AI Magazine Volume 14 Number 1 (1993) (© AAAI)

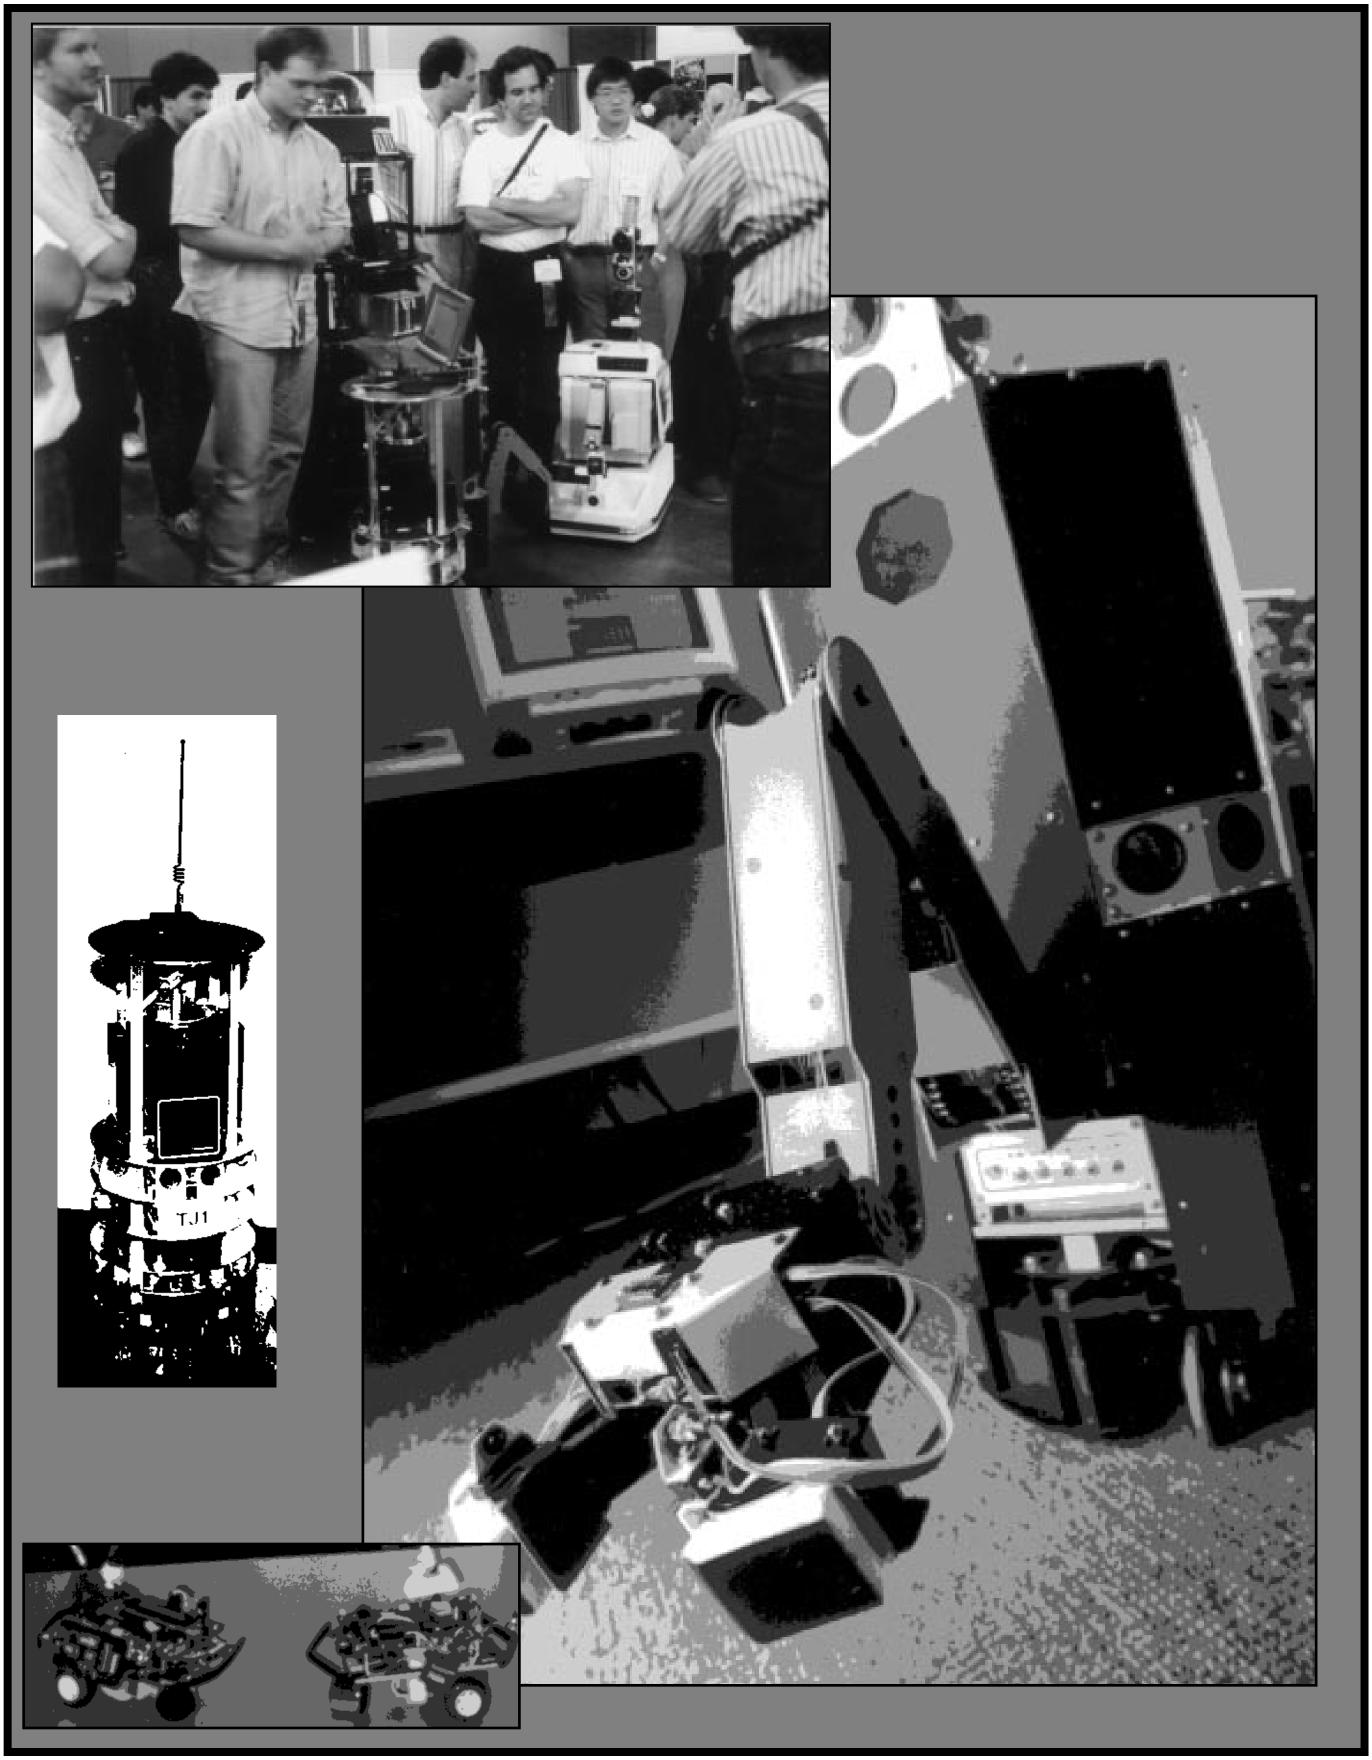

1992 AAAI Robot Exhibition and Competition

Thomas Dean and R. Peter Bonasso

■ The first Robotics Exhibition and Competition sponsored by the American Association for Artificial Intelligence was held in San Jose, California, on 14–16 July 1992 in conjunction with the Tenth National Conference on AI. This article describes the history behind the competition, the preparations leading to the competition, the threedays during which 12 teams competed in the three events making up the competition, and the prospects for other such competitions in the future.

Robot technology has progressed significantly in recent years. Advanced sensors and efficient actuators and power systems are now available for a wide range of applications. Related technology in vision, planning, and learning has also matured, and the time is ripe for a marriage of these technologies. Further, the growing economic incentives for robotic systems point the way to challenging research. Although the demand for high-precision industrial robots to paint, weld, assemble, and package products has leveled off, the demand for robots in less structured environments is expanding: Markets are opening up for robots that polish supermarket floors; deliver meals to hospital patients; perform remote sensing for military and government applications; and do a variety of tasks in hazardous environments, such as nuclear power plants, hazardous waste sites, and underground mines.

Robot applications, especially those in the service sector, require relatively sophisticated reasoning abilities because service robots must cope with a far wider and less controllable range of situations than robots that are confined to factories. The recent advances in technology and anticipated changes in the demand for advanced robots led the American Association for Artificial Intelligence (AAAI) to sponsor a combined exhibition and competition to call attention to applications and technological challenges that might profit from its members' expertise.

The exhibition-competition was designed to promote interaction among a wide range of industry and university researchers interested in these new applications of intelligent robotics. It featured robots, videos, and graphic images from university and industrial research laboratories around the world. The competition stressed the range of tasks that robots must master to move about in an unstructured environment and cope with interference, both deliberate and inadvertent, from the humans with whom they share this environment. These robotics applications focus on apparently mundane tasks that are, in fact, every bit as complicated as more traditional AI tasks (for example, playing master-level chess or providing decision support for air-traffic control) in terms of the interplay of behaviors and the physical interaction of the robot with the real world. These applications demand a degree of autonomy and robust execution that is unparalleled in prior commercial applications of AI technology.

The History of the Competition

The competition was to involve three stages, one each day, and two distinct phases. The first stage involved robots navigating in a cluttered environment and interacting with the people. This single stage made up the first phase, which served as a qualifying round to ensure that the robots were well enough behaved to participate in the subsequent stages. In the second stage, the robots were to explore their environment and find and identify 10 objects. In the third stage, the robots

Photo: Jim Wattenmaker

After an exhausting 2 days of preparation at the convention center interrupted rarely by sleep, 10 robots were ready to go on the morning of the first day…

