

Designing the 1993 Robot Competition

Kurt Konolige

■ The Second Annual Robotics Competition and Exhibition was held in July 1993 in conjunction with the National Conference on Artificial Intelligence. This article reports some of my experiences in helping to design and run the contest and some reflections, drawn from post mortem abstracts written by the competitors, on the relation of the contest to current research efforts in mobile robotics.

The Second Annual Robotics Competition and Exhibition was held in July 1993 in conjunction with the National Conference on Artificial Intelligence (NCAI) in Washington, D.C. The competition, which attracted teams from many of the top mobile robotics research laboratories in the United States (see side bar), was first proposed by Thomas Dean and held at the 1992 NCAI conference. Dean's concept was to further the research into the skills such robots need—sensing, interpretation, planning, and reacting—by bringing together interested parties in a cooperative and challenging environment. Ideas should be tested in the real world, not just the controlled conditions of the laboratory.

The Robot Competition

The first competition was an enormous success, thrilling spectators and generating camaraderie and shared experience among the teams. Dean and Peter Bonasso (1993) wrote

an excellent article on this competition; a companion article describes two of the top finishers (Congdon et al. 1993). The second competition attracted even more competitors, and the sophistication of the hardware and software systems seemed to have grown geometrically. An article by Ilah Nourbakhsh (Nourbakhsh et al. 1993) describes some of the top teams from the competition, the strategies they used, and their results.

As co-chairmen of the 1993 robotics competition, Terry Weymouth and I had the challenging and often chaotic job of designing the rules, coordinating the setup and administration of the contest, and trying to cope with the needs of the 15 teams that put so much energy into their entries. This article reports some of the experiences I had in helping to design and run the contest and some reflections, drawn from post mortem abstracts written by the competitors, on the relation of the contest to current research efforts in mobile robotics.

Most of the teams (see side bar) arrived at the cavernous exhibition hall by Sunday night. There was plenty of space for the robots and associated paraphernalia, and it looked like a true robot-builders convention, with computers and robot hardware in various stages of completion. As in the first contest, many of the teams were debugging hardware and coding major sections for the contest. Camaraderie and coopera-

tion were high as teams helped each other with missing or broken equipment and hardware problems. As usual, the teams were piqued by others' solutions: how had they solved the navigation problem, for example. One of the benefits of the competition is this free exchange of information, which is enhanced by seeing how well various techniques work out in practice.

To some extent, the hardware problems were lessened this year by the proliferation of commercially available mobile bases, sensors, and software systems, such as those from Real World Interfaces (GORT, ALFRED, JAMES, XAVIER), Nomadic Technologies (SCIMMER, ARGUS, XAVIER), Cybermotion (CARMEL, MARGE), Denning Robotics (CLEMENTINE, XAVIER), and Transitions Research Corporation (TJ, BORIS). However, some of the most interesting entries used home-brewed mechanisms, and many of these worked well and also proved to be highly entertaining. Continuing a tradition started in the first contest with SCARECROW, the TIN MAN team constructed its robot from off-the-shelf hardware components put together in inventive ways, including a weird Addams-family-like hand that guided the robot with a joystick. In addition, a commercial company, MacLeod Technologies, showed off a unique global-positioning system using a robot-mounted revolving laser and three or more stationary receivers.

Still, many teams suffered frustrating failures in hardware and especially software, leading to a general lack of sleep and noticeable exhaustion among the contestants by Monday night, the day before the contest. I know this from personal experience: FLAKEY was still being debugged for the contest, and I was wearing dual hats as a member of the three-person FLAKEY team as well as a co-chairman of the competition. Fortunately, the energy level was high, and most of the teams vowed that they would have their system problems under control "next time." All the teams managed to participate in at least one event, except GOPHER, which relied on a magnetic compass for its bearing; the convention hall had too

TEAMS IN THE AAAI-93 ROBOTICS COMPETITION

ARGUS	Lockheed Missiles and Space Company
REN	Georgia Institute of Technology
ISAAC	MacLeod Technologies, Inc.
TIN MAN	KISS Institute
CARMEL, BORIS	University of Michigan
GORT	Brown University
XAVIER	Carnegie Mellon University
FLAKEY	SRI International
GOPHER	Massachusetts Institute of Technology, Panasonic
CLEMENTINE	Colorado School of Mines
SCIMMER	Stanford University
TJ	Yale University
MARGE	North Carolina State University
ALFRED	California Institute of Technology
JAMES	Worcester Polytechnic Institute

many massive iron pipes, causing complete turnarounds in magnetic north across the arena.

