Sensing for Danger

Correlated sensor networks can help fight against nuclear terrorism and other threats.

PICTURE this scenario: A terrorist carefully negotiates city streets, moving ever closer to his target, an air force base on the outskirts of town. In the rear of his van, a homemade bomb—containing plutonium and high explosives—waits for the signal to explode. As one of the "good guys," you've received information that the attack is imminent, but your sources don't know its timing, the direction from which the vehicle will come, or what route it will take. What can you do to detect, identify, and track the van and its contents so that you can prevent the attack? At Lawrence Livermore, researchers in the Nonproliferation, Arms Control, and International Security (NAI) Directorate have been exploring responses to this threat and others like it.

The researchers are focusing on systems for detecting and tracking threats. The systems go by many names—correlated sensor networks, wide-area tracking systems, sensor or network fabrics—but the concept behind them is the same. Take a number of wireless sensors (for instance, seismic, magnetic, pressure, acoustic, nuclear, or particle-counting), tie them together with a communications network, develop a scheme for fusing the data (that is, converting the data into forms easily interpreted by users), and make the system easy to deploy.

Such correlated sensor networks can help detect a nuclear terrorist attack, track the movement and characteristics of a wildfire, assist military operations in taking out a target, determine earthquake damage to large structures such as bridges, and even protect the president.

The Power of Networking

The power of correlated sensor systems arises from their networked nature. "You could ask, 'Why not just use a bunch of stand-alone sensors?'" says Rob Hills, acting leader of the Tactical Systems Section in NAI. "Part of the problem is that many sensors, particularly those that detect nuclear signals such as gamma rays and neutrons, have a hard time differentiating between a 'hit' and normal variations in the background radiation. And to compound the challenge, the farther one moves away from a nuclear source, the weaker the signals become." Sid Niemeyer from the NAI directorate office agrees, saying that "Weapons-usable nuclear materials are difficult to detect under circumstances of unconventional nuclear warfare, as nuclear terrorism is sometimes called. As the distance between a detector and source increases, the radiation signature quickly fades into the background caused by other artificial and natural sources."

One solution is to network the sensors, that is, have them share the information they gather. "Networked sensors allow the user to 'see' more by creating a more complete picture of the situation, something that stand-alone sensors cannot do," says Niemeyer. For this article's opening scenario, for instance, a correlated sensor network

The Challenge—Smaller, Smarter, More Energy-Efficient Sensors

Most of today's wireless sensors are big and heavy. They have large power requirements and limited intelligence. Thus, large networks of such sensors are impractical. In the Nonproliferation, Arms Control, and International Security (NAI) Directorate, researchers are working to create sensors that use less energy, are more intelligent, and scale better to large networks.

The energy issue, notes engineer Rob Hills, is a big concern for sensors that are networked. "We have a saying that power is everything," he explains. "Power requirements make a network feasible—or not." For instance, the Joint Biological Remote Early Warning System (JBREWS) prototype network used 132 commercial sensors, each requiring two batteries (one in the sensor, one in the charger) to operate continuously, for a grand total of nine tons of batteries. To address this problem, the Laboratory developed a communication system that requires an average power of only 1 watt. "And we're pushing the power requirement down from there," says Hills.

In a back-to-basics project, Laboratory engineer Dave Harris is researching the underlying physics that is key to creating microsensors for seismic networks. "I believe that Dave's work—along with our data-fusion techniques—will allow us to create cheap and small sensors, which can be delivered from a remote platform such as unmanned aerial vehicles," says Hills. Harris has been been working with engineer Bruce Henderer, who has developed a prototype sensor about 3 centimeters thick and 6 centimeters square—small enough to hold in the palm of your hand—and containing a low-power communications device that allows the sensors to network and to configure themselves. "In other words," says Hills, "once laid down, the sensors would talk to each other and, by determining their neighbors, build a network and paths back to control." The data processing would take place out in the network, with the network sensors themselves being capable of pattern recognition, information fusion, and decision making.

The prototype sensor being developed at Livermore.

