

# The 1994 AAAI Robot Competition And Exhibition

*Reid Simmons*

■ The third annual AAAI Robot Competition and Exhibition was held in 1994 during the Twelfth National Conference on Artificial Intelligence in Seattle, Washington. The competition was designed to showcase and compare the state of the art in autonomous indoor mobile robots. The competition featured Office Delivery and Office Cleanup events, which demanded competence in navigation, object recognition, and manipulation. The competition was organized into four parts: (1) a preliminary set of trials, (2) the competition finals, (3) a public robot exhibition, and (4) a forum to discuss technical issues in AI and robotics. Over 15 robots participated in the competition and exhibition. This article describes the rationale behind the events and the rules for the competition. It also presents the results of the competition and related events and provides suggestions for the direction of future exhibitions.

This article describes the planning, organization, and results of the 1994 AAAI Robot Competition and Exhibition, which was held in Seattle, Washington, on 31 July to 3 August in conjunction with the Twelfth National Conference on Artificial Intelligence (AAAI-94).

The first such competition was held during AAAI-92 in San Jose, California. That event, which attracted much media attention, featured 10 robots literally searching for "tall poles" in a field of large boxes (Dean and Bonasso 1993). Its main objective was to demonstrate that mobile robots could perform interesting tasks safely and autonomously. The 1993 competition, held in Washington, D.C., shifted to the domain of office robotics and introduced more complex and realistic tasks and environments. The robots participating in

the three events had to maneuver in an office full of furniture, find and deliver an object in an arena of offices and corridors, and move large boxes around (Konolige 1994).

The three-day competition in 1993 was very successful but quite grueling for both participants and judges. The inevitable question was whether there was enough interest to sustain a third annual competition. An early e-mail call for participation was encouraging: Over 40 replies, from researchers all over the world, demonstrated that interest was still high for a competition to compare the state of the art in autonomous mobile robots. Let the games continue!

## Organization and Format

As chair of the 1994 competition, my first task was to determine its general direction. With the help of an able committee, consisting of David Kortenkamp, Erann Gat, and Jim Firby and past competition chairs Tom Dean, Pete Bonasso, Kurt Konolige, and Terry Weymouth, we decided to take an evolutionary approach—to make the 1994 events similar to those at the 1993 competition. The main reasons for adopting this approach were (1) it was felt that office robotics was an interesting task domain; (2) many researchers were already exploring the domain; (3) shifting domains would place a burden on participants who had put in much effort preparing their robots in 1993; and (4) many of the robots were unable to complete the 1993 events, so it did not make much sense to choose more difficult tasks for 1994.

The competition has three basic aims: (1) comparison of the state of the art in



*Figure 1. The Participating Robots in the Competition.*

autonomous mobile robots, (2) entertainment and publicity, and (3) technical exchange of information among robotics researchers. In the past, there was no clear distinction between these aims: A single competition served the first two, and the technical exchange of information was carried out informally while the participants worked at readying their robots for the competition.

I felt, however, that each of these aims could be achieved more effectively by separating them and designing events tailored for each. Comparison of the robots, which I considered to be the most important aim, would be achieved through a series of competitive events, each consisting of multiple trials. The trials would be held prior to the start of the conference technical sessions (during the tutorials and workshops). These events, although not exactly closed to the public, would be run primarily for the benefit of the participants. The idea was that with multiple trials, robots that did consistently well should be, in general, ranked higher than those that performed exceptionally well once but failed in subsequent trials. In addition, it gave the judges a chance to increase the complexity of the tasks from trial to trial, to find the performance limits of the robots (it also gave teams a chance to fine tune their algorithms between trials or even to try completely different strategies). For example, in the Office Delivery event, most teams were successful in the first trial, but no team performed the last

trial without some outside help. Although this part of the competition was the most grueling—two solid 8- to 10-hour days of trials—it was over in time to enable the participants to actually attend the conference (because of my participation in the previous competitions, I had attended a grand total of two conference sessions in two years!).

For the aim of entertainment and publicity, we decided to hold both a robot exhibition and the finals of the robot competition for the public. The events, which were professionally videotaped, were open to AAAI-94 participants, the public at large, and the media. The exhibition, which was organized by Firby, Jim Hendler, and Ian Horswill, was designed to showcase research that combined AI and robotics in interesting ways. Participants submitted abstracts describing what AI techniques their robots demonstrated, and those included in the exhibition were given a 20-minute time slot and an arena in which to “strut their stuff.”

The finals featured head-to-head competition between the robots that had performed best during the preliminary trials. A running commentary and visual aids helped the audience understand what was going on. Barbara Grosz, president of the American Association for Artificial Intelligence, presented certificates to all participants and prizes to the top teams. Afterward, all the robots lined up for a group photo, along with their proud owners (figure 1).

