AI Magazine Volume 14 Number 4 (1993) (© AAAI)

The Winning Robots from the 1993 Robot Competition

Illah Nourbakhsh, Sarah Morse, Craig Becker, Marko Balabanovic, Erann Gat, Reid Simmons, Steven Goodridge, Harsh Potlapalli, David Hinkle, Ken Jung, and David Van Vactor

■ The second annual Robot Competition and Exhibition sponsored by the American Association for Artificial Intelligence was held in Washington D.C. on 13–15 July 1993 in conjunction with the Eleventh National Conference on Artificial Intelligence. This article describes the robots that placed first and second in each event and compares their strategies and their resulting successes and difficulties.

The 1993 robot competition consisted of three distinct events: (1) Escape from the Office, (2) Office Delivery, and (3) Office Rearrangement. The unifying theme for these events was autonomous robotics in realistic office environments.

In the first event, Escape from the Office, the objective was to maneuver out of a confined space, simultaneously avoiding real-world office furniture, including difficult-to-see objects such as chairs and tables with thin legs, and then to quickly complete a slalom course and recognize the finish wall.

In the second event, Office Delivery, the objective was to self-locate using an office map, search an area for a given object (a coffeepot), and then navigate to a specified delivery area.

In the third event, Office Rearrangement, the objective was to modify the office environment by identifying and maneuvering appropriate boxes to create a specified pattern at the goal location.

These events present a significant challenge to the robotics community by testing many skills that a truly autonomous mobile robot must demonstrate. The results of the competition were as follows:

Event 1: Escape from the Office
Stanford University, First Place
California Institute of Technology, Second Place

Event 2: Office Delivery
California Institute of Technology, First Place
Carnegie Mellon University, Second Place

Event 3: Office Rearrangement
North Carolina State University, First Place
Lockheed Palo Alto Research Labs, Second Place

Event 1: Escape from the Office

Each robot started inside a four-meter by five-meter office with three doors; each door was marked with a large black-on-white cross. The office contained actual furniture, including chairs, a table, a file cabinet, and a bookcase. This realistic environment was a hurdle for conventional robotic sensory systems. Thin-legged tables and chairs are nearly invisible to sonars, as are black cabinets and bookcases to infrared sensors. The robots began in unknown locations facing the finish line. Between one and three minutes from the starting time, one of the three office doors opened, and the robot was to find a path out of the inner office, then across the outer arena, which contained a scattering of obstacles (boxes). The event was complete when the robot recognized that it was within two meters of the finish wall.

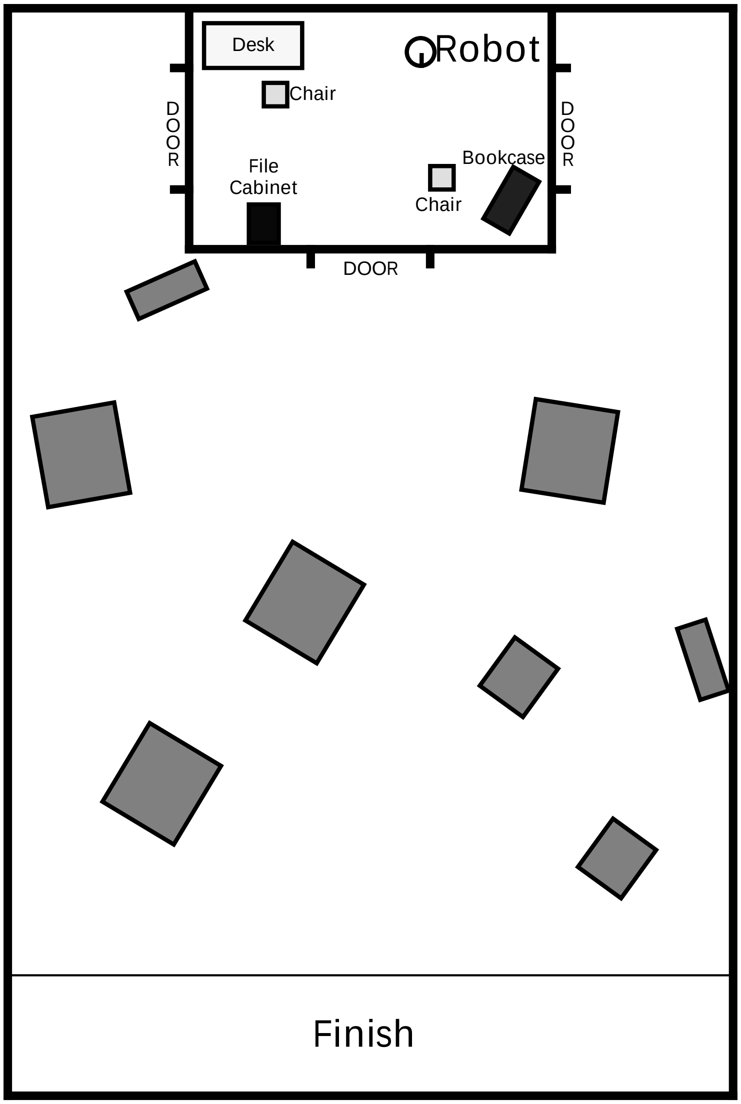

Figure 1. Diagram of Arena Layout for Event 1 (12 × 18 meters).

Robot Descriptions

The Stanford robot, SCIMMER (Sarah, Craig, Illah, and Marko's most excellent robot), and the CalTech robot, ALFRED, are remarkably similar. They are both 100-percent off-the-shelf synchrodrive robots. All software development for both robots was done on board, using MACINTOSH Common Lisp on ALFRED and TURBO C on SCIMMER. Both use a similar layered control structure.

SCIMMER is a NOMAD-200 mobile robot built by Nomadic Technologies. It detects range information using a bumper, a ring of infrared sensors, a ring of sonars, and a laser range-finding system. Other features include a speech synthesizer, an on-board hard drive, and a sensor turret that turns independently of the wheels. The entire robot is governed by a 386 running DOS and TURBO C.

