AI Magazine Volume 23 Number 1 (2002) (© AAAI)

The AAAI-2001 Robot Exhibition

Bryan P. Adams and Vandi Verma

■ The 2001 American Association for Artificial Intelligence Mobile Robot Exhibition provided an opportunity for AI researchers to interact and share ideas. Despite some difficulties with environment and timing, the primary objective of disseminating information was achieved. A short summary of each robot demonstrates the variety in form and function among the exhibitions.

The 2001 American Association for Artificial Intelligence (AAAI) Mobile Robot Exhibition provided a forum for robots to operate outside a competitive arena. With a wide variety in behavior and form, the robots in the exhibition created a sense of the broad range of function in the robotic community. As one of the only events with no scoring or judges, the exhibition takes on a very different feel and allows robots with different objectives to be shown and discussed. The exhibition has provided past AI researchers with new perspectives and ideas. Exhibitors have also benefited through hands-on exposure to real robots from other institutions.

The extreme diversity, however, presents numerous challenges. At the most basic level, finding the proper way to interact with an audience presents a challenge. Although most visitors are able to quickly and easily understand competitive robots and their tasks, the exhibited robots have no such context to quickly identify their work. The robots are often out of their natural environments, which can further obscure the purpose of a particular project. Finally, with no implicit timing or natural cycle, the exhibition takes on a completely amorphous quality that, in contrast to the natural scheduling of a competition, creates difficulty in finding or holding audiences.

This year's exhibitors were able to overcome many of these obstacles, though, and generate conversation about several important AI topics. Each of the following projects was presented at the exhibition as well as discussed in a more detailed talk on the final day of the workshop.

Air Hockey with Humanoid Robot DB

Programming a robot is often a complex and involved task. For many tasks, examples of successful runs are provided over and over by humans. By allowing the robot to learn and practice behaviors at run time, Darrin Bentivegna aims to create robots that can learn and adjust their behaviors without the traditional programming interface.

This question has been explored in different game domains, with the most recent work involving air hockey against a humanoid robot. After a human has specified some primitives (such as "defend the goal"), the system is able to break the primitives down, use the parts to create goals and desired outcomes, and generate motion as well as an error signal to learn from.

Both a video of the system and a laptop with a simulation were available at the exhibit. Together, they provided observers with both a sense of the difficulty of the task as well as some perspective on what the system has accomplished

–Darrin Bentivegna, Georgia Institute of Technology

Scout

The Center for Distributed Robotics focuses on creating small, inexpensive teams of mobile robots. Using wireless networks to both com-

Figure 1. SCOUT *from University of Minnesota.*

municate with each other and transmit pictures, these robots have been optimized for surveillance tasks. With two eight-bit processors and little sensing, the focus of this research has been on the development of a distributed control system that allows these robots to operate cooperatively using shared resources (figure 1).

In addition to their innovative control, these robots demonstrated several clever mechanical feats. A later revision of the cylinder robots is able to change the radius of their wheels, allowing them to gain greater speed and reliability without compromising their small package size. A spring-loaded winch mechanism also allowed the small robots to jump out of trouble spots and over some small obstacles that presented problems. They also demonstrated a wall-climbing robot that used a vacuum to hold itself to a vertical plywood surface. These robots were also able to interface, allowing the small robots to be carried up (and over) wall-like obstacles.

Because of their small size (a cylinder 11 centimeters long with a 4-centimeter radius), these robots had a great deal of success in the urban search and rescue course without any practice. As an unofficial entry, they were not scored; however, their success in navigating many of the trickier spots in the course with robots that were designed for a separate purpose was a testament to their robust design.

– Paul E. Rybski and Sascha A. Stoeter, University of Minnesota

LAS ABSURDAS MACHINAS

Although most of the robots in the exhibition had a clearly defined task space, other exhibitors displayed work that was more abstract. In an attempt to take the fluency with machine intelligence to other spheres, LAS ABSURDAS MACHINAS (the absurd machines) is an attempt to create a sculptural language using an intelligently controlled robot (figure 2).