The last part of the competition was a technical forum, held the day after the exhibition and finals. The idea (originally suggested by Mel Montemerlo) was to provide a formal opportunity for researchers to discuss techniques used and problems encountered in connection with the competition. The forum (organized by Kortenkamp) was divided into sessions: sensing, mapping, architectures, and hardware-manipulation. Each session featured short talks by several team members and commentary by an outside expert in the field (the participating experts were Dean, Ben Kuipers, Gat, and Blake Hannaford). The forum ended with a general discussion on the future of the competition and of autonomous robotics in general.

## The Competition Events

Given the general domain of office robotics, we wanted to design events that promoted reliability and robustness and that emphasized the skills of navigation, manipulation, and object recognition (primarily visual). We felt strongly that preparing for three separate events was too taxing (only two teams competed in all three events in 1993) and decided to limit the number and the diversity of events. Although this approach risked excluding participants whose research fell outside the more narrowly scoped events, we felt that there was sufficient current research in the areas of navigation, manipulation, and object recognition to sustain a competition. In addition, the robot exhibition provided a venue to demonstrate other research results.

Based on these considerations, we designed two events: In Office Delivery, the task was to navigate from one room to another within a large arena of offices, corridors, and foyers. This event was timed and consisted of three separate trials; in each trial, the environment was changed somewhat (by closing doors and blocking corridors) to see whether the robots could detect the changes and react appropriately. In the Office Cleanup event, the same arena was littered with cups, empty soda cans, and paper wads. The robots had to recognize and collect as many objects as possible in the allotted time and deposit them in nearby trash bins. This event also had several trials, with the main difference being the density of trash placed in the offices.

Once the basic outline was in place, the hard part began: designing a set of rules. A common complaint in previous competitions was that the rules were not explicit enough, leading to ambiguities in scoring and inter-

pretation. Although we knew that we could never devise a set of rules to satisfy everyone, we wanted to create rules that were tight enough to eliminate major ambiguities but were still flexible enough to enable teams to come up with creative solutions to the tasks.

The Rules Committee (Gat and Kortenkamp) decided on an elegant strategy: define a minimal set of rules and then add a set of penalties for robots that deviated from the rules. The idea was to allow any strategy for achieving the tasks but at a cost. Penalties were levied for either modifying the environment in some way or providing the robot with some additional a priori information. For example, in the Office Delivery event, robots could make use of a topological map of the arena for a small penalty, and for a much larger penalty, they could use a metric map. There were also penalties for enhancing the environment, such as by adding perceptual markers or shrouding the furniture to make it more detectable. The main problem was to make the penalties large enough so that there was incentive not to take the penalty but not so large that it was impossible to overcome. This set of penalties evolved over the months leading to the competition, mainly in response to questions and critiques from the participating teams, but the underlying event rules remained remarkably stable.

In general, the penalty system worked well, but of course, it was not perfect. For example, because most of the robots did not have arms, we added a penalty for using *virtual manipulation* (where the robots merely had to be facing the trash and announce their intentions to collect it). In retrospect, this penalty was much too small, providing a big disincentive for actually trying to grasp objects. Another problem arose from the penalty for using noncontest-supplied trash. The intention was to modestly penalize robots that might be able to recognize only certain shapes or colors, such as red soda cans. Two teams, however, chose to replace our small trash items with white cardboard boxes about a foot square and four feet high. This use of “big trash” placed them at a great advantage because the boxes were much easier to recognize than the small trash that the other teams used. Such examples merely illustrate the difficulty we faced in eliminating ambiguity from the rules.

Designing the office environment was also a major undertaking. It had to have sufficient variety to be challenging yet be simple enough to be navigated by the majority of the robots. We also wanted the environment