were to carry out commands issued by the judges, where a command consisted of a set of objects to visit in a specific order. The robots could make use of any information acquired in the second stage to expedite the execution of these commands. The second and third stages together made up the second and final phase of the competition. Prizes (first through fifth place) were awarded for each of the two phases as well as for a variety of special categories suggested by the judges.

The idea for the competition and its three-stage design originated with Tom Dean. Dean was inspired by a panel on household robots suggested by John Kender for the Ninth National Conference on AI (AAAI-91). As program co-chair for AAAI-91, Dean was forced to take over the organization of the panel on household robots when Kender had to choose between attending yet another conference and attending the birth of his child. The panel was really a set of presentations with accompanying videos of robots. Following the presentations, a startup company, Nomadic Technologies, Inc., Mountain View, California, demonstrated its robot wandering around in a crowd of fascinated attendees. The response of the attendees and the press was astonishing. It was clear that even the most cynical reporters and veteran researchers were not immune to the appeal of mobile robots.

At the time, AAAI was looking for ways to revitalize the national conference to regain some of the excitement and enthusiasm of the early years. Dean used the surge of enthusiasm generated by the household robot panel to convince the incoming president of AAAI, Pat Hayes, the conference chairs for the 1992 national conference, Paul Rosenbloom and Pete Szolovits, and the AAAI executive council that such a competition was feasible; they should authorize a group to organize it; and, perhaps most importantly, that they should allocate AAAI funds to cover the costs, which were estimated to run in the neighborhood of $25,000 to $30,000. Finding a group of dedicated individuals turned out to be easy. Pete Bonasso, Jim Firby, Dave Miller, and Reid Simmons had already begun discussions during AAAI-91 to refine the specifications.

Any number of robot competitions have been run in the past at other conferences. Most of these stressed hardware or low-level control issues that were not particularly appropriate for a conference on AI. It was certainly easy to imagine a competition that stressed more complicated forms of reasoning; the trick was to stage such a competition to occur within the temporal, environmental, and political constraints imposed by a national conference. Conference attendees were used to seeing videos of robots performing in the lab, but often, the performances were carefully orchestrated, used a variety of tricks to make the robot appear to be moving faster than it actually was, or simply documented one successful run out of many failed or less impressive runs. The competition would be held in public, in real time, and without the extensive controls possible in the laboratory. Most of those attending the competition would have high expectations.

Those organizing the competition were able to convince themselves and those sponsoring the competition that the research community had achieved the necessary level of robustness in hardware and software to stage such a public event and bring credit to the field and education and entertainment to the general community. There were plenty of occasions during the time leading to the competition that the competition organizers were to doubt themselves on this score.

The set of tasks had to be possible within the current state of the art, it had to exercise capabilities that were of interest to the AI community, and it had to be entertaining and fit in a three-day venue. We chose a set of tasks that involved interaction with people, navigation and exploration in an unknown environment, and path planning and command following. The idea was to encourage machines that exercise commonsense capabilities involving planning, control, and rudimentary learning and spatial reasoning.

