Lab Notes

NEWS FROM AROUND LINCOLN LABORATORY

AVIATION

Foul Weather Friend

Corridor Integrated Weather System now covers the continental United States, expanding air traffic planners' and airline dispatchers' view of the nation's weather

Weather prediction software developed at Lincoln Laboratory may help to ease delays in the nation's air traffic now that the Corridor Integrated Weather System (CIWS) has gone from a regional experiment to a continent-spanning demonstration project.

On June 3, CIWS, which provides FAA air traffic managers and airline dispatchers with predictions of the location and intensity of storms at airports and along airline routes, expanded its coverage to encompass the entire continental United States and the southern part of Canada.

"The fact that we've expanded the coverage of the CIWS products has helped greatly," says Michael McKinney, CIWS program manager at the Federal Aviation Administration (FAA), which sponsors the project. He says that, with partial coverage, traffic managers could only see how weather would affect flights in a particular region of the country. The updated system gives a more complete picture. "The CIWS expansion gives air traffic management personnel who had been outside of or on the boundary of the coverage area the ability to more fully assess how weather is going to affect their airspace," McKinney says.

As any frequent flier knows, passengers can be sitting in one airport on a perfectly sunny day waiting for a plane that has been delayed because of thunderstorms elsewhere in the country. "The nation's airspace is an enormous, interdependent network, so what happens in one location can affect operations in a very wide area," says Elizabeth Ducot, a senior staff member in the Laboratory's Weather Sensing Group and leader of the CIWS project. For example, when convective weather occurs in

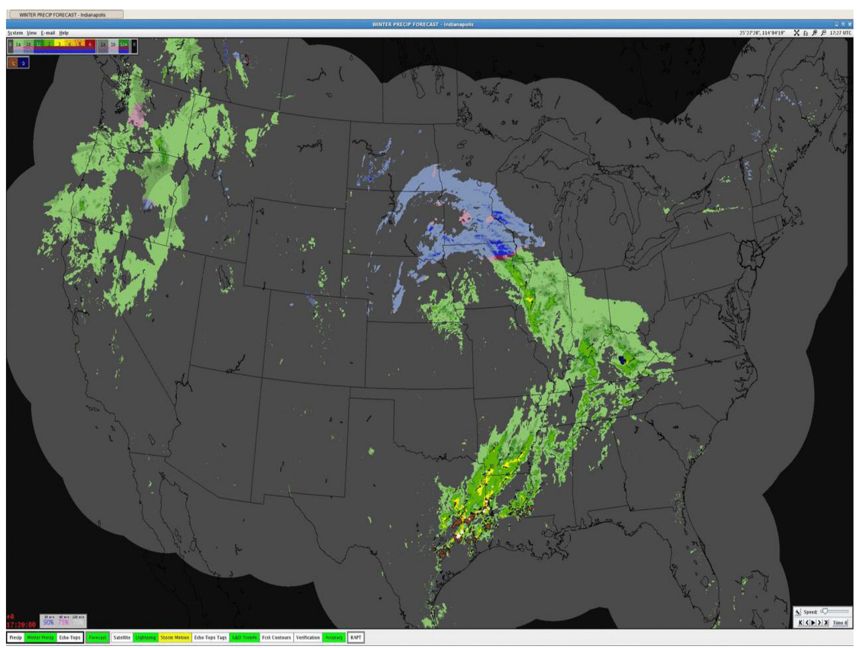

A November CIWS display shows snow (blue), ice (pink), and rain (green). This time-zero frame is the start of the winter weather forecast animation, which extends to two hours, looping in five-minute forecast increments; the entire forecast suite updates every five minutes. During the weather event shown, the CIWS forecasts enabled Northwest Airlines to collaborate with national and regional air traffic managers to develop a plan for when the passage of a heavy snow band would require a reduction in arrival rates at the Minneapolis-St. Paul airport. Knowledge of the snowfall's start time and duration allowed traffic managers to reduce airport demand, thus preventing excessive airborne holding and increased fuel usage.

Ohio, the air traffic from the western states headed for the New York, Boston, and Philadelphia airports may need to be rerouted north into Canada or south through Georgia, causing delays in Atlanta even though the entire Atlanta area is completely clear. "Severe weather can cripple the flow of traffic nationwide," Ducot says.

Delays have increased in recent years, as propeller airplanes that fly at lower altitudes have been replaced with jet planes, making the higher altitudes more crowded. An increase in corporate jets has also added to the congestion.

Daily planning for the nation's air traffic has traditionally relied on highly inaccurate weather forecasts that cover two to six hours in the future. So the CIWS project was started in 2001, with the aim of providing forecasts in the zero-to-two-hour range that are updated every five minutes. These forecasts provide air traffic personnel a much more detailed picture of how weather is changing in ways that could affect flight plans, not only for takeoff and landing, but also along the routes between airports.

CIWS provides forecasts of both precipitation and echo tops, which are the uppermost limits of the storms as measured by radar. Ducot says that previous forecast systems didn't look at echo tops, which are important in determining whether an airplane can fly over a storm. Additionally, CIWS forecasts whether the precipitation will be in the form of rain, snow, or a mixture, information that is particularly important in managing winter operations. Snow, for example, curtails visibility much more severely than the same amount of rain, so knowing when heavy bands of snow will arrive or depart in a terminal area is extremely valuable in trying to juggle arrivals and departures and in keeping the airport open through a winter storm.

The Lincoln Laboratory team began testing the system in the Northeast Corridor and along the Great Lakes in 2001 because those were the areas that are most congested, even in good weather. Over the years, the team collected data and refined the system on the basis of feedback from FAA air traffic managers and the airlines. A study in 2005 found that the use of CIWS reduced delays that year by a total of 90,000 hours, resulting in a savings in airline operations cost of more than $90 million.

Providing nationwide coverage required an overhaul of the system. The coverage area more than tripled, which Ducot says led to a significant increase in the number of computations required. To be able to create these larger products quickly, software designers broke up the grid into smaller sections, called tiles, so weather predictions could be run simultaneously on several sections and then stitched back together into a whole. Programmers also needed to change some underlying assumptions. While, on a smaller scale, they could treat the earth as if it were flat, there's enough curvature over the expanse of the United States that continuing to do that would introduce unacceptable errors into the predictions. So they adopted a nationwide standard mapping system that minimized the prediction errors and retested the software against it. "The reengineering was a fairly major undertaking," Ducot says. In the process, the team took the opportunity to remove redundant sections of code that were the result of the evolution of the prototype in support of various research ideas along the way and rewrite bits of old software that had been designed to run on the obsolete systems in the original CIWS prototype.

The Lincoln Laboratory group also had to upgrade the computer hardware. The old system consisted of different types of computers, mostly a mix of Sun computers purchased between 2001 and 2003. At the heart of the upgraded system is a new 150-node Linux cluster of processors. Each node is a Dell PowerEdge 1955 Blade with two dual-core 2.9 GHz processors, 8 gigabytes of memory, and dual gigabit Ethernet connections. The system is designed to easily scale up to accommodate advances in the forecasting science that will lead to new capabilities for the FAA. The team also upgraded the phone lines connecting the computing cluster to the airports where

The nation's airspace is an enormous, interdependent network, so what happens in one location can affect operations in a very wide area.

dedicated CIWS displays show air traffic managers their data.

When the nationwide system was unveiled, it did not include the precipitation phase part of the product suite since this wasn't needed in the summer. This feature was added in mid-November, just in time for winter. Dedicated CIWS displays have been installed in key air traffic facilities—primarily those in the heavily congested Northeast Corridor. But access to a website that provides the same information is available, upon request, to air traffic management personnel at facilities in other regions, as well as to the airlines. McKinney says air traffic managers around the country have been requesting accounts that give them access to the website, suggesting that they find the CIWS information useful.