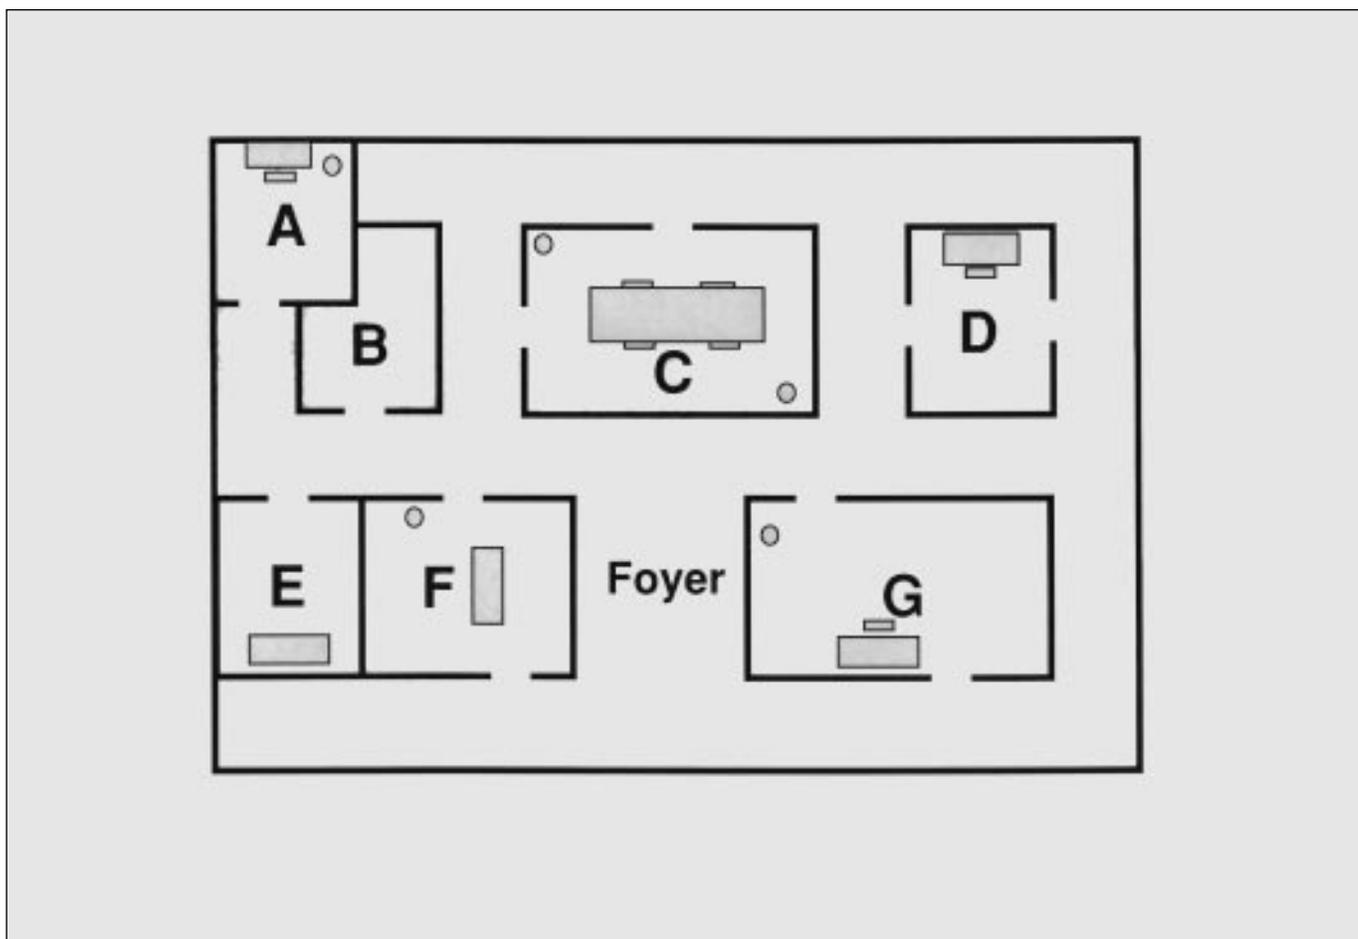


Figure 2. The Competition Arena.

to be fairly realistic, but it had to be affordable to construct and enable easy viewing by the judges and the audience. One compromise was to use a plastic material for the walls that, unfortunately, turned out to have poor sonar characteristics. This material caused problems for several of the robots, especially the entry from Carnegie Mellon University (CMU). We actually constructed two identical arenas side by side so that we could run the competition finals head to head and also allow one team to set up while a robot was running in the other arena during the preliminary trials.

Each 55- by 80-foot arena consisted of three major corridors, with two rows of offices and connecting corridors (figure 2). Several of the offices contained some furniture (indicated by squares and circles in figure 2) and had multiple doorways to provide for alternate pathways when we closed doors and blocked corridors. A large foyer was included because it was felt (correctly, as it turned out) that although autonomous corridor naviga-

tion was fairly well established, reliable navigation of large open spaces was pushing the state of the art.

One of the rules was that the exact layout of the arena would not be divulged until just before the competition. Participants did know that the arena would have a number of offices, corridors, and foyers and that, for example, corridors would be orthogonal, and each office would have at most one door on a side. However, it was not until the night before the competition that the participants were supplied with a topological map of the environment.

## The Participants

Fifteen robot teams showed up for the competition (although five were not actually able to compete, mainly because of hardware problems). In addition, there were five entries in the exhibition (described in a subsequent section). The robots varied from foot-long wheeled robots to a five-foot-wide hovercraft.

All the robots used sonar sensors, and some sported laser range finders, cameras (color and monochrome), and infrared and contact sensors. Teams came from as far away as Germany and represented universities, companies, and private individuals.

Many of the robots were commercially available products, which continues a trend seen in recent years; however, it is probably the first time in which over half the robots were off-the-shelf products.