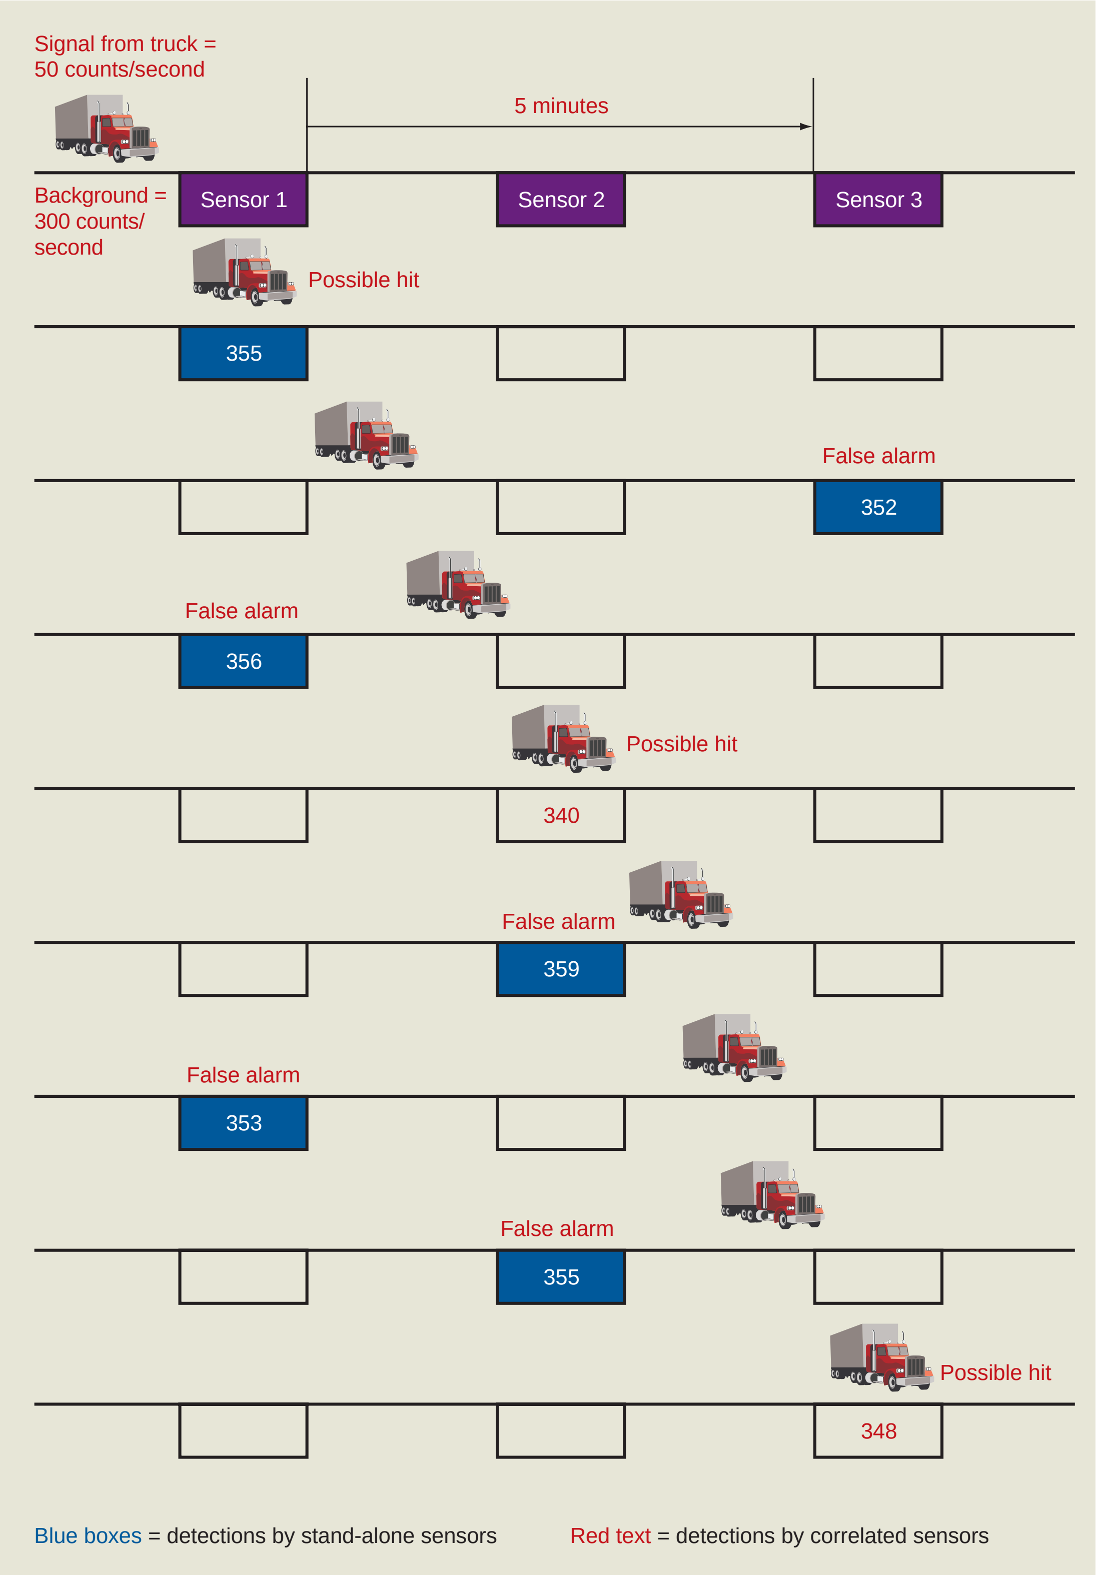
Signal from truck = 50 counts/second
5 minutes
Background = 300 counts/ second
Sensor 1
Sensor 2
Sensor 3
Possible hit
355
False alarm
352
False alarm
356
Possible hit
340
False alarm
359
False alarm
353
False alarm
355
Possible hit
348
Blue boxes = detections by stand-alone sensors
Red text = detections by correlated sensors

