S&TR June 2018

BUILDING A NETWORK OF COLLABORATIVE AUTONOMOUS MACHINES

A FIREFIGHTER enters a burning office building followed by several drones, searching for people in need of rescue. The drones scatter in different directions, moving down corridors and systematically scanning rooms. Each drone's onboard sensors measure smoke, temperature, sound, and motion to determine which places are too hot or smoky for the firefighter to approach and which spots show signs of a living human being.

No single drone acts as a central control device. Instead, the drones in this hypothetical squadron communicate with each other and use their powerful algorithms and ample computing power to pool and process their data and decide, as a group, whether a survivor has been detected and, if not, where to continue the search. The drones also "know" enough to avoid conditions that would disable them. The firefighter receives reports from her drone squadron and can command the devices, but she has no need to control each drone's every action—the devices have sufficient algorithmic might to operate on their own in real time, deciding collectively when their human partner needs to know about certain information. Such behavior by a group of networked devices is called collaborative autonomy, and Livermore researchers are working to make it a reality.

A human–machine firefighting team is just one example of collaborative autonomy's potential. "Any task that is dangerous, repetitive, or dirty would benefit from this capability," says Livermore's Reginald Beer, the leader of the research effort. The software and hardware for collaborative autonomy could be applied not only to drones but also to mobile surface or underwater robotic vehicles, self-driving cars, and even appliances in the Internet of things—in other words, machine agents of all kinds. Devices so equipped could perform search-and-rescue operations for missing persons in a wilderness area or after a disaster, look for a hidden nuclear device by detecting its radioactive signature, or measure the movement of a hazardous chemical release and assist in containing and cleaning up the substance, to name only a few possibilities.

Extending the Reach of Radar

Livermore's collaborative autonomy research is a natural outgrowth of a project to develop vehicle-mounted ground-penetrating radar (GPR) for detecting buried objects in real time. Beer's team developed an array of radar transmitters and receivers for the effort, which has applications such as detecting buried explosive devices. As the GPR-equipped vehicle moves forward, the array produces an image of what is located underground in a vertical plane. The system processes data much the same way that medical computerized tomography generates a planar image of the body's interior. Beer realized that the system could be given algorithms that not merely generate an image from the data but also interpret the image and relay that interpretation to the human operator. "From there, it was a natural step to imagining a team of sensors that can test for more types of targets and could even be airborne instead of mounted on a ground vehicle," he says. For GPR, he envisions a group of drones that fly ahead of a vehicle, autonomously select the areas on which to concentrate their scans, collaboratively interpret their data, and alert the vehicle's operator to the presence of a buried explosive device, which the operator then disables.

In short, the ultimate goal of this collaborative autonomy research is to develop software and hardware that allows a group of machine agents to collaboratively gather sensor data (observe), identify and interpret what the sensors detect (orient), make decisions about how to respond (decide), and implement those decisions (act). Furthermore, such a system must be able to decide

Livermore researchers are using simulations to improve communications between networked machine agents, such as drones. In these screen captures from video representing such a scenario, (below) a fleet of drones supports a vehicle using ground-penetrating radar to find buried explosive devices. The drones fly ahead of the vehicle, using their sensors to detect any buried metallic objects. (top) When a buried explosive is found, the drones notify the vehicle's human operator, who can then avoid the explosive. (Renderings by Adam Connell.)

and act quickly enough in real time to be useful to humans. In achieving the considerable technological advances required to realize this vision, the Livermore team began with the first step—developing and testing the algorithms for decentralized processing and communications in autonomous sensor networks.

Autonomy through Decentralized Computing

A project funded by the Laboratory Directed Research and Development (LDRD) Program and led by computation engineer Ryan Goldhahn focuses on creating decentralized processing and communications algorithms, while partners at the University of Texas at Austin are developing hardware for algorithm testing. Goldhahn says, "We wanted to move away from the model of nodes sending data to a central command center. That model is not scalable because of the vast amounts of data that must be centrally processed. In addition, the central node—such as an autonomous vehicle, a computing cluster at a command center, or any other type of lone machine agent—is a vulnerability because if the central node fails so does the entire network." With recent technological advances allowing engineers to equip each individual node with considerable processing power, the project team is working on a real-time capability in which each member of a large network possesses a high degree of autonomy.