Ducot says that while there are no measurements so far of the impact of nationwide coverage on air traffic operations, anecdotal evidence suggests that the upgraded system is helping to reduce delays. Over the next few years, Lincoln Laboratory will continue operating, monitoring, and upgrading CIWS, which is still considered a prototype.

In 2011, the FAA is planning to transition the system to the agency's William J. Hughes technical center in Atlantic City, N.J., with help from Lincoln Laboratory. At that point, CIWS will no longer be operated by the Laboratory as a prototype, but will become recognized as a formal part of the U.S. National Airspace System, which handles the flow of the nation's air traffic.

MICRO-OPTICS

Going Beyond Limits

New lens lets microscopes peer at much smaller objects

The optical microscope is a workhorse of biology, but it has its limits—specifically the diffraction limit, which says it can't resolve anything smaller than about half the wavelength of visible light. Now one Lincoln Laboratory researcher has developed a lens array that he says can increase the resolution of tiny objects by 400 percent, enough to study the innards of bacteria.

Such enhanced vision can allow people to quickly identify dangerous organisms, such as anthrax, by sight, instead of running cultures and waiting for the results. "You can see a lot more details, and you can better identify them," says Zong-Long Liau, a materials scientist and device physicist in the Laboratory's Electro-Optical Materials and Devices Group. "It will allow us to measure a biological sample in more like a three-dimensional way."

Liau has created tiny lenses that sit between the microscope objective and the object being studied. The lenses are made of gallium phosphide, which has a very high index of refraction of 3.5. If he places the sample directly against the lens at the microscope's focal point, that tiny area is magnified by a factor of the refractive index, in this case 3.5. As a result, the resolution of the image is now 350 percent of what it would be in air. Because the light is traveling through a high-index material, it slows down somewhat, essentially behaving as if it has a shorter wavelength and thus a smaller diffraction limit. Say you're looking at green light, which at roughly 550 nm is about in the middle of the visible spectrum. Its normal diffraction limit would be about 275 nm, but the gallium phosphide reduces that to about 80 nm—an improvement that makes smaller objects visible. Liau uses gallium phosphide because it is transparent at most wavelengths of visible light, whereas other high-index materials are transparent only in the infrared. Aside from the fact that optical microscopes rely on visible light, the longer wavelength infrared light would not achieve the same resolution, even when the wavelength is reduced to half its original value.

The concept is not entirely new. Scientists have long known that placing a tiny drop of oil between the microscope and the sample could increase resolution by 50 percent because of the oil's 1.5 index of refraction. In fact, it is by analogy with the oil drop method that this is called a "solid immersion lens." Though the sample is not truly immersed in the material, as it would be in the oil, it has to be pressed so close against the lens that it is as good as inside it.

Such microlenses have been considered before, but implemented in a reversed way. In the 1990s, Liau says, some researchers considered using such lenses to focus a laser spot down to a very small dimension; their goal was to

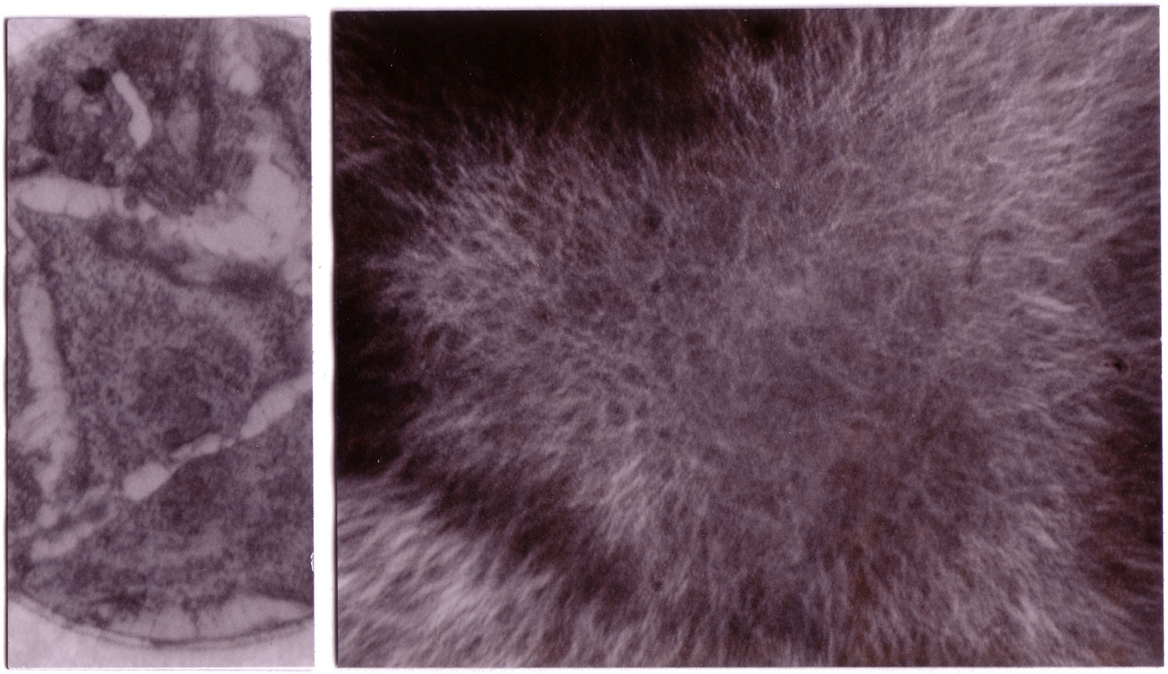

This comparison of micrographs of a cell shows the higher resolution made possible by the solid-immersion lens (right) over conventional microscopy (left). Both images were made using identical microscope objectives and magnifications. The cell is about 40 micrometers across.

increase the density of the markings that store data on a compact disc. That didn't work out very well, Liau says, because no one developed an easy way to manufacture the tiny lenses in quantity. Most microlenses are stamped out of low-index polymers, a process that is relatively simple. But these high-index semiconductor lenses have to be smooth regular hemispheres that measure only about 200 micrometers in diameter. Making them, Liau says, "is not entirely trivial."

James Leger, a professor of electrical and computer engineering at the University of Minnesota who worked with Liau at Lincoln Laboratory in the early 1990s, agrees that the fabrication process can be tough, especially if the maker wants to introduce any variation to the shapes. "If you want to make aspheric shapes or you want to make a bunch of hemispheres in a row, that's a lot harder to do," he says.

Liau's lab manufactures the lenses using photolithography and ion chemical etching, common procedures used for making computer microchips. But in chip making, the structures are flat. The lenses need to be hemispheric. One way to make spheres out of little bits of semiconductor material is to tumble them together, so that impact and friction smooth out their surfaces. But Liau says this process can create cracks within the material, rendering it useless as a lens. Moreover, he's not convinced that the grinding process produces a perfectly spherical shape.

Instead, Liau relies on a process he invented, and which the Laboratory patented. He uses photolithography and etching to carve gallium phosphide into an array of lens shapes, each very close to the final hemisphere shape he wants to end up with. Then he heats the array to 1050 C. That's not quite the melting point for the gallium phosphide, but it's enough to allow atoms on the surface to move around easily. Since a smooth surface represents a lower-energy state than a rough surface, the mobility lets the atoms rearrange themselves so the lens is smooth. "It sounds simple, but there is a lot of materials science, even physics, involved," Liau says. His plan is to put an array of thousands on a 1-centimeter chip. Such a chip could be used in a manner similar to a microscope slide, with a sample pressed against the backside of multiple lenses.

Leger says the processes Liau has developed might help the solid-immersion lens find wider use. "He has pioneered a lot of the fabrication techniques for some of these binary compounds," Leger says, referring to materials such as gallium phosphide that contain two elements.