This illustration shows the different results produced by stand-alone sensors versus a correlated sensor network. Here, the sensors are set to register signals of 350 or more counts per second from a truck carrying a signal-emitting device. The stand-alone sensor system simply detects six instances of over-350 signal counts (blue boxes). The networked system, having access to more information, correlates the information to discount all but the first detection as false alarms and to register two others that are under the 350-count threshold as likely “hits,” which are then correlated to the first hit.

could do double duty. First, it could provide a way to discard signals that are false alarms. Second, it could pick up on signals that might be real alarms but would have been ignored by stand-alone sensors because the signals were under a preset threshold of sensitivity.

The figure on p. 13 shows how correlated network and stand-alone systems differ. In the figure, the truck carries a device that emits signals averaging 50 counts per second; signals in the natural background are on average 300 counts per second. Three stand-alone sensors are set to register detections, or hits, of 350 or more counts, which would reduce false alarms caused by background variations.

The truck passes the first sensor, which detects 355 counts—a possible hit. It goes on its route and the stand-alone sensors detect five other instances of over 350-count signals. The detections provide no information to the person at the central command post as to whether they are real alarms or not.

However, if the sensors were networked and able to communicate with each other, a different picture would emerge. "For one thing," notes Hills, "you can include other information in the system, such as the approximate travel time of the vehicle. A reasonable assumption would be that the vehicle is traveling at the speed limit, because its driver probably would not want to attract attention to himself."

In the correlated sensor case, the 355-count signal at sensor one is noted as a possible hit. This information is shared with sensor two, and then the system clock starts to track travel time. The travel time between sensors one and three is assumed to be about 5 minutes. Sensor two is on the alert for signals above background that appear within a certain window of time centered on a predetermined time mark, say, 2.5 minutes. The closer a signal is detected to that 2.5-minute mark, the more weight is given to the probability that the signal is from a real source, rather than some random hiccup from background.

The ensuing signals at sensors one and three are discounted as false alarms because they are uncorrelated, that is, they show no relation to previously recorded data. If the truck is proceeding forward, it would not register at sensor one, which it just passed, and the signal at sensor three comes much too soon. However, the 340 counts detected at sensor two, even though a trifle low, is viewed as a possible hit because it falls within the allotted window of time and is considerably higher than background. This information is passed along to sensor three.