The piece is based around Goya's etchings *The Disparates* and *The Disaster of War*. Although the work is designed to recreate some of the same emotions as Goya's work, as the artist describes it, "[M]y work deals with the most fundamental aspect of artistic endeavor: the human condition embedded within a socio-historical framework."[1]

This robot, consisting of a long, flexible arm, attracted curious onlookers with its unique style and feel. The robot, named DIRTY RED, creates motion by changing the tension on cables connected to the arm, creating fluid and biological gesture, which invokes a different feel from the more typical robot arms. The use of old farm implements also creates a mood that stands in stark contrast to the other projects. This work will be on display at the Massachusetts Institute of Technology (MIT) in February 2002.

– Aaron Edsinger and Jeff Weber, MIT AI Laboratory

SLITHER

Using a snake as biological inspiration, SLITHER is an eight-segment robot that is capable of locomotion and navigation using only the movements of its body (figure 3). With each segment phase shifted by 45 degrees, a central controller propagates a sine wave through the body. With passive wheels on the bottom, SLITHER uses the resistance of perpendicular motion to generate forward movement.

By adding a parabola to the sine wave, the robot is able to turn corners. This behavior is used for object avoidance when two forward-mounted whiskers sense an imminent collision. SLITHER is also capable of reverse locomotion (accomplished by reversing the sign on the sine wave).

– Mark and Bill Sherman, Tennyson High School

FLYING GINSU

FLYING GINSU is a hovering robot that uses simple optical mice to regulate its position off the ground (figure 4). Using two wireless mice, the robot is able to estimate its distance from a (minimally reflective) floor. Two powerful downward-pointing fans keep the oval-shaped robot off the ground, and two orthogonal fans provide the robot with the ability to move across a surface.

Although the robot is able to successfully travel using easy-to-find components, the optical mice do not provide enough noise robustness for behavior in more severe environments. As sensors, they are serviceable but have two drawbacks. The first is a need for a surface that will provide sufficient reflection, which limits the surfaces that the hovercraft can maneuver. Second is a general need for stability because

