Lab Notes

NEWS FROM AROUND LINCOLN LABORATORY

CHEM-BIO SENSING

Threat detection

Cheap sensors plus mesh networking could yield an effective alarm system for biological and chemical attacks

We haven't heard much about biological warfare since envelopes containing anthrax spores were mailed to Congress and some news media in the fall of 2001. But it's still a concern for the military and the Department of Homeland Security. So researchers at Lincoln Laboratory are pursuing the use of commercially available mesh networking technology to make populations safer from airborne pathogens.

Sensor networks have existed for several years. Companies like Crossbow, Ember, and Dust Networks manufacture radio platforms that autonomously communicate with their neighbors, forming mesh networks that relay such data as temperature, moisture, and other environmental characteristics back to a central point. These systems haven't been used for biological or chemical defense applications, though, because sensors that can detect specific biological or chemical threats are not designed for low-power, low-cost deployments, says Adam Norige, a biomedical engineer in the Laboratory's Biodefense Systems group. State-of-the-art detectors incorporate sophisticated technologies that are usually geared toward improved performance rather than reduced cost, and scattering dozens throughout a city could be prohibitively expensive.

Norige's aim is to develop a sensing architecture focused on low-cost distributed detection, and to test these networks with various sensor prototypes. "Inexpensive biochemical sensors, which may even approach disposable in terms of cost, don't get a lot of attention from the R&D community," he says.

Although the sensitivities of inexpensive sensors may be a couple of orders of magnitude lower than those of the most advanced devices, linking them in a network can increase their overall detection performance. For one thing, no matter how sensitive a detector is, it works only if it's in the right spot to detect a threat—and given how unpredictably clouds of toxins can blow around, choosing the right location can be a difficult task. Scattering large numbers of sensors increases the odds that one of them will intercept the plume of threat particles that might be drifting toward a commercial district. "You don't have to be as sensitive, because the odds are increased that at least one or two sensors will be closer to the release point," Norige explains. And while a single sensor gives information at only one point, an array could provide spatial data about the plume's structure and propagation. "Once you have a network of sensors out there, you're gathering much more information than you would with a single, typical sensor."

Sensor at sunrise: a grid of high-fidelity smoke detectors and collocated anemometers was tested at Fort Devens, Mass., in August 2007. Such an array of sensor nodes can monitor the smoke plume's shape and movement—information that would be critical in a biological or chemical release.

For instance, with a network of sensors, officials could look at which sensor registered a signal at a given time and draw up a map of the plume's spread. That would help track the plume in real time—yielding clues as to where it came from and where it was headed.

"That would be a very helpful adjunct," says Robert Weiss, founder of Physical Sciences Inc., a company that develops sensors and other new technologies for government and industry. "The systems that are out there now are very large, very expensive, and few in number." He thinks the best arrangement may be to deploy one or two expensive sensors and then add Lincoln Laboratory's cheaper network, for a better mix of wide range and high sensitivity.

Weiss has discussed the work with Norige and thinks it shows promise for dealing with the threat of bioweapons, which he doesn't feel has been sufficiently addressed. "I don't think enough people are working on it," Weiss says.

The Lincoln Laboratory group has conducted proof-of-concept experiments using smoke as a threat cloud and a particle detector at each network node. In addition, each node included an anemometer to measure wind speed and direction. With such information as part of a real alert system, Norige says, officials could say, "We're seeing the threat here. Look at the way the wind's blowing. The people over here have so many minutes to get out of there." Norige and his colleagues are also working on algorithms to help distinguish between real detections and false alarms and track the plume's propagation. They are performing outdoor field experiments to test with plumes generated in an uncontrolled environment and supplementing their results with computational fluid dynamic analysis.

These networks could help detect chemical plumes as well as biological ones. Toward this end, Lincoln Laboratory is working with Timothy Swager, head of MIT's chemistry department, to develop an inexpensive chemical-agent sensor. Swager's device is based on carbon nanotubes and polymers with attached molecules that are designed to bind to specific chemical agents. When the agent binds to the polymer, it exerts pressure on the nanotubes, increasing the system's electrical resistance and thus signaling a detection. The Lincoln Laboratory team has built Swager's sensors into prototypes that incorporate commercially available mesh networking technology and is currently characterizing the prototypes in the laboratory.

Norige hopes to be able to build the sensor network for about $1000 per node; if it becomes commercialized, the price could drop to $100 to $200 per node, complete with sensor and networking equipment. At such a price, local officials could spread the sensors around to provide warning of biological threats. As part of its standard procedure in securing an area, the military could distribute sensors while setting up an operations base. The Environmental Protection Agency could place them near industrial plants to assure compliance with clean air rules or to warn if there's a chemical leak. It's just a matter of showing people that such a setup could work effectively and affordably, Norige explains.

"Low-cost, low-power, distributed detection is a hot field," Norige says. "We're trying to find it a home in biological and chemical defense."

ROBOTICS

Auto-mation

A robotic car bedecked with Lincoln Lab sensors takes on DARPA's "urban challenge"

From cruise control to antilock brakes to hybrid cars' fancy energy management programs, drivers have steadily ceded more and more control over their automobiles to computers. The logical endpoint of such developments—a car that drives itself—has long remained a futuristic mirage. But a team from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) and Lincoln Laboratory recently demonstrated that a self-driving vehicle is no mirage when they participated in the Defense Advanced Research Projects Agency (DARPA) Urban Challenge. The mission: design a fully automated, independent, self-thinking, self-correcting vehicle that could operate in the complex and cluttered environs of a modern city.

Earlier DARPA challenges required autonomous cars to make their way around the much simpler courses of deserts or mountains. MIT took a pass on these tests because they put relatively little

The TALOS autonomous vehicle begins a qualifying run at the DARPA Urban Challenge. Sensors scan and sweep across an intersection to show approaching and standing traffic that TALOS' algorithm must assess before continuing. Note the empty driver's seat.

demand on the situational-awareness technologies that Lincoln Laboratory and MIT specialize in. But DARPA realized that the true test of an automated vehicle was city traffic. "In the spring of 2006, there was chatter about an urban challenge," says MIT EECS and CSAIL professor Seth Teller. Urged by several students, Teller, CSAIL colleague and MechE professor John Leonard, MIT Aero/Astro professor Jonathan How, and Olin College professor David Barrett joined forces with Robert Galejs, Jonathan Williams, and Siddhartha Krisnamurthy of Lincoln Laboratory's Advanced Capabilities and Systems group and several other partners to develop a robotic vehicle. They named it TALOS, after the horseless golden chariot of Greek mythology.

The urban challenge was closely tied to dynamic real-world driving. The vehicles had to follow the roads, queue up at intersections, merge into traffic, park, and perform U-turns. Not only did TALOS have to sense its location on the map and travel a specific route, it also had to observe other vehicles (including both human-driven and robotic ones), lane markings, curbs, and obstructions. Consider the issue of a parked car that fills the driving lane on a road. Human drivers evaluate the problem, determine whether they have enough time to move into the opposite lane to pass the parked car, and proceed. The automated vehicles were expected to perform the same types of decisions and proceed in roughly the same time sequence as human drivers. (DARPA did simplify the task a bit by eliminating pedestrians and traffic sign sensing.)

Each vehicle was required to follow the curve in the road from one GPS point to the next. But because DARPA provided only minimal GPS information for the course, and the GPS waypoints were so scarce, TALOS relied on a tiling of radars and other sensors—all tightly positioned near one another and each pointing in a specific direction—to observe and interpret its immediate surroundings. Teller gathered collaborators in industry, academics, and Lincoln Laboratory with one fundamental goal in mind—a single algorithm that would handle all situations. A good driver, human or algorithm, "translates the laws of the road into a recipe for good driving," he says. After building the vehicle and testing the sensors, the team needed several months to encode approximate rules of the road into the vehicle.

