AI Magazine Volume 19 Number 3 (1998) (© AAAI)

The Home-Vacuum Event

Pete Bonasso and Karen Myers

■ This article discusses the setup and results from the Home-Vacuum event of the 1997 AAAI Robot Competition and Exhibition. After a summary of the rules, we outline the high and low points of the competition. Then we suggest ways such competitions could better accommodate new teams in the future.

When the original vacuuming contest committee—Pete Bonasso, Erann Gat, and Sebastian Thrun—began devising a contest, the big question in their minds was, "What does AI buy you in this task?" As early as 1993, it was apparent that industrial vacuuming robots were emerging as autonomous and well behaved at least in large industrial areas (Bonasso, Miller, and Kuipers 1993). Devising a sweep pattern on a bounded uncluttered surface to ensure complete coverage is a well-formed and solved problem. AI had nothing to contribute with regard to the basic cleaning task.

A domestic or small office venue offered more complexity. The areas were smaller and contained more furniture. However, all but the first AAAI Mobile Robot Competition had office-navigation tasks, and the top finishers showed that many efficient algorithms existed (for example, Nourbakhsh, Powers, and Birchfield [1995] and Yang et al. [1996]).

After some discussion, organizers agreed that the power of the intelligence in a home-cleanup task lay in (1) knowing how to clean efficiently, (2) knowing when there might be a need to touch up an area because of human intervention, and (3) being as unobtrusive as possible in the presence of humans. Efficient cleaning concerns deal with limited resources, for example, energy, time, and bag capacity, whereas touching up requires checking for messes in areas of recent human activity. Being unobtrusive involves adapting behavior to avoid interfering with both expected and unexpected activities of any humans in the environment.

The Rules

We combined our venue with that of the Find-the-Remote contest, which resulted in a five-room house (figure 1). The Find-the-Remote participants would use the kitchen and the living room, and the Home-Vacuum participants would use the hallway, the bedroom, the den, and the family room.

We designed three phases, the first of which required a one-time cleaning of all the rooms. We hoped phase one would be an easy navigation round where the entrants' robots would show competent navigation capabilities well established in past contests. The only wrinkle concerned bag capacity: if the robot encountered messes in the rooms (one-foot-diameter piles of white confetti), it had to return to a deposit area after it met two such messes. Points were awarded for cleaning the messes (or just moving over them) and making deposits. As in the past, points were deducted for colliding and for engineering the environment to accommodate the robot.

We planned for the defining part of the contest to be phases two and three. These phases were designed to highlight the intelligent aspects of a home-vacuuming task. Phase two, called Tidy Up, was intended to tease out efficiency issues. The robot was to station itself at the disposal area, face down the hall, and watch for humans entering and leaving the rooms. Whenever a human (one of the judges) entered a room, there was the possibility that a mess had been made; the robot had to check the room and, if necessary, clean it up. Again, after every two messes, the robot had to return to the disposal area before continuing. The objective was to obtain the highest score in a 15-minute run. We pushed for intelligence by

"The hardest part is still the sensing and acting."