One area of continued research will be on ways to speed up production, which right now requires very slow etching and uses special equipment built in Liau's lab. To commercialize the lenses, Liau will need to improve the throughput of the process.

Liau has come up with a further improvement to the lenses that increases resolution even more—to 400 percent better than in air. The key here is to alter the lenses' shape somewhat, in ways he doesn't want to describe until after he has filed for patent protection. To show how effective this is, Liau displays a pair of photographs of a cultured cell taken through a microscope. In the first, which used a conventional lens, the cell is a small triangular shape without much internal structure visible. In the second, it becomes a thick array of branches, the cytoskeleton that makes up the cell's structure. Organelles that perform functions inside the organism are visible.

There isn't much contrast between different structures within a cell, and even with Liau's lenses, the sample has to be stained to make the structures visible. But the contrast is improved enough by his lenses that he can see structures that wouldn't be visible through a conventional microscope lens without the use of complex fluorescence techniques.

Having demonstrated that the lenses work, Liau wants to start using them to see what he can learn about different bacteria, to show that valuable information can be gleaned. He also plans to apply to the Defense Advanced Research Projects Agency or the National Institutes of Health for funding for the research. It could take another two or three years to commercialize the lenses, he predicts. "This is a very strong lens, the strongest people ever produced," Liau says. "This is the first time that people have ever looked at things this way."

SENSING

Early-Warning Chemical Sensors Pass Indoor Tests

A relatively cheap sensor system can detect indoor chemical weapons attacks

Metal detectors and X-ray scanners may keep guns and bombs out of sensitive government buildings and military facilities, but they're powerless to stop the spread of airborne chemical weapons. So a team led by Juliette Seeley of Lincoln Laboratory's Sensor Technology and System Applications Group is working to convince people that a relatively simple and inexpensive sensing system can provide early warning of chemical attacks and allow building managers to respond appropriately.

"We're just trying to demonstrate that it's a good sensor that people should be using for indoor sensing," says Seeley, a physical chemist. Her method entails measuring the infrared light passing through a gas to identify its chemical makeup. Such long-wave infrared spectroscopy has been used outdoors by the Environmental Protection Agency (EPA) to monitor chemical plants and refineries for polluting emissions.

But Seeley's system differs from the EPA's. "If somebody's going to attack you with a chemical weapon, they're going to use a very high concentration," she points out. "So we don't have to be as sensitive and we can use cheaper technologies."

The system relies on an off-the-shelf infrared spectrometer, such as a laboratory might use for characterizing various substances. At the port where a container holding the test sample would normally go, Roshan Aggarwal of the Laser Technology and Applications Group built an optical system to pick up light coming from a source across the room. A collimated infrared beam from a filament travels along a straight line across one part of a room to a mirror, which directs it along another path to a periscope; the use of a periscope allows the positioning of the beam above the heads of people, so passers-by won't interrupt it. The telescope then collects the light and directs it to the spectrometer/detector coupled to a Fourier transform spectrometer. (This arrangement, with the sensor physically separated from a transmitter, is known as a bistatic system.)

By releasing sulfur hexafluoride gas in the Lincoln Laboratory lobby, Juliette Seeley (shown here with colleague Robert Ashworth) demonstrated that simple sensors could give early warning of chemical releases.

The resulting spectrum is used to identify any material that might have absorbed the light along its path, and a computer compares its findings to a library containing the chemical signatures of various materials, including chemical weapons agents. The signatures of all sorts of chemical weapons, such as the nerve agent sarin, blister agents such as mustard gas, as well as chemicals that cause bleeding and choking, are well known and can be stored in the computer. Chemical warfare agents have known signatures in the 8- to 12-micrometer wavelength range in which the system operates. The system is not as sensitive to biological agents, which have less distinctive signatures in that part of the spectrum and absorb less light. The system is also blind to a few toxic industrial chemicals. It's a simple matter, Seeley says, to tell the computer not to raise an alarm when it identifies a relatively benign substance such as that found in cleaning fluids. "We detect it, but we don't accuse it of being a threat," she says. (For details of the sensing system, see the article by Seeley and Jonathan Richardson, "Early Warning Chemical Sensing," *Lincoln Laboratory Journal*, vol. 17, no. 1, p. 85.)

To demonstrate their system, the researchers used Lincoln Laboratory itself as a testbed. Late one Saturday night, when they knew the lab would be nearly empty, they set up the system in the main lobby. As a test material, they used sulfur hexafluoride (SF_6), an inert gas used to check heating, ventilating, and air-conditioning (HVAC) systems for leaks. The team also rented point sensors that used the same technology (infrared spectroscopy) to detect chemicals, so they could compare the results. The spectrometer system detected every release the point sensors saw, and even picked up on one release that a point sensor missed. Similarly, the spectrometer detected the releases for much longer than did the point sensors, while the point sensors would have declared the area safe long before the area was actually clear.

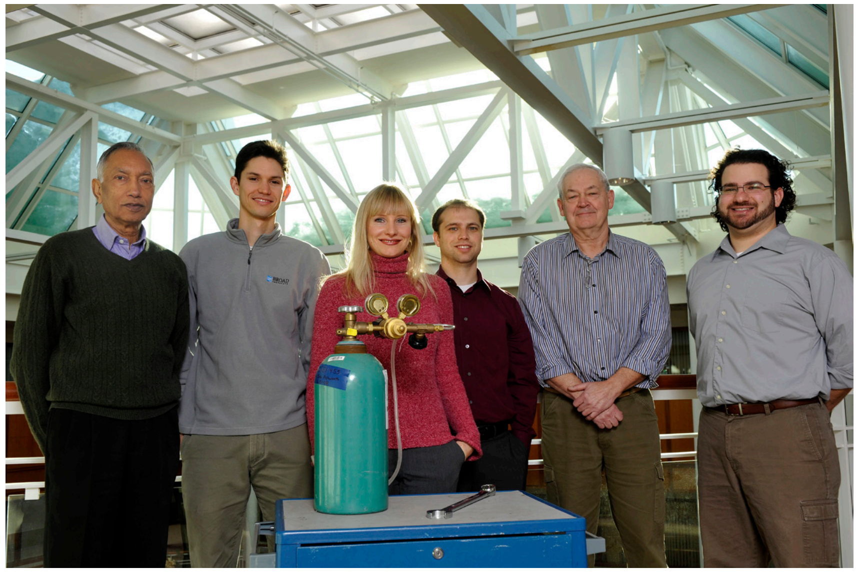

The chemical sensing experiment was conducted by (l. to r.) Roshan Aggarwal, John Aldridge, Juliette Seeley, Joseph Costa, Robert Ashworth, and Ira Mabel.

One worry was that the relatively heavy SF_6 might fall to the floor, below the infrared beam. As it turned out, that didn't happen. Seeley says normal air currents or diffusion was enough to stir up the chemical and send it into the path of the beam, even when it was released near the ground. She points out that a substance that sank to the floor and stayed there would not be a particularly effective weapon anyway.

The other question the researchers wanted to answer was whether detecting the chemical allowed them to respond effectively. "If you can't actually do anything by knowing about the attack, it doesn't help prevent injury," Seeley says.

So the team placed a point sensor in the auditorium on the level beneath the main lobby. They used the HVAC system to pump in extra air, so that air pressure inside the auditorium would be greater than outside. As long as they maintained that positive pressure, they detected no SF_6. When they turned off the HVAC, the chemical began to seep in. When they turned it back on, the level of the chemical dropped again. Those results, Seeley says, show that a building manager could provide some protection from a chemical attack through simple manipulation of the HVAC system. Another possibility would be to use the sensors to identify clear and contaminated parts of a building and thus prescribe an evacuation route.

If you can't actually do anything by knowing about the attack, it doesn't help prevent injury.