The TALOS team started with a Land Rover, and automated the gas, brake, steering, and shifter controls. Sensors included one short-range 360° lidar, a dozen longer-range planar lidars with 180° fields of view (seven were oriented downward as if they were push brooms and five were oriented horizontally), 15 radars each with an 18° fields of view, five wide field-of-view cameras, and one narrow field-of-view camera. The massive computer cluster processing the stream of sensor data required an air condi-

The long term goal is to finish the assigned route. The short term goal is to avoid the pothole in front of the vehicle by tweaking the steering.

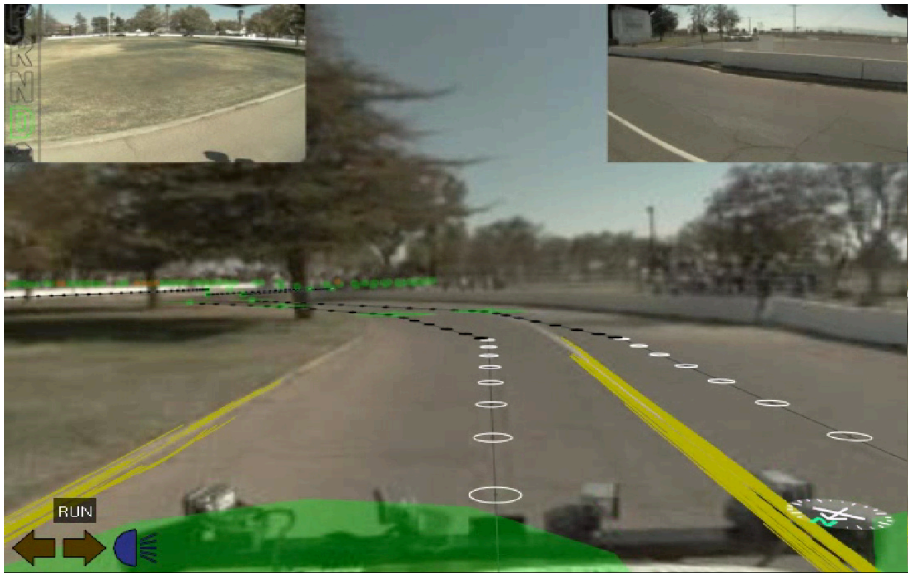

TALOS is preparing to negotiate a curve in the road to reach it next destination, the green marker, which is a DARPA-supplied GPS waypoint. The optical image is overlaid with a proposed path of circled destination points (white indicates that the lane-marker detection and lane-center estimation are functioning, and black indicates subsequent steps without current estimations), directional lines of motion to those points, and yellow lane and curb markers.

tioner, a gas-powered generator, and a battery supply (to keep the computer running whenever the generator had to be shut down—for example, for refueling).

Combining all the sensor's inputs, the vehicle's algorithm defined a series of routes, selected the best route that complied with the rules of the road, and reevaluated its situation about ten times per second. The algorithm's long-term goal was to complete the several-mile "mission" defined by DARPA. Its intermediate goal was to get the vehicle to the next GPS point on the route, tens of meters ahead. And its short-term goal? "Avoid the pothole in front of the vehicle," Teller says.

With new data arriving many times every second, potentially confusing the algorithms, the vehicle might stop, think for a while, and recover—or not. "I can't go in and help TALOS: it has to figure it out by itself," says CSAIL doctoral student Edwin Olson, a member of the TALOS team. During the Urban Challenge, the MIT entry distinguished itself as the highest finisher of all those that had not been involved in the first two challenges. TALOS made further news by being involved in two bot-to-bot collisions, but it was absolved of fault in both cases. Still, "human drivers probably would have avoided those accidents, so there is still work to be done," says Olson.

Teller relied on Lincoln Laboratory to evaluate and calibrate the radars and lidars. Galejs and his associates needed to develop a long-range radar sensing capability—identify radar options; characterize radar for accuracy, multi-radar interference, and clutter rejection; and develop software algorithms to merge with the other sensors and the controls of the vehicle. "Our leverage," says Galejs, "was our test facilities and the knowledge of how to test radars." And the collaboration with MIT was a success story: TALOS was a truly autonomous vehicle.

The challenge showed some of TALOS's strengths and weaknesses. On the positive side, TALOS correctly overrode the "do not go over the center line" rule when the alternative—staying in the driving lane—would have resulted in a collision with a parked car. TALOS paused, waited until it determined that no vehicle was in the passing lane, and went around the obstruction—exactly as a human driver would be expected to behave. Two difficulties that arose were primarily in response to the camera sensors. At one point in the course, several trees cast shadows across the roadway. These shadows confused TALOS's vision-based lane detection system, and made the vehicle stop until it sorted things out. TALOS also struggled to navigate a section of the course consisting of a dirt road with no curb or centerline. "Getting TALOS to drive more quickly and smoothly in perceptually difficult environments is something that the team will continue to work on," says Olson.

People might appreciate some advantages that would come with robotic cars, says Leonard: "You won't need parking lots next to strip malls. After getting out of your car, you just tell it to go park itself."

REMOTE SENSING

Good Vibrations

A different take on terahertz radiation can measure moving motors or beating hearts

Terahertz radiation—the part of the electromagnetic spectrum that lies between infrared and microwaves—has recently become a hot technology for security applications. T-rays, as they're called, can penetrate such barriers as clothing, paper, plastic, and cardboard, and can identify the chemical make-up and physical shapes of substances that they find. And they do so without the hazardous ionizing effects of X-rays. Now Lincoln Laboratory researchers Jerry Chen and Sumanth Kaushik have taken T-rays in a new direction: using this radiation to listen for vibrations. "As far as I know," Chen says, "this is the first time anybody has used T-ray technology for vibration sensing."

The interferometric technique starts by splitting the T-ray into two separate beams. Next Chen and Kaushik aim one beam at the object they want to examine. The beam hits the object and bounces back to a detector while a reference beam is routed directly to the detector. A beam reflected from a vibrating object will be out of phase with the reference beam, causing an interference pattern on the detector. Taking a Fourier transform of these time-varying patterns reveals the object's vibrational frequency.

Chen, an electrical engineer in the Laboratory's Active Optical Systems group, tested his system by placing an ordinary stereo speaker behind a cardboard barrier. Using the T-ray interferometer, he measured the peak velocity of the speaker as it vibrated. The results told him not only that the speaker was moving, but also which pitches it was producing. To check the validity of his approach, he performed the same test without the barrier, using a helium-neon laser beam in conventional optical interferometry, and got the same results.

The advantage of T-rays over a laser-based vibrometer is their ability to penetrate many nonmetallic barriers. For example, T-rays could detect the ticking of a time bomb inside a leather briefcase. Other ways of measuring vibrations often involve physical contact with the object being measured, an intrusion that can throw off the measurement.

T-ray vibration sensors could check the efficiency of motors, whether they're inside huge aircraft or tiny microelectromechanical systems such as those used in projectors. The technology, Chen says, could thus provide a method of testing industrial machinery without having to take it apart. An automobile designer might use a T-ray system to trace the source of a particular frequency to see just which part is causing an unwanted sound.

The device could also detect a beating human heart or vibrating vocal cords. That ability might prove useful in, say, examining a battlefield to quickly separate the injured from the dead. T-rays could work well in such a setting because they can penetrate smoke and dust as well as cotton and Kevlar. T-rays' speed and accuracy in triage for both military and civilian emergencies (such as a major traffic accident or building collapse) could save lives.

Chen filed a patent application on the system early last year. He says he'd welcome the chance to refine his detector, and sees no reason it couldn't be commercialized.

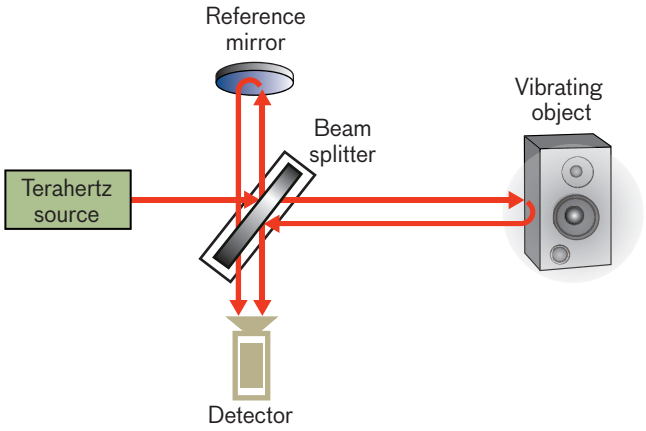

Terahertz (T-ray) interferometry can detect vibrations. T-rays that bounce off a vibrating object (such as a loudspeaker) are combined with a reference T-ray beam; the resulting interference between the beams yields information about the vibration.