— David Kortenkamp

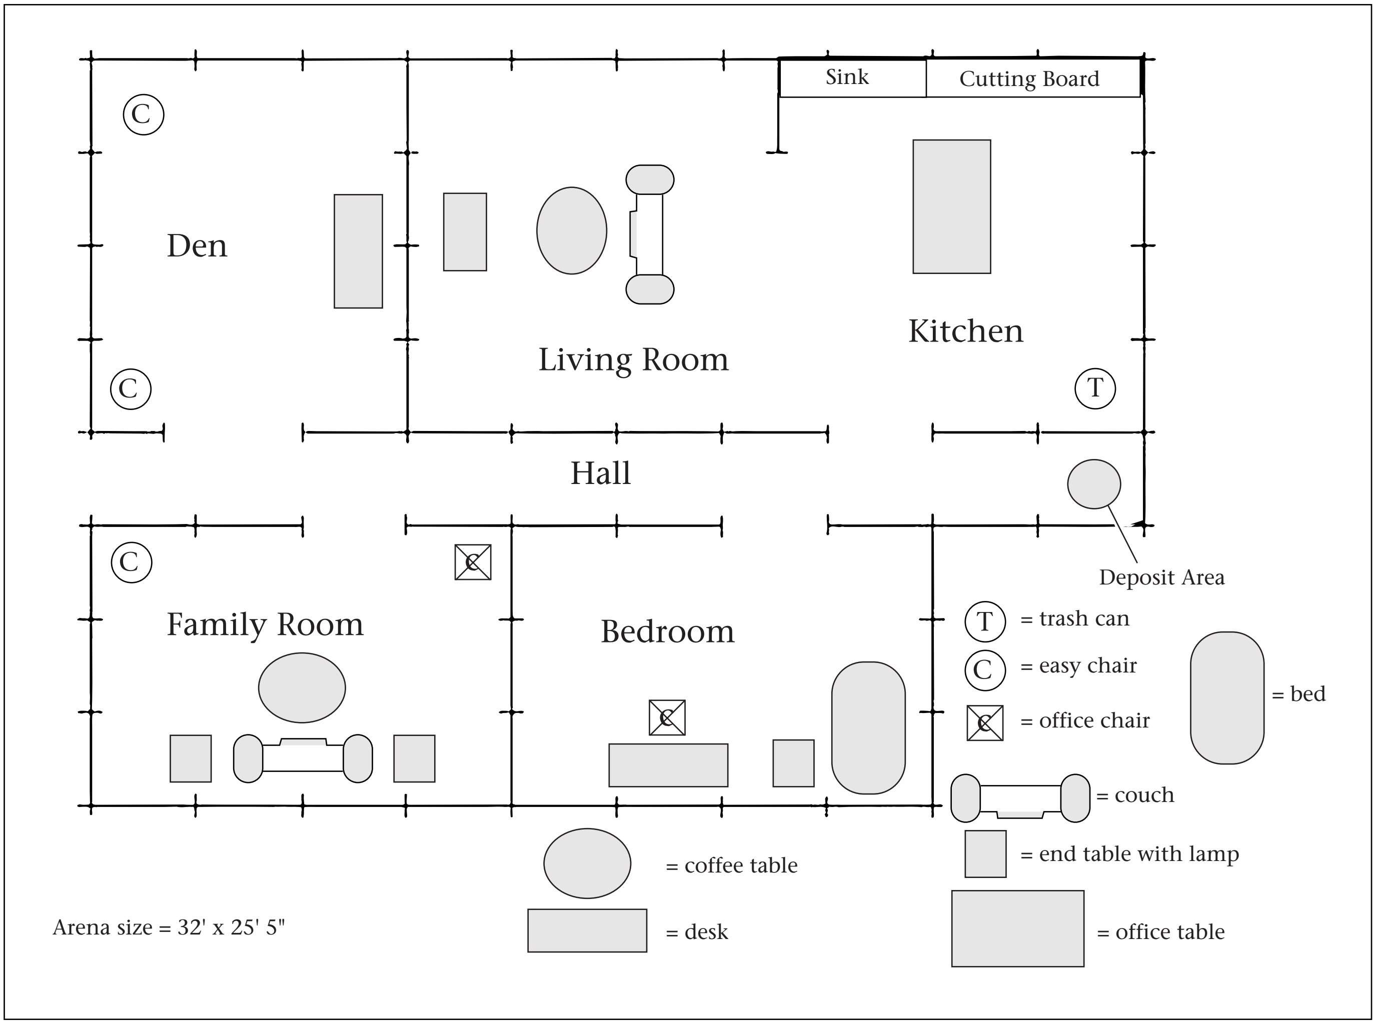

Figure 1. The Home-Vacuum Venue.

making the family room twice as profitable as the other rooms and giving the teams the expected frequency of human visits for each room. Thus, the robot had to concern itself with missing a family room opportunity while cleaning a lesser room or emptying the bag. In this phase and the next, discretionary points were awarded for innovative vacuuming mechanisms.

Phase three, Clean My Room, was designed to emphasize unobtrusive activity in the presence of humans. The robot was to start at the disposal station. A human (one of the judges) would come into the hallway and indicate to the robot that a room needed to be cleaned. The robot was to then move to the room, clean any messes there, and return to base to deposit the trash. While the robot was cleaning the room, one or more humans would enter. The robots were not to clean while a human was in the room but were to move to the nearest wall and wait until the human left. This phase allowed discretionary points for innovative human-robot interaction.

What Happened

It is important to note that these rules reflected two major assumptions made by the original organizers and were agreed to by Karen Myers who joined the team when Gat and Thrun were unable to attend the Fourteenth National Conference on Artificial Intelligence: First, robust navigation is a solved problem for the home layout shown in figure 1. Second, autonomous operation of robovacs is both feasible and essential.

As it turned out, these assumptions were far from the contest reality. Of the five teams that entered, none exhibited fully autonomous, safe

Figure 2. LOBOTOMOUS.

Figure 3. DIABLO.

navigation in the venue. Two teams, those from Swarthmore College and the Georgia Institute of Technology, were snake bitten with repair problems from the outset. Swarthmore finished the first phase but was too disheartened to continue. Georgia Tech recovered enough to make one run of a much simplified (see later discussion) second phase on the last day.

The rest of the teams had different kinds of problems: Dartmouth College's navigation approach depended on too few sensors to do the job (although the team assured us that if only that part worked, it would have had a dynamite planning algorithm to demonstrate). The University of New Mexico team ran its robot, LOBOTOMOUS, completely open loop (that is, without sensor feedback), depending on furniture that was exactly placed and then never disturbed. This team's setup required teleoperating the robot to each room and recording the coordinates of various way points. Then the robot moved to each coordinate and, running the vacuum, hoped there was trash to pick up at the way points. The robot from the University of Texas at El Paso, DIABLO, had a navigation algorithm that worked well initially but was overly dependent on dead reckoning, causing it to lose its way after running the length of the hallway.