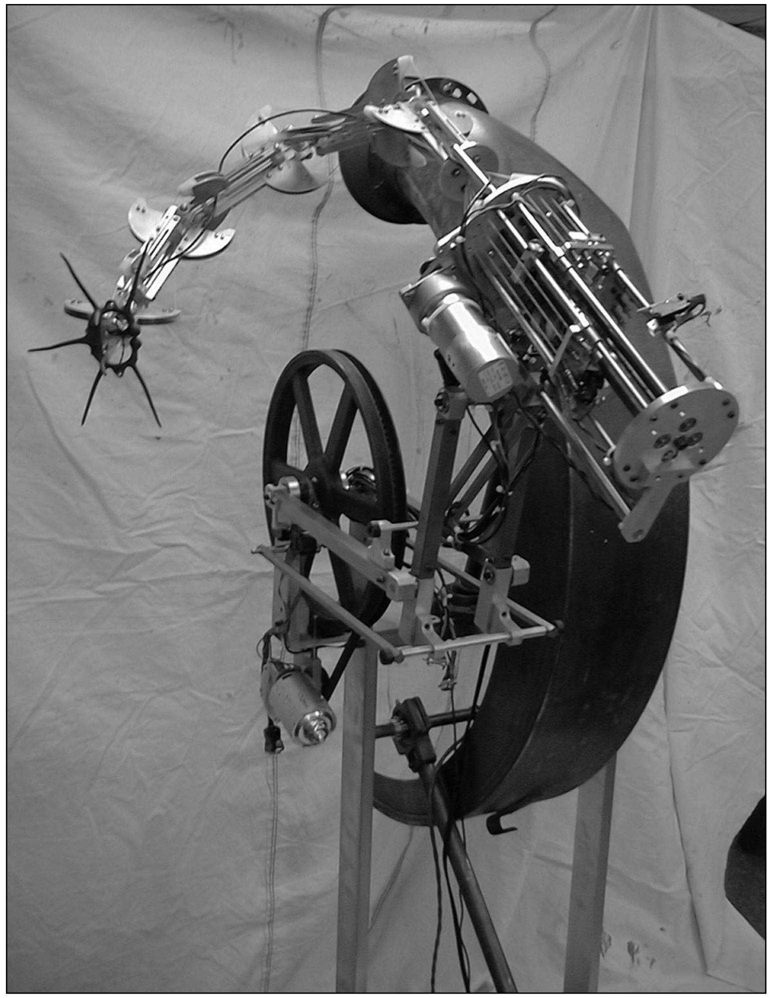

Figure 2. LAS ABSURDAS MACHINAS from Massachusetts Institute of Technology.

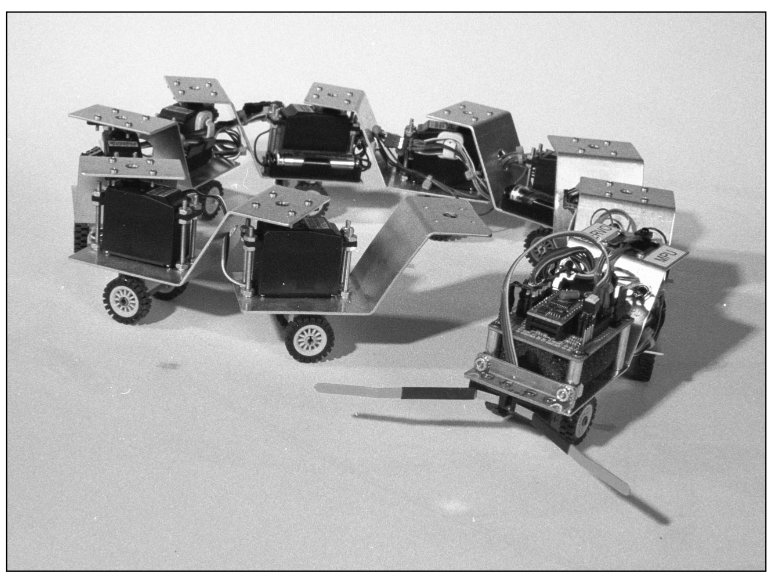

Figure 3. SLITHER from Tennyson High School.