NETWORK SECURITY

Plugging the Right Holes

NetSPA software maps computer networks to find paths most vulnerable to hacking

On the night of November 1, 2004, according to published reports, hackers in the Chinese province of Guangdong broke into computers at the Army Information Systems Engineering Command in Arizona, the Defense Information Systems Agency in Virginia, the Naval Ocean Systems Center in California, and the Army Space and Strategic Defense Installation in Alabama. The attack, *Time* magazine and the *Washington Post* wrote, was part of Titan Rain, a series of breaches of U.S. government computers that occurred between 2003 and 2005 and may have captured sensitive information about military readiness.

In fact, says electrical engineer Richard Lippmann, a senior staff member in Lincoln Laboratory's Information Systems Technology group, U.S. government and defense computer networks are attacked all the time. In response to this chronic cyber threat, he and his colleagues developed NetSPA, a software tool to identify potential avenues of attack in computer networks. NetSPA (for Network Security Planning Architecture) uses information about networks and the individual machines and programs running on them to create a graph that shows how hackers could infiltrate them. Although system administrators can examine visualizations of the graph themselves to decide what action to take, NetSPA analyzes the graph and offers recommendations about how to quickly fix the most important weaknesses.

NetSPA relies on vulnerability scanners, such as Nessus, to identify known vulnerabilities in network-accessible programs that might allow an unauthorized person access to a machine. Fast-spreading worms, for instance, often take advantage of weaknesses in servers or operating systems to spread from one machine to another. But simply being aware of vulnerabilities is not sufficient; NetSPA also has to analyze complex firewall and router rules to determine which vulner-

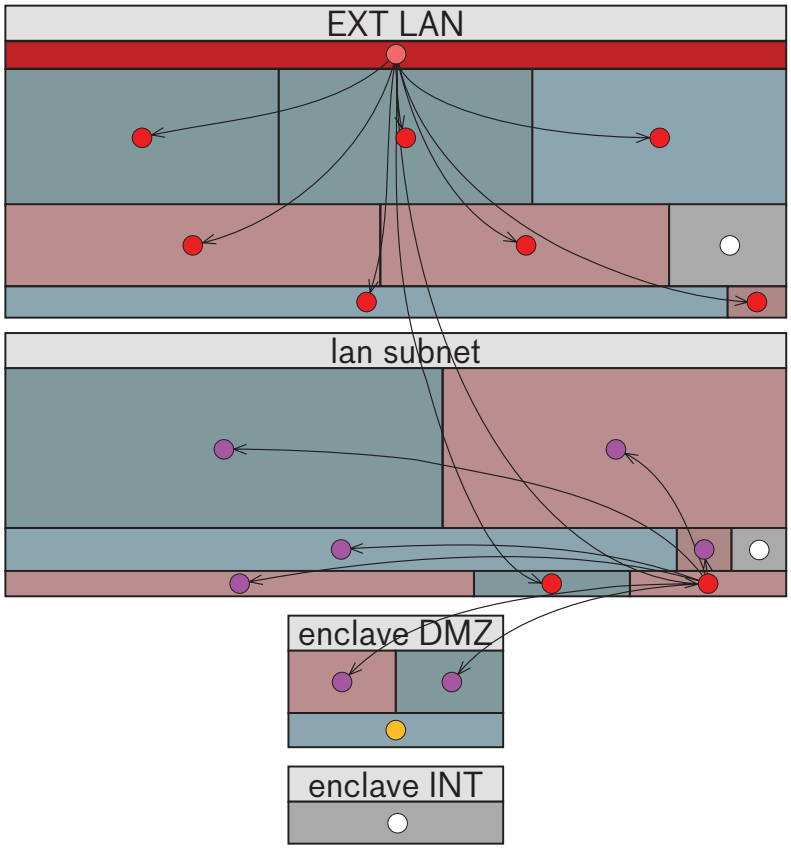

A screen shot shows an attack graph cascade. Each of the four large rectangular regions represents one subnet in a larger network. Within each subnet, the smaller rectangular regions represent groups of hosts that are treated identically by all firewalls and that are compromised by an attacker to the same level. The dot at the center of each region signifies all hosts in that region. The attacker starts at the upper subnet ("EXTLAN") on a single host (topmost dark rectangle). Lines connecting hosts represent vulnerabilities that the attacker uses to progressively compromise more hosts. After one hop, the attacker compromises all vulnerable hosts in the upper subnet and jumps to two hosts in the next subnet ("lansubnet"). On the next hop the attacker compromises all vulnerable hosts in the second subnet and jumps to two hosts in the third subnet ("enclave DMZ"). On the third hop the attacker compromises one more host in the third subnet and cannot reach the fourth subnet at the bottom of the display.

abilities can actually be reached and exploited by attackers and how attackers can spread through a network by jumping from one vulnerable host to another.

"It's a matter of what the attacker can get to and in what order," says Kyle Ingols, a computer scientist in Lippmann's group who is working on NetSPA, along with Seth Webster (who is focusing on ways to make the system more automated) and MIT graduate student Leevar Williams (whose master's thesis is on visualizing attack graph data). It takes a long time to patch all hosts in a network. "If you spend time patching vulnerabilities the attacker can't get to first," Ingols says, "you've left your network exposed longer."

NetSPA aims to solve that problem. "Instead of patching or fixing or blocking a thousand hosts," Lippmann explains, "we could say: There are ten critical hosts. Patch those first."

The software finds the most critical weaknesses by combining information from vulnerability scanners with firewall rules used to allow and block access and information about the physical structure of the network. For instance, if a firewall allows a certain kind of access, hackers could use that access to reach a vulnerable machine on the inside of the network. That might grant them access to only one machine, but once they take over that machine inside the firewall, they then gain access to many more. Thus a route through the firewall to a vulnerability on a single "stepping stone" host is much more critical than the potentially many other vulnerabilities on the network.

This insight sounds obvious, but applying it to real systems can be a huge challenge. A network comprising thousands of computers may have dozens of filtering devices such as firewalls and routers, and each device may have 200 or more different filtering rules. The multitudinous combinations of possibilities are far too many to track down by hand, and even very complex for a computer algorithm to compute. The original version of NetSPA, in fact, could handle networks of only about 17 machines before the modeling complexities made it too slow to be useful.

Since then, however, the Lincoln Laboratory researchers have developed ways to speed NetSPA up. For instance, firewalls may have rules that treat a number of different machines on the same network in the same way. Rather than modeling each of those machines individually, the software uses the same model for all of them, saving significant computing time. The researchers have also developed new types of attack graphs and efficient algorithms to compute these graphs.

In examining firewall rules, NetSPA also has the potential to discover unforeseen avenues of attack. For instance, a network might have had to share data with an outside vendor several years ago, so the system administrator would have added a rule to allow access from that vendor's IP address. That long-forgotten permission could be exploited by someone forging that address.

Lincoln Laboratory researchers have received one patent for the first type of attack graph they developed, called a "predictive" graph, and have one patent pending for a much more efficient and recurrent type called a "multiple prerequisite" attack graph. They're testing NetSPA on different networks and developing ways to make it easier to use. Eventually they hope to see it commercialized. "Ideally we would like to transfer this to a security company that could deal with all the details," Lippmann says. This tool would provide a protective umbrella in case anything like Titan Rain were to fall again.

AVIATION

Untangling the Friendly Skies

Computer recommendations could clear up some weather-caused airline delays

You've just finished a day of meetings in Washington, D.C. You arrive at Washington National airport in time for the 4:45 p.m. shuttle to Boston, only to discover that your flight is delayed—indefinitely—because of thunderstorms. The storms, however, are not in Boston or Washington, but in New York, and they've left the airplane that was supposed to take you to Boston stranded at LaGuardia.