By the end of the first day of competition, only LOBOTOMOUS—the teleoperated robot—had successfully completed the once-a-week cleanup. LOBOTOMOUS was fun to watch because it used a noisy suction motor in its head to suck up the confetti from the floor (figure 2). We extended the first phase for another half day, which allowed Swarthmore and Dartmouth to try again. With the exception of LOBOTOMOUS, all the robots needed to have a team member present to help them navigate without becoming hung up on walls or furniture. None of the robots cleaned up all the rooms in the time allotted, but each encountered enough messes in its best run to allow us to give win, place, and show awards. The New Mexico team did not intend to enter the other two phases of the contest, and as mentioned previously, Swarthmore, plagued with equipment problems, left the contest early. Thus by the third day, with only three teams entering and those still struggling to overcome nonrobust navigation capabilities, we knew we had to drastically revise the contest.

We decided for the second phase to have the robots check all the rooms for messes and, if there were any, to clean them up. To limit the amount of navigation required, we informed the teams that only one of two rooms would have messes. Also, in our scoring, we leaned toward giving points for getting any part of a mess and

It is important to note that these rules reflected two major assumptions. …

First, robust navigation is a solved problem for the home layout. …

Second, autonomous operation of robovacs is both feasible and essential.

Once-a-Week-Cleaning Contest	
First Place: LOBOTOMOUS, University of New Mexico	
Second Place: SK, Dartmouth College	
Third Place: Swarthmore College	
Certificate of Participation: DIABLO, University of Texas at El Paso	
Tidy-Up Contest	
First Place: DIABLO, University of Texas at El Paso	
Second Place: AMADEUS, Georgia Institute of Technology	
Certificate of Participation: SK, Dartmouth College	
Against All Odds Award	
Swarthmore College	
The Murphy Award	
Georgia Institute of Technology	

Table 1. Final Awards in the Home-Vacuum Event.

going into and out of rooms safely, avoiding deducting points for all but the most egregious violations. During this phase, as in the previous phase, all the teams had to restart their robots numerous times during their run, losing points in the process. The University of Texas at El Paso (figure 3) won this phase using an infrared sensor to detect white confetti on a dark floor.

The final awards given are shown in table 1. The last two awards reflected the continued honoring of the tradition to recognize teams encountering bad luck with their equipment. Marc Slack was the first recipient of such an award at the first competition in 1992 (Dean and Bonasso 1993).

Conclusions

With the exception of the Georgia Tech team, which had severe hardware problems until the last day, the participants were relative newcomers to the autonomous robot field. As such, they weren't yet experienced enough to put together robust, autonomous mobile robot systems in a short period of time. Robot engineering remains an art to a large degree: it still requires a lot of experience to make these systems operate reliably, notwithstanding the various conceptual advances that have been made in recent years. Despite the number of advanced practitioners of this art—as evidenced by the more sophisticated robots that have participated in previous years or participated successfully in other events at the 1997 competition—it takes a lot of time and effort to attain this level of expertise.

We recommend that future organizers make their assumptions about the expected state of robot engineering and the level of autonomy explicit in the contest announcements. This could possibly include recounting what the successes of past contests indicate about the expected level of capability of any robots entering the current year's events.

Finally, perhaps we should make room for new, young teams to compete in events of rudimentary autonomous navigation. They would thus be involved productively in the experience of the competition and not come away discouraged at their robot's poor showing in the more difficult contests.

References

Bonasso, P.; Miller, D., and Kuipers, B. 1993. Instantiating Real-World Agents, Working Notes, AAAI Fall Symposium Series, 22–24 October, Raleigh, North Carolina.

Dean, T., and Bonasso, R. P. 1993. The 1992 AAAI Robot Exhibition and Competition. *AI Magazine* 14(1): 34–48.

Nourbakhsh, I; Powers, R.; and Birchfield, S. 1995. DERVISH: An Office-Navigating Robot. *AI Magazine* 16(2): 53–60.

Yang, H. S.; Chung, J.; Ryu, B. S.; and Lee, J. 1996. CAIR-2: Intelligent Mobot Robot for Guidance and Delivery. *AI Magazine* 17(1): 47–54.

Pete Bonasso is a senior scientist for AI and robotics at TRACLabs at Johnson Space Center. He supports the Automation, Robotics, and Simulation Division investigations of intelligent robot control using layered architectures. He is the codeveloper of the 3T robot control architecture and has applied this architecture to the control of robotic and life-support machines. He received his B.S. in engineering from the U.S. Military Academy at West Point in 1968 and master's degrees in operation research and computer utilization from Stanford University in 1974. His e-mail address is bonasso@mickey.jsc.nasa.gov.

Karen Myers has been a member of the research staff at SRI International since completing her Ph.D. in computer science at Stanford University in 1991. Myers participated in the first two AAAI robot competitions, where she helped to design and build the SAPHIRA architecture for SRI's robot FLAKEY. More recently, her research has focused on the integration of reactive planning and execution systems, multiagent problem-solving architectures, and advisable problem-solving systems. Her e-mail address is myers@ai.sri.com.