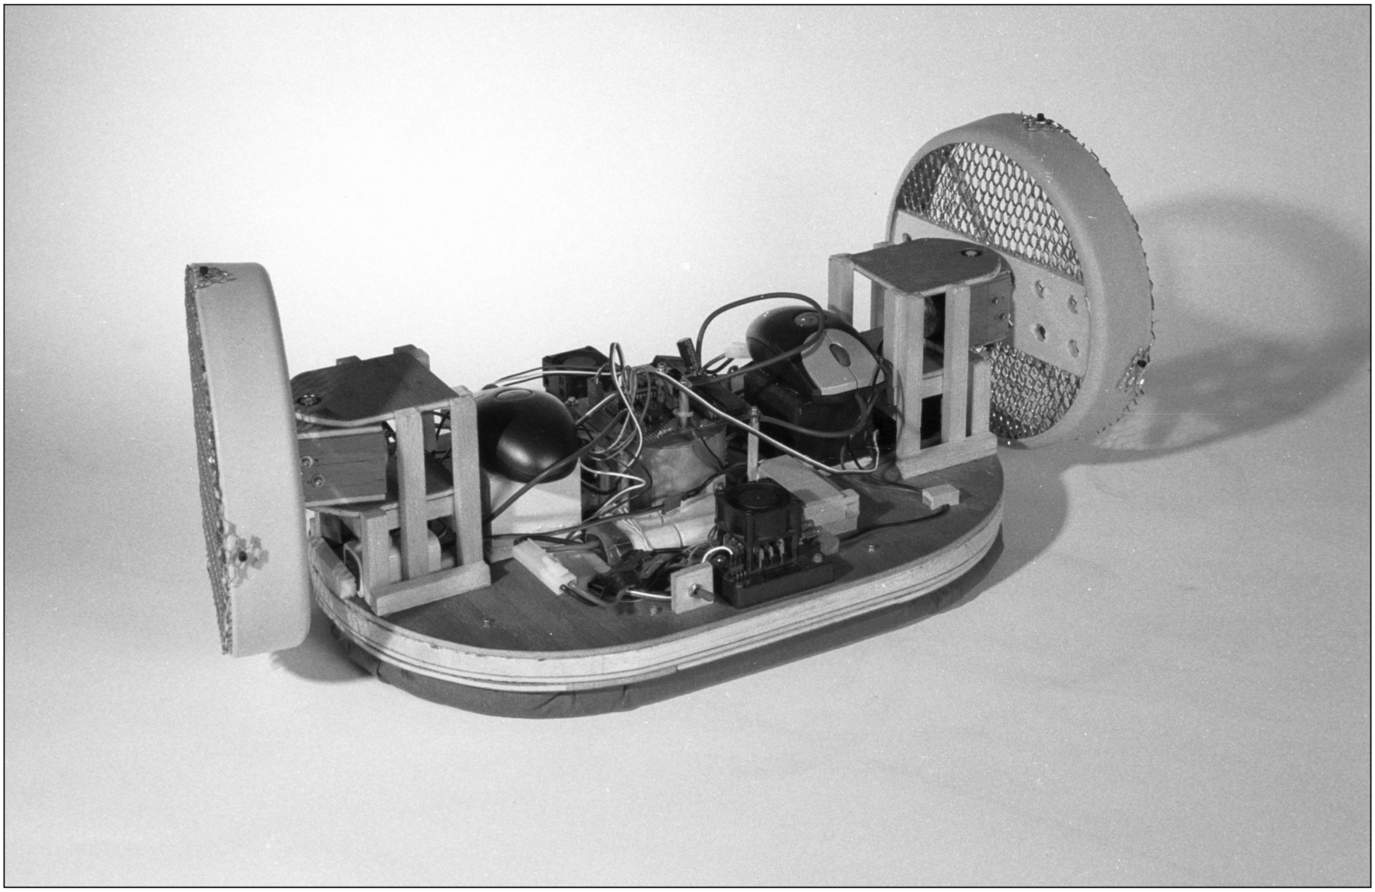

Figure 4. FLYING GINSU *by the Carnegie Mellon University Undergraduate Robotics Club.*

the optical mice can easily be thrown off when operating at a great distance. This robot probably represents the furthest extent that these devices can be used as sensors.

– Brian Kirby andAnthony Rowe, CMU Undergraduate Robotics Club

CEREBUS

CEREBUS is a simple, low-cost, mobile robot built from a commercial robot base (a Real-World Interface MAGELLAN); a low-end laptop computer; and a USB video camera (figure 5). Although CEREBUS can perform a number of tasks, ranging from visual navigation to the following of simple natural language commands, its primary purpose is to be a "self-demoing" robot, in the sense that it can use reflective knowledge of its own capabilities to answer questions about itself and give interactive demonstrations on command.

CEREBUS's demonstration task for the AAAI-2001 Mobile Robot Exhibition was to give an interactive technical talk on itself; it launches a PowerPoint presentation (within its own laptop), reads the text of each bulleted point of each slide, and uses keyword matching to determine the node in its semantic net that corresponds to the topic of the bulleted point. It resolves match ambiguities using a discourse stack to determine which matching node is most closely linked to the previous topic. Having selected a node, it generates text by walking the subtree rooted at this node. When a node is tagged with a behavior, it also runs the behavior to demonstrate it. When the subtree is complete, it moves on to the next bulleted point or slide.