Studies at Lincoln Laboratory suggest that getting even two or three flights per hour out of otherwise closed airports in highly congested areas can significantly

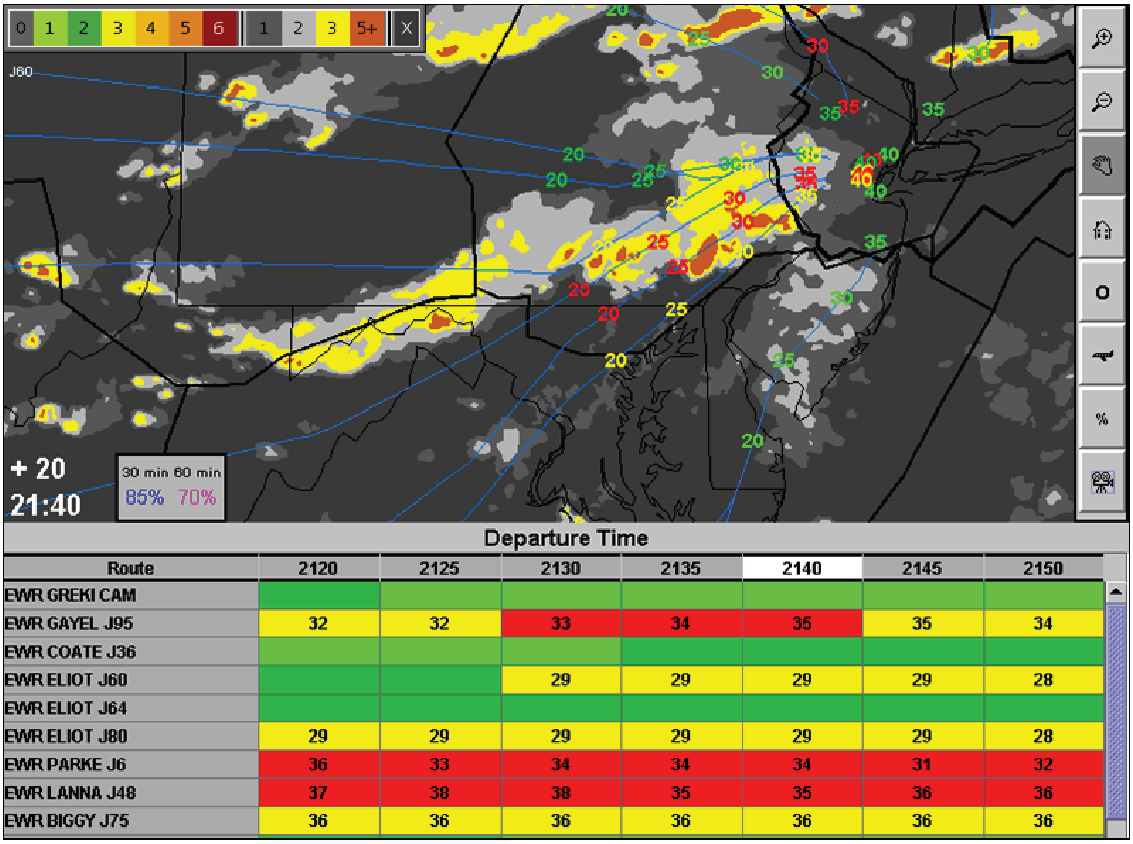

Departure Time

Route	2120	2125	2130	2135	2140	2145	2150
EWR GREKI CAM							
EWR GAYEL J95	32	32	33	34	35	35	34
EWR COATE J36							
EWR ELIOT J60			29	29	29	29	28
EWR ELIOT J64							
EWR ELIOT J80	29	29	29	29	29	29	28
EWR PARKE J6	36	33	34	34	34	31	32
EWR LANNA J48	37	38	38	35	35	36	36
EWR BIGGY J75	36	36	36	36	36	36	36

A screen shot of RAPT user interface shows the departure status timelines for Newark, N.J., departure routes and the weather forecast and projected departure trajectory animation screen. Red and yellow departures in the timeline include the height of the echo tops encountered, to help users understand the RAPT guidance. Specific departures in the animation frame are color coded to match their departure status.

reduce the weather-related delays that ripple across the nation's air travel system. And Laboratory researchers are developing a computerized prediction model and graphic display that can increase the odds of sneaking a few jets out between thunderheads.

Richard DeLaura of the Laboratory's Weather Sensing group is working on the Route Availability Planning Tool (RAPT) to give air traffic managers assistance in deciding whether to allow planes to take off during inclement weather. The computerized tool takes weather information from satellites and radar systems, makes predictions about whether a pilot would choose to fly through such conditions, and displays the information graphically to enable an air traffic controller to make a quick decision.

The display shows a map of the airport with lines radiating outward to indicate the various departure routes. A grid below the map lists departure times in rows, which are divided into columns of five minutes running from the present to half an hour in the future. Each rectangle on the grid displays a color that tells whether departure at that time along that route seems feasible. Red means the route is blocked. Yellow means there's some heavy weather that might pose problems. Dark green says there's weather, but that it shouldn't be an issue. Light green represents clear sailing.

Generally, air traffic managers get weather information and have to come up with a picture in their heads like the one RAPT displays, then make decisions based on that mental image. If the weather is changing rapidly and there are a lot of flights in the air, the process of conjuring such a picture can become so time consuming that controllers decide not to let any flights out. Instead they concentrate on landing the ones in the air. But

if too many departures are stuck at their gates, the arriving aircraft have no place to go once they land. The result is a major traffic jam. DeLaura hopes that RAPT will take away some of the manager's burden, making more departures possible and thus minimizing delays.

"What they really need to know is the following. When can I start moving departures along this route? At what rate?" DeLaura says, "We can say, we think you can start limited departures in about 15 minutes and go to full capacity in 30." RAPT bases its guidance on a computer model that combines the departure route geometry, forecasts for precipitation intensity and the height of radar echo tops (a measure of storm height), and a model for pilot behavior in convective weather (e.g., thunderstorms). The model estimates the probability that pilots will deviate significantly to avoid the weather along their routes and assigns the departure route status color based on that probability. The pilot model is based on studies of pilot behavior and other data gathered by DeLaura's group. "We're building the models from what we observe," DeLaura says.

A prototype of the system has been used in the New York City region—including LaGuardia, JFK, and Newark airports, several regional air traffic control centers, and commercial airline dispatch operations—for about four years, with modest funding from the Port Authority of New York and New Jersey. With 10,000 flights a day, it's among the busiest areas of the country. But this past year, the FAA began funding RAPT and asked for a major field assessment of the system. DeLaura and his colleagues spent the summer looking at how the system was performing. While they made some adjustments to the computer algorithms, their biggest discovery had to do with human factors. The air traffic managers using the system wanted to be sure they could trust what RAPT was telling them, and needed to know why it was making certain predictions. For instance, managers trying to decide whether to release departures along a yellow route looked at estimated heights of cloud tops in the RAPT timeline display to help them decide what to do. When they noticed that those heights were low enough that pilots could safely fly over the storm, they were able to make an informed judgment about the best course of action. RAPT also includes an animation of how storm systems intersect with flights. Managers can glance at that to confirm that they're getting a green or yellow light because the storm is moving away from their flight routes.

In their study, the researchers found that RAPT reduced flight delays in 2007 by 2300 hours. In terms of the costs of operating aircraft, plus the value of passengers' time, that delay reduction saved $7.5 million. They estimate that fully implementing RAPT in the New York region could save 8800 hours per year, which translates to $28 million in costs saved. "It certainly provides us with exceptional benefits in most scenarios with severe weather," says Leo Prusak, the FAA district manager for the New York area. "I think it's a fabulous product."

The researchers are also adjusting the model to take more account of the impact of arrivals on the system. Because incoming planes often avoid storms by deviating into the airspace used for departures, too many additional arrivals can disrupt or even stop departures altogether. And they're working on picking a site for the deployment of a second prototype system.

