Lab Notes

NEWS FROM AROUND LINCOLN LABORATORY

MICROFLUIDICS

Pump on a Chip

Electronically manipulated fluids could lead to implantable silicon devices

Jakub Kedzierski twists a valve on a tiny syringe holding some ordinary water. On a nearby computer screen, an image of a droplet forms, captured by a camera pointed at the syringe. The droplet balances between the needle's tip and a thin layer of a polymer underneath, a sphere of water. Kedzierski flips a switch and the bottom of the droplet, where it touches the polymer, instantly flattens like a balloon pressed against the table. When he adjusts some knobs and flips the switch again, the droplet flattens even more.

Kedzierski, an electrical engineer in Lincoln Laboratory's Advanced Silicon group, is demonstrating electrocapillary action. In this process (also known as electrowetting), changing the voltage in a capacitor changes how liquid interacts with a surface. He's hoping to use the phenomenon to build microfluidic labs-on-a-chip, in which tiny volumes of fluids can be moved around on a silicon chip without external pumps and valves. Such devices, he says, could be used for chemical sensing in environmental monitoring—say, as a safety check on a municipal water supply. They might also be useful for medical applications, perhaps in an implantable device that would provide continued monitoring of blood levels of hormones or insulin, or in rapid bedside diagnostics that would eliminate the wait for lab workups of blood samples.

A microfluidic lab-on-a-chip could move tiny volumes of fluids around on a silicon chip without the use of external pumps and valves.

Without some sort of on-chip pumping system, microfluidics chips need a system of tubes connected to outside pumps that vary pressure in the chip's fluid channels to move liquid around. That makes the device bulky and power hungry and unsuitable for implantation or for long-term remote use. The electrowetting technique should allow the creation of devices that are much smaller, and therefore don't use as much expensive reagent for chemical tests. They could also work much faster. Kedzierski says some researchers have attempted to use techniques developed for microelectromechanical systems to build on-chip pumps, but without much success: at the scale of tens of microns, surface tension is too powerful.

"They try to take a classic pump that uses discrete parts and do it on a small scale, but that's very difficult," he says.

The concept is fairly straightforward. An electrode is topped with a thin layer of an insulating polymer—Teflon works well—to form a capacitor. Place a droplet of some liquid on top of the polymer, and the droplet will press against the surface, flattening slightly where it touches. The flattening, or contact angle, depends on factors such as surface tension in the droplet and the electrochemical nature of the surface. Apply a voltage, and the droplet deforms in a repeatable, predictable way. "It will seek a different physical shape that will try to minimize its electrical free energy," Kedzierski says. He compares the action to squeezing a balloon between your hands: apply energy and the balloon flattens; let go and it reverts to its original roundness.

With this deformable droplet, Kedzierski now has a shape-changing material that can act as a miniature piston. One device he designed has a shallow channel to carry liquid, with deeper cavities inscribed in the channel's floor at regular intervals. Water sits on the bottom of the channel, with oil floating above it; electrodes sit beneath the floor of the channel. As Kedzierski switches the electrodes on and off, droplets of water balloon upward and shrink back, first on one side of the cavity and then on the other, while a flat lake of water in the cavity rises and falls. The result is a set of alternating pistons, creating a pump that propels the oil along. One fluid—in this case, the water—acts as a mechanical system pumping the other fluid—oil—around the chip. But the fluid to be moved could be something else, a blood sample, for instance, or drinking water, a variation that might call for a change in the pumping fluid as well. If the device designers want to keep the fluids from mixing, they can ensure that by using additives.

Kedzierski and fellow researcher Shaun Berry are studying the properties of various fluid mixes. The best way to control the deformation of droplets, Kedzierski says, is to add a surfactant, which changes surface tension. Plenty of surfactants are cheap and readily available; the researchers found several that work for their purposes, notably sodium dodecyl sulfate, used in many soaps and shampoos. Increasing or decreasing a concentration of salt also provides a way to manipulate the droplet's properties.