CEREBUS is unusual because it is implemented entirely using feed-forward networks and finite-state machines. Long-term memory is represented using a marker-passing semantic network, and working memory is distributed through the marker bindings of its semantic net, the attentional state of its active vision system, and the activation levels of its logic network. CEREBUS's reasoning and planning components support circuit semantics: Most decisions in the system are continually recomputed from sensor input and the state of work-

CEREBUS: A Higher-Order Behavior-Based System

CEREBUS is an attempt to scale behavior-based robots directly to higher-level cognitive tasks without adjoining a traditional planning system. CEREBUS combines a set of behavior-based sensory-motor systems with a marker-passing semantic network, a simple parser, and an inference network to form an integrated system that can both perform tasks and answer questions about its own ability to perform these tasks. For the 2001 American Association for Artificial Intelligence Mobile Robot Exhibition, CEREBUS's demonstration task was to give a formal technical talk on itself, complete with PowerPoint slides and interactive demonstrations.

CEREBUS is structured as a parallel network of logic gates and finite-state machines. Inference rules in CEREBUS are compiled into a feed-forward logic network, giving CEREBUS's knowledge base circuit semantics: The input of the network monitor the truth values of premises as generated by the sensory systems, and the output of the network track the truth values of conclusions in real time as the premises change. In effect, the entire rule base is rerun from scratch to deductive closure at sensory frame rates. Although this approach sounds inefficient, the CEREBUS rule engine can run a base of 1000 Horn rules with 10 conjuncts each, updating at 100 hertz (100 complete reevaluations of the knowledge base a second), using less than 1 percent of the central processing unit. Using a generalization of deictic representation called *role passing*, the network is able to implement a limited form of quantified inference—a problem for previous behavior-based systems. Rules can be quantified over the set of objects in short-term memory, provided they are restricted to unary predicates (predicates of one argument).

CEREBUS implements *reflective knowledge*—knowledge of its own structure and capabilities—through two mechanisms: (1) a marker-passing semantic network provides a simple mechanism for long-term declarative memory and (2) role passing allows variables within inference rules to be bound to behaviors and signals within the system. The network allows the system to answer questions about its own capabilities, and the role passing allows it to answer questions about its current state and control processes.

CEREBUS can follow simple textual instructions. When a human types a command such as "drive until the turn," its simple parser, which is formed as a cascade of simple finite-state machines, examines each individual word as it is typed, binding the appropriate words to the appropriate roles. In this case, the parser binds the drive behavior to the role activity and the turn sensory signal to the role destination. When it detects a stop (a pause or period), it triggers the handle-imperative behavior, which implements the rules:

> If the signal bound to destination is false, activate the behavior bound to activity.
>
> If destination is bound to a sensory signal and that signal is true, deactivate activity and myself.
>
> If activity deactivates itself, also deactivate myself.

Because this behavior is parameterized by other behaviors, we call it a higher-order behavior, compared to the higher-order procedures of functional programming languages. Other examples are the *explain behavior*, which walks the subtree of the semantic network that describes its argument behavior to produce a natural language explanation of the behavior, and the *demo behavior*, which both explains and runs the behavior. Role passing and higher-order behaviors are easily implemented using parallel networks of gates and finite-state machines, making them a natural choice for the kind of distributed, parallel processing environments often found on mobile robots. They are implemented in GRL, a functional programming language for behavior-based systems that provides many of the amenities of LISP, and it statically compiles programs to a network of parallel finite-state machines.

When CEREBUS gives a talk, it uses a component object model interface to open the specified PowerPoint presentation, then reads the text of each bulleted point and keyword matches it to an appropriate node in its semantic network. It uses a novel distributed representation of a discourse stack to resolve ambiguities using only single-instruction multiple-data marker-passing operations. Having determined the node to which the bulleted point refers, it uses spreading activation to mark the subtree rooted at the selected node as being relevant. It then improvises banter about the topic by continually selecting and explaining the "highest-priority" relevant, unexplained, node. Priorities are computed offline using a topological sort so that if topic *A* is required to understand topic *B, A* will always have higher priority.