Designing the Events

What is the robot competition supposed to measure, and how do we design appropriate events? The distinguishing feature of the NCAI competition is the emphasis on autonomous action in natural environments. The operative word here is *emphasis*; mobile robots today do not display anything near the flexibility and adaptability that humans have in performing tasks. Each year we expect to make the environment a little more natural and less engineered for robots. In the first NCAI competition, three events tested the robots' abilities to move among obstacles and to find and return to a distinguished set of objects in the arena. This format was a good base to start from because we had an idea of how well the robots could perform and what parts of the competition we could push to stretch their performance. Over the eight months before the contest, in consultation with a mailing list of interested parties, we developed a new set of events based on the following guidelines:

First, the tasks should be described in terms of performance, for exam-

ple, delivering a coffeepot or moving from one place to another. Any method of accomplishing the task that fell within the constraints of the contest would be considered equivalent. There would be no penalty for using or not using a certain strategy or capability, such as dead reckoning or map learning. One of the most interesting facets of the competition is to see the creative strategies that the teams find to accomplish a task.

Second, the tasks should be doable but demanding enough so that simple engineering solutions would not work. The key to accomplishing this goal is to broaden the class of environments that the robots are expected to work in and to keep the environment natural (no modifications by the teams). In the first competition, SCARECROW was engineered and built specifically for the contest and had very little utility in even slightly different environments. We were looking to reward more general-purpose robots that could act autonomously in widely varying circumstances.

Third, given the ubiquity of visual cues in human environments, there was a conscious effort to encourage the use of visual perception. Special visual markers were designed and placed in key areas of the events. This part of the competition was one of

the most successful because even though image understanding is difficult and computationally expensive, many teams integrated it into at least some of their work (see discussion later).

Fourth, manipulation of the environment should be part of some task. Although still difficult for current mobile robots, we should push this aspect of the competition because it is essential for truly useful robots.

Fifth, because robots must often interact with humans, we tried to emphasize communication between man and machine. With a few exceptions, this aspect of the competition is still disappointing, and it is difficult to design tasks that reward appropriate communication. As many of the team members noted, there is a tension between autonomy and communication: One can view teleoperation as an extreme form of communication. This area is one that future competitions need to address.

Sixth, robot-robot cooperation was also encouraged. It is clear that there is still a long road to travel here, but CARMEL and BORIS made a promising start as the only multiple-robot entry, acting in a master-slave configuration.

Because a typical application area for mobile robotics is office automation, we picked an office domain for the competition. The first event, called Escape from the Office, emphasized reactivity to the environment and the ability to navigate among unmapped obstacles. Robots were positioned somewhere in a 4- x 5-meter office replete with real office furniture. After a few minutes, one of three doors would be opened; the robots had to find the door, exit, and make their way among obstacles to the end of the arena. Almost all the robots competed in and finished this event.

The second event was Office Delivery: Given a map of an office environment, find a coffeepot located somewhere in one quadrant of the map and deliver it to a designated room. This event tested navigation and map-registration skills. However, there was a twist that made this event one of the most difficult: Robots were

started at an arbitrary location in one quadrant of the map. Initial self-localization with respect to the map turned out to be one of the most difficult problems for the robots, and only four managed to do it successfully (ALFRED, XAVIER, SCIMMER, and FLAKEY). ISAAC, by virtue of its global-positioning system, always knew where it was after an initial calibration.

The final event, called Rearranging the Office, was the only one involving manipulation of the environment: The robots had to position four marked blocks into a pattern in a designated area of the arena. This event was highly entertaining, with creative home-brewed devices for manipulating the boxes in a way that would leave the front of the robot clear to see for navigation. XAVIER lifted the boxes over its head using a grappling arm; MARGE had a unique suction-tail device that could pull a box along behind it. CARMEL and BORIS cooperated, with CARMEL first plotting a route and then BORIS pushing a box along the route.