The competition organization was broken down into three committees: the rules committee (chaired by Bonasso) to set up the detailed specifications for the robots and the tasks; the hardware committee (chaired by Firby) to deal with the myriad of details involving computing equipment, such as communications for robots that require access to remote computing resources; and the local arrangements committee (chaired by Rick Skalsky) to deal with the construction of the competition arena, displays, and shipping and to work with the people at the convention center. Dean served as coordinator and dealt with sponsors, exhibitors, press, and the many other details that arose during the course of the preparations. Dean also solicited and organized the judging for the competi-

PHOTO: JIM WATTENMAKER

CARMEL
MARC I (Michigan Autonomous
Robot Competition, Team I)

Carmel is a Cybermotion K2A mobile platform with a ring of 24 sonar sensors. Motor control and sensor firings are controlled by 2 on-board computers. A third, 486-based, personal computer (PC)–compatible on-board computer runs all the competition software and communicates with the other on-board computers. Object detection is accomplished using a color camera connected to a frame grabber that occupies a slot on the 486 PC. All these systems, plus the robot, are powered by 2 on-board 12-volt batteries, thus eliminating any need for a tether and allowing for fully autonomous operation. CARMEL *uses error-eliminating rapid ultrasonic firing to accomplish fast obstacle avoidance while it navigates. It uses a certainty grid technique for global path planning. The objects in the competition were recognized by using bar-code tubes that were detected by the color camera.*

"Perhaps my greatest surprise in judging the robot competition was how much I learned about an area I thought I knew something about."

– Jim Hendler

Photo: SRI International

FLAKEY
SRI International Mobile Robot Team

Flakey is a custom-built octagonal robot approximately 2.5 feet in diameter and 3 feet high, weighing 300 pounds. Locomotion is achieved with 2 independently controlled wheels located on the sides; maximum speed is about 5 feet a second. Sensors include a bottom-mounted ring of touch sensors, a circular array of 12 Polaroid sonar sensors, and a structured-light system using an infrared laser and a charge coupled device (CCD) camera. Internal computers include microprocessors that control the wheel motors and sonars and a sun-3 master that coordinates the other controllers, collects and processes the structured light, and communicates with an off-board controller through a 200-kilobyte wireless ethernet bridge. The controller is written in Lisp, with c subroutines where necessary for speed. Multiple real-time processes are implemented as a software round-robin queue. There are processes for basic communications and screen display, motion control, sensor interpretation, and mapping. Behaviors written in fuzzy-control rules direct Flakey's motion. Object modeling uses a surface-based representation, extracting coherent surface patches from the sonars and structured-light sensors.

tion. As the competition neared, Bonasso and Dean found themselves mired more and more in these many details, to the extent that their efforts were directed almost exclusively on the competition. Our comments regarding the work involved in such a competition are meant to educate those considering similar projects and prepare them (and their employers) for the effort involved.

The initial specifications for the competition were out in September 1991, distributed widely on the electronic mail networks and published in the fall issue of *AI Magazine*. Several labs made tentative commitments and participated in subsequent refinements of the specifications in the following months. On the basis of these commitments, we were able to secure funding from the AAAI executive council. Part of the funding would be used to actually stage the competition and part to provide scholarships to support student participants wanting to attend the competition and help defray the cost of shipping the larger robots.

Later, we were able to obtain additional funds for scholarships from the National Aeronautics and Space Administration (NASA) through the efforts of Mel Montemerlo and Peter Friedland. General Motors and the MITRE Corporation also made donations to the scholarship fund. These scholarships were essential to the success of the project. Some of the teams involved several students, and the larger robots required in excess of $500 to ship to San Jose. Participation in the event was certainly an important educational experience for the students. In addition, given the effort required in repairing hardware damaged in transit and modifying software to adapt to the conditions in the convention hall (for example, radio frequency [RF] noise and high-intensity lighting), many of the competitors would not have been able to complete the events without the effort that the multistudent teams were able to muster.

In the winter and spring of 1992, some changes were made in the list of competitors, but most of the early entries stuck with it to the end. The complete list of participants appears in the sidebar to this article. At the 1992 Spring Symposium Series at Stanford University, a number of those involved with the competition met to design the arena layout and finalize the specifications. The final design required materials that were relatively inexpensive, and, at the same time, not too difficult to detect with available sensors. Foam-core sheets were used for the walls of the two rings, cardboard boxes for the obstacles, and polyvinyl chloride (PVC) drainage pipe for the objects. The rest of the arena was constructed from standard components used elsewhere at the conference.

The last months leading to the competition involved a flurry of activity that included finding additional sponsors, cajoling judges into spending a good portion of the conference involved with the competition, convincing computer manufacturers to lend equipment, compiling a library of videos and graphic materials, and finding additional exhibitors to supplement the competitive events. Nobody was sure how many teams would show up for the competition or whether they would be able to compete even if they did manage to show up. The videos and additional exhibitors were a form of insurance to help salvage the event should severe problems cripple the competitors.

AAAI was also concerned that the press not misrepresent the event; so, there was considerable effort to prime the press. Jim Wattenmaker handled press relations for AAAI, contacting the Cable News Network (CNN), *Popular Science*, and a host of other television and print news services before the convention. We tried hard to make it clear that the aims were to push the state of the art in autonomous agents and mobile robotics, educate the press and the AI community in the recent progress in the field, and provide an exciting educational opportunity for the participants that would result in a useful transfer of technology through a concerted effort on a common task.

The Competition

For the participants and organizers, the competition began Sunday, 12 July, when the shipping crates were unpacked and the arena constructed in the San Jose Convention Center. Inevitably, equipment was damaged in transit, tools and parts were left behind, software complications arose for those using borrowed computing machinery, and a wide variety of problems existed that had to do with the arena and the convention hall. With regard to this last point, the two competition rings were larger than expected, dwarfing the smaller robots; the obstacles (white cardboard boxes) were not what we ordered; the lighting posed problems to vision systems; and perhaps the most important factor for many of the competitors, the hall was a veritable soup of RF noise: portable microphones, transmitters used by the press, two-way radios used by

PHOTO: HOLLY YANCO, MIT

ODYSSEUS

CMU Mobile Robot Team

ODYSSEUS is a small, wheeled robot equipped with an arm, sonar sensors, and a camera system. It is connected by radio links to a pool of computers that control the robot. ODYSSEUS combines low-level reactive mechanisms with global planning and exception handling, using a wide variety of control and AI techniques, ranging from A planning and hidden Markov model–based speech recognition to artificial neural networks and reinforcement learning. On the lowest level of behavior, the robot uses several fast on-board obstacle-detection and obstacle-avoidance mechanisms for safely operating in unpredictable, dynamic environments. ODYSSEUS's global navigation is map based. The sonar sensor is used for incrementally constructing a model of its environment. The camera is used for detecting target objects. ODYSSEUS is able to identify and navigate to particular objects as well as explore its environment autonomously to gain knowledge. The robot is operated using a speaker-independent speech-recognition and speech-generation system. In addition, a graphic interface is used for monitoring the operation of the robot.*

> *"Robots bring people back to the core of what AI is aimed at: learning how to build intelligent creatures."*
>
> *– Ben Kuipers*

the convention center employees, and dimmers and starters for the powerful halogen lighting, not to mention that the competitors used a wide range of incompatible devices for transferring data between robots and remote computing equipment. Problems with communications and equipment damaged in transit would figure prominently in the outcome of the competition.