The largest contingent was robots built by Nomadics. DERVISH, from Stanford University, is a NOMAD-100 robot, augmented with 2 POWERBOOKS to provide for higher-level planning and reasoning. ARGUS, from Lockheed Missiles and Space Company, ESCUINCLE, from ITESM/Stanford, and WILLIE, from Kansas State University, are all the larger NOMAD-200 robots. Although all three relied heavily on sonar, in addition, ARGUS used a laser range finder, and ESCUINCLE used a camera to find visual landmarks.

RHINO, the entry from University of Bonn, and AMADAUS, from Worcester Polytechnic Institute, are both B21 robots from Real-World Interfaces. The robots feature a sonar ring, bump panels, and a color camera on a pan-tilt head. The other two commercial robots were CLEMENTINE, a Denning MRV3 from Colorado School of Mines, and EGOR, a TRC LABMATE from University of Utah. CLEMENTINE, which was damaged in transit and could not actually compete, had a sonar ring, two video cameras, and a laser-navigation system. EGOR sported sonars, infrared sensors, and a camera for visual object recognition. Much of the computation was performed on a 486 notebook computer strapped to the robot.

The other seven robots were all home brews. Two—CHIP, from University of Chicago, and XAVIER, from CMU—used bases from Real-World Interfaces but customized them substantially. CHIP, built on a 12-inch-diameter base, relied heavily on stereo color cameras and an off-board DataCube image-processing board. It also had grafted on a five-degree-of-freedom HERO II arm. XAVIER, built on a 24-inch base, has a sonar ring, laser range finder, and color camera on a pan-tilt head. Unfortunately, XAVIER's sonars did poorly in detecting arena walls, which forced it to withdraw from the Office Delivery event.

The most unusual entry was brought by the University of Maryland. AIRS, a Hovercraft robot, was originally developed at Brigham Young University to investigate control of space-faring robots. The robot is about five feet wide and seven feet tall, and its fans use



*EGOR Gets Some Last-Minute Instructions.*

lots of power to keep it suspended a few inches above the ground (one of the main difficulties was supplying sufficient battery power and dealing with the high-power electronics). Unfortunately, AIRS was badly damaged in transit, and the Maryland team spent most of the week trying to repair the damage. Eventually, AIRS was fixed enough to be exhibited but did not take part in the competition itself.

The remaining four teams had small robots, on the order of 12 to 18 inches long, that were built entirely from scratch. ERRATIC, from SRI International, used the same type of fuzzy-control rules as the SRI FLAKEY robot but in a much smaller design (see the article by Kurt Konolige, also in this issue). The Georgia Institute of Technology entered a team of three robots named IO, CALLISTO, and GANYMEDE. These bright-green robots featured a small color camera and a simple manipulator (see the article by Tucker Balch and his colleagues, also in this issue). The other multirobot entry was EULER and ARCHIMEDES, the work of two individuals. Rounding out the entries was CUJO, a lightweight, modular robot from Simon Fraser University. Unfortunately, neither CUJO, EULER, nor ARCHIMEDES ended up competing.



*ARGUS Is Led to the Starting Gate.*

## The Competition Trials

Robot teams began arriving on Friday, 29 July, and quickly got to work setting up computers, uncrating robots, and putting the robots through their paces. Although some of the robots were damaged during shipping (the most severe was AIRS, which was shipped in a wooden crate but somehow arrived, in pieces, in a battered cardboard box!), most of the problems were soon corrected.

By the late afternoon, construction of the arenas was completed, and teams had begun practicing for real. At this point, although we reserved the right to make minor modifications to the arena before the competition, for all intents and purposes, the participants knew what they would be up against. As mentioned, the wall material caused some problems, and several teams spent many hours improving their sonar-processing algorithms (and, in at least one case, modifying the sonar ring itself) to better detect the walls.

Thus began nearly 40 straight hours of testing, revising, and retesting algorithms that had worked perfectly in the lab but now did not quite perform in this foreign environ-

ment. In some cases, strategies, such as those for detecting objects and exiting from rooms, were completely rewritten during the night. Although such revisions could be viewed as failures of the initial software, many of the participants indicated that they were, in fact, testaments to the flexibility of the software architectures used by their robots.

Seven teams successfully participated in the Office Delivery trials, which were held on Sunday, 31 July. To ensure consistency, all trials were judged by the same two individuals, Gat and Kortenkamp (with help from Marcus Huber and Dean in the Office Cleanup event). The Office Delivery event consisted of three trials: In the first trial, the arena was identical to the topological map given out; in the remaining trials, doors were closed, obstacles were added in the corridors, and some corridors were blocked altogether. Teams had 10 minutes to complete each trial, and the final score was a combination of the time and penalties taken. Teams were allowed to restart their robots, at no penalty, by moving them back to the starting point (however, the clock kept running during the restart).