ALFRED is a B12 from Real World Interface, with a development enclosure housing a GESPACK MPL-4080 68000-based single-board computer. A MACINTOSH DUO 230 sits on top of the development enclosure. The only sensor modality on ALFRED is the canonical 12-sonar ring. A dedicated Motorola 68HC11 controls the sonars. ALFRED detects collisions by monitoring the drive motor current. The robot has no cameras. The control architecture used on ALFRED is a stripped-down implementation of the ATLANTIS control architecture (Gat 1992). The controller runs on the MPL-4080, and the sequencer (along with the development system) runs in Lisp on the MACINTOSH. Communications between the MACINTOSH and the MPL-4080 and between the MPL-4080 and the base are through a 9600-baud RS-232 link.

Strategies

Both robots used a similar three-part strategy in the first event. During the first 60 seconds (when the doors were guaranteed closed), the robots explored the office to find its boundaries. The second phase consisted of searching for an open door and escaping from the inner office. Finally, the robots used different strategies to speed through the outer slalom and across the finish line (figure 1).

SCIMMER constructed its map by exploring the office while turning the laser range finder at a constant speed. At the end of the first minute, the robot fit a rectangle to the resulting laser range data. This computation gave SCIMMER an idea of the extent of the inner room as well as a rough idea (+/-2 meters) of its distance to the finish line.

ALFRED's mapping strategy was considerably simpler. Because the robot's initial orientation was known, it just wandered randomly around the office and noted the minimum and maximum *X* and *Y* coordinates that it reached. ALFRED then combined this information with knowledge of the size of the office to compute the robot's position with associated uncertainty bounds. Each robot's strategy for finding the open door was somewhat different from the other. SCIMMER searched for the open door by following the perimeter of the office, but ALFRED tried each door in turn.

The challenge for the final stage was to create a control program that would guarantee completeness in complex outer-room configurations and would still be fast in simple cases. SCIMMER's solution allowed it to speed toward the goal at top speed while it looked ahead 10 feet. If the robot detected an obstacle on the horizon, it initiated early evasive action. If these gentle turns failed to provide a clear path, then SCIMMER would actively search for a clear path by following the wall around the obstacle. If this exploration violated a preset maximum-allowable backtracking distance, then SCIMMER would change direction to find a path around the other side of the obstacle. By incrementally relaxing this backtracking restriction, the robot is able to find the shortest path around an obstacle quickly in the easy case and eventually in the difficult case.

ALFRED's navigation strategy was essentially identical, except that instead of backtracking on distance limits, ALFRED backtracked on angle limits. If the robot had to turn farther than a certain angle to avoid an obstacle, it would stop and turn the other way. This turn limit would be relaxed every time the robot changed direction until it managed to make some forward progress.

ALFRED recognized the goal wall by dead reckoning within two meters of the wall and then moving toward the wall until it encountered an obstacle. This strategy was a safe one because the last two meters were guaranteed to be free of obstacles. SCIMMER used two strategies to recognize the finish line. If its encoders indicated that it must have reached the finish line in spite of the arena size uncertainty, then the robot stopped when it detected the goal wall. If SCIMMER saw any wall at the expected orientation that was as wide as the arena, then it would also reason that it had reached the finish line. In both the regular contest and the playoffs, the judges opened the particular door that was the last to be searched, resulting in the longest escape time for both robots. The outer arena was easy in both cases, however, and both robots found the finish line without backtracking. SCIMMER won the playoffs because of an extremely fast time to the finish line (after the door opened): 30 seconds.

In a shameless attempt to influence the judges, ALFRED entered and left the arena autonomously (the only robot to do so). The ATLANTIS sequencer made it simple to add this capability.

Lessons Learned

Both SCIMMER and ALFRED used a multilevel control architecture with low-level reactive obstacle avoidance guided by high-level strategy planning. The Stanford and CalTech teams agree that developing this type of architecture in a bottom-up fashion is the correct development paradigm for mobile robot programming.

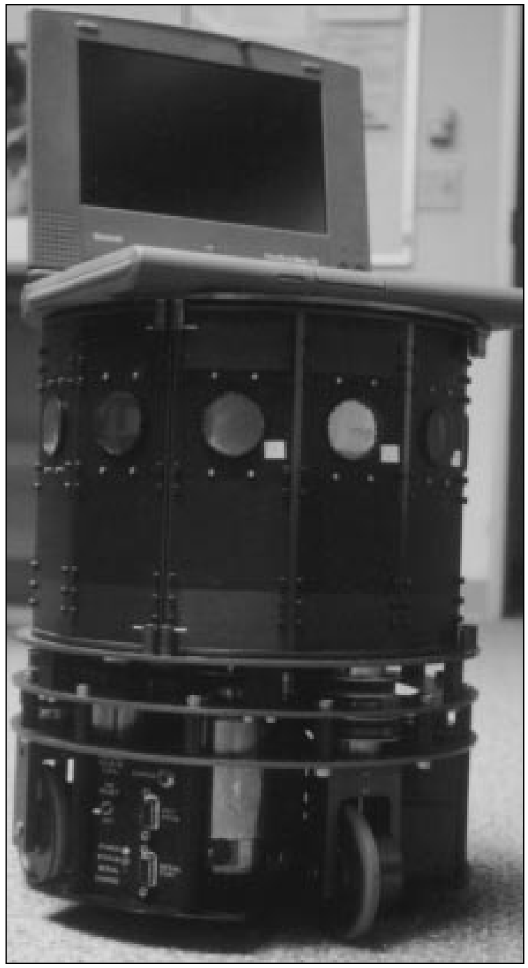

Alfred

SCIMMER and ALFRED did not use vision in the first contest and were thus unable to recognize the doors using the perceptual markers. SRI International's FLAKEY and Carnegie Mellon's XAVIER used vision to locate the door markers. They determined that a door was open when the perceptual marker on the door disappeared from view.

