

Herding Robots: Enabling Fleets of Multiagent Systems to Collaborate in Unprecedented Ways

Larry Hardesty, MIT



Writing a program to control a single autonomous robot navigating an uncertain environment with an erratic communication link is hard enough; write one for multiple robots that may or may not have to work in tandem, depending on the task, is even harder.

As a consequence, engineers designing control programs for “multiagent systems” — whether teams of robots or networks of devices with different functions — have generally restricted themselves to special cases, where reliable information about the environment can be assumed or a relatively simple collaborative task can be clearly specified in advance.

In May 2014, at the International Conference on Autonomous Agents and Multiagent Systems, researchers from MIT’s Computer Science and Artificial Intelligence Laboratory (CSAIL) will present a new system that stitches existing control programs together to allow multiagent systems to collaborate in much more complex ways. The system factors in uncertainty — the odds, for instance, that a communication link will drop, or that a particular algorithm will inadvertently steer a robot into a dead end — and automatically plans around it.

For small collaborative tasks, the system can guarantee that its combination of programs is optimal — that it will yield the best possible results, given the uncertainty of the environment and the limitations of the programs themselves.

Working together with Jon How, the Richard Cockburn Maclaurin Professor of Aeronautics and Astronautics, and his student Chris Maynor, the researchers are currently testing their system in a simulation of a warehousing application, where teams of robots would be required to retrieve arbitrary objects from indeterminate locations, collaborating as needed to transport heavy loads. The simulations involve small groups of iRobot Create, programmable robots that have the same chassis as the Roomba vacuum cleaner.

Reasonable doubt

“In [multiagent] systems, in general, in the real world, it’s very hard for them to communicate effectively,” says Christopher Amato, a postdoc in CSAIL and first author on the new paper. “If you have a camera, it’s impossible for the camera to be constantly streaming all of its information to all the other cameras. Similarly, robots are on networks that are imperfect, so it takes some amount of time to get messages to other robots, and maybe they can’t communicate in certain situations around obstacles.”

An agent may not even have perfect information about its own location, Amato says — which aisle of the warehouse it’s actually in, for instance. Moreover, “When you try to make a decision, there’s some uncertainty about how that’s going to unfold,” he says. “Maybe you try to move in a certain direction, and there’s wind or wheel slippage, or there’s uncertainty across networks due to packet loss. So in these real-world domains with all this communication noise and uncertainty about what’s happening, it’s hard to make decisions.”

The new MIT system, which Amato developed with co-authors Leslie Kaelbling, the Panasonic Professor of Computer Science and Engineering, and George Konidaris, a fellow postdoc, takes three inputs. One is a set of low-level control algorithms — which the MIT researchers refer to as “macro-actions” — which may govern agents’ behaviors collectively or individually. The second is a set of statistics about those programs’ execution in a particular environment. And the third is a scheme for valuing different outcomes: Accomplishing a task accrues a high positive valuation, but consuming energy accrues a negative valuation.

School of hard knocks

Amato envisions that the statistics could be gathered automatically, by simply letting a multiagent system run for a while — whether in the real world or in simulations. In the warehousing application, for instance, the robots would be left to execute various macro-actions, and the system would collect data on results. Robots trying to move from point A to point B within the warehouse might end up down a blind alley some percentage of the time, and their communication bandwidth might drop some other percentage of the time; those percentages might vary for robots moving from point B to point C.

The MIT system takes these inputs and then decides how best to combine macro-actions to maximize the system’s value function. It might use all the macro-actions; it might use only a tiny subset. And it might use them in ways that a human designer wouldn’t have thought of.

Suppose, for instance, that each robot has a small bank of colored lights that it can use to communicate with its counterparts if their wireless links are down. “What typically happens is, the programmer decides that red light means go to this room and help somebody, green light means go to that room and help somebody,” Amato says. “In our case, we can just say that there are three lights, and the algorithm spits out whether or not to use them and what each color means.”

The MIT researchers' work frames the problem of multiagent control as something called a partially observable Markov decision process, or POMDP. "POMDPs, and especially Dec-POMDPs, which are the decentralized version, are basically intractable for real multirobot problems because they're so complex and computationally expensive to solve that they just explode when you increase the number of robots," says Nora Ayanian, an assistant professor of computer science at the University of Southern California who specializes in multirobot systems. "So they're not really very popular in the multirobot world."

"Normally, when you're using these Dec-POMDPs, you work at a very low level of granularity," she explains. "The interesting thing about this paper is that they take these very complex tools and kind of decrease the resolution."

"This will definitely get these POMDPs on the radar of multirobot-systems people," Ayanian adds. "It's something that really makes it way more capable to be applied to complex problems."

Source URL (retrieved on 05/29/2016 - 11:20am):

<http://www.scientificcomputing.com/news/2014/02/herding-robots-enabling-fleets-multiagent-systems-collaborate-unprecedented-ways>