Nearly all the teams took the small (5-point) penalty for using a topological map (each point was equivalent to 10 seconds). The exceptions were ERRATIC, which used a metric map (see the Konolige article in this issue), and RHINO, which learned its own map by exploring the environment (see the article by J. Buhmann and his colleagues, also in this issue). In addition, several teams took penalties that facilitated their exiting from the start room, either by adding visual markers near the doors or by starting at a known position and orientation in the room. The other robots used fairly simple strategies for exiting the rooms, typically by finding and following the office walls until a door-sized opening was detected.

Most of the teams succeeded easily in the first trial run, in which the robots had to exit from room A (figure 2) and travel to room D. Several teams also did fairly well in the second trial, where the judges shut the closer door to goal room D. The robots that succeeded either detected that the door was closed by noticing an offset from the wall or realized that they had missed the door when they came to the end of the corridor without seeing an opening. They then replanned a path to the other side of the room and quickly found their way into the office. In addition, few robots had problems with a chair placed in the corridor, which was meant to test the capability to avoid obstacles and still track the corridor.

This capability was further tested in the third trial, where a column of chairs was placed in the foyer near point X (figure 2). The idea was to force the robots into the open, away from the corridor walls. The results were instructive: No robot was able to complete this trial totally autonomously. *DERVISH* came closest, but it needed some help to avoid “decapitating” itself by driving underneath one of the chairs.

At the end of the Office Delivery trials, *ESCUINCLE* was ranked first, and *DERVISH* and *ARGUS* were tied for second, followed by *WILLIE*, *ERRATIC*, *RHINO*, and *EGOR*. The results generally showed the advantage of using a topological map and the ability to exit a room quickly (both *ESCUINCLE* and *ARGUS* used specialized strategies for this part of the event). *RHINO*, the one robot that learned its own map from scratch, had troubles using the map in subsequent trials and did not fare well overall. In general, the faster robots did better in the earlier trials, but as the complexity of the environment increased, the slower, more robust robots tended to perform better.

Not everything ran smoothly, however. One particular incident occurred when a TV cameraman showed up. We rolled out the robots that had consistently been performing well that day for him to film them navigating around the arena. In 20 minutes of filming, not one robot made a successful run. We tried to convince him that it was because of the radio frequency noise from his camera, but the robots performed just as poorly when the camera was off. Finally, he left, and everything mysteriously went back to normal. The TV news report, however, featured some prominent crashes of the robots—I guess the audience learned that a bit more research remains....

The Office Cleanup trials were held the next day. This event featured a much greater diversity; although most teams used vision to detect the trash and trash bins, some used monochrome, and others used color; some robots used laser light stripers; some robots actually had manipulators (*CHIP* and the team from the Georgia Institute of Technology); and Georgia Tech used a team of three cooperating robots (see the article by Tucker Balch and colleagues, this issue). In addition, the strategies for finding trash differed sharply. Robots such as *ARGUS* and *ESCUINCLE* used methodical approaches, scanning for trash first before picking it up, but the Georgia Tech robots used random walks to obtain coverage of the office.

The robots were started in the “large con-



*ERRATIC Races Toward the Finish Line.*

ference” room (room C), which was filled with trash and had two trash bins on opposite sides of the room. The robots without manipulators had to stop and announce when they were in front of a piece of trash. If they actually were facing some trash, a human would pick it up for the robot. Similarly, when the robot announced it was at a trash bin, the trash would be put down (either in the bin or back on the floor if the robot was not, in fact, near a bin). Several of the robots had problems with this second part, often confusing walls and the conference table for a trash bin. Because points were acquired only for successfully depositing the trash, this mistake was a particularly costly one.

Probably the most effective strategy was employed by *ESCUINCLE*: It detected large visual markers attached to the big trash boxes it substituted for the contest-supplied trash (*ARGUS* employed a similar strategy but used a shorter-range laser range finder to detect trash boxes). *ESCUINCLE* would move to the center of a room; turn in place; scan for the markers; and visit each object it detected, in turn, occasionally dropping off the trash in



*AIRS Gets Ready to Take a Spin.*

the nearest bin. After clearing one room, ESCUINCLE would navigate to another room (using the navigation strategies from the Office Delivery event) and continue. Although the judges were concerned that using this type of trash violated the spirit of the rules, the robot's performance was still quite impressive.

Other notable performances were turned in by the Georgia Tech team and by CHIP. The small Georgia Tech robots used a simple gripper and computationally inexpensive color-vision algorithms for detecting trash, trash bins, and each other. Their cooperative strategy for clearing the room was highly effective, especially because the robots would occasionally break down or become stuck. CHIP was the only robot to perform the task with no penalties: It used the contest-supplied trash and actually found, picked up, and deposited one piece of trash during the allotted time (although the Georgia Tech robots could pick up trash, they could not reach high enough to place it in the bins).