Neither CalTech nor Stanford used a simulator during any stage of the development or debugging process. The teams differ in the degree of vehemence with which they eschew simulation. The Stanford team is strongly against simulations, claiming that the time and efficiency gains of simulation do not outweigh its tendency to mislead the programmer and misrepresent the real world. The CalTech position is somewhat less extreme. Although simulations can be misleading, they can also be useful time savers if used properly. However, proper use does seem to be rare.

In a shameless attempt to influence the judges, ALFRED entered and left the arena autonomously …

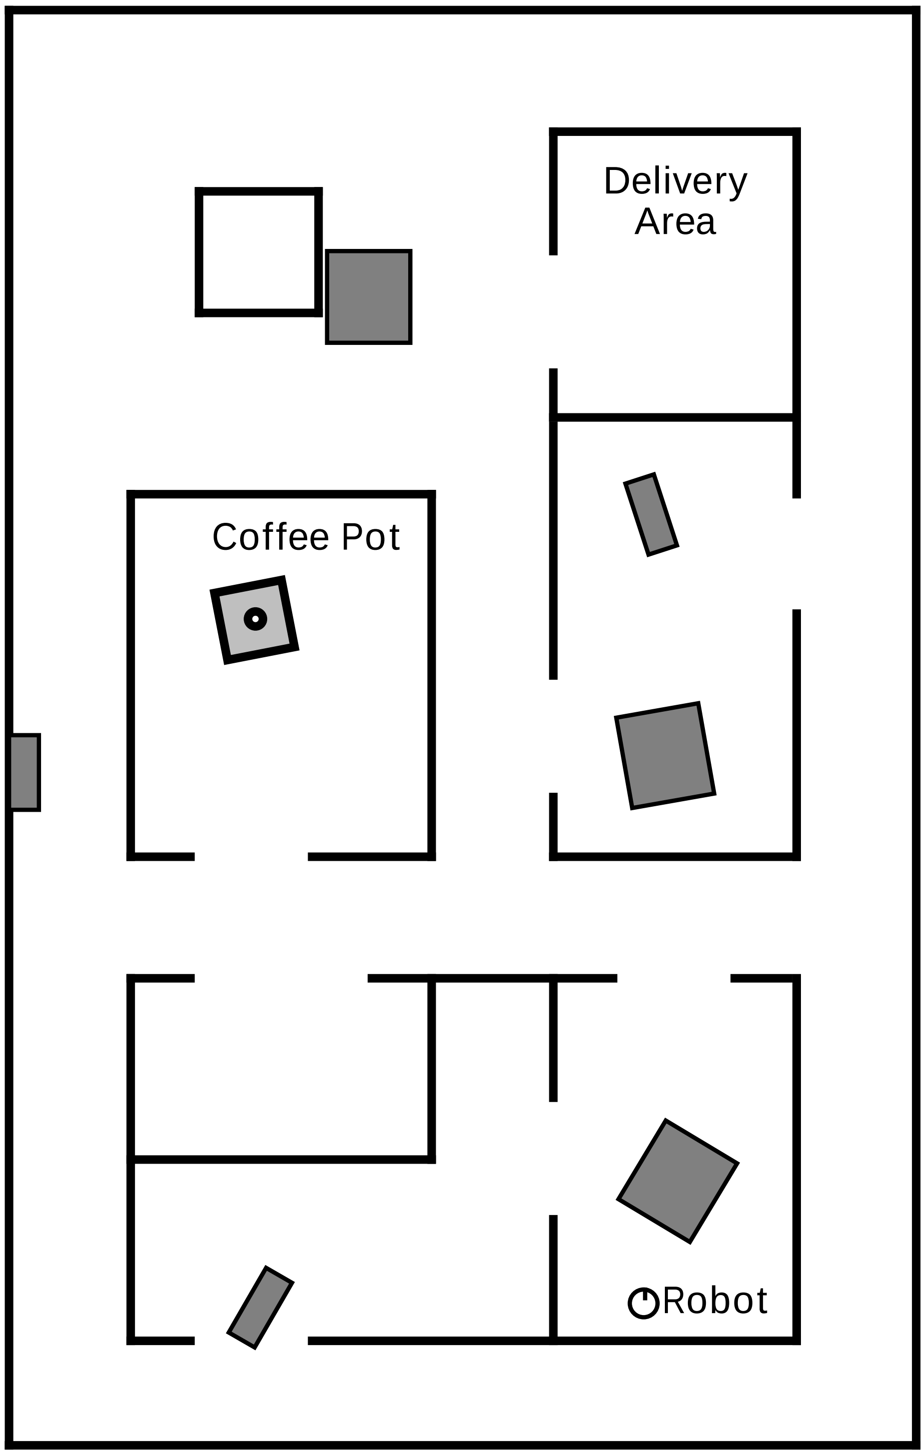

Figure 2. Diagram of Arena Layout for Event 2 (15 × 24 meters).

Event 2: Office Delivery

The office delivery arena was a realistic office floor plan, including offices, corridors, and open areas. Several of the doors were perceptually marked with large plus signs (+) and bar codes. Teams began with an approximate map of the arena layout and dimensions. Robots began in unknown positions and orientations within the arena. The task was to find a marked coffeepot at an unspecified location and deliver the coffeepot to a particular office. The robots received hints about which quadrant they were starting in as well as which quadrant the coffeepot was in (figure 2).

Robot Descriptions

CalTech's ALFRED is described in Event 1: Escape from the Office. Of all the entries, ALFRED came the closest to completing the event, doing everything except actually locating the coffeepot (because it had no sensors capable of detecting it).