Three signals follow and are discounted as false alarms because of their location and timing. However, the

Bayesian Statistics at Work

While developing the computer algorithms to perform distributed decision making for a sensor network, a team of researchers, including physicist Chris Cunningham, came up with an approach based on Bayesian algorithms. As Cunningham explains it, the Bayesian approach has a couple of pluses. First, it is energy-efficient because communication only occurs when there is a sufficient probability that a target has been detected. Second, each sensor independently extracts features from its raw sensor signals, compares these features with the targets, calculates the likelihood of detection, fuses the likelihoods received from neighboring nodes, and communicates only the new likelihoods to its neighboring nodes. This statistical data fusion can allow each sensor platform to make decisions based on the total information in the network, while reducing the volume of communications among sensors.

The method is based upon the work of an English mathematician, the Reverend Thomas Bayes. Bayes developed a mathematical formula that allows scientists to combine new data with prior conditions. In a sense, it addresses the question, "Given that an event has occurred that may have been the result of any of two or more causes, what is the probability that the event was the result of a particular cause?" The answer lies not in an absolute yes or no, but in the set of probabilities that the various causes are at play. Bayesian methods allow scientists to combine prior information about a population parameter with information contained in a sample to guide a statistical inference process. A prior probability distribution for a parameter of interest is specified first. Sample information is then obtained and combined through an application of Bayes's theorem to confirm the prior assumptions. Bayesian methods are used extensively in statistical decision theory.

Livermore's Wide-Area Tracking System (WATS) is one example of a correlated sensor network that uses algorithms based on Bayesian constructs. In WATS, each sensor computes and exchanges information with its near neighbors in the form of Bayesian probabilities for possible sources. Algorithms reduce the sensor data to probability estimates and then fuse the estimates among the multiple sensors.

348-count signal at sensor three is recorded and its probability of being real is calculated and correlated with the preceding hits.

"What's happening here is that we're actually correlating signatures in different domains," explains Hills. "For this example, we're correlating data from both temporal and spatial domains: correlating whether the appropriate sensor gets the hit—which is the spatial domain—and whether that hit may be due to the source based on the time of travel between sensors—which is the temporal domain. We then perform some statistical calculations to determine how probable it is that the hit is real, based on the number of counts detected and when—within the allowable window of time—the counts are detected."

Performing these kinds of calculations for three networked sensors is one thing, but widen a network to include 100 sensors and it becomes extraordinarily challenging. The computer algorithms needed to track and follow more than one likely pattern and calculate all of the probabilities are extremely complex (see box on p. 14) and are only now possible with the increases in computing power.

Military and Other Applications

Livermore researchers have been working on many applications of correlated sensor networks. For instance, the Laboratory has developed a prototype correlated sensor network for detecting and tracking a ground-delivered nuclear material. The Wide-Area Tracking System (WATS) is a network of gamma and neutron detectors and communications links, with information continuously evaluated by Laboratory-developed data-fusion algorithms. The sensors can be permanently deployed at chosen locations or mounted in vans for deployment on demand to protect specific areas for specific situations or events.

The individual sensors share their data with neighboring sensors, process the data, integrate and combine them with other available information (for example, data gathered previously; observed radiation signatures, spectra, and backgrounds; road maps), and finally determine the probability that the signal comes from a real source—all while the system is in the field. In this way, a WATS sensor network can drastically reduce false alarms and detect the entry of a nuclear device or radioactive material into the protected area and track its movement.

The analysis could be performed by a centralized computer at, say, command headquarters, but researchers have found that communications