Eventually, DeLaura would like to see the system deployed at other large airports where congestion and convection cause major problems. Placing it at key spots could reduce delays at both large and small airports all over the country, he says. "You wouldn't need to have it in Elmira in order for passengers in Elmira to see some benefit.

BIOTECHNOLOGY

Looking for a Reaction

A new DNA amplification method is better, faster, and cheaper than traditional tests

When the polymerase chain reaction (PCR) was developed in the early 1980s, it revolutionized DNA cloning and analysis by allowing investigators to take a few strands of DNA and replicate them until there were so many that it increased the odds of detecting them. Now researchers at Lincoln Laboratory say they've brought PCR to a new level by developing reagents that make the process

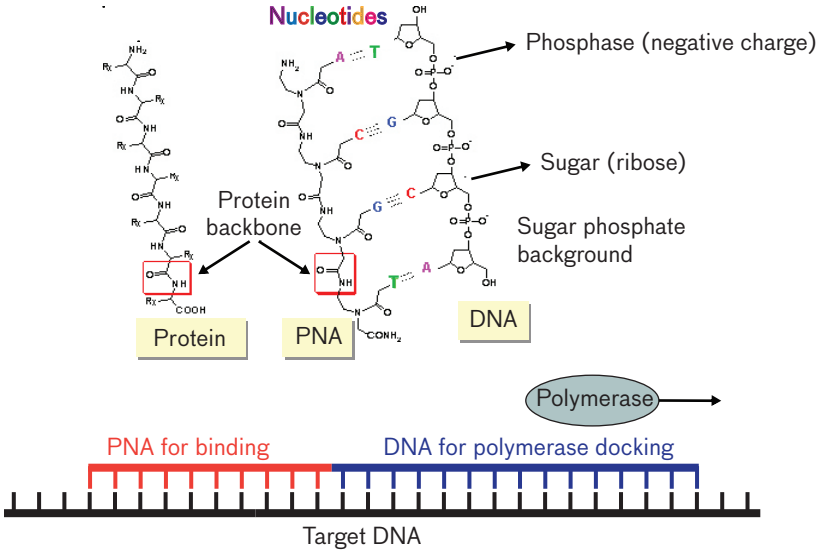

DNA and PNA are both composed of the same nucleotides (top). But while the DNA backbone is made up of sugar phosphate molecules, the PNA has a polyamide backbone similar to that found in proteins. A chimeric primer (below) starts with a few units of PNA (red) followed by DNA (blue), which together bind to a target string of DNA (black).

faster, cheaper, and more sensitive. This innovation opens up new possibilities for detecting pathogens in military uses, food safety, and clinical diagnostics.

Chemist Christina Rudzinski, molecular biologist Laura Bortolin, and applied biologist Amanda Stephens of the Biodefense Systems group have developed an artificial molecule that enhances PCR, which demands less sample preparation and provides heightened sequence specificity. The molecules designed by the trio are a modification of traditional PCR primers, which initiate the DNA amplification reaction. The new primers rely on a synthetic component, a peptide nucleic acid (PNA). PNA is made up of the same four nucleic acids as DNA.

Instead of the sugar phosphate molecules that make up the backbone to which the nucleic acids are attached, PNA has a backbone based on peptide bonds. Because DNA backbones have a negative electrical charge, opposing strands tend to repel each other, although the tendency of complementary nucleic acids to bind overcomes that. PNA backbones, however, lack charge, so they bind much more strongly to DNA. Because of this strong binding, PNA has been used for more than a decade to block certain DNA reactions in assays. But it has never been used for DNA replication in real-time PCR before. As it turns out, the lack of charge gives the molecule new properties that improve the workings of PCR, which up to now has been the gold standard for gene amplification.

The group's chimeric primer is a strand of nucleic acids specific to the DNA they're trying to replicate. Say they want to detect anthrax. Using a computer program Stephens developed to design the primers, they pick a portion of a DNA sequence that's unique to anthrax and have an outside company synthesize a matching primer. Like the mythical Chimera, with the head of a lion and the body of a goat, the primer starts with one to seven units of PNA, followed by a string of DNA. The PNA, which is more strongly attracted to the target DNA, snaps onto the beginning of the target sequence, and the DNA follows, just as in natural replication. Rudzinski says the presence of PNAs essentially jumpstarts the reaction.

On rare occasions, all-DNA primers sometimes bind to the wrong spot on the target molecule. For a number of complex reasons, the PNA-DNA primers generate fewer binding errors. Such specificity is important if you're trying to tell, say, whether a sample contains anthrax or *Bacillus thuringiensis kurstaki*, a common insecticide that is often mistaken for anthrax.

Because the PNA at the head of the primer isn't charged, it's much less sensitive to salt levels in the sample, whereas PCR won't work unless the salt concentration is just

The chimeric primer opens up new possibilities for detecting pathogens in military uses, food safety, and clinical diagnostics.

right. The charge difference also makes the reaction less sensitive to pH levels. That means that preparatory steps to get salt and pH levels just right are no longer needed, and the samples can be much less pure. A sample requiring PCR amplification is commonly treated by using DNA purification kits, which remove extraneous proteins and other material that might interfere with the amplification process. The chimeric primers don't need that step. The group tested samples with 100 milligrams of soil to a milliliter of water, in which all-DNA primers don't work at all, and found their chimeric primer produced results almost as good as in a pristine sample. This insensitivity of the process to sample purity increases testing speed while cutting costs. It also worked in a drop of blood only slightly diluted with water; such a sample would stymie normal PCR because of salts and other components in the blood. Moreover, the PNA renders the primer unrecognizable to the DNases and proteinases—that is, the enzymes that break down natural DNA and proteins. Again, it's the non-standard backbone that accounts for the difference. With fewer molecules attacking it in the blood sample, more of the primer survives to find the target DNA, making the test more sensitive.

Reducing the need for sample preparation cuts down on time and cost. In tests, the group found that, depending on the type of sample, their PNA primers cut prep time approximately in half, and cut costs by two thirds.

The PNA system has obvious appeal for military and homeland security uses, but it could also be developed for food safety and environmental testing, or even medical diagnostics, the researchers say. They're hoping for funding to test it for other types of targets. For instance, they haven't yet shown that it can work in the reverse-transcription PCR assays required to detect RNA. But they see no reason that a chimeric primer wouldn't work there, too. "If you could do that, you could do HIV detection in blood without sample preparation," Bortolin says. "That would be tremendously useful." It could mean a test in a doctor's office could detect actual virus in a few minutes, whereas today's quick tests can find only antibodies.

SPACE SURVEILLANCE

A Big Eye Sees Small Things

An upgrade to the Millstone radar antenna will ensure uninterrupted tracking of the ever more crowded geosynchronous orbit.

It takes a really big antenna to locate small, faraway objects. At 26 meters across, Lincoln Laboratory's Millstone Hill Radar (MHR) antenna certainly fits this description. And thanks to a recent renovation, the antenna is now easier to maintain and much less likely to suffer from downtime.

MHR is one of the principal tools for maintaining the Deep Space Catalog—the listing of the more than 3000 objects that are circling the earth 40,000 km away in geosynchronous earth orbit (GEO). The antenna consists of a 84-foot-diameter Cassegrain-fed parabolic reflector, together with counterweights; its motion is controlled by a set of motors mounted to the azimuth deck. In an elevation-over-azimuth axis configuration, the entire structure is supported on a 26 m tower. With its wide beam, MHR can view large regions of space and can locate objects for the first time or after they have been lost because of a planned or unplanned shift in orbit.