Carnegie Mellon's XAVIER is built on an RWI B24 synchrodrive base. Its sensors include bump panels, a Denning sonar ring, a Nomadics laser scanner, and a color camera mounted on a Directed Perception pan-tilt head. On-board computation consists of two 66-megahertz Intel 486 computers connected to one another by ETHERNET and connected to the outside world by a Telesystems 1-megabyte radio ETHERNET. XAVIER runs a distributed, concurrent software system under the MACH operating system. All software development was done in C using the task control architecture (TCA) (Simmons 1990, 1992), which handles interprocess communication and synchronization. Communication with XAVIER is primarily speech driven, using an off-board NEXT computer running the SPHINX real-time, speaker-independent speech-recognition system (Alleva 1993).

Strategies

CalTech, ALFRED: The most difficult part of this contest for ALFRED was self-localization. The contest was designed to allow robots to self-localize through the use of vision, but ALFRED had no camera. Instead, the robot used a strategy of building a map of its surroundings until it was able to unambiguously match the map it was constructing against a portion of the a priori map, thus determining the robot's position.

Unlike most map-matching strategies, ALFRED's map representation was procedural rather than geometric, which greatly simplified the matching process. The robot began searching for a wall by wandering until an obstacle caused it to move in a straight line for more than one meter. It then followed the wall in the opposite direction to verify that it was, in fact, a wall. Once verified, it then continued to follow the wall, turning at corners and at wall ends and recording the pattern of left and right turns. It turned out that this pattern was unique for each wall assembly in the test course, allowing ALFRED to uniquely determine its position. Thus localized, the robot used a simple network-based path planner to explore for the coffeepot and go to the target room. Without vision, of course, the robot was unable to actually find the coffeepot. Instead, it systematically explored the appropriate quadrant until it was told by a virtual coffeepot sensor that it had entered the correct room. The virtual coffeepot sensor actually failed because of an obscure bug in MACINTOSH Common Lisp, and human intervention was required at this point. It was suggested after the contest was over that the vision system of TIN MAN (KISS Institute) could have been mounted on ALFRED, allowing it to complete the event without human intervention.

One problem that plagued many of the other entries was dead-reckoning errors introduced by rotational drift. ALFRED avoided this cumulative error by periodically aligning the robot to the walls. The infrastructure provided by the ATLANTIS sequencer made it easy to add this capability.

Carnegie Mellon, XAVIER: The Carnegie Mellon team began event 2 by using the speech-recognition system to input the map into XAVIER. The team described the sizes and locations of the rooms and corridors in natural language. XAVIER acknowledged verbally that it understood, and it displayed the map graphically. Speech input was also used to indicate the quadrant where the coffeepot would be found and where to deliver it. The robot was then told to find and deliver the coffeepot.

XAVIER localized itself by first traveling forward until it found a wall, then following walls until its sonars detected a corridor (consisting of two straight, parallel sides). Once in a corridor, the robot navigated in the direction of the corridor, turning only when it found that the corridor ended. While navigating, XAVIER analyzed the sonar readings for evidence of doorways. To compensate for noise in the sonars, evidence was accumulated using Bayesian updating until the robot had enough confidence that a door was actually present. At this point, it would stop in front of the doorway, pan its camera toward the two sides of the doorway, and look for the perceptual markers that were at specified locations on the map. Once XAVIER found a marker (along with the corresponding bar code), the robot assumed it had successfully localized itself on the map. It then was to plan a path to one of the rooms in the quadrant that contained the coffeepot and navigate through the corridors to the room. The corridor navigation used a combination of dead reckoning and feature detection to find junctions and doors that corresponded to locations on the map. Once in a room, XAVIER would visually search for the coffeepot. If it did not find the pot, it would go to and search another room in the quadrant. When the coffeepot was found, it would navigate to the delivery room.