Goldhahn explains, "Each node must be able to decide which data matter and communicate those data to the rest of the network. We are using what are called gossip and consensus algorithms, in which one node sends a measurement deemed relevant to its neighboring nodes. As they exchange more and more data, the nodes agree on what they are sensing." The algorithms also eliminate the problem of "Byzantine data"—poor measurements or deliberate misinformation—by first agreeing what data are significant and then deciding where additional measurements are needed to be more certain of their interpretation. One capability required for such a system is determining which nodes in the group will make the subsequent measurements and which nodes will relay those measurements to the rest of the group. To this end, each node must autonomously decide where to position itself relative to the other nodes and whether to investigate a potentially important event. The entire group must agree on who does what without any one node being in overall command. A significant challenge is to achieve a decision making capability that can be scaled up to hundreds of nodes without overwhelming the network with computational complexity or communications volume.

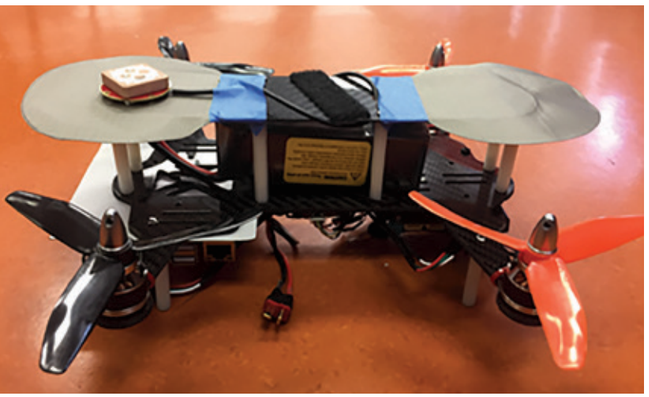

Lawrence Livermore collaborators at the University of Texas at Austin have developed this drone prototype along with a network simulation platform. Together, the hardware and software test the ability of individual nodes in a mobile network to collect and share sensor data.

Simulating Network Communication

An important aspect of this effort is to achieve maximum efficiency in communication among the nodes. To this end, the researchers are studying simulations of sensor networks run on Livermore's high-performance computing resources. Computer scientist Peter Barnes leads Livermore's network simulations team, which has developed the capability to simulate realistic networks—on the order of 10,000 to 500 million computers—using the open-source software program ns-3.

Barnes and Anton Yen are developing the ability to computationally model communications and other behavior among nodes using a simulated set of nodes. "Communications technology faces constraints on bandwidth and range, such as the time it takes for information to travel," says Barnes. "In addition, a specific node may need to communicate only to one other node but not all the others, or may send and receive messages to and from multiple nodes nearby." The simulations are therefore examining internodal communications in great detail, from parameters such as range, latency, and bit rate to the impact that the parameters have on the collaborative autonomy algorithms. The ultimate goal is to optimize the process of gathering, sharing, and interpreting the data and calculating the next steps.

A Belief Network

Gerald Friedland, a computation scientist at Livermore and an adjunct professor at the University of California at Berkeley, is leading an LDRD project to apply Bayesian belief propagation to collaborative autonomy. Specifically, Friedland and his collaborators Kannan Ramchandran and Maya Gokhale are developing software and hardware for a network that uses Bayesian statistics to arrive at a belief, that is, to compute the probability of a particular outcome. For instance, as each node in the network collects data from its sensors, the nodes "vote" to arrive at a consensus interpretation of the data.

Using YFCC100M—an open-access database of more than 100 million photos and videos for artificial intelligence research—the team is developing algorithms that will be distributed over a network so that the nodes can collectively determine where a specific video was filmed. To solve this challenging problem, nodes will use visual, audio, and even text data to guess at a location and then vote repeatedly on the location. "Expert" nodes will emerge from the group by virtue of being closer to the correct answer than others, leading eventually to a majority vote that represents a collective consensus on the video's location. "In some ways, the nodes behave like humans," states Friedland. "One node has a certain belief, and another may have a different belief, and the two nodes can agree or disagree." The power of the belief-network approach is its ability to be generalized to any problem, exploiting and combining available information to arrive at the best possible answer. As part of this LDRD project, Friedland's team will install the software in specially developed hardware—field-programmable gate arrays designed to be lightweight, fast, and low power. These collaborative autonomy efforts are joining forces to take the first steps toward what Beer calls a "networked machine intelligence that is capable of autonomy of action."

—Allan Chen

Key Words: Bayesian network, belief network, byzantine data, collaborative autonomy, consensus algorithm, decentralized processing, field-programmable gate array, gossip algorithm, ground-penetrating radar (GPR), Laboratory Directed Research and Development (LDRD) Program, network simulation, node, ns-3, sensor network, YFCC100M.