Perceptual Markers

In an attempt to promote the use of vision systems, a set of uniform perceptual markers was added to the environment. These markers consisted of black crosses, bars, X's, and O's on a white background. They were placed on key areas in the events: in event 1, on each door; in event 2, near selected doorways and the coffee pot; and in event 3, on the boxes. Using the markers was not obligatory, and all the events had alternative means for identifying objects with markers, such as their shape (in event 1, the doorways could be identified after they were opened).

It was gratifying to see teams take advantage of the perceptual markers with a wide variety of visual systems. GOPHER used miniature, low-resolution video cameras to find the markers and also to find free space. FLAKEY had a stereo head on one of its cameras with the capability of real-time stereo for obstacle detection; unfortunately, the camera's autogain interacted badly with the stereo algorithms, and it was never used. TJ's

vision system was based on a theory of finding invariants in special bar-code-like marks printed on standard paper; it was fast and robust, taking 15 milliseconds to recognize the marks in cluttered scenes.

Most of the other vision systems used standard video cameras and processed the information with general-purpose computers. A variety of pattern-recognition methods were used in detecting the marks; their speed varied from milliseconds to several minutes. Almost all the vision algorithms were bothered to some extent by perspective distortion of the marks, especially in the Office Delivery event, where the robot tended to view the marks at oblique angles in the narrow corridors.

Integrating vision with other sensing and action was one of the most interesting challenges for those teams using vision. For example, in event 1, although the ultimate winner, SCIMMER, did not use vision, two of the top challengers (FLAKEY and XAVIER) did. Unfortunately, they both experienced strategic programming problems in the finals. Here is XAVIER's post mortem account:

The problem was that it [XAVIER] got stuck in a corner for most of the first minute, and so did not have a chance to circumnavigate the room before heading to the center to begin visually scanning for markers. As such, it thought the room was smaller than it really was. The problem was that, wanting to eliminate spurious marker sightings, we programmed XAVIER to ignore perceived markers that were far outside its model of the room. Thus, XAVIER mistakenly thought that the wrong door was open (since it ignored a marker it had actually recognized), and headed for that "door." When it was found to be closed, the robot headed back and looked for markers again. This cycle repeated until time expired. In retrospect, we should have trusted our vision more and explicitly encoded the probable office dimensions (as many of the other contestants did). In any case,

alternative strategies to avoid looping behavior would have been helpful, to say the least.

This is what the competition is all about: integrating perception, planning, action, and error recovery to achieve a task.

The use of perceptual markers was not without controversy. Many teams protested that using the common markers made it difficult for those who were without vision systems but who had invested in other technologies, such as reflective bar-code readers. Their main point was that the perceptual markers were no more natural than bar codes. This issue is bound to resurface in future competitions.

Natural Environments

To be accepted as cooperative workers, mobile robots must be able to function in the same environments that humans use, without excessive engineering to make it easy on them. Current robots have some severe limitations, of course: stairways, for example. Still, we hope to move to more and more natural environments, creating harder challenges for the robots as the competition continues.

We considered three dimensions of naturalness. The first dimension is the shape of the environment and the objects that are in it. To many at the contest, the office arena did not look natural: Obstacles were big white cardboard blocks with contrasting tape on their sides; doors were mostly just openings; and there were no potted plants, moving people, rugs, or the thousand other common objects one finds in an office environment. In event 1, in which the robot had to escape from an office when the door opened, we did use real office furniture such as chairs and file cabinets; the other two events did not. We hope that the move toward natural space will progress in future competitions.

A second dimension of naturalness is the a priori information that the robot is given. Consider event 1; three pieces of information would be useful: (1) a map showing the dimen-

sions and position of the office and arena, (2) the location of the doors, and (3) the initial position of the robot.

Given this information, a simple robot equipped with a proximity sensor and with good wheel encoders for dead-reckoning movement should be able to do well using a simple strategy: Move to each door by dead reckoning (perhaps avoiding obstacles), check if the door is open using the proximity sensor, and go through the door to the finish line if the door is open. This strategy won't work if we eliminate knowledge of the door positions rather than put perceptual markers of some kind on the doors. Now the robot must be able to find the doors using perception. Even harder would be to eliminate information pieces 1 and 3 so that the robot has to figure out the dimensions of the office, localize itself within the office, and find the doors. We compromised by giving the following limited information: (1) approximate dimensions of the office and arena, (2) perceptual markers on the doors, and (3) the initial orientation of the robot but not its position within the office.