During the two days prior to the competition, hardware was repaired and replacement equipment flown in. Software underwent major revisions throughout the competition. There was an enormous amount of sharing of expertise and equipment; the competitors helped one another in a myriad of ways. The fact that they understood the task in intimate detail made communication almost telepathic; all the competitors wanted to know exactly how the others had approached the competition. Some used general approaches that could deal with a wide variety of tasks; others adopted a task-specific approach. Some had primitive hardware but sophisticated software, and some relied on hardware specifically designed for the task at hand. Indeed, the experience was exhilarating for all the competitors because for five days, there was assembled under one roof some of the brightest, most intelligent, and incredibly clever hardware and software hackers in the nation.

After an exhausting 2 days of preparation at the convention center interrupted rarely by sleep, 10 robots were ready to go on the morning of the first day. Not knowing exactly how long each entry would take, we decided to run robots in both rings with some overlap. The judges were partitioned into 2 teams of three, 1 for each of the 2 rings. One ring had been reconfigured for the smaller robots by dividing it in half using a wall of large boxes. Before each robot was to compete, the robot team briefed its judging team on the technical approach to be used. This briefing allowed the judges to better understand the behavior of the robot and adjust their scoring accordingly. Indeed, some of the judges felt the whole experience was quite educational.

To wring out whatever orchestration bugs there might be, we decided to start the first day's competition an hour before the general public and press were to be allowed in the arena. This approach proved useful because the first teams to start (based on a random ordering) had several difficulties that caused us to restart with a different ordering by the time the public arrived.

The first day was meant to exercise the

"Personally, I got a concrete understanding of many of the issues which are debated in the robotics world; I suspect that many of the people watching did as well."

– Howard Shrobe

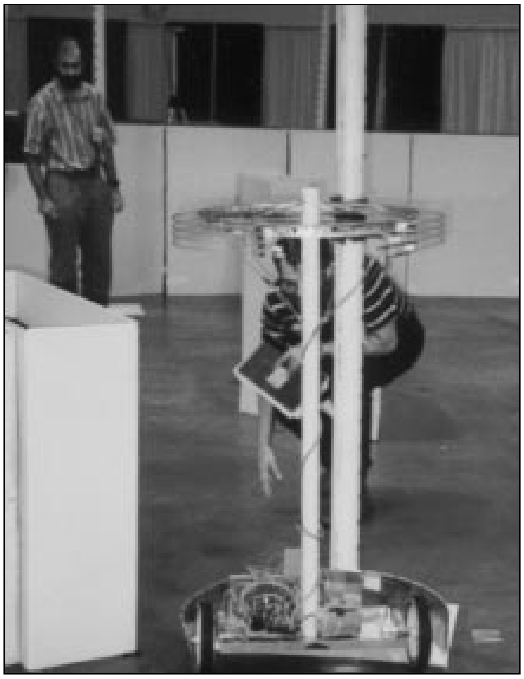

Photo: Holly Yanco, MIT

SCARECROW

MIT/JPL/ISR Mobile Robot Team

SCARECROW is a completely self-contained, totally autonomous mobile robot that was specifically designed and built for the 1992 mobile robot exhibition at the American Association for Artificial Intelligence convention. The robot is approximately 4 feet tall, roughly cylindrical in shape, with a diameter of approximately 2-1/2 feet. The robot moves using differentially driven wheels and a set of casters. All batteries and motors are located within 6 inches of the ground to keep the robot stable. The robot is equipped with tilt sensors to disable the actuators should the robot start to tip. The robot has a programmable top speed that can be adjusted to less than 2 feet a second. The robot is ringed with soft bump sensors along its lower base. The robot has a sensor ring on its top that can read a conductive bar code of any object that is labeled and that it comes in contact with. The robot detects obstacles and objects by collision and distinguishes objects from obstacles using the head sensor. SCARECROW maintains almost no state information. For it to visit the objects in the correct order, it is programmed by the operator with object numbers in the desired order. It uses a simple finite-state automata to track which object it last saw and which it wants to see next.

robots' abilities to navigate autonomously, avoiding obstacles both fixed and moving. The judges were encouraged to interact with the robots. To the teams' consternation, the judges decided to push their capabilities to the limit, both to determine some criteria for distinguishing among the competitors and to better understand the limits of the technical approach used. Ken Forbus, Ben Kuipers, Larry Matthies, Stan Rosenschein, Yoav Shoham, and Beverly Woolf were the judges for the first day, dashing in front of the robots; moving boxes; and, even in some cases, lying down in front of the robots. A number of the robots were cheered for their graceful movements, their agility, and their various means of communicating with the judges and their handlers. Computer-generated sounds and voice synthesizers were particularly appreciated by the crowd.

Although on the first day, there were some problems involving robots communicating with remote computing devices through radio modems, several of the robots were able to perform in the first stage using only on-board computing. The University of Michigan robot did all its computing on board during all three stages. Another robot, FLAKEY, ran into communications problems in the preliminary trials, but the team (from SRI International) was able to solve its problems by strapping a portable computer to the top of the robot with duct tape.

The attendees were captivated by one of the competitors, Jacob Milstein, who with his father, David Miller, built an intriguing special-purpose machine designed just for this competition. Miller and Milstein's entry, SCARECROW, was a real crowd pleaser but did not fare particularly well in the first stage because its primary means of sensing involved crashing into things and people (deemed obstacle avoidance by collision detection by one observer) and then dashing off in the opposite direction. However, SCARECROW brought added excitement to the competition with its clanking and buzzing and whistling and wobbling. The crowd became attached to the robot, groaning when it got stuck and cheering when it recovered.

The winner of the first phase was IBM Watson's TJ. TJ combined ultrasonic sensors (sonar) for long-range obstacle detection with near infrared sensors for short-range obstacle detection, creating a robust, agile navigation and obstacle-avoidance system. SRI's FLAKEY was a close second, followed by Michigan's CARMEL, Georgia Institute of Technology's BUZZ, and University of Chicago's CHIP to make

PHOTO: JIM WATTENMAKER

BUZZ

GEORGIA INSTITUTE OF TECHNOLOGY / DENNING MOBILE ROBOTICS TEAM

BUZZ is an MRV3 Denning robot, a three-wheeled holonomic robot, 27 inches in diameter and 36 inches high. buzz has 24 ultrasonic sensors mounted in a ring 22 inches above the floor. A 68000 microprocessor running os/9 acts as the robot interface, controlling other microprocessors in charge of the motors and ultrasonics. High-level commands are transmitted to the robot over a 9600-baud serial link. One black-and-white CCD video camera is used to find the objects. A small light source is mounted near the camera to illuminate bar codes made of reflective tape placed on the objects. Our primary hardware system is a Sparc IPC computer with a digitizer board. Both video and serial data are transmitted through radio frequency (RF) data links. BUZZ utilizes a schema-based reactive control system to make its way around. This control system is a subset of the hybrid hierarchical-reactive autonomous robot architecture (aura) that controls a set of processes. There is no explicit hierarchy of processes, but the processes can be grouped under the headings "planning," "motor control," and "perception."