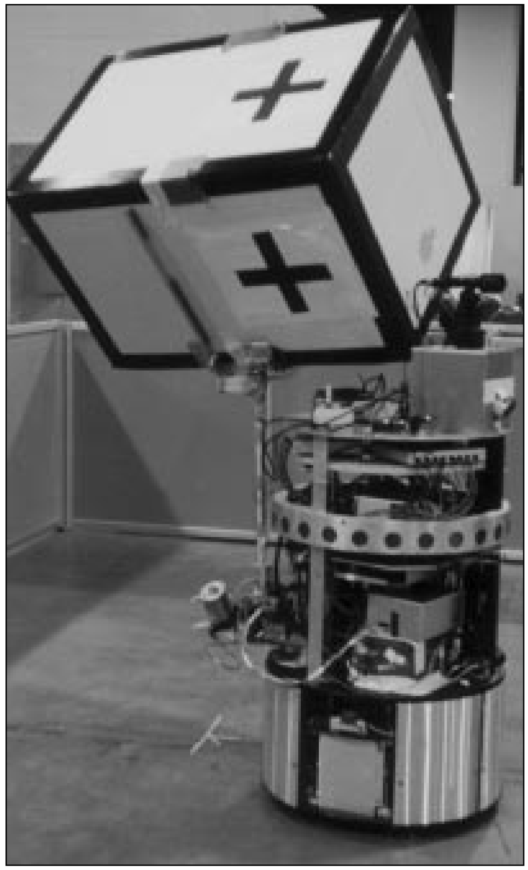

Xavier

In the actual competition, XAVIER quickly

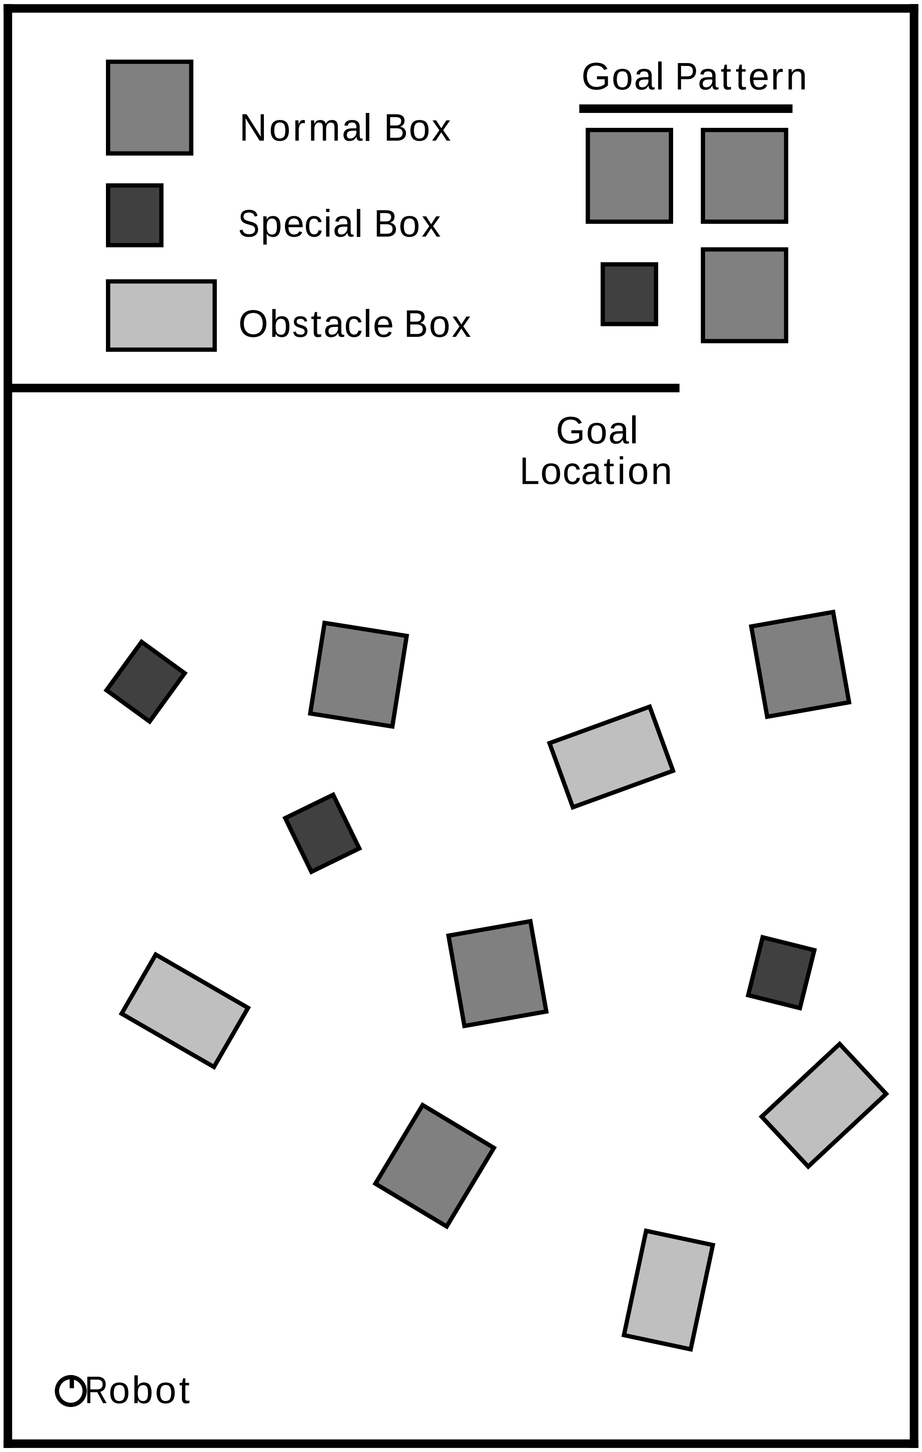

Figure 3. Diagram of Arena Layout for Event 3 (15 × 24 meters).

got out of the starting room but wandered the corridors a bit before finally finding a door with a marker. Its corridor navigation was fairly robust, but its door detector found many false positives, causing the robot to stop unnecessarily several times in search of markers. Unfortunately, after succeeding to self-localize, an obscure software bug (which had never shown up during testing) caused the path planner to crash, ending the run and Carnegie Mellon's chances to deliver the coffeepot.

Lessons Learned

The outcome of event 2 is fairly strong evidence of the power of cognizant failure (Firby 1989; Gat 1992). *Cognizant failure* is a design principle that states that the control system should be designed to detect failures whenever they occur (as they inevitably will) so that corrective action can be taken. By detecting and recovering from failures, the overall system reliability can be made high, even though component failures can be common. Both the ATLANTIS and TCA control architectures provide exception-handling capabilities, but because of time constraints, XAVIER did not use the TCA exception-handling capabilities. To some extent, XAVIER compensated for this by using redundant sensors (in particular, sonar and vision) in self-localizing and combining potentially unreliable sensor readings over time.

The ATLANTIS design methodology, by contrast, calls for exception-handling capabilities at every step of the design process rather than simply as an afterthought at the end of the process. ATLANTIS assumes that things will fail. By designing the system to detect failures and recover from them, the final product can achieve robust overall behavior. It turns out that approaching the problem in this way can dramatically reduce the amount of effort required to program a robot reliably. When the CalTech team left for the conference, the only software that had been written was the low-level control routines. One programmer developed all the contest-specific code in the three-day period immediately preceding the contest. All the code specific to the second event was developed in about 12 hours. This achievement is attributable in large part to the disciplined application of a bottom-up development methodology based on failure detection and recovery. The Carnegie Mellon team is confident that when it is able to make full use of TCA's similar contingency facilities at next year's contest that XAVIER will be a formidable entry.

Event 3: Office Rearrangement

The final event was unique in that the task required the robots to actually touch objects and manipulate their environment. This task presented a unique challenge: How can a robot maneuver objects and complete the event safely? If the robot is pulling a box, its tail can collide with other objects as it goes around obstacles. If the robot is pushing a box, it is blind in front and must use an internal map for guidance (figure 3).