By creating uncertainty about the environment, the problem becomes much more challenging because the robot designer must use more general strategies for dealing with the uncertainty. In the case of event 1, the robot must have at least rudimentary map-making, door-sensing, and navigational capabilities.

A third type of environmental naturalness is whether engineering is allowed or not. *Engineering* the environment takes the form of adding perceptual markers or other modifications that make it easy for the robot to perform a particular task. For example, in event 1, some teams put cardboard edging on the office furniture so that it could be seen more easily by the sonars. In event 3, teams added special hooks or other mechanical devices to boxes so that they could grab them (the most spectacular were XAVIER's box-lifting arm and MARGE's vacuum-suction tail) as well as perceptual markers such as reflective bar codes. Perhaps the most

obvious example of such environmental engineering is ISAAC's global-positioning system. This system had remarkably good results in event 2, where the robot could navigate by tracking its position on a map of the offices.

This aspect of naturalness engendered the most controversy. There is no doubt that engineering is good practice in many applications. In fact, humans constantly do such engineering: roads, street signs, color-coded buttons, and so on. In the factory and workplace, the best (cheapest, most reliable) solutions for mobile robot tasks can involve special markers or other aids. Still, we were not trying to promote just good engineering for a specific task but also the ability to operate in more unconstrained and unengineered environments. Penalties for modifying the environment were high, although it was clear that many teams were not happy with this arrangement and considered the perceptual markers that were used (see discussion later) no more natural than the ones they wanted.

Finally, many teams wanted prior information about the statistical nature of the domain for help in designing their strategies. For example, SCIMMER took advantage of the fact that the walls were all on a rectangular grid when self-localizing during event 2. Other teams wanted to know the range of sizes of blocks, the approximate density of the obstacles, or the range of door sizes. This last point was a good example of how difficult it is to find compromises that suit all the teams. TIN MAN wanted doorways that were human size so that it could tell by its touch sensors that it had gone through; CARMEL needed doorways nearly a meter and a half wide so that it could go through without bumping.

In my opinion, the need for precise a priori information about the environment, although helpful in designing special strategies for robot tasks, is symptomatic of nonautonomy in mobile robot systems. The designers are obviously bright, and given enough information, they can code strategies that work on their

machines in particular environments. What we need to do is put some of these smarts into the robots, so that they can derive the necessary strategies for solving a task in the class of environments they are likely to encounter. Such smart robots are still years away, and we have to compromise to make the tasks both interesting and doable.

Strategies and Results

Event 1 was the simplest, and most of the robots were able to successfully compete. There were two basic strategies: Map the office using dead reckoning and wall sensing (either sonar or laser range finders), or use vision to find the door markers. XAVIER and FLAKEY successfully used vision in their strategies, and their ability to sense door markers at a distance and move purposefully toward an open door was impressive. CLEMENTINE, CARMEL, MARGE, and GOPHER also were prepared to use a vision-based strategy in this event but had to forego it when their systems developed hardware and software problems.

Mapping the office using sonars and dead reckoning was straightforward, although the real office furniture caused some problems with sonar-based obstacle avoidance. SCIMMER, ALFRED, GORT, JAMES, CARMEL, CLEMENTINE, and REN all turned in credible performances using this strategy with no vision. However, to find the open door, they would have to circumnavigate the office, which potentially put them at a disadvantage with the vision-based robots.

Event 2 was difficult because of the lack of initial positioning information. The two crucial problems were self-localizing at the start and finding the coffeepot. Again there were two types of strategies for self-localization, one based on landmarks and the other on wall signatures. *Landmarks* are markers that uniquely identify a given place. Some teams (TJ, REN, and ISAAC) used their own landmarks (at a penalty) and were able to easily find their location with respect to the map. XAVIER was the only robot to try to find the unique visual markers posted on some of the doorways;