Installed in 1957, Millstone was the first radar system to do space surveillance (it observed the Sputnik satellite) and satellite launch tracking. But the venerable system was showing its age. The motors and motor generators replaced in this renovation were original 1950s era equipment. "They were past their end of life. The motors were worn from years of use and regular rebuilds, and the inefficient motor generators were failing frequently," says Paula Ward of the Control Systems Engineering group. Each failure would shut down the antenna for a significant period of time. Now the system is easier to troubleshoot, and downtime can be kept to a minimum. Jeff Dominick, the site manager of the Lincoln Space Surveillance Complex (LSSC) that includes MHR, stresses the importance of MHR to LSSC and how important LSSC is to the Air

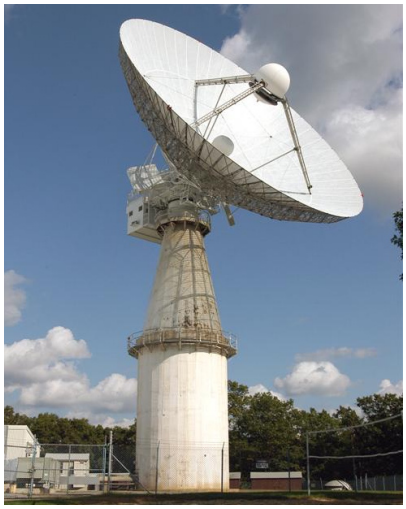

Thanks to a critical upgrade, the Millstone Hill Radar in Westford, Mass., rejoins ALTAIR (in the Western Pacific) and Globus II (in Norway) to monitor the increasingly cluttered geosynchronous orbit.

Force Space Command. "Losing MHR for any period of time would impact our ability to track in this region," he says, pointing to the arc of GEO above the United States that isn't covered by the other two surveillance radars—ARPA Long-Range Tracking and Instrumentation Radar (ALTAIR), located on the Kwajalein Atoll in the Marshall Islands, and Globus II in Norway.

As project lead of the recent upgrade, Ward was working with two extremes—very heavy and bulky motors and gear boxes, and new software controls running in a real-time embedded environment. The components that were replaced included the drive motors, solid-state silicon-controlled-rectifier power amplifiers, some of the gear boxes, the servo control unit, the programmable logic controller, and all the position and velocity sensors. Installing and aligning the new motors and gear boxes was challenging, since no mechanical computer-aided-design models existed. At the other end of the spectrum, while many upgrades had been done to the radar system and associated computers over the years, the antenna control system had been upgraded only twice in the last fifty years. "The previous antenna control upgrade, completed 21 years ago, used a microprocessor for position-loop control, but most other functions were still done in hardware," says Ward. "Development for controls to drive new motors using the obsolete 386-based microprocessor was out of the question." The new servo control software, written in C, runs on a VME-based Motorola 6100 processor by using the VxWorks real-time operating system. MatLab and Simulink were used to create a detailed model of the antenna and drive train. An improved algorithm for control of the motor pairs in each axis to mitigate effects of drive train compliance and backlash was also implemented.

"The safety interlock controls were obsolete as well," Ward says. "And safety at every level was of prime importance." Special precautions were taken for personnel safety during integration. Equipment safety was also critical, since the new motors and gearboxes were interfacing with one-of-a-kind 50-year-old equipment in several places. The legacy programmable logic controller, an obsolete 286-based computer, was replaced with a new Allen-Bradley programmable logic controller to implement many of the safety functions. Extensive testing of all safety functions was performed during integration.

Prior to the upgrade, an operator needing to move the antenna for maintenance had to turn mechanical knobs to rotate the antenna and read meters on a panel indicating positions. Now the interface is more intuitive and is done on a laptop running a custom application written in LabVIEW. The operator simply sets the desired positions in azimuth and elevation, and clicks Run. A complete suite of servo and safety data is recorded automatically on this laptop whenever the motors are enabled. In addition, the upgraded system provides, for the first time, remote access to the antenna's local displays. The new maintenance laptop, combined with the servo control unit, gives unparalleled ease of use for maintenance and complete insight into antenna operations.

The Millstone team is very satisfied with the results of the renovation. With its broad beamwidth Millstone surveys about a 400 km circle at GEO altitude, and its high angle accuracy provides excellent data for the other antennas such as Haystack to track objects. Today, MHR and its partners, ALTAIR and Globus II, cover the entire 360 degrees of GEO as well as monitor satellite and spacecraft launches. With more than 12,000 known objects in earth orbit, mankind is quickly filling the vacuum of space. MHR searches for new objects and reacquires the essential information on drifting or unstable objects. Dominick concludes, "We're trying to reduce the probability of collisions. We have to rely on Millstone and Haystack to keep track of everything. This upgrade has sig-

nificantly reduced downtime and maintenance tasks associated with the MHR antenna control system."

ERGONOMICS

Order from Chaos

Human factors engineering adds value to complex systems by making them seem simple to the user.

Sometimes too much information is a burden. Say, for instance, you were searching all the satellites in Earth's orbit. Advanced space surveillance technologies provide enormous amounts of raw data—but none of it does much good if it's too difficult to make sense of.

Enter Ann-Marie Lind of Lincoln Laboratory's Surveillance Systems group. Lind's work on the Optical Processing Architecture at Lincoln (OPAL) program ensures that users will be able to easily and efficiently work with the information provided by the technology the Laboratory creates. Lind is a human factors engineer—an expert in a discipline that aims to understand and improve the interactions among humans, the tools and systems they use, and the environments in which they interact.

There are thousands of satellites already in space, and the number of additional satellites is growing rapidly because of increased military, governmental, and commercial use. Newer satellites are becoming smaller and more difficult to detect. The OPAL software helps monitor, log, and track these space objects and presents the data through a graphical user interface (GUI).

Systems prototyped at Lincoln Laboratory are typically developed by a team of specialists with expertise in electrical engineering or physics. But as Lind points out, the end users of these systems are often non-technical, entry-level laypeople, such as soldiers or air traffic controllers. Human factors engineers translate user needs into design specifications. At the Laboratory, human factors engineers focus on the most efficient and effective ways for the user to view and access vital information for making a correct and timely decision. Lind says that her main task is to "turn chaos into order."

Early user interface prototypes for OPAL had been designed by four separate software developers—each creating screens with a distinctly different look and feel. The result was a confusing mishmash that was an invitation to frustration and human error. To address the situation, Lind and Donna Anastasi—a human factors engineer in the Surveillance Systems group—became part of the OPAL team. Their mission: ensure that the technology provided by OPAL could be used to its fullest potential by the human operator.

Lind and Anastasi redesigned the OPAL GUIs to reflect best practices in usability engineering. They worked with the software developers and provided written specifications and graphical mock-ups to convey how the GUIs should look and perform to ensure their usefulness and usability.

Principles applied to the OPAL GUIs include:

- **Provide a consistent look and feel.** Lind and Anastasi developed the GUI Style Guide for Developers—a quick reference for each of the four developers to follow in implementing the GUI. Applying the conventions listed in the guide led to a usable interface, as well as cost savings in developer time. Guidance was provided in areas such as format (fonts, use of acronyms and abbreviations, labeling, spacing), windows (content, format, management), widgets (proper use and labeling of buttons, checkboxes, and tables), visualization (use and layout of graphs, maps, colors, and symbols), and feedback (status information, progress bars, tool tips).
- **Organize the layout** of the display to support the job the user needs to do. The human factors engineers arranged widgets in an efficient, easy-to-find sequence that supported the task flow. Buttons were placed according to the usage sequence and reading style (top to bottom, left to right).
- **Keep users informed.** Too many software systems leave users wondering whether their requests went through. The OPAL GUIs provide a steady and meaningful stream of informative feedback messages. In addition, the GUI provides error checking where feasible and informs the user of the proper range of values to be entered.
- **Use familiar language** rather than computer jargon. In OPAL, this principle was as straightfor-