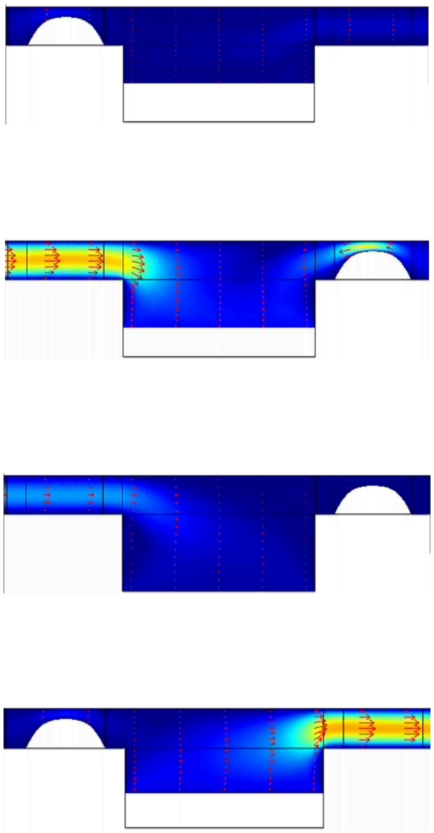

Alternately expanding and shrinking bubbles create a pumping action to move fluid around on a chip.

A few other researchers have taken similar approaches to his and have already started companies based on electrowetting. Varioptic, of Lyon, France, is using a similar idea to make tiny liquid lenses for cell phones. These lenses zoom and focus by deforming a liquid—and can do so much more quickly than is possible by physically moving a piece of glass in conventional optics. Advanced Liquid Logic of Research Triangle Park, N.C., is trying to develop handheld medical diagnostic devices that incorporate an electrowetting microfluidics chip developed at Duke University. Kedzierski says his device is much smaller and uses significantly less voltage—4 or 5 volts as compared to 15 to 100 for other devices. The lower power consumption that results means that a device can run for much longer on a small battery, making it suitable for implant or for constant environmental monitoring. His device can be as small as 20 micrometers on a side, compared to a few millimeters, and thus use much smaller samples—as little as 200 femtoliters—and less of the expensive reagents to test them with. Smaller samples can also be moved faster.

Having the basic concept in place means the next step will be to come up with a specific project and figure out what design and chemistry works best for that goal. For instance, if Kedzierski wants to make a remote sensor to monitor drinking water, he'll need to enlist a chemist to figure out what oil or surfactant can be used as the engine to move the water along without mixing with the water or interacting with possible contaminants. Moving blood would pose a similar challenge, but one with a different solution, because of blood's specific chemistry. He'll also have to incorporate detection equipment that's small enough to work in, say, a palm-sized device, perhaps using a tiny photodetector and laser. But he is confident that research dedicated to a specific project will pay off. "You just need to figure out what the different issues are," he says, "and engineer around them."

SEMICONDUCTORS

Small Packages

A novel way to squeeze compound semiconductors onto silicon chips could lead to smaller, more efficient microelectronics

---

If engineers could pack circuits made of normally incompatible materials together on the same silicon chip, it could lead to lower-cost radar modules, less power-hungry cell phones or other wireless devices, and smaller, cheaper radio frequency identification (RFID) tags. One Lincoln Laboratory researcher thinks he has come up with a way to integrate transistors made using complementary metal oxide semiconductor (CMOS) processes with other devices built from compound semiconductors, such as the III-V materials gallium arsenide or indium phosphide.

"CMOS is very good, especially for digital-logic, low-power applications, whereas a lot of III-V devices can have superior transistor speed and higher gain," says Jeremy Muldavin, an electrical engineer in the Laboratory's Advanced Silicon Technology group. "You'd like to be able to pick and choose which transistor technology you use for which function."