"I think it's good to have people outside robotics realize how hard the problems are, at least by seeing what the realistic state-of-the-art robots are like today. We're just not close to R2D2s and C3POs yet, and people should realize why not."

– Maja Mataric

the top five ranking robots. The other competitors finished in order from sixth through last: Miller and Milstein's SCARECROW, NASA Johnson Space Center's (JSC) SODA-PUP, Brown University's HUEY, Carnegie Mellon University's (CMU) ODYSSEUS, and MITRE's UNCLE BOB.

There was remarkable agreement on the basic form of the robot architectures. Almost every architecture consisted of a multilevel control system, coupling low-level routines for sensing, primitive navigation (for example, wall following), and obstacle avoidance that operate continuously in highly responsive feedback loops with high-level routines for planning and map generation that rely on the low-level routines to keep the robot moving synchronously with real-time events. There were systems based on variants of Rod Brooks's subsumption architecture (IBM Watson) and Stan Rosenschein and Leslie Kaelbling's situated automata approach (MITRE) as well as systems based on fuzzy control theory (SRI). There were also systems that did not adopt any particular architecture but incorporated the lessons that such architectures have taught us over the years. For example, the Brown entry was a modular, object-oriented software system, utilizing multiple processors to manage low-level and high-level routines asynchronously. In general, the software developed for this competition was extraordinary for its sophistication and ease of modification. In some cases, code was considerably revised during the competition to cope with hardware failures and complications introduced by conditions in the conference hall.

The first day was fraught with small problems orchestrating the event; happily, the audience was both patient and enthusiastic, and many people volunteered their services. In particular, Lonnie Chrisman from CMU demonstrated his talent as a robotics announcer by giving a blow-by-blow account of the events, providing background on the competition and the participants, and interviewing competitors to obtain details of the robots' performance during the trials. His announcing turned out to be especially important on the second day when most of the robots took much longer than expected to accomplish their exploration task. Terry Weymouth and Holly Yanco took turns as announcer during periods when Chrisman was otherwise occupied.

During times in which there were no robots competing in the rings, various robots roamed about randomly; Mark Gordon from Georgia Tech showed off a flying robot

PHOTO: NOMADIC TECHNOLOGIES

SODA-PUP

NASA-JSC Automation and Robotics Division
Mobile Robot Team

The SODA-PUP robot uses as its mobile base a nomad 200 mobile robotic system built by Nomadic Technologies, Inc., Mountain View, California. The Nomad 200 has a cylindrical structure, consisting of a lower base and an upper turret. The drive system provides synchronous translational and rotational motion to nomad's 3 wheels as well as independent rotation to the turret. The turret functions as the superstructure of the mobile robot, housing both the on-board computer systems and the majority of the sensors. Sensors include a contact bumper, 16 ultrasonic and 16 infrared sensors, and a color CCD camera. The SODA-PUP robot communicates through a 9600-baud modem with off-board os/2- and unix-based computers that provide the primary processing power and serve as the operator interface. The SODA-PUP robot architecture is built on individual processes, which communicate with each other using an in-house–developed interprocess communications package. These processes are grouped into sense, perception, knowledge, motivation, planning, and action modules, and data progress from low-level data-acquisition processes through higher-level data-interpretation functions and, eventually, to the planning and action processes.

(which could not fly for safety reasons); Yanco from the Massachusetts Institute of Technology showed off two small robots, BERT and ERNIE, that performed in tandem; Jim Slater and David Zhu showed off robots from Nomadic Technologies; Jeff Kerr from Zebra Robotics demonstrated a small robotic arm; and the San Francisco Robotics Society showed off a prototype for a robot vacuum that it is developing.

In the hours between the end of the first day's activities and the beginning of the second day's, there were many minor and major hardware disasters. Several teams made frantic calls to manufacturers and colleagues back home to ship replacement parts by overnight shipping companies. There were still more discoveries of damage in transit. One particular poignant story involved MITRE's UNCLE BOB, which was discovered to have a wheel assembly badly bent out of alignment. The team was able to compensate by modifying the software to some extent; however, in their exhaustion later that night, one of the team members accidentally shorted out a part that resulted in freezing the top rotating sensor platform. UNCLE BOB missed the second day but managed to return on the last day to finish the competition despite the damaged drive assembly and the incapacitated head. MITRE 's performance was a tribute to good software design.

On the second day, the two rings were cluttered with obstacles constructed from cardboard boxes, as in the first stage. In addition, there were also 10 "objects," which were 8-foot poles constructed of PVC drainage pipe. The different teams were allowed to rig the poles with any sort of sensor stimuli that they wanted: Bar codes, colored rings, reflective tape, and infrared beacons were all used by one team or another. The robots were to explore the ring and identify all the objects. Robots that were not able to identify labels on objects were forced to first differentiate a pole from a box—which is actually pretty tricky using sonar alone—and then distinguish each object from the others using object location. Robots that used sonar exclusively for pole recognition generally ran slower than robots using specific stimuli attached to the poles. FLAKEY performed well using a hybrid strategy, identifying potential pole candidates by their sonar signature at long range, then verifying with a structured light sensor, which was extremely reliable, at a shorter range. FLAKEY found 8 of the 10 poles in the allotted time by circumnavigating the ring and making occasional forays into the

PHOTO: IBM

TJ

IBM AI Mobile Robot Team

TJ is built on a real-world interface (RWI) B12 three-wheeled omnidirectional synchrodrive mobile platform. The robot is 12 inches in diameter, stands 3 feet high, and weighs 55 pounds. TJ has an array of 12 short-range infrared proximity sensors, 4 longer-range sensors to aid in local navigation, and 8 ultrasonic sensors for obstacle avoidance. Object detection is accomplished through the use of a planar, rotating, infrared, phase-based range sensor. Object identification is achieved using a small video camera connected to an on-board low-bandwidth vision system that subsamples the image for analysis. Object identities are announced by a speech synthesizer. Local navigation and reflexive goal-seeking behaviors are controlled by an on-board network of 8 microprocessors connected by a high-speed multidrop serial network. These microprocessors run a distributed behavior-based subsumption controller that reevaluates the robot's situation every 70 milliseconds. All these systems are self-contained and powered by an on-board battery system. Supervisory control and a symbol-based human interface are provided by an off-board IBM PS/2 Model 70 workstation connected to the robot over a bidirectional 915-megahertz, 9600-baud, spread-spectrum Arlan radio link.