it successfully self-localized after searching several corridors. FLAKEY, ALFRED, SCIMMER, and CARMEL all used various wall-signature techniques to try to self-localize. *Wall signatures* are features extracted from the sonars or laser sensors, representing a part of the local wall surface. Such signatures are not unique, but by gathering a pattern over a larger area and comparing it to the map, a unique match can eventually be made. This strategy is difficult to follow, however, because it involves a directed exploration of the environment, looking for new information. The matching techniques varied, from extremely simple (rectangular overlay search for SCIMMER) to sophisticated (belief network in CARMEL). Some of the teams took advantage of unique features of the map (which was handed out the night before the contest) in hand coding their strategies. The FLAKEY team noted that the perimeter corridors had unique patterns of doorways, and ALFRED depended on a unique signature of turns in navigating from one office to the next. This kind of hard-coded strategy relative to a given map should probably be discouraged in future competitions, perhaps by handing out the map in machine-readable form just before the event. Self-localization was not a problem for ISAAC because its global-positioning system gives accurate, real-time information about its position and orientation.

Finding the coffeepot involved looking for the specially marked box that the coffeepot was sitting on. A hint about the quadrant was given; as it turned out, there was only one possible room, so exploration was kept to a minimum. Of those robots that successfully self-localized, not one was able to find the coffeepot, either because they lacked the crucial vision sense (ALFRED and ISAAC), or they experienced software failures after localizing (SCIMMER, FLAKEY, XAVIER). We look forward to next year for the first successful completion of this event.

Event 3 was a novel one because the robots had to change their environment. As initially planned, robots would have only imprecise knowledge of the arena walls and would

have to perform some self-localization. In practice, the teams rebelled because the rest of the event was so difficult, and we distributed a precise map of the arena along with the robot's initial position.

Surprisingly, even with this information, keeping the robot registered with respect to the map turned out to be a major problem. A spectacular example was provided by XAVIER: After lifting a box over its head, it misjudged the deposit location and dropped the box on the arena wall. Robots that relied too much on dead-reckoning movement and did not reregister their position with respect to the arena were subject to this problem. MARGE sidestepped the problem: Initially, it dropped a large landmark beacon at the deposit position, making it possible to home in on the position without using the arena map or having to map the position of any boxes.

Another difficult part of this event was finding the correct boxes to move into position: Some of the boxes were obstacles, and some were movable. CARMEL, XAVIER, and MARGE used vision to detect the markers and home in on movable boxes. ARGUS had a laser range finder that could recognize the shape of the boxes. All these teams were successful in coupling complex sensing with manipulation. XAVIER and MARGE engaged their special couplers, and ARGUS performed a balancing act as it pushed a box ahead of itself. CARMEL wanted to use a special coupling mechanism, but when it failed, the team switched to a multirobot strategy, with CARMEL finding a clear path for BORIS to push a box. The audience clearly appreciated the importance of the tight coupling between sensing and manipulation, and these teams solved some of the most difficult tasks of the competition. Given that this task was new and that the teams were so successful, we can look forward to increasing the complexity of the manipulation task in future competitions.

Some Open Issues

I want to comment briefly on some of the research ideas that have suc-

cessfully worked in the competition and some of the research areas that still need work. The most obvious success is in the use of reactive systems. For avoiding obstacles and accomplishing goals in tightly constrained contexts, reactive architectures have proven to be a winner. Rather than plan a complex series of moves, reactive architectures rely on sets of behaviors, or skills, that each implement small tasks, such as keeping the robot from bumping into an obstacle. Using reactive behaviors guarantees that the robot will exhibit at least a modicum of intelligent behavior, staying out of trouble and performing well when the situation is favorable. Still, there are problems with reactive behaviors. One of them is that they can take a long time to program and debug because the robot designer must test them. It would be better if the robot could be given a description of its goal and learn the requisite skill.

Another problem and one that is being addressed actively in the competition is the integration of low-level behaviors with planned activity. For example, it is relatively easy to write behaviors that keep the robot from running into obstacles; it is much harder to keep the robot headed for a particular goal, such as escaping from an office, at the same time. Architectures that blend reactivity and planning are emerging from the research arena into the real world: FLAKEY and MARGE's fuzzy control sets, XAVIER's TCA, and ARGUS's ATLANTIS. Error recovery using these architectures can be much more robust than using just a reactive system because the reason for the failure of low-level behaviors can be analyzed and corrected.