By continually reselecting the highest-priority relevant unexplained node using circuit semantics, CEREBUS can respond instantly to changes in relevance when, for example, an unexpected contingency during a demonstration opens up an opportunity to explain a feature. It also allows CEREBUS to cleanly respond to, and return from, interruptions without replanning. However, such topic shifts require the generation of transition cues such as "but first …" or "getting back to …." CEREBUS detects these abrupt topic shifts by tracking the current semantic net node, its parent node, and the previous node and parent. By comparing these nodes, CEREBUS can determine whether it has moved locally up, down, or laterally in the hierarchy or whether it has made a nonlocal jump to an unrelated node. It then generates the appropriate transition phrase.

CEREBUS is far from fluent. It is not intended to demonstrate that behavior-based systems should be the implementation technique of choice for natural language generation. Instead, it shows that parallel, finite-state networks are much more powerful than previously believed. Moreover, by implementing as much of a robot's control program as possible with these techniques, we get efficiency, easy parallelization, and flawless synchronization of the knowledge base with the environment—and all for free.

– Ian Horswill, Northwestern University
Computer Science Department

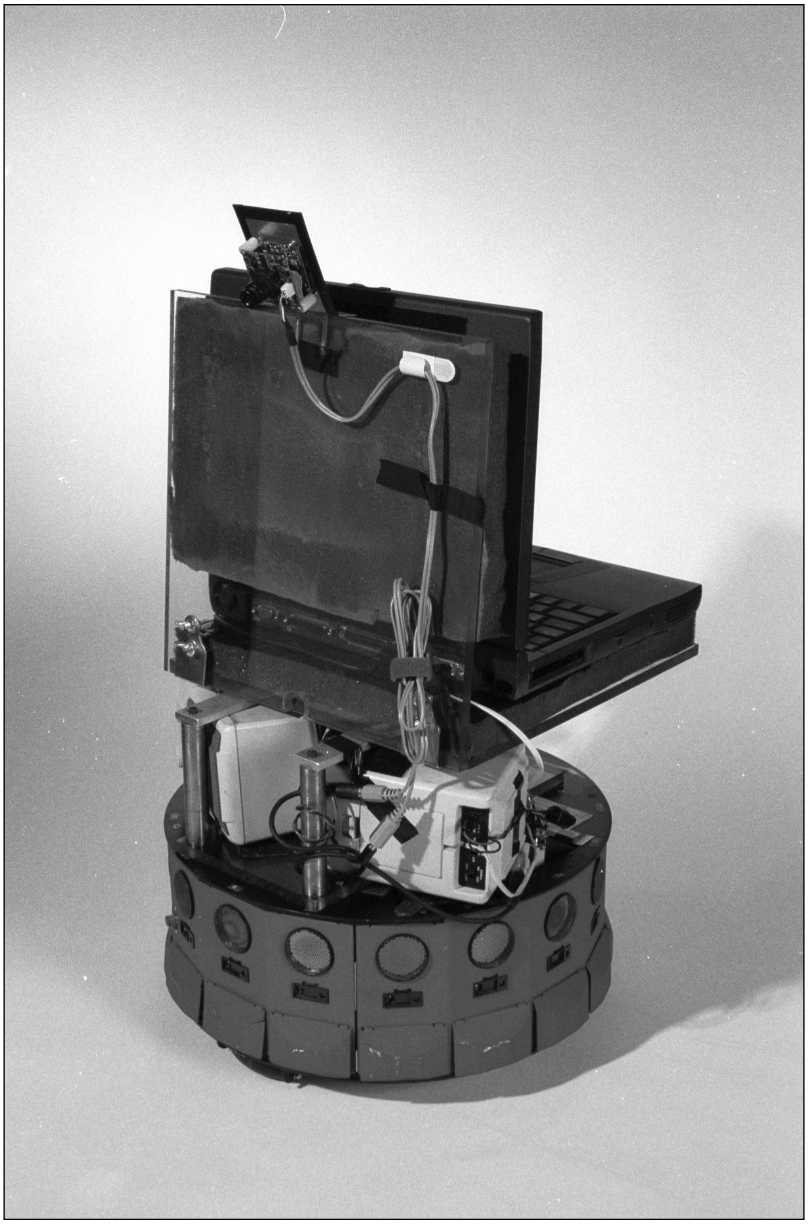

Figure 5. CEREBUS, Northwestern University.