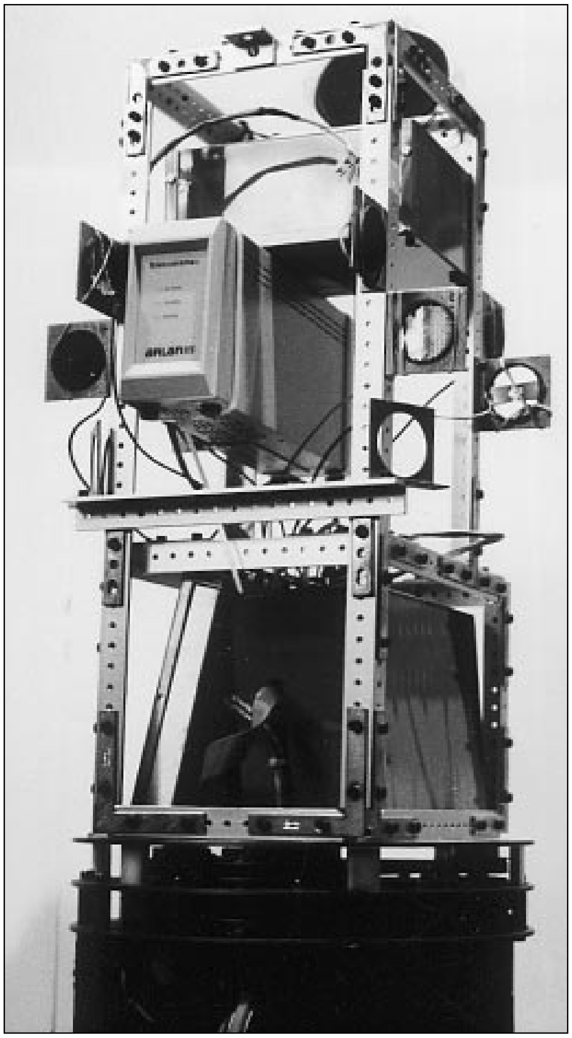

Photo: Brown University

HUEY

Brown University Undergraduate Artificial Intelligence Group

The robot consists of a RWI B12 mobile platform and a superstructure housing on-board computers, sensors, and radio communications equipment. The robot is 12 inches in diameter, stands 18 inches high, and weighs 40 pounds. It is powered by two 144–watt-hour batteries. HUEY's sensors consist of 6 sonar transducers arrayed in pairs (2 forward pointing, 2 each right and left pointing), 2 additional sonars, and a pair of forward-pointing infrared sensors. There are several small on-board computers with a Motorola MC6811 handling most of the on-board coordination and communications. An Alan radio modem, operating at 9600 baud, provides communications with an off-board Sun sparcstation. The system interprets the robot's sonar readings to generate probabilistic descriptions of interesting spatial locations and links these descriptions into a network whose arcs are described by the operations the robot must follow to get from one location to another. Because of the limitations of HUEY's sensors, the robot can only identify objects by their sonar characteristics.

interior.

This stage was meant to encourage teams to construct some sort of internal representation of the ring, encoding the location of the objects for use in the final stage of the contest when the judges would ask the robots to visit a set of objects in some particular order. Most of the robots did indeed build some sort of map of the ring, and during the competition, the judges were shown graphic displays depicting these maps.

At the suggestion of some of the competitors, the teams competed in the reverse order of their first-day finish. This idea turned out to be a good one because it built something of a crescendo to the end of the day. Many observers began to have favorites among the entries, and on the final day, the high-scoring robots drew the largest crowds.

The robots with cameras or other more sophisticated sensors tended to do better in this stage. In particular, robots with cameras on stalks or tall robots that could see over the boxes excelled, and those relying on local sensing were at a disadvantage. CMU, Georgia Tech, JSC, and Michigan all used cameras that allowed them to detect objects over the boxes from some distance. FLAKEY, with its hybrid strategy using sonar and structured light, did well without a long-range camera.

TJ was hampered by a broken analog-to-digital converter that crippled its pole-recognition system. Nevertheless, the IBM team members were able to modify their software so that TJ found and identified three objects. The team members were somewhat disappointed; however, they were able to perform well enough to compete in the final stage of the competition, in which they did well.

Brown's was the only entry to rely entirely on sonar for sensing the environment. Despite a sophisticated technique for interpreting sonar data and positioning to obtain multiple views for pole recognition, their robot was not able to recognize any of the poles during the allotted time of 20 minutes. The Brown entry, HUEY, proved far too slow because of both the method of sensing and the delays resulting from communications problems. The undergraduate students responsible for this robot were the youngest competitors in the contest (with the exception of Jacob Milstein) and were sorely disappointed because HUEY's performance was not sufficient to allow them to compete in the final stage of the competition.

The surprise of the second day was SCARECROW, which managed to find and detect 7 of the 10 objects in the allotted time. As could be expected by its pseudorandom walk, the robot found the first 4 objects in 5 minutes and spent the remainder of the 20-minute period looking for the rest. SCARECROW was enormously popular with the crowd, and Jacob Milstein kept the crowd entertained with his antics throughout SCARECROW's run. SCARECROW lurched about the ring always seeming on the brink of crashing to the floor and smashing into boxes and occasionally brushing against a pole with its tall, circular antennae that were used to recognize poles. Miller and Milstein had rigged each pole with a steel wool ring at exactly the height of the circular antennae. If SCARECROW brushed against a pole in the right way, the steel wool served to close an electric connection involving the antennae, allowing SCARECROW to register that it had located a pole. SCARECROW did not use a systematic method for exploring its environment. Instead, it performed a good approximation to a random walk. SCARECROW offered a dramatic illustration of the theoretical result that a short random walk in an undirected graph can produce visits to every location in the graph with high probability.

SCARECROW did not learn anything during its exploration of the ring. This fact, coupled with the fact that several other robots found all 10 objects, put it at somewhat of a disadvantage in the final stage of the competition in which it would be asked to visit specific locations. SCARECROW would have to use the same random walk strategy to carry out the judges' request. However, what SCARECROW lacked in intelligence (like its namesake in *The Wizard of Oz*, SCARECROW had no brains), it made up for in raw speed, and as a result, SCARECROW was still a contender going into the final stage.