The arena contained three types of boxes: (1) obstacle boxes (nonmovable), (2) normal boxes (movable), and (3) special boxes (movable). The boxes were scattered in the arena and could be distinguished by both size and perceptual markers on their sides. The task was to maneuver the boxes into a specified pattern consisting of three normal boxes and one special box in a square pattern. Any type of maneuvering was allowed. The boxes could be pushed, pulled, or even carried.

Robot Descriptions

MARGE (mobile autonomous robot for guidance experiments) was recently configured by North Carolina State University students to serve as a test bed for new approaches to the field of autonomous navigation. Based on a Cybermotion NAVMASTER robot, MARGE carries an on-board computer and sensor payload built at the university. Three 68040 processor cards networked on a VME bus process data from 2 charge-coupled device cameras, 19 ultrasonic range finders, and 28 tactile whiskers. A colony of fuzzy controllers compete and cooperate to determine the emergent behavior of the robot at the reactive level.

Lockheed's robot, ARGUS, is a NOMAD 200 from Nomadic Technologies. It has a laser range finder, a ring of 16 sonar sensors, a ring of 16 infrared sensors, and 2 rings of tactile sensing bumpers. Computation was done on a SPARC 2, and actions were communicated to the robot through a 9600-baud radio modem.

Strategies

N.C. State, MARGE: The office-rearrangement task required that the robot not only navigate through a random obstacle field but also modify its environment in a specific manner. This level of complexity poses a tremendous challenge to classical world-modeling techniques. Inspired by the work of Rodney Brooks at the Massachusetts Institute of Technology, the N.C. State team chose instead to embark on the "road less traveled" and design a reactive architecture.

Ants, wasps, and honeybees are adept at finding food and returning to a goal point. These capabilities are clearly analogous to the task of moving marked boxes. Therefore, an office rearranger does not require intelligence so much as it requires reactive competence. The N.C. State team abandoned the idea of a world model and instead configured MARGE as a complete creature adapted to the contest environment. The solution was to develop an artificial nervous system that mixes neuroethology with fuzzy logic.

Prior to the competition, MARGE was being used for research involving visual landmark recognition, ultrasonic sensing, and fuzzy systems. Obstacle-avoidance and goal-seeking behaviors were developed using fuzzy control rules. MARGE's fuzzy development environment allows independent behaviors to be fused using additional qualitative rules. As new behaviors were added for the competition, the resulting network of controllers became analogous to a simple nervous system, complete with motivational states such as landmark attraction and frustration. Finally, the adaptive qualities of the control scheme ensured that MARGE would not get stuck in endless repetitive behaviors.

Because MARGE could not see over the boxes with most of its sensors, the robot would not be able to push the boxes without a map. The solution was to drag the boxes behind it. The robot pulled boxes using a vacuum gripper, which required little modification to the target boxes. A single 70-centimeter-diameter suction cup at the end of a compliant foam tail allowed MARGE to drag the boxes robustly. An on-board, modified air compressor provided the necessary vacuum. Tactile whiskers surrounded the suction cup, allowing MARGE to sense a box in its grip. Using a landmark recognition algorithm that finds the + and × signs in the lower camera image, MARGE locat-

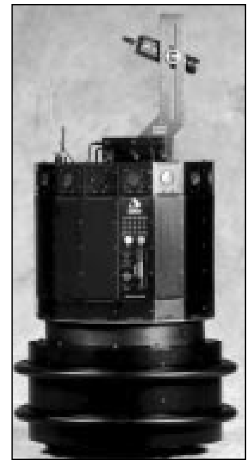

Scimmer

Marge

Call for Participation

Third Annual AAAI Mobile Robot Contest and Exhibition

The third annual AAAI mobile robot contest and exhibition will be held at the AAAI-94 conference in Seattle, Washington, July 31-August 4. The format will be similar in spirit to last year's contest, with events designed to allow participants as many opportunities as possible to showcase their robots' capabilities in a spirit of friendly competition. In addition, this year we will provide the opportunity for participants to showcase their robots' capabilities in a public exhibition.

The organizing committee hopes to have a draft set of rules distributed very early this year in order to allow participants as much time as possible to prepare for the contest. In order that we may meet the needs of as many participants as possible the committee requests that anyone interested in participating in the contest contact us so that we may start a mailing list for obtaining feedback about the rules and other aspects of the contest.

What follows is a the current strawman proposal for the contest format. Although it is couched in imperative terms, it may be subject to some change in the next few months.
The competitive phase of the contest will be held July 31 and August 1, concurrent with the tutorials and workshops. The contest, which will be judged, will consist of two events, one focusing on navigation and the other on an integrated delivery-type task. Both events will take place in the same arena, which will be a hallway-and-office environment similar to that used in the 1993 contest.

A publically attended robot exhibition will be held during the main conference. The exhibition will include the final rounds of the competition, plus "freestyle" events that showcase robot capabilities. Participants are encouraged to enter the exhibition, regardless of whether they participate in the competition.

In addition, a symposium is being planned, to be held during the conference, that will discuss strategies, tactics and open problems that arise in the context of the competition tasks. Speakers at the symposium will include participants in the competition and exhibition.

Interested participants are asked to contact Reid Simmons at the School of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213 (reids@cs.cmu.edu).

ed the boxes. It then turned around and backed up to dock with them using the vacuum gripper.

The final problem involved finding the drop-off location. Dead reckoning would be accurate enough for a single trip, but the cumulative effects of successive trips would cause too much positioning uncertainty. Because MARGE did not have a map of the arena, relocalization presented a problem. The solution was to place a special landmark at the drop-off point. It was against the rules for team members to modify the environment in such a way; however, it was perfectly legal for the robot itself to place a landmark in the arena! The N.C. State team used a tall marker with large H (Home) signs on all sides and wheels on its base that allowed MARGE to tow the special marker to the goal location at the start of the event. This special marker was taller than all arena obstacles, so MARGE was able to see it from any part of the arena with its upper camera.