Seeley says the system her team developed can predict how the sensors would respond to an actual chemical attack. The government has data about chemical warfare agents' physical properties, which affect how the substances move through air, and their infrared absorption coefficients. Comparing those to the test data on SF_6 allows Seeley and her colleagues to create computer models of the system reacting to other chemicals. Sarin, for instance, is not as easy to detect as SF_6, but would still be captured by the sensor, the model predicts.

There are other types of chemical sensors based on other technologies that can also provide protection. But Seeley emphasizes some of the advantages of the bistatic infrared system. For instance, ion mobility spectrometry, a competing technology, identifies chemicals by measuring how long it takes them to move through an electrical field after ionization; as currently implemented, however, this technique is more susceptible to false positives. Cost is another issue. The point sensors Seeley rented can each cost about the same as her system, but it takes five or six of them to cover a whole room. Maintenance has also been a major obstacle with point sensors. Since the bistatic system is noncontact, the need for maintenance is greatly reduced.

Since completing these tests, Seeley's team has been looking for a company interested in developing the system commercially. Seeley hopes the extra measure of security the system could bring to government and military buildings makes it appealing. "We're trying to bring awareness that this is something useful," she says.

OPTICAL NETWORKS

A Fast Switch

Work on a high-speed optical switch could provide one component for future optical communications networks

Jade Wang's office is nothing fancy, just a cubicle in one corner of a lab within Lincoln Laboratory's Optical Communications Technology Group. But it's here that Wang, who just this year earned her Ph.D., is collaborating with electrical engineering professor Leslie Kolodziejski from MIT's Research Laboratory of Electronics on finding ways to speed up optical communications. Wang's particular contribution to this complex problem is a method that may help make optical switches practical.

As more and more information streams along the Internet, it strains the capacity of the fiber-optic network. One way to get more data through the system, of course, is to increase the bit rate—and indeed, the latest communications hardware can handle 40 gigabits per second, up from the current standard of 10 Gb/s.

But there's a problem. After coursing speedily through the optical fibers as pulses of infrared light, data transmission comes to a screeching halt when signals get to network nodes, where they have to be rerouted into a different fiber. The routers and switches at these nodes read the data packets' information to figure out their destination and send them along the right paths. To get that information, these devices need to perform logical operations on the data stream, and they must do it electronically—optical computing devices are not yet practical. Converting the optical data to an electronic signal, examining and rerouting it, and then converting it back into light takes time and requires additional equipment, slowing networks and driving up costs. "There's a bottleneck there," Wang says.

The optical switch that Wang is developing is based on interferometry: the incoming optical signal is split in two, and each component travels along a separate channel before recombining. If the signal in one of the channels is processed, it can speed up or slow down. Depending on how much the signal was changed, when it meets the signal from the other channel, it can interfere constructively, making the combined signal stronger, or destructively, making the combined signal weaker. If this merged signal is stronger, that counts as a digital "one;" if it is weaker, it's a "zero."

The signal-altering devices in these switches are semiconductor optical amplifiers (SOAs)—lengths of a semiconductor material, usually indium-gallium-arsenide-phosphide, that work much like a laser. Pumping an electrical current through the SOA produces charge carriers (negative electrons and positive holes). An optical pulse traveling through the electrically energized device causes the electrons and holes to recombine, emitting photons with the same characteristics as the original pulse, thereby amplifying it.

The SOA is made of nonlinear material—that is, the optical signal passing through it alters the material's index of refraction. So if a second pulse follows right on the heels of the first, it is slowed down or sped up by an amount proportional to the change in index. Therefore, Wang says, "I can control how much index change I get by controlling how much power I send in."

To create optical logic, the researchers build an interferometer with an SOA in each channel. A control pulse passes through one, altering its index of refraction. A signal pulse follows, passing through both channels and then recombining at the other end. The power of the pulses controls how much out of phase with one another the signal pulses are.

We've gone from a 12-by-18-inch board to an inch-long device that performs optical logic.

With the pulses adding up to ones and zeroes, the routers can now read the headers of data packets and decide where to send them without having to go through the electronic conversion step. Previous researchers have shown this design can perform switching at 80 Gb/s.

An SOA/interferometer optical switch was first demonstrated about a decade ago, but it was built out of discrete components on a

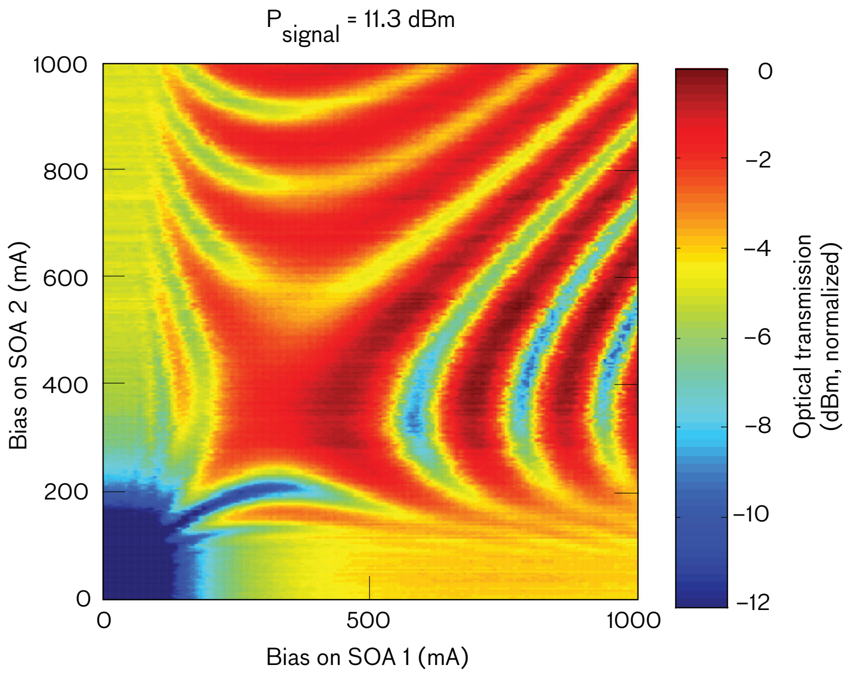

The color-coded map shows constructive and destructive interference as a function of bias current on the two semiconductor optical amplifiers (SOA 1 and SOA 2). Constructive interference between the two interfering light beams, shown in deep red, results in high optical transmission (a digital 1), while destructive interference, shown in blue, produces low optical transmission (a digital 0).

circuit board—much too large for integration into optical networking equipment. Wang and company are working on integrating the same design into a much smaller package, making it easier to use and less expensive to produce. "We've gone from a 12-by-18-inch board to an inch-long device that performs optical logic," Wang says. The next step is to put multiple logic gates on a single chip. One issue the researchers are trying to tackle is that the optical pulses lose power as they pass through the device. The amplification effect of the SOAs can compensate for that, but using too many SOAs adds complexity and introduces noise that can interfere with the signal. "There may be a way to reduce that loss by being clever about fabricating the chip," Wang says.

For instance, light moves around the device through a passive waveguide; the SOAs, by contrast, are active waveguides, altering the signal. Chip makers can fabricate the passive and active waveguides out of similar materials, and determine which is which by changing the mix of materials slightly. Wang uses Kolodziejski's design, which moves the light from the passive to the active waveguides by slowly tapering the end of the waveguide; the change in shape forces the light to move from one waveguide to the other.