Near the other end of the complexity spectrum was Michigan's CARMEL. CARMEL managed to find and visit all 10 objects in under 10 minutes, which was significantly faster than any other robot. The key to CARMEL's impressive performance in stage 2 was that it could zip about the ring at speeds in excess of 300 millimeters a second and still avoid obstacles. This speed was more than twice as fast as its speediest competitors. Controlling CARMEL at these high speeds was a unique obstacle-avoidance method developed by researchers at the University of Michigan Mobile Robotics Lab. This system has two major components: (1) a method for detecting and rejecting noise and crosstalk with ultrasonic sensors, called error-eliminating rapid ultrasonic firing (EERUF) and (2) an obstacle-avoidance method called the vector field histogram


PHOTO: WATTENMAKER

UNCLE BOB

MITRE Autonomous Systems Laboratory
Mobile Robot Team

UNCLE BOB is a Denning MRV with a ring of 24 sonars mounted approximately 0.7 meters from the floor and 6 sonars mounted approximately 0.1 meters from the floor, with baffles to increase the dispersion angle of the sonars. The mission sensor is a self-contained laser target reading system on loan from Denning that can detect bar codes made of reflective tape. The robot's onboard 68000 microprocessor controls the actuators and sensors except the laser, which communicates its data through the RS232. The robot's intelligence is located in a MACINTOSH QUADRA that is on board and powered by a 100–ampere-hour DC source. The MACINTOSH is connected to the base and the laser target system through the rs232 and communicates to an off-board MACINTOSH IIFX through an RS232 9600-baud RF link. The MACINTOSH IIFX is used to display robotic telemetry and start and stop the robot's autonomous activity. Robot behaviors are coded in REX and then sequenced with reaction plans coded in GAPPS. All navigation on the robot is orchestrated through the use of navigation templates, which are similar to potential fields except for the addition of knowledge about the goals, which avoids the traps of local minima.

> *"The biggest controversy [in] judging involved the question of whether judging should be purely observational—in other words, whether a robot should be rewarded or penalized on the basis of the judge's opinions about the kind of program running it or whether the score should be based on its observable performance."*
>
> *– Martha Pollack*

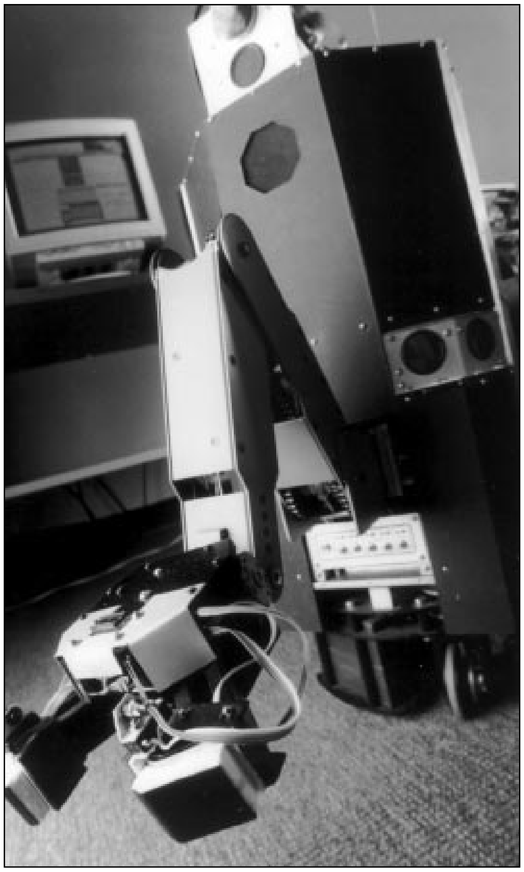

Photo: Matt Gilson / University of Chicago

CHIP

University of Chicago Mobile Robot Team

CHIP is roughly cylindrical, with a diameter of about 18 inches and a height of about 3 feet. It rides on a synchrodrive base from RWI and is equipped with a bumper near floor level to detect collisions. The robot has 8 long-range sonar range sensors and 16 infrared short-range sensors arranged around the robot 18 inches above the floor. The robot is also equipped with a simple HERO II arm that can reach the floor. Sitting on top of the robot is a color camera on a custom pan-tilt head. The image from this camera is broadcast from the robot by radio to a Sun SPARCstation equipped with several DataCube image-processing boards. All motors and sensors on the robot are managed by on-board microcontrollers that communicate with a MACINTOSH computer off the robot using a 9600-baud radio modem. CHIP uses a potential field-based approach for navigation. Should obstacles appear suddenly close to the robot, low-level routines stop or move the robot away. Objects are perceived using color-histogram visual identification. Each object to be identified is marked with a color-coded sign that can be seen from all directions.

(VFH). The combination of EERUF and VFH proved to make CARMEL well suited to high-speed obstacle avoidance.

The judges on the second day were Jerry Dejong, Mike Georgeff, Jim Hendler, Ian Horswill, Matt Mason, and Martha Pollack. The judges worked hard to deal with disparities in the size, the speed, and the sensor capabilities of the various robots. A large part of the judges' discussions centered on what constituted an object classification. Some robots clearly identified the objects at a distance and then navigated to them to establish the object location, but SCARECROW didn't identify an object until it ran into it. After the tallies of the second stage, CARMEL (Michigan) was on top, followed by BUZZ (Georgia Tech), FLAKEY (SRI), ODYSSEUS (CMU), and SCARECROW (Miller) to make the top five. SODA-PUP (JSC) was sixth, followed by TJ (IBM) and HUEY (Brown).

The hours between the end of day two's activities and the beginning of day three's activities were long, with many of the competitors working to prepare for the final stage. Lessons learned from days one and two were being used to tweak parameters and revise code. Chicago and MITRE encountered severe hardware problems and were hacking much of the night. A critical vision board on Chicago's entry burned out, and the team was recoding furiously in an effort to get around the problem. Team members were also awaiting the delivery of a new board in the hope that a replacement board would correct their problems. The next day would be even more frustrating for the Chicago team when the replacement board arrived and promptly burned out just like the earlier one. After an extraordinary effort, Chicago had to withdraw from the second phase entirely.

For the final day of the competition, eight robots that had completed the earlier stages were still in good enough shape to continue. MITRE's robot had not run on the second day because of hardware problems; however, it managed to complete enough of the course early in the morning before the events of the final day to compete in the final stage. UNCLE BOB was a crowd favorite because of the extraordinary effort made by the MITRE team in coping with the robot's disabilities and because of the sounds that it generated to punctuate its exploits (for example, "I swear I will not kill anyone" from Arnold Schwartzenegger in *Terminator II* and "I've got to rest before I fall apart" from C3PO in *Star Wars*).

The final stage would be run in two heats, with the robots performing in reverse order of their standings on the previous day. In the morning, we saw the entries from JSC, MITRE, and IBM Watson in addition to SCARECROW. TJ recovered from its low ranking on the previous day with a great run using the map that it built on day 2. SCARECROW ran a close second to TJ, randomly visiting many poles but using a mechanical state counter when it found a desired pole in the desired order. SODA-PUP and UNCLE BOB put in credible performances but did not appear to be in contention for the top places.

The judges on the final day were Erann Gat, Steve Hanks, James Crawford, and Henry Kautz, all sharing duties because of other conference conflicts, and Maja Mataric, Nils Nilsson, and Howie Shrobe. Again, the judges had the extraordinarily difficult job of comparing these very different robots that were performing in different rings and, in some cases, performing subtly different tasks.