MARGE performed well on the day of the event, moving four boxes into place well within the time limit. In the process, MARGE demonstrated its ability to wander through the obstacle field, capture boxes, and circumnavigate obstacles with its cargo in tow. This team stuck to the original competition rules in terms of the robot's starting location, the obstacle locations, and minimal environment modifications.

The success of MARGE's biologically inspired design demonstrates the importance of concurrent engineering. Hardware and software development are not independent events; rather, the creation of a reliable autonomous creature requires a close coupling of these two disciplines. Inspection of living organisms gives testimony to this relationship between intelligence and physiology. As the sciences become more specialized, scientists and engineers are in danger of overlooking many such useful multidiscipline solutions to problems. The N.C. State team hopes that its emphasis on real-world competence will promote more sophisticated and robust applications of mobile robots.

Lockheed, ARGUS: The Lockheed strategy for this event was to use existing hardware and compete on the basis of innovative software. ARGUS used its laser range finder to scan the environment, distinguishing potential boxes from walls by virtue of line-segment length. When it detected a potential box, ARGUS maneuvered to the opposite corner of the box, focusing its attention on the sides of the box to better map the box's position. During this maneuver, ARGUS would see two adjacent sides of the box. After analyzing its sensor map of the box region, ARGUS could identify the box's type by virtue of its length and width.

Argus

Once ARGUS identified the box, it tried to get into the best position from which to push it. ARGUS did not use its bumper sensors to detect when it was pushing the box because the boxes were not heavy enough to trip these tactile sensors. When it was pushing the boxes, ARGUS did not use the sonar or laser sensors because the boxes were too close to the robot to get accurate range readings. However, the infrared sensors were accurate from the surface of the robot to approximately three feet; so, ARGUS used them to detect both the position and the orientation of the box.

After the robot was approximately in the pushing position, it attempted to push the box to the goal location. If the robot was not in the correct starting position when it attempted to push the box, it was sometimes unable to sense the box in front of it, and it would retry to get into the correct position. While pushing a box, ARGUS tried to keep it balanced in front of it. Because ARGUS was a round robot pushing a square box, the box had a tendency to shift to one side or the other. As ARGUS was pushing, it was constantly checking the orientation of the box and adjusting its trajectory. For example, if the box shifted 20 percent to the right, ARGUS maneuvered 60 percent to the right in an effort to get the box back on the correct trajectory. If the box got too far out of balance, the robot would briefly stop pushing the box and maneuver to get a more centered pushing position. Through this fairly tight feedback loop, ARGUS was generally able to keep the box correctly balanced in front of it.

ARGUS performed well in this event. It operated both safely and completely autonomously, moving all but one box into the goal position. ARGUS safely pushed three of the four boxes into the correct position, successfully completing 3/4 of the task, including getting the special box in the correct position. ARGUS then maneuvered into position and was starting to push the final box to the goal location when the communication link between the robot and the computer failed. By the time the robot was restarted and ARGUS rediscovered the box, the time limit was up. The judges said they were impressed with ARGUS's ability to keep a box balanced in front of it.

Lockheed has been working in autonomous systems and robotics for several years, acquiring the NOMAD 200 early this year. Much of the preparation time was spent developing general movement, obstacle-avoidance, and mapping capabilities for the robot as well as working on the competition-specific routines. Currently, Lockheed is expanding its replanning and mission-analysis tools, using the robot to demonstrate and validate the tools and capabilities.

Lessons Learned

MARGE and ARGUS used different approaches to complete this task, but both performed well. Although MARGE demonstrated the advantages of reactive sensor-based behaviors, ARGUS showed how a reliable world model can be constructed and used to perform robustly without the benefit of immediate sensor data. Vision allowed MARGE to detect signs on the boxes from a long distance, which gave it one advantage over ARGUS. However, infrared and laser range finders provided ARGUS with more precise information about the position and orientation of the boxes, allowing ARGUS to make an accurate map and move the boxes in a precise manner. ARGUS used a radio modem and an off-board computer system, but MARGE used entirely on-board control. Its on-board control meant that MARGE used up battery power quickly; however, it was not susceptible to the radio modem problems that ARGUS encountered.

The other teams that competed in event 3 were not successful in completing the task but demonstrated innovative solutions. Carnegie Mellon's robot, XAVIER, featured a large V-shaped arm that used two electromagnets to hook onto metal plates mounted on the corners of the boxes. The arm was custom built for the task of picking boxes up over the robot's head. The mechanism turned out to be a reliable means of moving boxes as well as a real crowd pleaser. The University of Michigan team used two robots, CARMEL and BORIS, to share the task of navigation and box manipulation. While BORIS pushed the box, CARMEL moved ahead of BORIS and guided it to the goal position.

Unstructured and unknown environments have always posed a challenge to

… the field of autonomous mobile robots has made tremendous strides in the past five years.

autonomous robots. The office-rearrangement task required a robot not only to navigate safely in such a world but also to actively modify it. Even if the environment were mapped at the beginning of the event, once the robot moves a box, the environment might be modified in an unpredictable way. One point that was demonstrated by this competition is the importance of reactive competence and the ability to sense the state of the world. Another point is that high-level planning can be a powerful tool but only if the system can detect its mistakes and recover from them. Finally, it is important to keep solutions simple and task oriented. A robust solution to a real-world problem does not always have to be expensive or complicated.

Conclusions

The overwhelming lesson of this contest is that the field of autonomous mobile robots has made tremendous strides in the past five years. These strides can be attributed at least in part to the availability of off-the-shelf systems that are capable of performing complex tasks without requiring custom hardware development. Three of this year's entries were 100-percent off-the-shelf hardware, and almost all used an off-the-shelf mobility platform. There has been a developing consensus that mobile robot control architectures must combine features of traditional sense-plan-act architectures with those of the more recent behavior-based architectures. Four of the five robots described in this article use such a hybrid architecture. However, the debate continues about whether such architectures should be developed top down or bottom up. This year's contest appears to provide a data point to support the bottom-up position, but the jury is still out.