The results of the Office Cleanup trials had Georgia Tech in first, followed by ESCUINCLE, RHINO, CHIP, ARGUS, EGOR, and XAVIER. The top six teams in this event and the top five teams in the Office Delivery event were selected to participate in the finals.

## The Robot Exhibition

The Robot Exhibition and Competition finals were held on the afternoon of the first day of the technical sessions. During the afternoon, robot videos from about a half dozen institutions were shown continually. For two hours, the exhibitors had center stage—actually, two 30-foot circular arenas. The various exhibitions proved to be great crowd pleasers.

The University of British Columbia exhibited a team of small soccer-playing robots running on the DYNAMITE test bed. The test bed included several radio-controlled vehicles that used a common vision system to monitor the positions of the ball and each other. The fast-paced action was so popular that it was demonstrated several times during the afternoon. An exhibition of machine learning in robotics was provided by the University of Minnesota. Their PBMIN system used an unsupervised neural net-type approach to learn to balance a pole on a small moving vehicle.

The AIRS hovercraft robot, from the University of Maryland, demonstrated several skills, including autonomous station-keeping while turning in place. Given the size of the robot and its novel means of locomotion (plus the noise it made), it was quite an impressive demonstration (and one filled with an air of danger). An interactive exhibition was provided by RHINO, from the University of Bonn, which followed moving objects by tracking a colored visual marker. It would speed up and slow down depending on its perceived distance from the marker. People held the marker and walked around, and RHINO followed, keeping a safe distance.

XAVIER, from CMU, demonstrated its prowess at finding and picking up large, meter-sized boxes. XAVIER used a speaker-independent speech-recognition system to accept commands. It used a camera to locate visual markers on the boxes, navigated to the boxes, and used its laser range finder to verify the size and position of the box. XAVIER then picked up the box and lifted it over its "head," before wandering off to place the box at some designated location.

In all, the exhibitions presented the audience with a variety of research being performed in the field of robotics and AI. From the various reactions, it is clear that they had a definite impact.

## The Competition Finals

The competition finals were held immediately following the exhibition. Judging were the Rules Committee members, Erann Gat and



*XAVIER Demonstrates Its Box-Picking Prowess.*

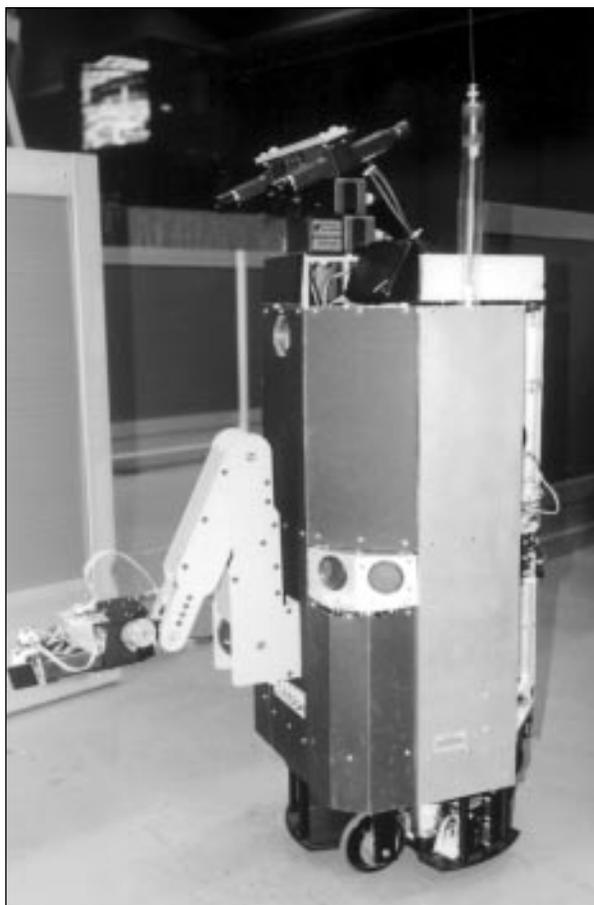
David Kortenkamp, and “celebrity” judges Kuipers, Bonasso, Leslie Kaelbling, Austin Tate, Martha Pollack, Stan Rosenschein, Montemero, and Dean. Each robot had one opportunity (not counting restarts) to perform the task. To increase excitement, we had wanted to run all the trials head to head; however, because two-thirds of the robots used radio communication to off-board computers, we were forced to do some of the runs separately. To aid audience understanding, we had a “play-by-play” commentary and visuals showing the current time and score to beat. In some cases, team members described their robots as they were competing, but most teams were too busy hunched over computer screens to provide this commentary. However, most of the robots had some speech-generation capability, so they provided their own commentary. Although this speech capability was invaluable for the judges, it was not quite as effective for the audience because it was difficult to hear in the large hall (one exception was EGOR’s pronouncements, which were loud and clear, although unmistakably accented).