In the afternoon, two of the top four contenders from the previous day, ODYSSEUS and FLAKEY, turned in impressive performances, both of them completing the specified tasks although a little slower than TJ and SCARECROW.

Georgia Tech faltered after its second-ranked performance on the previous day. Most likely, a number of factors were involved, but RF interference was suspected. As excitement over the competition grew during the three days, the press gathered in increasing numbers with their array of electronic devices. It was almost impossible to control all those who were trying to record the event. Because Georgia Tech relied on communications with a remote computer for this stage, it is suspected that communications problems led to the downfall of its robot in the final stage.

An additional unforeseen problem was the unabashed intrusion of the press into the rings during all stages of the competition. The reporters were obviously unaware that they were recognized as obstacles to be maneuvered around and might have changed the performance of some of the slower robots (although the top five finalists all coped well with moving obstacles in stages 2 and 3).

The last robot to run, Michigan's CARMEL, and the leading contender from the previous day started with some problems. There was a false start in which one of the team members entered the wrong command, and the robot confused everyone by flawlessly executing a sequence quite different from that given by the judges. On its second attempt, CARMEL seemed to get disoriented and at one point stood by the wall of the ring scanning the crowd for the obstacle it was looking for. It was as though an Olympic champion had advanced to the finals, leading in the standings, only to fail inexplicably in the last seconds. In this case, however, CARMEL recovered. The judges gave the robot one final try in which it executed the sequence flawlessly, turning in the fastest time. This performance, coupled with that of the previous day, led the judges to award CARMEL first place in the second phase.

Giving the robot additional tries within the 20-minute period in any of the stages was not part of the original rules but was unanimously adopted by all the judges on all 3 days. As Nilsson put it, "We've all given demos before."

In the closing ceremonies, Hayes presented the awards. In addition to the top five places in the two phases, there were a variety of other awards. MITRE was awarded a prize for its dramatic recovery, JSC for its impressive performance with the shortest development time, and Miller and Milstein for their innovative design.

Future Competitions

The competition drew a great deal of interest, and many people commented that they would like to see additional competitions staged at the national convention in future years. Many people volunteered to help out with these competitions.

Whether future competitions take place depends on a number of factors, including the continued support of sponsors, the continued efforts of competitors and organizers, and the degree to which such competitions are seen as forwarding science. Our purpose was to encourage approaches that relied on techniques that borrow from and extend current research in planning, learning, and spatial reasoning.

For the most part, entries that used simple sensors and sophisticated methods for dealing with uncertainty in sensing and movement found it difficult to compete with those robots equipped with more sophisticated sensors. FLAKEY, however, was a notable exception, performing well in both phases and relying exclusively on sonar for long-range sensing. As was apparent from SCARECROW's performance, sophisticated reasoning was not necessary to perform passably in some of the stages. CARMEL's fast obstacle avoidance, long-range sensing, and the fact that all its computing was done on board were significant factors in the Michigan team's success.

It is our hope that those participating in

"I think it was a big win and really helped to make the AAAI special this year."

– Jim Hendler

"Nearly everyone I spoke with at AAAI thought it was the highlight of the conference."

– Martha Pollack

future competitions can build on the work of those competing in the 1992 competition. In particular, we would like to see robots that combine the best of the 1992 robots. We would also like to see the robots challenged by a wider range of tasks, encouraging more general reasoning capabilities. It is clear from the performance of the robots in this competition that robust mobile platforms capable of sophisticated reasoning are technologically feasible; we would like to see future competitions stretch this technology to its limits.

That the specially designed robots did particularly well was neither a surprise nor a disappointment. We believe that the tension between approaches tailored to a particular problem and those striving for generality is extremely healthful for the field. Future competitions should continue to play these different approaches against one another in an effort to better understand the basic tradeoffs. FLAKEY's performance demonstrates that it is possible to perform well without the use of extremely accurate dead reckoning or sophisticated sensing. The lesson is, as Kurt Konolige of the SRI team pointed out, that sophisticated processing can compensate for less sophisticated sensing.

There was some fear that the competition would draw attendees away from the talks or the paid exhibitors. In fact, because one had to go through the exhibition hall to get to the competition, the competition seemed to draw people to the exhibition hall. Also, attendees tended to use the competition arena as an interesting place to meet and spend some time between talks that they wanted to attend.

Many people commented that the competition had made them think about what problems there are in robotics and about how AI might further contribute in this area. Others thought that the general idea of a competition might be applied to other areas, such as natural language understanding, in much the same way as it was applied to robotics.

Finally, we might have taught the press a thing or two about autonomous robots, as evident from an anecdote presented by Hayes about the CNN coverage. CNN had apparently decided not to cover the event, but the local news stories got the network interested again. When CNN arrived on the last day, the reporters were predisposed to thinking of the event as just another mechanical olympics. However, when team members began showing them that the robots were running on their own (no one behind the curtain), the CNN team began scrambling to make a story of the event. The final clip that showed on the CNN news hour the week following the competition indeed discussed the possibilities of intelligent robots assisting in hazardous environments and with space exploration.

Whatever the prospects for future events, the competition staged at the Tenth National Conference on AI (AAAI-92) was a resounding success in the eyes of the competitors and the attendees. Several of those participating mentioned that the pressure of the competition had enabled them to compress a year or more of research and development into just a few months. It now seems likely that other such competitions will be held, if not at the next conference, then the following. Your comments and suggestions are welcome.

Acknowledgments

We want to thank some of the people who made the competition possible: the past and current presidents of AAAI, Danny Bobrow and Pat Hayes; the executive director of AAAI, Carol Hamilton; the program chairs for AAAI-92, Paul Rosenbloom and Pete Szolovits; Mel Montemerlo and Peter Friedland of NASA; Steve Holland of General Motors; and the executive council of AAAI. We also want to thank Howard Moraff of the National Science Foundation and Erik Mettala of the Defense Advanced Research Projects Agency for their encouragement during the early stages of organizing the competition.

Thomas Dean received a B.A. in mathematics from Virginia Polytechnic Institute and State University in 1982 and his Ph.D. in computer science from Yale University in 1986. He is currently an associate professor on the faculty of the Department of Computer Science at Brown University in Providence, Rhode Island. His general research interests include temporal and spatial reasoning, planning, robotics, learning, and probabilistic inference.

Pete Bonasso is the division assistant for AI and robotics in the Information Systems Division of the MITRE Corporation, Washington C3I operations. Currently, he is the principle investigator for MITRE's research in intelligent control of autonomous vehicles. Besides robot intelligence, his research interests include qualitative reasoning, planning, and AI system integration. He received his B.S. from the United States Military Academy at West Point in 1968 and Master's degrees in operations research and computer utilization from Stanford University in 1974.