One conspicuously missing aspect of the current work in robotics is a strong tie between theory and practice. There is a large body of advanced mathematical results in the robotics literature that was little in evidence in this year's contest. Although there is too little data to draw strong conclusions, it does appear that system engineering issues are the main limiting factors in the current state of the art. Thus, there appears to be some support for Herb Simon's (1993) position that we should spend a bit more effort improving our engineering skills and leave the math to the mathematicians for a while.

Nevertheless, it is an exciting time to be working with mobile robots. The field of robotics appears to be at about the same stage of development as the microcomputer field was in the early 1970s, when off-the-shelf computers were first mass produced for education and hobbyists. Although we are still far from having a general-purpose autonomous robot, the dream does seem to be more within reach now than at any time in the past.

References

Alleva, F.; Huang, X.; and Hwang, M-Y. 1993. An Improved Search Algorithm Using Incremental Knowledge for Continuous Speech Recognition. In Proceedings of the 1993 IEEE International Conference on Acoustics, II, 307–310. Washington, D.C.: Institute of Electrical and Electronics Engineers.

Firby, R. 1989. Adaptive Execution in Dynamic Domains, Ph.D. diss., Dept. of Computer Science, Yale Univ.

Gat, E. 1992. Integrating Planning and Reacting in a Heterogeneous Asynchronous Architecture for Controlling Real-World Mobile Robots. In Proceedings of the Tenth National Conference on Artificial Intelligence, 809–815. Menlo Park, Calif.: American Association for Artificial Intelligence.

Simmons, R. 1992. Monitoring and Error Recovery for Autonomous Walking. In Proceedings of the IEEE International Workshop on Intelligent Robots and Systems, 1407–1412, Washington, D.C.: Institute of Electrical and Electronics Engineers.

Simmons, R. 1990. An Architecture for Coordinating Planning, Sensing, and Action. In *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control,* 292–297. Los Altos, Calif.: Morgan Kaufmann.

Simon, H. 1993. Artificial Intelligence as an Experimental Science. In Proceedings of the Twelfth National Conference on Artificial Intelligence, 853. Menlo Park, Calif.: American Association for Artificial Intelligence.

Illah Nourbakhsh is a graduate student in the Stanford University computer science department. His research has included protein structure prediction under the GENOME Project, interleaving planning and execution, software reuse, and mobile robot navigation. He is a member of the Logic Group at Stanford.

Sarah Morse is a senior in computer science at Stanford University. Her research with the Logic Group has focused mainly on mobile robotics and computer vision.

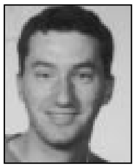

Craig Becker is a graduate student in the computer science department at Stanford University. His research interests include planning, visual recognition, and visual navigation for autonomous mobile robots.

Marko Balabanovic holds a B.A. from Cambridge University and is currently a Ph.D. student in the Stanford University computer science department. His research is in architectures for the control of mobile robots and the automatic learning of maps.

Erann Gat is currently a visiting associate in the Department of Mechanical Engineering at the California Institute of Technology as well as a member of the technical staff at the CalTech Jet Propulsion Laboratory. He has been working on autonomous mobile robots since 1988. He received a Ph.D. from the Virginia Polytechnic Institute and State University in 1991. Gat's research interests include the combining of strategic planning and reactive control, robot control languages and compilation techniques, and theoretical analysis of systems embedded in unmodeled or unmodelable environments.

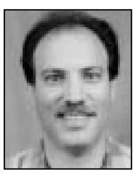

Reid Simmons is a research scientist in the Department of Computer Science and Robotics Institute at Carnegie Mellon University. He earned his B.A. in computer science in 1979 from the State University of New York at Buffalo and his M.S. and Ph.D. from the Massachusetts Institute of Technology in 1983 and 1988, respectively, in the field of AI. 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. This work involves issues of robot control architectures that combine deliberative and reactive control, selective perception, and robust error detection and recovery. The ideas are currently being applied to the XAVIER robot, a six-legged planetary rover, and an autonomous excavator.

Steven G. Goodridge is a research assistant in the Department of Electrical and Computer Engineering at North Carolina State University. His interests include autonomous robots, space vehicles, fuzzy control systems, and biological models of intelligence. Goodridge received his B.S. in electrical engineering from the University of New Hampshire in 1991. He is a member of the Institute of Electrical and Electronics Engineers and Tau Beta Pi.

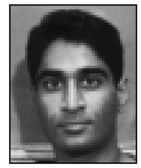

Harsh Potlapalli is a research assistant in the Department of Electrical and Computer Engineering at North Carolina State University. His interests include autonomous robots, neural networks, texture analysis, and natural scene understanding. Potlapalli received his B.E. in electrical engineering from Osmania University in 1987 and his M.S. in electrical engineering from Tulane University in 1989. He is a member of the Institute of Electrical and Electronics Engineers and Eta Kappa Nu.

David Hinkle is a research scientist in the Artificial Intelligence Center at Lockheed Palo Alto Research Labs. He received his B.S. in computer science from Syracuse University in 1987 and his M.S. in computer science from Northeastern University in 1989. His research interests include robotics, autonomous systems, planning and scheduling, machine learning, and case-based reasoning. He is a member of the American Association for Artificial Intelligence.

Ken Jung is a research scientist in the Artificial Intelligence Center at Lockheed Palo Alto Research Labs. He earned his B.S. in 1982 from Arizona State University and his M.S. from the University of Illinois at Urbana-Champaign in 1983, both in electrical engineering. His interests include robotics, signal and image processing, parallel computing, data compression, and very large-scale integrated architecture.

David Van Vactor received his B.S. in mathematics and physics (independent program of study) from Harvey Mudd College in 1972. He has worked at Lockheed since this time in a variety of areas, including thermodynamics, real-time satellite control, AI, and automatic image interpretation. He has developed a model-based image-interpretation system for robots and is responsible for incorporating this design into the Lockheed test bed for autonomous robotics.