The final Office Delivery task was to start in room B and travel to room F across the hall—simple except that the closer door on room F was closed, and the corridor at point Y (figure 2) was completely blocked off. This task was about the same complexity as the third task in the preliminary trials, and four of the five robots failed to successfully complete the task. Three robots (ESCUINCLE, ARGUS,

and WILLIE), which used similar path-planning software, had the exact same problem: Although they noticed the closed door correctly, they got trapped in the cul-de-sac formed by the blocked corridor and kept circling. Unfortunately, even after they were restarted, they quickly got back in the same trap, and this continued until time ran out.

The other two robots turned in very different runs. DERVISH, which was the only robot to complete the task, moved slowly but steadily toward the goal. DERVISH’s use of probabilistic methods and its unique sonar arrangement enabled it to deal robustly with sensor noise and changes in the environment (see article by Illah Nourbakhsh and colleagues, also in this issue). ERRATIC moved much faster but also took many more penalty points (350 seconds worth) because of its use of a metric map and known starting position and orientation. Even with this handicap, ERRATIC would have won the event except for one small mistake: It lost track of its exact position and turned into the foyer, thinking that it was in the goal room. For want of about five extra feet of travel, ERRATIC placed second, and DERVISH won, with ESCUINCLE, ARGUS, and WILLIE tying for third place.

Although the outcome of the Office Delivery finals was quite different from the trials, the Office Cleanup finals proceeded more according to prior expectations. Again, the Georgia Tech team of robots was effective in cleaning up the conference room, although



*CHIP Prepares to Find Some Trash.*

they occasionally mistook the blue siding around the tables for trash bins. ESCUINCLE and ARGUS did well finding their large trash boxes. RHINO was the most improved robot in the finals: its color vision object-detection algorithm worked exceptionally well, and RHINO was not only able to clean up the conference room, it also had time to venture out and start cleaning up other rooms. In the end, first place again went to Georgia Tech, with RHINO finishing second. ESCUINCLE was a close third place, followed by ARGUS, CHIP, and EGOR.

## The Robot Forum

The next day, a forum was held to discuss technical issues arising from the competition. There were many interesting presentations and discussions, but I mention just a few points that I found particularly noteworthy.

After one set of presentations, Dean asked, "Where is the science?" The teams demonstrated that their systems had great flexibility by modifying software at the last minute, but is it AI? Firby responded, "The science is what

you do and throw away the night before the competition." Tom Henderson countered that the "competition creates problems that motivate people to work on the science"; what one sees that does and does not work is what motivates science. In general, it was felt that basic research provides the tools to solve specific problems but that the competition, by revealing the inadequacies of these tools, provides the impetus for further research (to do better at the next competition!).

It was interesting to see the apparent convergence in corridor-navigation techniques. For one thing, performance has improved significantly in the last several years: In 1992, the fastest robot in the competition traveled at 30 centimeters/second (Dean and Bonasso 1993); this year, that rate was about the average speed, and RHINO demonstrated corridor navigation at speeds as high as 90 centimeters/second (although not during the competition itself). Most teams used topological maps, typically without metric information. In addition, researchers are specifically addressing the issue of uncertainty in navigation. In particular, the fuzzy control rules of ERRATIC and the probabilistic methods used by DERVISH proved to be fairly robust to sensor noise.

However, there is much room for improvement in the area of navigating within crowded offices and open spaces. The robots that effectively exited rooms all used specialized strategies that would not necessarily be effective in more realistic environments. The robots also tended to get confused within the foyer area, and some, such as DERVISH and ERRATIC, explicitly avoided traversing such areas at all (preferring instead to always stay in sensor contact with known entities, such as the walls).

Similar observations hold for the state of the art in mobile manipulation. Only CHIP successfully completed the Office Cleanup task, as originally formulated. As Blake Hanaford aptly expressed it: He had nothing but "virtual praise" for those robots doing virtual manipulation. Several participants felt that this area was the next big one to tackle and that future competitions should be structured to encourage research in mobile manipulation.

The state of the art in visual object recognition, however, has noticeably improved over past competitions. In particular, many of the teams are now using color vision—from the "cheesy vision" of the Georgia Tech team (their term) to the sophisticated, multiple-strategy approach of CHIP. One potential problem, however, is that most of the teams

that used vision (except for the Georgia Tech team, EGOR, and XAVIER) needed off-board processing as well. In several cases, reliance on radio communication to off-board computers proved catastrophic, causing robots to operate open loop and crash into walls. I expect that more teams will move to all on-board computation in the future, which might imply the use of simpler, more specialized vision algorithms.

The most successful robots often relied on fairly simple strategies; these strategies usually succeeded because the environment was relatively well structured and benign. For example, the offices were quite tidy, making it easier to do wall following to find the exit to the room. It is also interesting to note that the teams with fewer members tended to be more successful than the larger teams. One possible explanation is that systems integration, which is critical to the success of these systems, is easier with just one or two people in charge.