ing memory with each successive clock cycle (100 milliseconds), allowing CEREBUS to seamlessly add or rearrange talk text during a demonstration to take advantage of unforeseen opportunities. CEREBUS is part of an overall research agenda intended to show that behavior-based systems supporting circuit semantics can cleanly implement both sensory-motor tasks and everyday cognitive tasks in a single efficient, integrated architecture.

– Ian Horswill, Northwestern University

GROWBOTS and JUNIOR

The objective of this project has been to develop and evaluate command and control architectures that permit deployment and operational tasking of many small- to mid-sized robots. For large numbers of robots to be deployed as a viable force, human users must be able to interact with functional units rather than issue commands to each individual robot. Rather than exert global, centralized control from above, this project has developed individual robot behaviors that can promote the emergence of swarm intelligence.

The Idaho National Engineering and Environmental Laboratory (INEEL) has worked at an individual robot level toward highly robust suites of sensors and behaviors. To enable useful group behavior, INEEL has developed an innovative, multimodal communication architecture consisting of acoustical chirping and infrared and radio frequency communications. The objective of the command and control system is not only to provide a means of task dissemination but also to facilitate mission planning, retasking, and operator understanding.

The GROWBOTS demonstrated these principles at the exhibition in the form of a spill-detection task. After spilling water on the floor of the exhibit, the robots were able to use a series of chirps and beeps to successfully navigate much of the perimeter of the water. Although the lack of a context made the task somewhat abstruse to the casual observer, the robot behavior was clear and robust.

- David Breummer, Idaho National Engineering and Environmental Laboratory

CYE

The Minnow Project is designed to study intelligent, efficient teams of robots. The goals of the project are to develop and implement teams of robots that effectively cooperate and collaborate to achieve complex tasks in complex environments. Issues of multirobot interactive behavior, communication, cooperative sensing and localization, behavioral and hardware heterogeneity, and learning are key elements of this program. By collaborating behaviors for all aspects of a task (sensing, goal achievement, support), the performance of the team can vastly be improved.

The demonstration presented at the exposition was designed to show a particular element of interactive robot behavior. In the robot soccer domain, which these robots used in the exhibition, it is sometimes required that the goalie leave its home area in front of the goal to move the ball away from it (figure 6). In these cases, the goal is left unguarded and vulnerable in the case that the goalie loses the ball. In this situation, the goalie uses communication to announce its intention to leave the goal. The defensive player, hearing the goalie's announcement, moves back to cover the goal during the goalie's absence, preventing interference with the goalie and providing defense for the goal.

—Tucker Balch, Ashley Stroupe, Rosemery Emery, and Kevin Sikorski, Carnegie Mellon University

Conclusion

As the exhibition looks toward next year, there is a great deal of room for change and improvement. The goal is to increase and improve conversation among roboticists in an environment that stimulates education and discussion. With this in mind, next year's exhibition will have more attention to structure. It is clear now that a great variety in robots will require a great variety in environment in addition to a change in how the demonstrations are scheduled.

However, the concluding workshop confirmed that this year provided an effective exchange of ideas. Groups from different backgrounds provided insight into both the ideas and research behind the robots as well as the diversity of missions and approaches that result. Maintaining this free and open communication channel while creating new opportunities for more individual exchange, both among exhibitors and between exhibitors and the public, will be critical for success in the following years.

Note

1. www.ai.mit.edu/people/edsinger/Absurdas/absurdas_statement.htm.