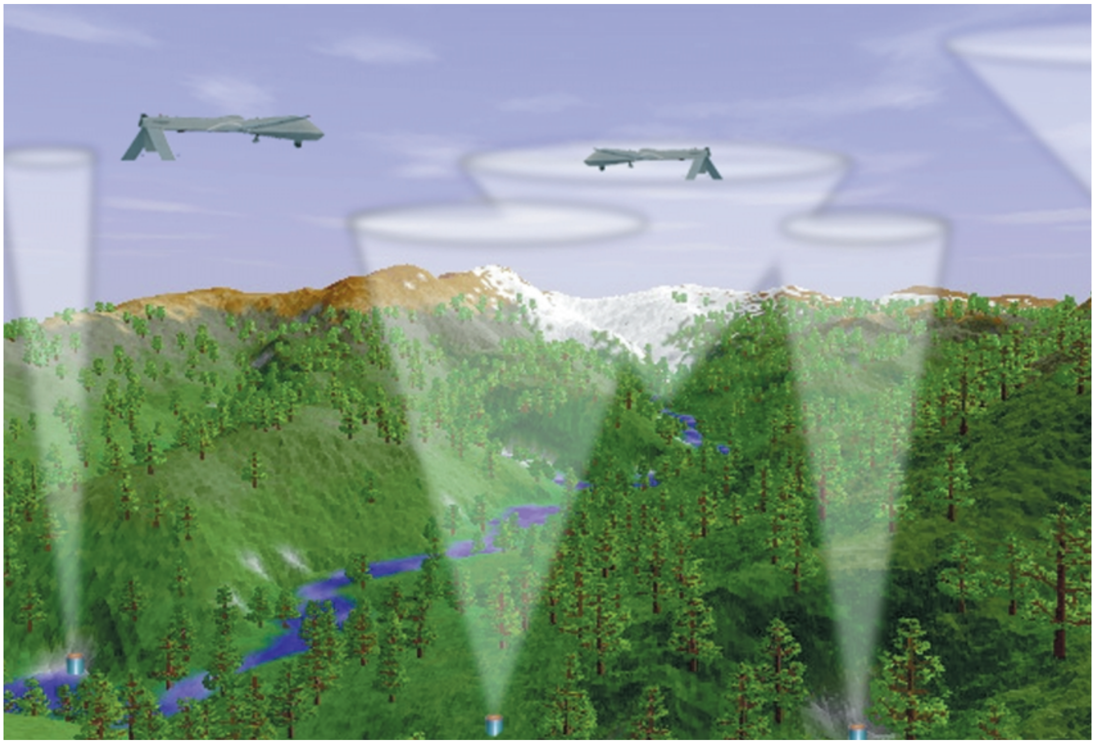

Researchers at Livermore are exploring the possibility of using unmanned aerial vehicles (UAVs) to place, operate, and maintain sensor networks in rugged terrain. In the figure, the sensors in the network are shown sending their information.

limitations—latencies, available bandwidth, and so on—can be a significant bottleneck for these types of networks. When data are processed in the field, it is necessary to send only bits of information between neighbors, with the final result going to the human user. This type of operation makes the network much more scalable.

Another example of correlated sensor network development involves a recently concluded project called Joint Biological Remote Early Warning System (JBREWS). For JBREWS, the Laboratory was responsible for developing the command, control, communications, computers, and intelligence systems for a network of biodetectors that could provide U.S. field troops with early warning of a biological attack. Although the project is not continuing, it has allowed the Laboratory to make important progress in developing data-fusion solutions that could be applied to any type of correlated sensor network. The communications paradigms that were developed in JBREWS let Laboratory researchers take a big step toward solving one part of the data-fusion problem—that is, how to quickly and automatically establish a communications fabric for data fusion to work within.

In this communications scheme, the array of sensors forms an automatically reconfiguring, or self-healing, network, as follows. Once the sensors are in place, they communicate with each other via radio frequencies so each sensor can map where its neighboring sensors are. The sensors then radio-test each other and develop an efficient communications path back to the central command post. If, for example, one sensor can't communicate directly with the command post on the other side of a hill, it passes its data to its neighbors, to be relayed with the neighbors' data to other units, and so on, until the information reaches its destination. If a unit is knocked out by a malfunction or hostile action, its communication relay functions are picked up by surrounding units and a secondary path is formed. In short, the system quickly recognizes and adjusts to the absence of any sensor units. A big plus for this type of network and others like it, Hills notes, is that there are no single-point failures.

Another military application would connect these sensor networks with other systems, such as the Laboratory-developed Counterproliferation Analysis and Planning System (CAPS). CAPS can model the various processes (chemical, biological, metallurgical) used by proliferators to build weapons of mass destruction and their delivery systems. CAPS helps users identify critical processing steps or production facilities that, if disabled or destroyed, would prevent that country from producing weapons of mass destruction. "Now imagine adding correlated sensor networks to the mix," says Hills. "Sensors on the ground and in the air could track processes in real time. A user could click on the Web-based CAPS page and find out what's going on right then at such-and-such a facility."