Several participants touted the availability of complete mobile robot systems and standardized components as a boon to research in mobile robotics. The buy-in factor is no longer so prohibitive, as evidenced by the diversity of organizations and teams that participated in the competition. It is encouraging that over half the robots competing this year were commercially available, off-the-shelf systems. Although this availability of components is somewhat driven by the robotics community, in large part, it is driven by the consumer market, which produces inexpensive personal computers and laptops, camera chips (camcorders), servo-control motors (model airplanes), and so on. Masayuki Inaba, from the University of Tokyo, showed a videotape of some of the remarkable robotic creatures his lab has produced from these standard components, the most impressive being an apelike robot that can sit down, stand up, and walk around on its own.

Finally, there was a long discussion on the direction that future competitions should take. It was generally agreed that although competitions, as such, are useful for gauging the state of the art, they have the potential for stifling innovation by too narrowly scoping the tasks to be performed. In addition, the tasks are often not terribly exciting to watch, especially because the robots move at subwalking speed. Montemerlo likened it to the pre-VISICALC days of the PC, where a powerful tool existed without any real applications to drive it. The suggestion is that we need to find (or create) compelling applications that will need autonomous—or at least

semiautonomous—mobile robots. As yet, few such applications have been identified (some include security, delivery, automated wheelchairs, and automated vacuum cleaners), and fewer still have been marketed (none, so far, in the mass consumer market). Several participants contended that a way to encourage the development of such applications is to make the exhibition portion more prominent, perhaps by having a freestyle-type competition. The trick is to do so without merely showcasing “cool demos.” One suggestion was to have the robots perform “routines” that include both mandatory and optional portions, along with predefined difficulty factors (such as navigating in a static environment versus navigating in a peopled environment). It was felt that although judging such a competition might be fairly subjective, it could also encourage wider participation.

## Conclusion

The 1994 AAAI Robot Competition and Exhibition was successful in achieving its three aims: (1) comparing the state of the art in autonomous mobile robots, (2) providing an entertaining view of robotics research, and (3) promoting the exchange of technical information. Participants liked the overall format, which was less intense than in previous years. The events, although challenging, were still within reach of many of the robots, and they provide good benchmarks for the future.

I would recommend that future competitions continue holding multiple preliminary trials to promote reliability and robustness and to provide an opportunity to test the robots' limits away from the glare of the public. The competitions should continue to feature navigation and to create more realistic environments, but they should also move more strongly into the areas of visual recognition and manipulation—two areas that I believe will be key to the widespread application of mobile robots. The events should also be designed to encourage cooperation—both multirobot and human-robot. Another key area is learning. Little machine-learning technology was used by the teams, but this area must be nurtured and developed to achieve truly reliable and competent robots.

Preparing for a competition such as this involves a tremendous amount of effort (I probably should have heeded Konolige's advice not to chair and participate in the competition simultaneously!). Much credit goes to all the robot team members for their tireless

efforts and to all those who helped make this competition a success. In the end, is it worth it? The participants think so—many produce an invaluable infrastructure of software and hardware that aids in future research, most have fun, and all come away with a heightened appreciation of the difficulty and the promise in creating truly autonomous agents.

#### Acknowledgments

Many thanks go to my capable Organizing Committee: David Kortenkamp, Erann Gat, Jim Firby, Tom Dean, Pete Bonasso, Kurt Konolige, Terry Weymouth, Jim Hendler, and Ian Horswill. In particular, Kortenkamp and Gat did a great job with the competition rules and the judging. Financial contributions by the National Aeronautics and Space Administration, Microsoft Corporation, and AAI provided generous travel grants for students and the robots. The AAI staff, in particular Rick Skalsky and Carol Hamilton, made my job much easier by taking care of a myriad of details and logistics. Sara Hedberg shepherded the publicity aspects of the event, and Elsevier Publishers and Network Cybernetics

Corporation provided prizes for the winners of the competition finals and trials.

#### References

- Dean, T., and Bonasso, R. P. 1993. 1992 AAI Robot Exhibition and Competition. *AI Magazine* 14(1): 35-48.
- Konolige, K. 1994. Designing the 1993 Robot Competition. *AI Magazine* 15(1): 57-62.



Reid Simmons is a research computer scientist in the Department of Computer Science and the Robotics Institute at Carnegie Mellon University. His Ph.D. from the Massachusetts Institute of Technology is in the field of AI, where his dissertation dealt with algorithms for planning and data interpretation. Since coming to Carnegie Mellon in 1988, Simmons's research has focused on developing self-reliant robots that can autonomously operate over extended periods of time in unknown, unstructured environments. Simmons has been involved with the robot competition since its inception as both a participant and an organizer.