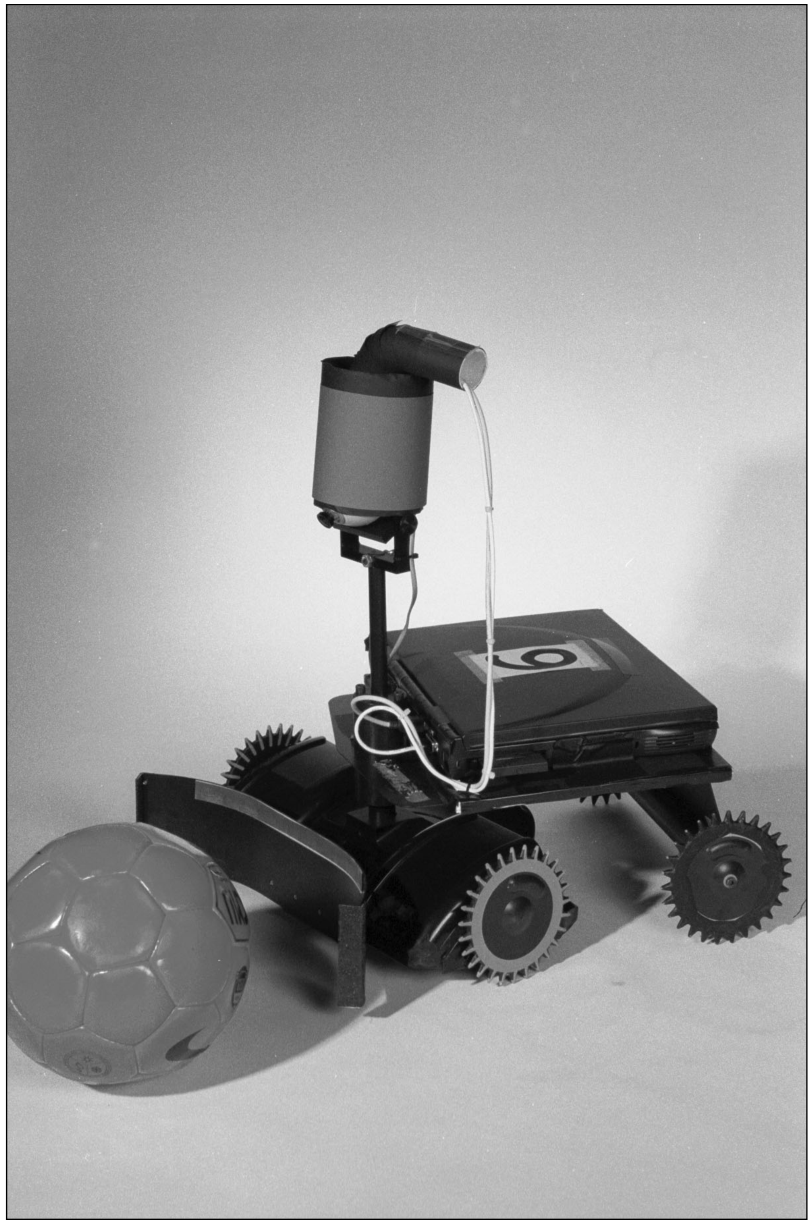

Figure 6. CYE*, Carnegie Mellon University.*

Bryan Adams is completing his Ph.D. with Rodney Brooks's Living Machines Group and is interested in

how the homeostatic principles of living creatures can be applied to robots. He received his B.S. and M.Eng. in electrical engineering and computer science from the Massachusetts Institute of Technology. His e-mail address is bpadams@ai.mit.edu.

Vandi Verma is a doctoral student in the Robotics Institute at Carnegie Mellon University (CMU). She is interested in using machine learning and AI to enable adaptive behavior in robotic systems. She is currently working on fault detection and diagnosis

for the HYPERION rover at CMU. Some of the earlier projects that she has been involved with are autonomous navigation for Mars rovers at CMU, the K-9 rover for Mars exploration at NASA Ames, the EVE rover for lunar exploration at LAAS-France, and the NOMAD rover for Antarctic meteorite search at CMU. Her e-mail address is Vandi_Verma@gs93.sp.cs.cmu.edu.