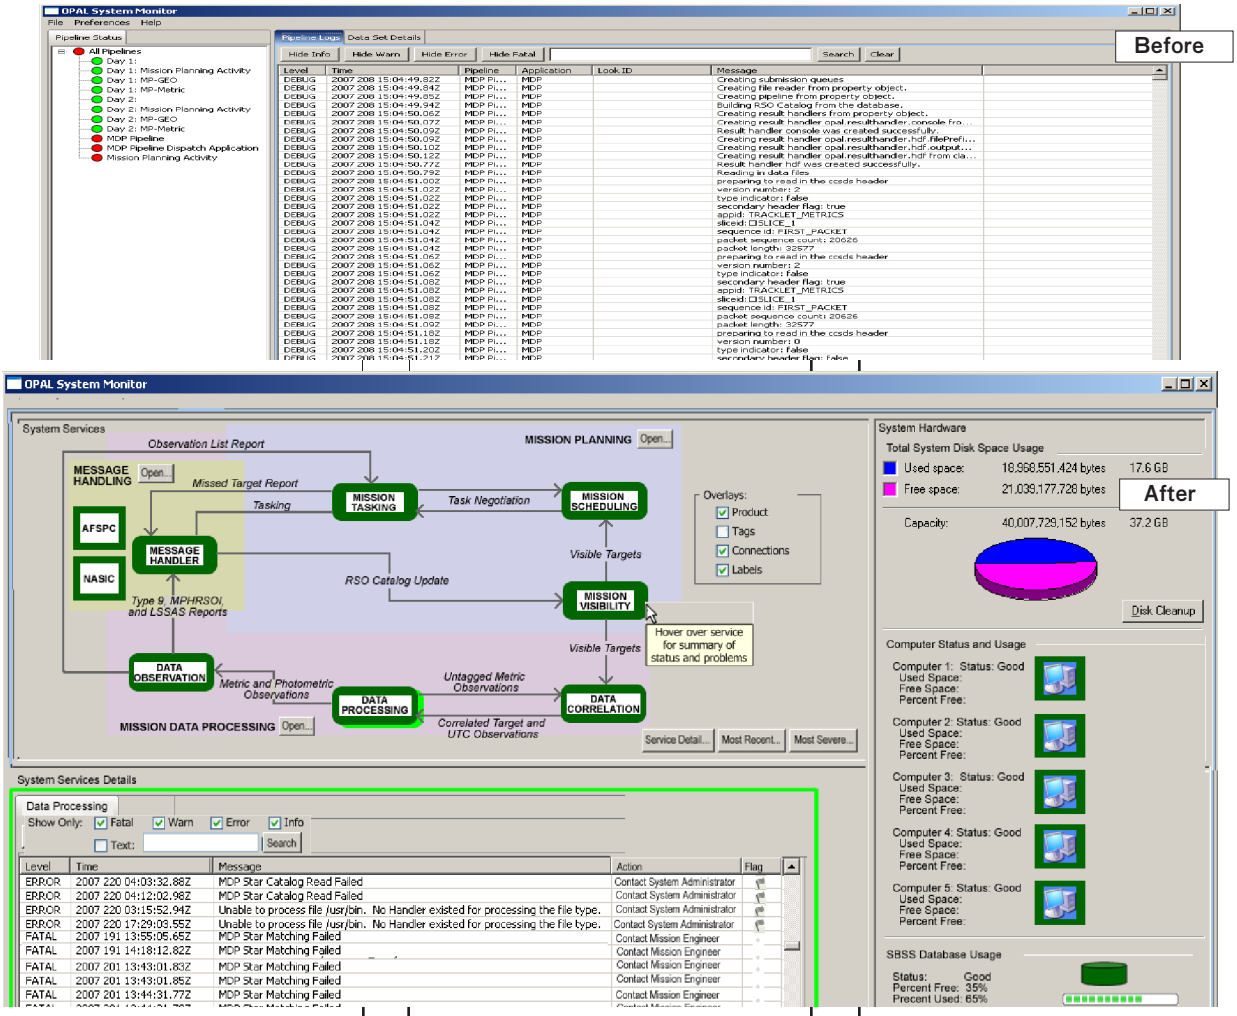

A before-and-after comparison of software recommendations from human factors engineers. The top of the figure (before) shows the system monitoring display prototype. Primarily textual, it requires heavy information interpretation. This GUI lacks "at-a-glance" diagnostics and indication of system status. Filter settings are unclear, since push-button labels change depending on message filter status. The lower half of the diagram (after) organizes the system monitor GUI into three sections. The system services (upper left) are shown in a high-level schematic, with overlays for user control of information complexity. You can hover the pointer over a service to see a summary of the status and problems for that service. Hardware status (right side) is represented textually and graphically. System service details (lower left) implement checkboxes rather than push buttons for filtering messages to clearly identify which filters are selected.

ward as consistently referring to messages as "messages" rather than "files"—a term that reflects a programmer-centric rather than user-centric mindset.

- **Provide shortcuts.** Instead of requiring users to burrow through menus within menus within still other menus, the OPAL GUIs provide quick ways to select frequently performed actions.

"I think the biggest improvement made to a number of our displays was work flow," says George Zollinger of Lincoln Laboratory's Space Control Systems group and program manager of the OPAL program, who worked with Lind and Anastasi. "Many of our GUIs are functionality driven, so developers tended to build screens full of features. But they didn't think as much about how a user would interact with the tool. Ann-Marie made the tools more intuitive and improved many displays." Managers and users alike, he says, agreed that applying human factors principles and techniques results in increased efficiency and effectiveness in system development and in end-system use. Zollinger adds that the contributions of Lind and Anastasi "demonstrate that Lincoln Laboratory is taking diligence in delivering a quality product."

PROCESS CONTROL SYSTEMS

Standing Guard

Q&A with Robert Cunningham

The process control computers that are in charge of such functions as distribution of electricity along the nation's power grids, the flow of natural gas through pipelines, and the operation of water treatment plants are becoming more and more accessible through the Internet—and thus potentially vulnerable to terrorist attack. Robert Cunningham, associate leader of Lincoln Laboratory's Information Systems Technology group, leads a project for a consortium of universities, national laboratories, and federally funded R&D centers called the Institute for Information Infrastructure Protection (I3P). The aim of I3P is to make sure such process control systems are secure. Lincoln Laboratory Journal contributing writer Neil Savage spoke with Cunningham about the consortium and the problems it aims to solve.

Lincoln Laboratory Journal: *When you began to look into the state of this infrastructure, what surprised you most?*

Robert Cunningham: My biggest surprise was when I got to see a process control plant. What I saw was a 30-year-old DEC machine—Digital Equipment Corporation, which no longer exists—as the primary process control system for this particular plant. And it was operating side by side with a Windows 2000 system that's also many years old and that has lots of well-known vulnerabilities. Then there's a modern laptop sitting right next to the other two systems. So it's like 30 years of the history of computing equipment, all nearby, some just recently connected to the Internet. All these things sitting in one room, sometimes connected together, are pretty worrisome.

LLJ: *Why are process control systems so vulnerable?*

Cunningham: Most of these systems are RTUs and PLCs. An RTU is a remote terminal unit and a PLC is a programmable logic controller. They were originally designed to talk on a dedicated wire that connected them to a human–machine interface—so there was no sort of authentication ever built into the protocols they use. There is nothing that says, "Are you really my PLC?" And there's nothing in the protocol to ensure that the data can't be changed either, because it was assumed that faults would occur very rarely. Finally, the network software stack was tested in very limited ways—making assumptions about what would talk to it and how.

The second thing to know is that the commodity systems that these are being built on stay in service for exceptionally long times and, unlike an enterprise computer system—like, say, the Macintosh that I have back in the office—they aren't patched at the same rate. It's easy to see why. After all, these things are the process control for the plant. Taking them down and upgrading them is a big deal, resulting in lost production and revenue. So if they are not broken, they are not usually fixed. After a decade of unpatched service, an enterprise operating system typically will have many well-known vulnerabilities. These vulnerabilities have particular importance in process control systems, which are where computers systems touch the physical world. If the systems aren't carefully designed and operated, human beings can die.

After a decade of unpatched service, an enterprise operating system will typically have many well-known vulnerabilities.