Yet another application for such networks is in tactical engagement systems. With sensor networks as part of these systems, a soldier would never be alone in the field. The sensor network could supply information not just to people in the field, but to those who are out of harm's way as well. They would all be tied together in a collaborative environment. With such a system, the electronic network would be displayed in a chest-top system so that a soldier could "see" the environment and watch his back—all from one small device.

Correlated sensor networks could also be used in nonmilitary applications to provide temporary communication infrastructures after a destructive earthquake or to provide information during large firestorms. For example, there are microclimates within a large fire. A correlated sensor network could track temperatures, humidity, and wind in three dimensions, providing valuable information to firefighters.

David McCallen, director for Livermore's Engineering Center for Complex Distributed Systems, notes that current research to develop self-healing, self-configuring networks of seismic sensors would be useful in studying how large structures respond in earthquakes. "Once these networks are developed, it's a small step to apply them to large structures, such as bridges, to gather data on how these structures vibrate and respond under various circumstances," he explains. "When you consider that to densely instrument a structure like the Golden Gate Bridge takes hundreds of sensors, having a system that's wireless and self-configuring is very attractive." He adds that the California Department of Transportation is also interested in using such networks to monitor steep hillsides for possible landslides.

Putting Sensors in Their Place

One of the challenges to using these networks is getting them in place, in real terrain. "In a battlefield scenario, for instance, or during a wildfire, you can't have people tromping in to set down sensors," says Hills. One answer is to use unmanned aerial vehicles (UAVs), such as the U.S. Air Force's Predator or even smaller, 2-meter-wingspan UAVs. In one project, researchers are evaluating the use of UAVs to rapidly place, operate, and maintain sensor networks in rugged terrain. Such vehicles could drop the sensors in predetermined locations and then act as airborne routers. Once in place, the sensors would form a network, communicate with each other,

and send information skyward to be collected and transmitted by the planes.

Using Laboratory-designed software, researchers could create self-configuring and self-healing networks made up of small, low-power sensors. If the sensors are cheap enough, the result is a ready-to-use network—a wireless “network on demand.” In this kind of setup, the UAVs become part of the system, sharing information about locations of all the sensors and other UAVs, sensor data requirements, connectivity maps, and UAV-sensor assignments; leveling the workload; and backing each other up in case one or another UAV is put out of commission. “This is just one of the directions in which we’re moving to position ourselves for the future,” says Hills.

Looking toward the Future

The idea of correlated sensor networks is not Livermore’s alone. Other organizations and commercial companies are exploring applications and, like the Laboratory, pushing on what’s possible in the laboratory to get to what’s feasible in the field. “The key,” says Hills, “is to find ways of gathering all those data together and turning them into usable, real-time information to let the user make decisions. Here at the Laboratory, we’ve got the key in hand and are turning it in the lock. It’s only a matter of time before the door opens.”

—Ann Parker

Key Words: Bayesian statistics, correlated sensor networks, Counterproliferation Analysis and Planning System (CAPS), gamma detector, Joint Biological Remote Early Warning System (JBREWS), neutron detector, nuclear terrorism, seismic detector, sensor or network fabrics, tactical engagement systems, unmanned aerial vehicle (UAV), Wide-Area Tracking System (WATS).

For further information contact Rob Hills (925) 423-7344 (hills1@llnl.gov).

About the Scientist

ROB HILLS is the acting associate division leader for the Tactical Systems Section in the Nonproliferation, Arms Control, and International Security Directorate. He leads a variety of projects that include research and development for sensors and sensor networks, military systems analysis, and computer-based battlefield conflict simulation models. Hills received a B.S. in electrical engineering from the Michigan Technological University in 1983. He joined Livermore in 1988 to perform research that involved automating the transfer of existing digital designs to new implementation technologies. Thereafter, he participated in several projects to develop sensor systems and image-processing technologies for astronomical telescopes. For example, he was a member of the team that developed the camera system to detect dark matter; the system won an R&D 100 Award for being one of the most technologically significant new products in 1993. Hills has led research and development efforts for microtechnology tools, such as a polymerase chain reaction system, used both in medical and national security applications. And he has engineered optical interconnects for parallel computer systems as well as overall architectures for self-configuring and self-healing communications networks.