Wang is also focusing on creating a streamlined method for checking the devices' performance. Because variations in fabrication produce devices that are slightly different from each other, each logic circuit has to be tested to determine what combination of current and signal timing gives the best match for constructive interference—the strongest "on" signal—and the best match for destructive interference—the strongest "off" signal. "The better we can turn things off compared to on, the better the switch operates," Wang explains. Wang has devised a way to measure the output power of the device at each current setting. Her technique produces a color-coded "bias map" of the device that shows the optimal current setting, or bias. Her method takes about 10 minutes, as opposed to the hours needed for earlier testing techniques. Better testing and better fabrication should eventually yield a reduction in costs that will help make these switches practical, she says. Wang's work, under the direction of her former thesis advisors Scott Hamilton (currently associate leader of the Optical Communications Technology Group) and Erich Ippen (a professor of both physics and electrical engineering at MIT) is part of the Defense Advanced Research Projects Agency's (DARPA) Data in the Optical Domain program.

According to DARPA, this program has, for the first time, shown a path toward a network that eliminates electrical-to-optical-to-electrical conversion at each router—a transformation that could potentially speed up transmission tenfold. Efforts like Wang's will help to "enable new optical networks that can meet the growing Department of Defense need for bandwidth while also minimizing latency," says DARPA spokesperson Jan Walker.

Wang doesn't want to project how long it might be before optical switches actually make their way into communications systems. For one thing, researchers still have to figure out how to get multigate logic, which would allow data to pass through the gates sequentially, on a single chip. With such a serial approach, it's especially important to reduce signal loss in each gate. "Things with optical logic have been limited to a few gates, so the functionality's not really there yet," she says. And complex optical computation will require some sort of optical memory, which does not yet exist. But she's hopeful that one day an all-optical network will eliminate the electronic bottlenecks and data will whizz around the Internet at speeds much greater than are now possible.

SPEECH PROCESSING

Dialect Detectives

Systems that distinguish among variants in spoken languages could enhance automated machine translation

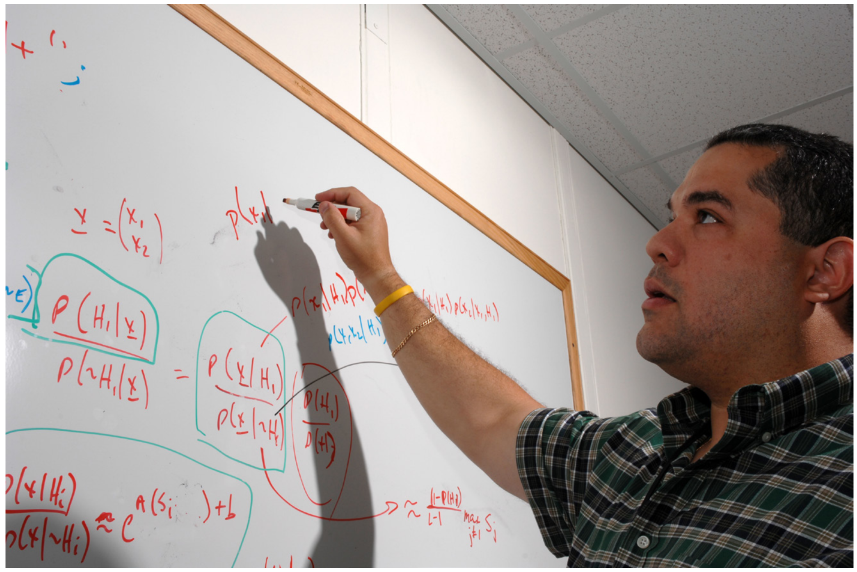

Pedro Torres-Carrasquillo is working on techniques for machine-based identification of dialects in a spoken language.

A law enforcement agency intercepts an international phone call alerting a suspected drug dealer to a new shipment. While the translator listening to the message is confident the caller's Spanish carries a South American accent, he cannot pinpoint a more specific region for agents to put under surveillance. But technology under development by Pedro Torres-Carrasquillo and his colleagues at Lincoln Laboratory may lead to a dialect identification system that compensates for a translator's inexperience with multiple variants of a spoken language.

Language identification systems that can recognize as many as 29 languages from written text are already marketed, and systems that can identify a spoken language from a prescribed range of choices also exist. So far, however, no system that automatically discriminates one spoken dialect from another is available.

Lincoln Laboratory's earlier work on dialect identification focused on building models that mapped the audiowave frequencies of phonemes—the individual sounds of a spoken language. Torres-Carrasquillo, an electrical engineer specializing in speech processing in the Laboratory's Information Systems Technology Group, says his group has recently moved from this phonetic-based approach to lower-level acoustic systems that use the basic spectral similarities of small pieces of spoken utterances. "We are not looking for the types of data linguists deal with—larger units such as phonemes and words," he says. "We're looking at the statistical distributions of basic frequency spectra of small pieces of sounds."

The Laboratory researchers are building a model that classifies the training data, finding markers that discriminate the frequency characteristics of the data. Previously, Torres-Carrasquillo says, the approach was to "get a lot of examples, and then build a model that looks like your examples." But he is tackling the problem in a different way. "Our group's idea is that we don't need a model that looks like our data—we need a model that can classify our data," he explains. "We take very small pieces—snippets of speech—turn them into frequencies, add up all these contributions, and make a model that can tell them apart. We're looking for patterns from just milliseconds of speech."

The researchers are using pattern recognition and classification methods known as support vector machines (SVMs) and Gaussian Mixture Models (GMMs) that use models trained to emphasize the more distinctive tiny features seen in the frequency patterns of small pieces of the dialects in question. The trained GMMs have the edge in accuracy, but SVMs are "an order of magnitude faster than GMMs," according to Torres-Carrasquillo. Even more effective than either SVMs or GMMs alone, he says, is the combination of the two techniques. In a test to discriminate general American English from Indian-accented English, for example, the error rate was 10% when GMM was used alone, 15% for SVM alone, and only 7% for a fusion of GMM and SVM.

To be incorporated into an automatic machine translation system, a dialect identification system would have to be able to recognize a dialect without having to process lengthy strings of speech data. Torres-Carrasquillo's goal is to be able to determine a speaker's dialect by categorizing discrete, characteristic markers in the snippet, and then create a model without having to use large sets of training data. "We'd love to see a short-term spectrum characteristic that is a strong discriminator, is very pervasive in the dialect, and that could be reliably detected in a sample," he says.

Finding this characteristic is a tall order. "You're not going to have a single spectrum characteristic that gives away the identification," Torres-Carrasquillo says. The linguistic differences between dialects of a language are often small; for example, vowel sounds in Cuban Spanish are slightly longer than those of Puerto Rican Spanish. The subtle differences between the spectral pictures of dialects are difficult to detect, especially in the milliseconds of speech used in the Laboratory experiments. "But as you look at the data" says Torres-Carrasquillo, "the differences start to pile up and you have a profile." The Laboratory's work to classify dialect differences may lead to the discovery of a strategy for any dialect problem—a global approach that could be exploited for various classes of dialects instead of a method that works only for specific dialects.

The Lincoln Laboratory research on dialect identification may contribute to approaches for language identification more generally, but Torres-Carrasquillo offers a caveat: "The differences one can exploit within two dialects are very specific—maybe too specific to be applicable to language ID." Still, when a universal machine translation system arrives on the scene in some future decade, it may well depend on Lincoln Laboratory research to ensure that nuances of meaning conveyed in dialects are not lost in translation.

Our group's idea is that we don't need a model that looks like our data—we need a model that can classify our data.

TACTICAL NETWORKING

Radio Wave of the Future

Modeling of communications protocols helps the military design next-generation software-defined radios

In 1997, the U.S. military set up the Joint Tactical Radio System (JTRS) to design the next generation of military communications. The core technology: radios that could simultaneously handle both voice and data transmissions and that could be upgraded to new communications capabilities with the simple addition of software. The multibillion-dollar program is now on the verge of fielding these radios throughout the military, thanks in part to work in Lincoln Laboratory's Communications and Information Technology Division, which is giving developers of the system early insight into how the radios will work in the field prior to their full deployment. The aim is to help the military develop practices and tactics for using these radios in combat and peacekeeping operations, as well as to pave the way for improvements in future generations of JTRS radios.