LLJ: *What sort of mischief might someone cause by hacking into these systems?*

Cunningham: I have a couple of examples. The first is the Bellingham, Washington, incident where there was a 16-inch pipeline that ruptured and poured 237,000 gallons of gasoline into a creek. Two

10-year-old boys and an 18-year-old fisherman died when the gas ignited and sent a fireball down the creek. This was not an example of terrorism. It's simply a case where a system administrator was changing the records of a database to monitor a few extra things. The admin committed the changes and went to the bathroom for 15 minutes. While he was gone, the people who were in charge of controlling the system had opened up the input to this long pipeline but were unable to control the output. More and more and more gas was flowing into the pipe with nowhere to go. Because their monitoring at the far end wasn't working either, they couldn't see that it hadn't opened and they couldn't tell that there was a pressure problem as well. As a result, the pipeline ruptured. Here's a case where it pretty clearly demonstrates that if you're not careful and your process control system fails, you can end up killing people.

The other example that I like to use is from the wastewater treatment plant in Maroochy Shire, Australia. A disgruntled former employee of the company that made the facility's process control system applied for but was turned down for a job at the wastewater plant. Monitoring and control for these systems was communicated wirelessly, and he knew how to access the network via the wireless network. So he would drive to a site near a receiver, connect in as pumping station #4, and reprogram the PCS to cause the control systems to dump sewage into the nearby rivers. Then he would call up the operating company and say, "You know, I would serve as a consultant for you if you need somebody to help fix your problem." Eventually he was caught and sentenced for two years in prison.

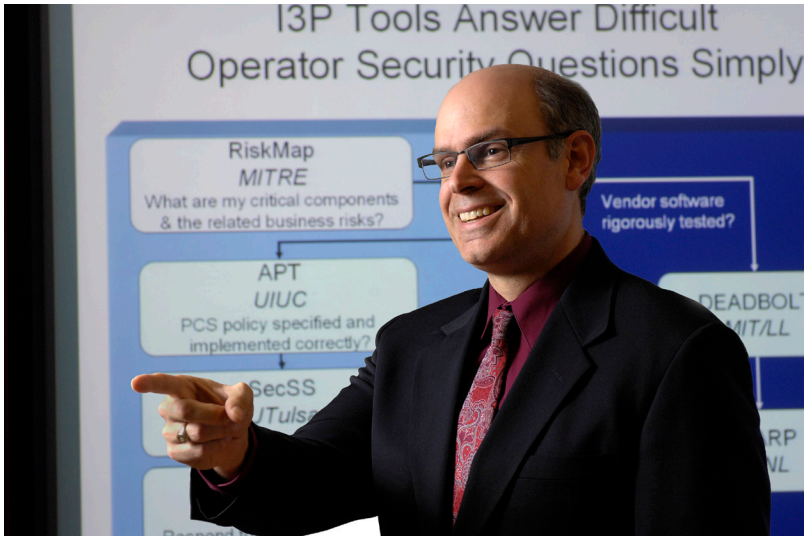

Lincoln Laboratory's Robert Cunningham leads the Institute for Information Infrastructure Protection (I3P), a consortium dedicated to the security of process control systems in an age of increasing vulnerability.

Process control systems are where computers touch the physical world. If the systems aren't carefully designed and operated, people can die.

LLJ: *Could he have done this from the Internet?*

Cunningham: I'm not sure. But the employee knew the PCS network better than business network, and he may simply have exploited what was easiest for him. I do know that it's increasingly common to have the PCS network connect to a DMZ network, which connects to a business network, which connects to the Internet. Sometimes companies don't think this is so—a few years back Paul Dorey, then chief security officer at BP, asked if the company's process control systems were connected to the Internet. Dorey was told that none were. He was skeptical, though, and so he did a careful study that discovered that in fact 89% of BP's process control systems were connected to the Internet. If those system connections are carefully designed, there are at least one and maybe more firewalls. A common mistake is to think only of outside attackers. But if attackers can get to the PCS networks, then they can often reach back as far as into the business network.

LLJ: *Are the system operators not aware these problems exist?*

Cunningham: Not until recently. With some of the problems the operators hadn't really thought

through the process. There was an argument until relatively recently that said, "Well, I don't think we have a problem here," because we didn't have examples of real cases of where people have successfully attacked systems. So that's why the Maroochy Shire example is a very nice one.

Even when the IT folks in a water treatment plant started to notice that a problem existed, little changed at first. In some cases they were responsible for the business network but not the process control network, and in other cases their suggestions to improve the security of the system were falling on deaf ears.

Furthermore, the IT personnel at the plants didn't have the tools to make the business case for better security. In the market space, most operators were not asking for security features to be built into products. That needs to happen—there needs to be a market pull. You also need to have a market push—vendors should be making equipment available that is more secure at about the same price. The program that I've been working on in collaboration with lots of other folks from other labs and other universities has been trying to build both pieces of this, working with operators to build the pull for buying new equipment, and working with the vendors to improve the quality of the offerings and the security of their products.

LLJ: *What is Lincoln Laboratory doing about the problem?*
Cunningham: We helped fashion the research program, and we have one element of the solution. Our piece in all of this is trying to secure software that vendors are making as a part of the PLCs and the RTUs. We're building a tool that will allow vendors to automatically test for certain vulnerabilities. Input is fed into the system, the system runs, and our instrumentation keeps track, for example, of how much memory is allocated or free. It can also tell you if you write beyond a range of memory dedicated for that information. If memory use starts to grow without bound or is improperly written to, we can tell you exactly where—at what line of code—the potential vulnerability is occurring. Then vendors can go back and fix their software to make sure that the vulnerability doesn't continue. I hope that this tool will become useful elsewhere at Lincoln Laboratory, too. The Laboratory has lots of long-lived embedded systems, like satellites and radar systems, that could benefit from a tool like this.

LLJ: *What other steps is the consortium taking?*
Cunningham: You have to think about this from the point of view of an operator, because the marketplace is ultimately going to have to buy this capability. An extremely high percentage of the nation's critical infrastructure is owned by independent systems operators. So we have to talk to them and help them build a business case for why it's important to include security in their systems. Then we tell them some of the questions that they need to be able to ask their vendors.

LLJ: *What sort of questions?*
Cunningham: Are they using secure protocols? Have they had their software systems tested? Has the platform they're working on been hardened against sets of attacks? The first couple of questions about rigorous testing are being answered by the work we do at Lincoln Laboratory. The folks at the Pacific Northwest National Laboratory are working on platform hardening. Then operators need to configure their systems so that only certain people, coming from certain locations, are able to access and control the components and processes running there. Researchers at the University of Illinois at Urbana-Champaign are developing a tool to make sure that firewalls are being configured correctly.

Once operators have everything configured in a way that's secure, they still need to monitor use, because they usually don't want to prevent all access, and not all the attackers come from the outside. So we've got a mechanism to monitor use, which is being worked on at

A high percentage of the nation's critical infrastructure is owned by independent systems operators. We help them understand why it's important to include security in their systems.

the University of Tulsa. And then because we think even that may someday fail, we need the system to be able to be recovered as quickly as possible. Our team members at Sandia National Laboratories are trying to build a tool to automatically recover and restore a system that's been broken.

The last thing we are trying to do is help out with technology transfer. Our first tool is being made commercially available this year. This is MITRE's RiskMAP tool, which connects business objectives to network nodes and indicates where investment should be made.

LLJ: *What else do you have to do to get the message out about securing the nation's infrastructure?*
Cunningham: The program has been going on for two and half years now, and it's got about another year

and half left. We usually say that there are three sets of customers. One is the government, which has asked us to help participate on running a number of workshops for them. We've done that. In fact, we just had a very successful workshop in Houston, where the attendees raved about our approach. And I've participated in webcasts for the SANS (SysAdmin, Audit, Network, Security) Institute. Another customer is academia. We've held a couple of conferences and published one book covering our and others' research, and we'll be working on a second book this coming year. I've given several invited talks to various universities. Finally, there is industry. We've got a couple of patents pending, and we hope to file for additional ones within the next year. We'd like to have at least one more commercial system on the market in addition to the MITRE tool. We've also put together an advisory board made up of vendors and operators in the oil, gas, and chemical businesses. Some companies have asked us to come in and help work with them, which has two benefits—it's an opportunity for them to make their systems more secure and for us to make sure that the systems that we're building actually work. In fact, we get more requests than we can handle.