability, multi-agent robotic systems, and cognitive models of perception. Dr. Arkin has authored over 40 papers in these areas. He is a senior member of the IEEE and a member of ACM and AAAI. He is included in several Who's Who publications including Who's Who in the world.