The idea behind JTRS was to replace the 25 to 30 different radio systems in use by the various branches of the military with fewer software-based radios that would operate over all of the military's radio-frequency spectrum allocations below 2 gigahertz. The radios

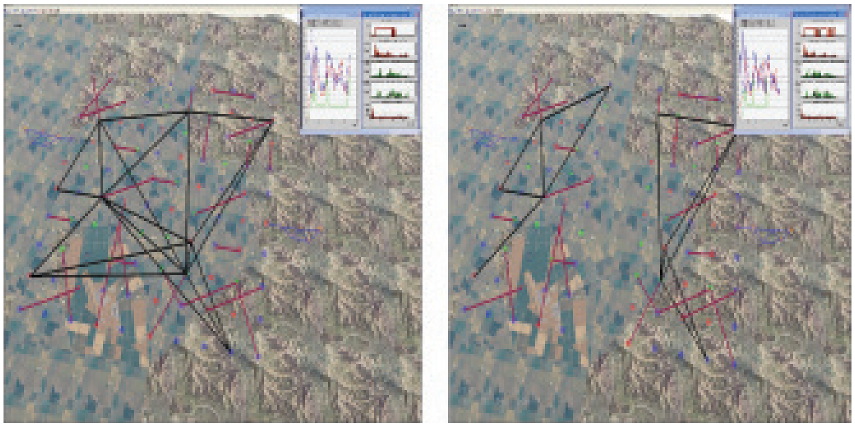

JTRS waveform protocols automatically adapt to changes in user locations and needs. The network on the right is partitioned into two halves where each segment is fully operational by itself. Once the two network partitions move toward each other, various JTRS protocols, including medium-access-control (MAC) and routing protocols, dynamically collaborate to form one fully connected network. The live simulation statistics depicted on the upper-right corner of the network topology screens illustrate various JTRS performance metrics (such as routing convergence, routing overhead, throughput, and latency) during the simulation.

would provide voice, video, or data communications. Each radio employs a complex set of protocol standards that govern their operation. In the parlance of JTRS, all the protocols that each radio must execute are called a waveform. The waveform software on each radio must run protocols to form networks, without access to fixed infrastructure such as cell towers or fiber-optic cables. Moreover, protocols must form these networks whether the radios are carried by planes in the air, vehicles on the road, or soldiers on foot. “The military needs anywhere, anytime connectivity,” says Tom Macdonald, leader of the Laboratory’s Wideband Tactical Networking Group.

In order to provide useful communication, each waveform’s core protocols perform many complex tasks without involving the radio operators. This ability of the waveform to hide the complexity of creating and maintaining communication links allows operators to focus on their military mission. The core protocols in each waveform spell out the rules of the electronic “handshake” by which one device connects to another. They predict rates of errors and apply correction or mitigation techniques to combat errors introduced by noise in the radio channel or from adversary jammers. They figure out how to form networks with a changing number of other moving radios. They determine the best path to reach each other radio and which data are the most important to send immediately. Because these are military communications, the protocols have to be secure—a requirement that can involve encrypting the transmissions. To optimize performance, the system’s designers need to know not only how well each individual protocol works, but also how all the different protocols interact. Complicating matters, the protocols’ behaviors affect the design of the devices themselves, from the shapes and size of the antennas, to the number of computer chips required, and to the type and life span of batteries.

One difficulty with figuring all this out, says Siamak Dastangoo, a specialist in communications and networking and the Laboratory’s principal investigator on the project, is that a variety of government contractors are responsible for different parts of the system. A large system integrator like Boeing might design the overall architecture and oversee contributions from other industrial partners, like ITT Corporation, which might build the devices; a specialty company like BBN Technologies might design some of the waveform’s networking protocols. And all this work may be going on simultaneously. “A lot of times the device is being designed as the protocols are being developed,” says Dastangoo. This concurrent development can be challenging both for the protocol designers who may not know what the hardware is capable of and for the hardware designers who may not have all the details of the protocols that their device must implement.

This is where Lincoln Laboratory comes in. Under the direction of the JTRS Program Office, Dastangoo and his colleagues have built software models of all the protocols used in one of the JTRS key waveforms, so they—as well as device manufacturers and protocol designers—can run simulations and see how they function. These models allow the development of a system-wide perspective of how all the protocols interact. The models facilitate easy capture of diagnostic information during the design process. They also provide an ability to upgrade each component or protocol separately.

It may seem as if such models wouldn't be needed and that it would be ideal to test the actual protocols on the actual hardware. However, that would take a room filled with racks of routers, computers, and other equipment—a setup that not only would be expensive but also would be difficult to transport, making it hard to share among the different contractors that are developing the different constituent protocols and hardware components. Furthermore, it would take many years before all the protocols and hardware were sufficiently mature to build up this capability, and, at that point, it would be more difficult to go back and change each individual protocol. Lincoln Laboratory's work helps look at large-scale end-to-end network performance much earlier than waiting until all the hardware and software are complete. The models distill the very complex protocols to their very essence, in a compact form that could be run on a laptop. Dastangoo compares the work to taking 1000 pages of text and cutting it to 50 pages that still tell the whole story.

Using a commercially available network-modeling environment called OPNET, Dastangoo and his team create streamlined behavioral models of the protocols to simulate how a network performs under different conditions. They can see how protocols written by different companies interact with others and identify potential trouble spots. They can add or disconnect communications links to mimic moving radios in complex environments like city streets or dense forests. They can try different types of network traffic—voice, then text, then video, for example. "We can go in and pinpoint where the problems are," Dastangoo says.

The military needs anywhere, anytime connectivity.... We are ideally positioned to bring together all the information from different partners.

To make sure their model is accurate, Dastangoo and colleagues use results from the actual protocols running on a few real radios to assess the fidelity of their computer models and tweak them where necessary. After this confirmation at small scale with real equipment, the simulations offer a cost-effective and simple way to investigate the performance of the much larger networks required in the field. Chris Burns of the Space and Naval Warfare Systems Center Pacific says this capability is an important outcome of the project. "That's the unique part where JTRS will really benefit going forward, to use models and simulation to save the government lots of money," he says. Burns is the modeling and simulation lead for the Airborne and Maritime/Fixed Station arm of JTRS, the group in charge of the radios that will be deployed on Army and Air Force aircraft and Navy ships.

A key benefit of the Lincoln Laboratory–generated models is that the government can freely distribute them to all the individual protocol and hardware designers so that everyone involved in the process can obtain more complete and more accurate information on the work being done by the other organizations. This early knowledge helps avoid problems when the different protocols and hardware devices are integrated. Dastangoo points out that Lincoln Laboratory was well suited to this role in providing early technical bridges between the different contractors because the Laboratory is an independent organization with a systems perspective. "We are ideally positioned to bring together all the information from different partners," he says.

The original project focused on one waveform, but having accomplished the modeling with the first set of protocols for that waveform, the team has been asked to tackle other waveforms. Dastangoo says it will take about two more years to develop and refine models for these additional JTRS waveforms. Lincoln Laboratory will then hand the models off to the government to maintain and update as needed.

Plotting the Laboratory's Technology Future

Q&A with Zachary Lemnios

Former Chief Technology Officer Zachary Lemnios spoke this spring with Lincoln Laboratory Journal *editor in chief Herb Brody about the Laboratory's strengths and challenges. Since this interview, Zach Lemnios has been appointed Director of Defense Research and Engineering.*

Lincoln Laboratory Journal: *How have the changes in the nature of the threats to U.S. security affected what Lincoln Laboratory does?*

Zach Lemnios: As the 9/11 attacks showed, a new set of asymmetric threats has emerged—threats that aren't well characterized and that are not necessarily driven by nation-states so it's hard to assign where they've come from. The United States doesn't have the tools or the experience and understanding of how to work through that.