For further information, contact Reginald Beer, (925) 424-2232 (beer2@llnl.gov).

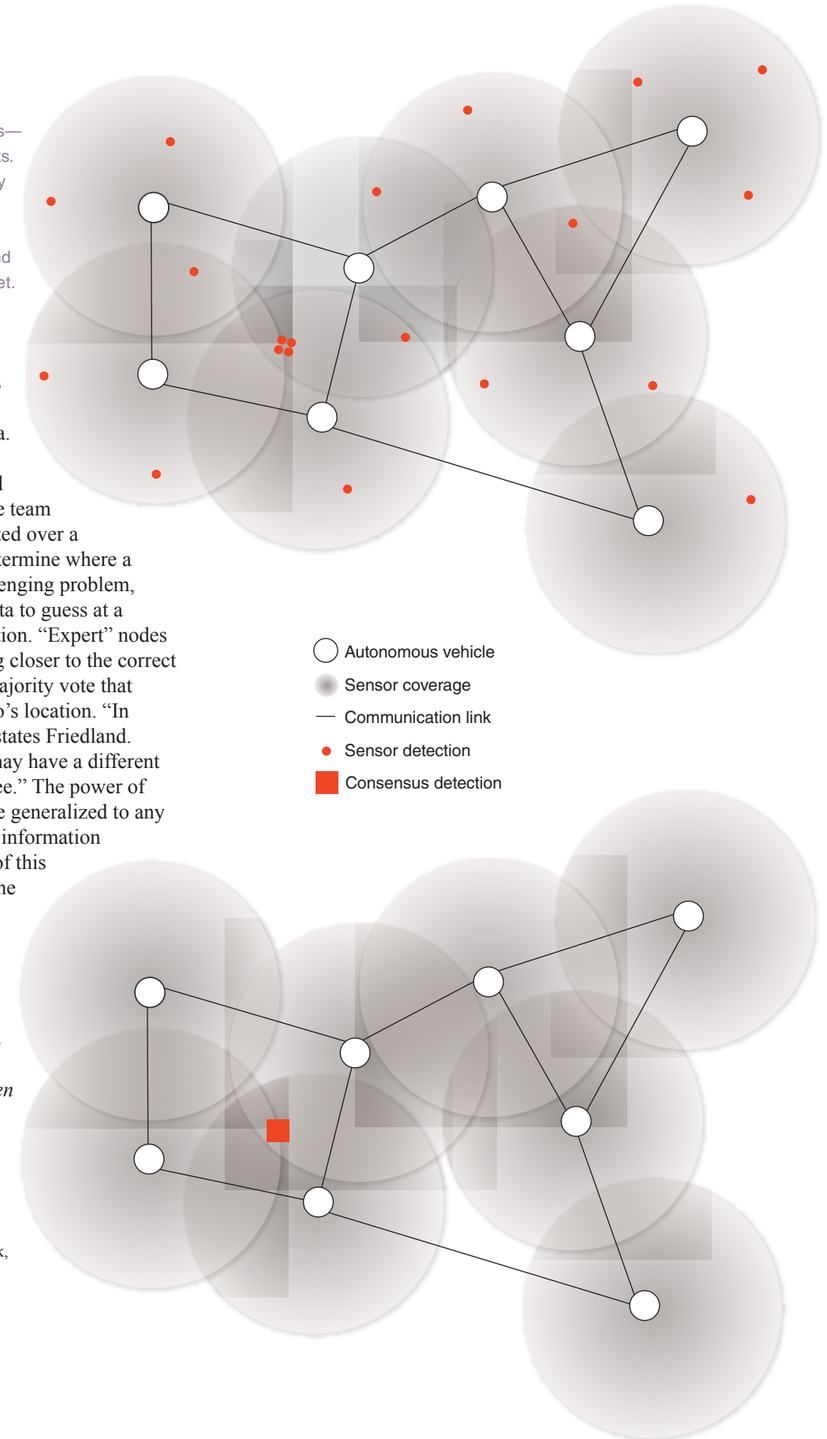

In this network simulation, a fleet of machine agents—drones or other devices—is deployed to detect targets. Individual agents communicate with others nearby but not necessarily with every agent in the fleet. Gossip and consensus algorithms allow the agents to (top) share sensor data about possible targets and (bottom) arrive at a consensus about a likely target.