Perhaps the most challenging part of the robot competition, especially with the push toward more natural environments, is understanding the local environment through the robot sensors. Sensors such as the sonars are data poor but easy to use: They give a direct reading of the distance to surfaces at a fairly low rate. Constructing an internal model of the environment from these sensors is difficult; so, one must rely on strong modeling assumptions about the

New from AAAI Press:

Proceedings of the First International Conference on Intelligent Systems for Molecular Biology

*Edited by Lawrence
Hunter, David Searls, and
Jude Shavlik*

The interdisciplinary work in this proceedings represents original biological results as well as pragmatically-inclined applications of computational research, including work in statistics and databases.

460 pp., index.
ISBN 0-929280-47-4
\$45.00 softcover

To order, call
(415) 328-3123

or fax to
(415) 321-4457.

MasterCard and Visa accepted.

AAAI Press
445 Burgess Drive
Menlo Park, CA 94025
USA

world. For example, given the robot's initial position in a global map, it is possible to keep registered with respect to the map by using sonars to find door openings and wall segments and matching them to the map. As event 2 showed, it is a much harder problem to use these data-poor sensors to localize oneself when the initial position is unknown. In addition, finding and manipulating objects is also difficult because data-poor sensors generally do not provide adequate information to distinguish objects from each other.

However, data-rich sensors, such as passive vision, require much more complicated processing to be useful. Still, when it is possible to do this processing, the advantages can be tremendous. In event 3, for example, some vision systems could find the movable boxes as far as 10 meters away, so their searching time was considerably shortened. Also, vision can be used to give a much more complete picture of the local environment quickly, so that instead of traveling around a room to map its walls with short-range sensors, the robot can simply scan from several vantage points.

The sensor-interpretation systems in use are far from exploiting all the data that are available. With a single glance down an office corridor, we can tell the shape of the corridor and the positions of offices, water fountains, and countless other objects. Vision systems are nowhere close to matching this capability. Further, the sensor interpretation used in the competition was, to a large extent, hand coded, with relatively little autonomous learning. This is another area where learning techniques could make a big difference and take a large burden off the shoulders of the robot designer. We look forward to seeing learning algorithms being applied to a greater and greater extent in future competitions.

Acknowledgments

I want to thank the teams, who contributed untold person-hours to making their entries work. Their enthusiasm and creativity made the contest a success. Terry Weymouth contributed

his time generously in setting up and administering the competition, especially during the running of the events. Reid Simmons worked hard getting funding for the competition. The staff of the American Association for Artificial Intelligence, especially Rick Skalsky, Carol Hamilton, and Mary Livingston, made the competition possible and worked overtime helping us to get needed materials and services. Thomas Dean and Peter Bonasso, who ran the first competition, provided invaluable insight and encouragement, along with Benjamin Kuipers, Jim Firby, and David Payton. Dean generously loaned me his notes from the competition, and I borrowed heavily from them. Of course, the competition was heavily dependent on the abilities of the fine group judges who contributed their time.

References

- Congdon, C.; Huber, M.; Kortencamp, D.; Konolige, K.; Myers, K.; and Saffiotti, A. 1993. CARMEL versus FLAKEY: A Comparison of Two Winners. *AI Magazine* 14(1): 49-57.
- Dean, T., and Bonasso, R. P. 1993. 1992 AAAI Robot Exhibition and Competition. *AI Magazine* 14(1): 35-48.
- Nourbakhsh, I.; Morse, S.; Becker, C.; Balabanovic, M.; Gat, E.; Simmons, R.; Goodridge, S.; Potlapalli, H.; Hinkle, D.; Jung, K.; and Van Vactor, D. 1993. The Winning Robots from the 1993 Robot Competition. *AI Magazine* 14(4): 51-66.



**K u r t
Konolige**
is a senior
computer
scientist in
the AI
Center at
SRI Inter-
national.
He re-
ceived his
Ph.D. in

computer science from Stanford University in 1984; his thesis, "A Deduction Model of Belief and Its Logics," developed a model of belief based on the resource-bounded inferential capabilities of agents. His research interests are broadly based on issues of commonsense reasoning, including introspective reasoning, defeasible reasoning, and reasoning about cognitive state. More recently, he has been working on making moving robots more intelligent.