LLJ: *What are the ramifications of that shift for technology development?*

Lemnios: The first issue is thinking through how an adversary uses commercially available technology. After all, we are a country that publishes freely and puts an awful lot up on the Web for anyone to read. We see that today in IEDs [improvised explosive devices]. The triggering mechanisms are exploitations of what's available commercially. We need to be thinking through how an adversary exploits open source understanding of technology and adapts it to their means.

LLJ: *Broadly speaking, what does Lincoln Laboratory contribute to this effort?*

Lemnios: The Laboratory has a keen role here—probably a unique role—in understanding how an adversary adapts and works in an environment that is entirely open, where the barriers to entry are very low, where the consequence is very high, and where the fingerprints are hidden. How do we understand that environment technically? How might we build countermeasures—tools to identify these threats or to mitigate their effects? The Lab does that very well because we are one of the few places where field testing, state-of-the-art development, prototyping, and an understanding of the application space and threat space come together.

LLJ: *What's the most important area of technology that you think the Lab ought to become more deeply involved in?*

Lemnios: One thing that worries me a lot is this: how do we understand the enormous amounts of data that the Department of Defense now collects on emerging and operating threats? How do we find salient features, such as targets and signatures that aren't well characterized? The second is how do we make this information more clearly accessible to the end user.

LLJ: *So this would be kind of data mining?*

Lemnios: Yes. And, in particular, social network analysis, which means not just looking for a single target but looking for networks of targets—people, places, or things that have some relationship. Rather than finding the single IED, for example, you want to look for the network of the bomb maker, the emplacement team, the trigger team, and the team that broadcasts the attack. This is a particularly urgent problem as our young men and women in service are under attack this very moment.

LLJ: *How might you go about doing that?*

Lemnios: There are many signatures—communications is one, imagery is another, and there may be cultural cues. Our understanding

Former Lincoln Laboratory CTO Zachary Lemnios says the Laboratory's expertise in language processing and social network analysis is especially valuable in a world of asymmetric threats.

of the Soviet Union was based upon a published set of operations and a published set of capabilities. They published it, we verified it, and we trained against it. We had a very-well-understood schema of how the Soviets would employ forces, what their strategic approach would be, what their tactical engagements might look like. This is a different world. The real challenge now is finding the indications and warnings and signatures of events when you don't have a lot of history on what that person or what that group might do. One of the challenges is trying to understand from this enormous data collection these very diverse set of sensors, how we cohere a picture of what's going on strategically and tactically. The Laboratory is developing the fundamentals in that space.

LLJ: *You also mentioned user interfaces. Why is that so important?*

Lemnios: We need to do a better job at presenting information so that the end user understands things in a natural way. A person looking at a photograph can pick out the salient features. We need computer systems that do that with sensor data. The human-machine interface hasn't yet scaled to the level of complexity of the systems that we're building. We need to be able to convey the information so that the user really understands it.

LLJ: *It seems you have defined two separate but interrelated problems: one, how to abstract out salient features from huge amounts of data; and two, how to then convey this higher level of abstraction to the end user. What are the specific approaches that Lincoln Laboratory is pursuing?*

Lemnios: The first question is, how do you build algorithms to look for salient features? This is all about building networked sensors, building a schema for storing data, data retrieval, putting in place Bayesian filters or hidden Markov model filters. There's a body of science that's helping us understand how to do that. The second piece is, once you've done that, how do you convey lots of information in a way that doesn't overload the user? While there are a few examples of good human-machine interface, the fundamentals are not well understood.

LLJ: *What specific capabilities would you like to see developed?*

Lemnios: Well, right now users have to learn about their computers. I'd like to see it the other way around: computers should learn about their users and adapt themselves accordingly. Why is it that when I sit down at my keyboard, I have to type the same things over and over again—login, password, pulling up certain screens? I do that every day. Why doesn't my system just do that for me? It knows that when I come in, the first thing I do is check e-mail, and I look at a particular website. The operating system should build a model of a user and tailor its interaction to the user. We do this socially all the time. When you have friends over for dinner and they are interested in a topic—gardening, say, or fine wine—you'll end up having a discussion about that.

LLJ: *Advanced information technology can also be used against us, though, right?*

Lemnios: Yes. In fact, the Laboratory has been working for many years to understand the cyber threat. A team of four can do enormous damage. An individual can basically shut down a city. It's a problem set that the Lab has resources to address, and we're starting to do so.

***LLJ:** On this, and other national security issues, what do you think sets Lincoln Laboratory apart from the many other institutions focusing on similar work?*
Lemnios: Well, we have very bright people, of course. But there are bright people everywhere. More uniquely, Lincoln Laboratory has access to the right data sets—that's the heritage of Lincoln Lab. A lot of other organizations are trying to solve these problems in the abstract. But here, we have access to user data and an understanding of threats that have been identified. We have access to not only the worm or virus itself but also to information as to how it was inserted into the system.

One thing that worries me a lot is this: how do we understand the enormous amounts of data that the Department of Defense now collects on emerging and operating threats?

***LLJ:** How would you characterize the Laboratory's place in the defense technology ecosystem?*
Lemnios: We have a dual role. On the one hand, we are a gold standard for data and analysis; that is, we provide the independent assessment of ideas and technologies developed by others. In other cases, we are a pioneer of new concepts. The Laboratory takes on projects that have enormous technology risk, far more than the private sector would take on.

***LLJ:** You are on R&D management side of things and spent several years at the Defense Advanced Research Projects Agency [DARPA], which funds a lot of cutting-edge technology development. If you could have all the scientists and engineers at the Laboratory understand something better than they do now, conceptually, what would it be?*
Lemnios: I'd like to see a more consistent ability to think beyond technology into the application space and into the countermeasures space. The Laboratory does that really well. That's because we live in field sites, we collect data, we do the data analysis, and we own the data sets. We have many people here doing analysis not just on the data that we collect, but also on systems that the United States operates and on systems that will counter our systems. So we're really thinking all the time about countermeasures.

***LLJ:** So what tends to guide the Laboratory's activities more—"technology push" or "applications pull"?*
Lemnios: There's both. Part of our job in this office is to balance the two—that is, to find the right connections between innovative technology and compelling applications. In some cases, a systems engineer will pose an application that absolutely demands a technology and we connect those up. But it's that intersection that is so rich. And there are a lot of people at the Lab who live in that intersection.

***LLJ:** What Laboratory programs do you think exemplify the Laboratory's main strengths in this regard?*
Lemnios: There are a few. One is the Air Force red team. Threat assessment, technology development, countermeasures assessment—that's exactly their game. The same is true for the work that we're doing in counter IED work. And then there's the ERSA [Enhanced Regional Situation Awareness] activity, which is all about threat assessment. You build a core capability, drive it with what's available commercially, add where you need to, but don't be overly aggressive about technology development. Another good example is a program we have for analyzing patterns of communications to help predict what actions a terrorist group is going to take. It's called CT-SNAIR, for counterterror social network analysis and intent recognition.

***LLJ:** How did this effort originate?*
Lemnios: Several years ago, when the Laboratory started building a counterterrorism program, the first thing that we did was take a close look at the threat network and how it might evolve, and try to understand where the vulnerabilities were in the kill chain. Take the threat of improvised explosive devices, or IEDs, for instance. While many organizations are looking at defeat-

ing the trigger of the IED, we have been doing a lot of work to try to understand the operation of the network of people that are responsible for IEDs. How are these networks structured, and how do they evolve? Where are the vulnerabilities in the process of designing, building, and placing IEDs, as well as in the filming and subsequent reporting of the attacks?

LLJ: *Isn't this what intelligence organizations do all the time? How is CT-SNAIR different?*
Lemnios: Conventional intelligence operations certainly do help identify people and where they are on the ground. The capability that is emerging—which CT-SNAIR is a part of—is to analyze the communication network amongst those who pose the threat so that we can better understand what they might be doing.

LLJ: *Even though the information could be in any number of languages?*
Lemnios: Yes. In fact, the ability to do natural language translation is an important piece. The Laboratory has a lot of strength in that technology. So language understanding, coupled with social networking analysis, is really what the CT-SNAIR effort is all about. CT-SNAIR uses tools that build upon the language processing work the Laboratory's been doing for many years. This effort couples well to work on MIT campus. We're also collaborating with Carnegie Mellon and with the University of Massachusetts, Amherst, in particular, on the front-end social network analysis tools. This effort entails more than just word spotting; we are also trying to identify phrases, disambiguate aliases, and understand cultural context, all of which are really hard research problems.

LLJ: *How far along is CT-SNAIR toward being a deployable system?*
Lemnios: It's not anywhere near close to that stage yet. At this point, we're trying to understand the underlying science of the problem and validate the algorithms that are being used.

LLJ: *What are some key areas with this system that need to be worked on?*
Lemnios: Well, right now the false-alarm rate is a big problem. Getting false-alarm rates down to acceptable levels is key. It won't be very useful if the system keeps triggering alerts in situations when the communications being analyzed are benign. There's a huge area here of research on how to get a system like this to do machine learning in very dynamic environments with unstructured data. This is an area that the Laboratory is very interested in.

LLJ: *Is the goal of CT-SNAIR more to monitor the activity of a known group or to identify new groups that we didn't previously know existed?*
Lemnios: For now, we are focusing on improving our understanding of existing networks. Eventually you'd like to be able to identify networks that aren't yet known, but this would give rise to enormous false-alarm rates. So at best I think these will be tools that cue an analyst rather than replace an analyst. And if we could provide some insight that the analyst hasn't yet seen, that would be a win.

LLJ: *One common theme to describe work at Lincoln Laboratory is that it would seem important for there to be in place a culture and a structure that encourages interaction among different disciplines.*
Lemnios: Yes, it is all about that. In fact, most of problems that we take on are so interdisciplinary that it's far more important to have a team with technical depth across a number of domains than a single individual researcher with twice as much depth in his or her own specialty.

LLJ: *How does the Laboratory foster this kind of innovation?*
Lemnios: We have a couple of mechanisms to help launch new ideas. The first is the Advanced Concepts Committee, or ACC. The ACC is all about finding a way to get the young, bright staff member with a glowing idea enough money to prove his or her thesis. This program gives a staff member enough money to run an experiment and collect the first data set. All the ACC work is internally funded, and the consequence of failure is low. We're basically betting on people and betting on good ideas.

LLJ: *So that's more at the component level. What about systems level thinking?*
Lemnios: That's where the other mechanism—the New Technology Initiatives Program [NTIP]—comes in. The NTIP is the first opportunity to start integrating a number of concepts into a systems picture. The

level of funding is higher than that of the ACC and the time horizon is closer. With an NTIP project, we assume that the enabling technologies are mature enough. The bet that we are placing on these efforts is the integration concept. We want to make sure that technical developments hang together from the system perspective. It's also the first time that we start thinking about measures and countermeasures. If somebody comes up with a new radio frequency tag, say, and wants to demonstrate it against some sort of aircraft, the NTIP might also look at what the countermeasures to that tag might be. The NTIP might think through what's the concept of operations. That is, how would you actually use this, and how does that compare with the existing concept?

LLJ: *What are some examples of NTIP successes?*
Lemnios: Well, one is CT-SNAIR, which we've already talked about. Another is MASIVS [Multi-Aperture Sparse Imager Video System], a system concept to build a very-wide-area imager from a number of commercial imagers. Whereas the ACC would be funding the actual device, the NTIP is looking at how we take that and start integrating it. It really takes a different mindset; we have different types of people in each of these two groups. The ACC largely funds work by people who love to spend time at the lab bench, whereas the NTIP projects are more for system thinkers. Both are absolutely vital.

LLJ: *Lincoln Laboratory has traditionally shunned the spotlight. When you talk to people in government, academia, and industry about the Laboratory, what's the most common misperception that you encounter?*
Lemnios: A lot of people confuse us with the Department of Energy labs—the national labs such as Los Alamos and Sandia. Of course, we are not a DOE lab. And in fact, I would argue that while those national labs are good in certain selective areas, Lincoln Laboratory has a deeper foundation in a broader set of topics that are of interest to the DoD—as well as the agility to take on new problems in lots of areas. Another way a lot of people have the wrong idea about Lincoln Laboratory is that they still regard us as a radar house. While that was the heritage of the Laboratory, we really are now an information technology organization in the sense that much of the work here is really about target identification, discrimination, and understanding, and the flow of that information across very complex systems. Radar is one sensor. We work with many other sensors.

LLJ: *What effect do you think the Laboratory will feel as the new administration takes the reins in Washington?*
Lemnios: We just don't know. I suspect that missile defense work might scale back a bit, but there will be other areas that increase. Homeland protection, for instance, is an area that I think we'll be seeing a lot more interest in. Counterterrorism work will grow. And we're starting a small effort and trying to understand where the Laboratory can make contributions in the energy space. Energy, after all, is now correctly viewed as a national security problem and so it is entirely consistent with the larger mission of Lincoln Laboratory.

LLJ: *What aspects of energy technology do you see the Laboratory focusing its efforts on?*
Lemnios: Well, you can divide up the energy space into two major pieces. First, there's a supply side that deals with fuels, distribution, and power conversion at the megawatt level. Then there is the consumer piece, which pertains to making sure that appliances and other energy-consuming systems are efficient in their use of electricity. I suspect most of the work we do will be on the usage piece—building low-power electronics, say. I could see our involvement in building systems that have power management schemes that go beyond what is available commercially, and maybe new technologies that allow for more efficient dc-to-dc conversion. We have also talked

Computers should learn about their users and adapt themselves accordingly.... The operating system should build a model of a user and tailor its interactions to the user.

with some professors at the University of California, Davis, about developing a smarter electrical grid that could take advantage of the energy storage and load-management that could be provided by a large fleet of plug-in electric hybrid vehicles. We are primarily interested in what the system architecture might look like. And we have a significant effort in collaboration with MIT campus to build a thermoelectric power generator (see "Power to Go," *Lincoln Laboratory Journal*, vol. 17, no. 1, p. 9). That's a good example of one of the first projects that puts the Lab in this space. The Solid State Division is doing a lot of work to extend the low-power operation of conventional CMOS and CMOS SOI [silicon-on-insulator] technology. These efforts could lead to much lower power consumption as we build systems with large numbers of processors, and processors with higher and higher transistor counts.

LLJ: *What impact do you expect the present economic crisis to have on Lincoln Laboratory?*
Lemnios: There's going to be a lot of budget pressure on all elements of DoD, including the Laboratory, to propose the best ideas.

LLJ: *The Laboratory is also less reliant on DoD than in the past, right?*
Lemnios: Yes, we've got a portfolio that's richer than we had a few years ago. We have a large number of sponsors and more of a diverse assortment of sponsors. We're moving into some new areas. We may thin down other areas within the Laboratory; missile defense, as I said, may be an area that has some sponsor pressure. During past downturns, we've been fortunate and we've had sponsors continue the work here. I suspect that's going to happen this time as well. Generally, we weather these kinds of things out; we try to reduce the overhead expenses as best we can. Now, as has been the case historically, we have had an over-commitment of sponsors—that is, we have more work that we can possibly do. But fundamentally, technology remains key to many of the problems that DoD has, and the Laboratory is in a unique position to provide a lot of those key technologies. That's not going to change.