Muldavin has come up with a packaging scheme that he says will let engineers pick the very best materials for every application, perhaps using silicon to provide digital processing and III-V devices for amplification or radio frequency transmission. Such flexibility could lead to power efficiencies and smaller overall devices. He has developed a way to isolate individual transistors so that they can be placed in the most convenient location without creating electrical interference with neighboring transistors. The packaging makes it much more practical to integrate CMOS and III-V devices on the same silicon chip, which until now has been a difficult challenge.

A key problem to overcome is electromagnetic interference. An active component, such as a field-effect transistor, gives off an electromagnetic field and needs to be placed far enough away from similar components that the two fields won't interfere with each other. That limits how many active components a chip can have and how small the chip can get. "You could have a huge number of things in that area if you could find a way to isolate them," Muldavin says.

He achieves that isolation by using micromachining techniques to build cavities coated with a conductive metal, essentially creating little cans into which he can place active components. That allows him to add a few dozen III-V transistors made of different compound semiconductors onto a chip without wasted space and without electromagnetic interference. The design, he says, is good for systems with a fairly high density of CMOS devices and a low number of III-V devices.

He starts with a layer of bulk silicon, about 100 μm thick, on top of an insulator, and etches into this silicon a series of small, flat-bottomed wells about 90 μm deep. He then makes another layer with its own series of wells that he places of top of the first layer. He coats the surfaces of each piece with a conductive metal, preferably gold, although copper is also a possibility. He places the active component within the wells, then bonds the two pieces together, separated by a thin layer, perhaps 10 μm thick, of high-resistivity silicon. Now each component sits in its own little cavity, completely surrounded by metal. In certain spots, a tiny via pokes all the way through the silicon to the insulator, providing a conduit for the electrical connection. The layers of metal create what is essentially

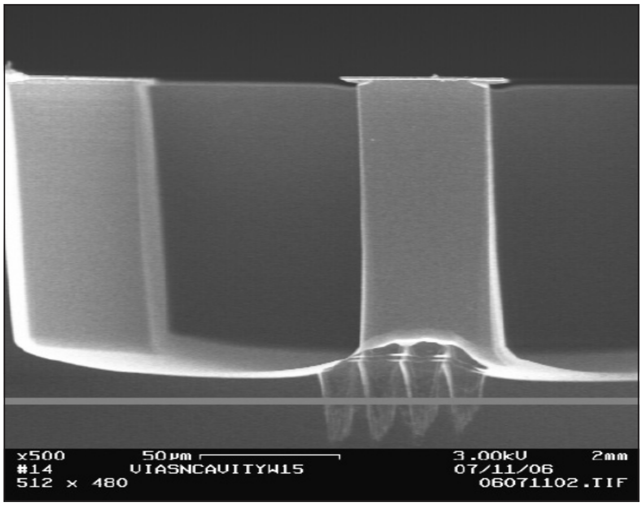

SEM cross section shows vias at the base of a cavity wafer. The drawn line represents buried oxide of 100-μm-thick silicon-on-insulator wafer.

a coaxial transmission line, carrying power to the components with little insertion loss while isolating the electromagnetic fields. The bulk silicon helps the III-V devices operate better by carrying away the excess heat they generate.

One previous approach some engineers have taken to mix and match dissimilar materials is to grow the III-V crystals on the same chip where the CMOS devices are built. Unfortunately, the lattice structures of the III-V crystals don't match very well with those of the silicon crystals, limiting the quality of the material. Also, many compound semiconductors can withstand temperatures up to only about 350°C, whereas silicon manufacturing processes require temperatures up to about 1000°C. Even if it works, this approach still requires placing active devices far enough apart to avoid electromagnetic interference.

Another approach is wafer-scale bonding: growing the III-V and CMOS devices on different chips, then bonding the two. But say the circuit design requires one III-V device on one end of the CMOS die and a second on the far side. That means the die with only two III-V devices needs to be the same size as the CMOS die. That way the devices line up in the right spot, even though the rest of the die is empty. "It's a little bit wasteful," says Muldavin. And that still doesn't solve the problem of combining, say, a gallium-nitride and a gallium-arsenide device on the same chip. Trying to grow them together leads to lower yields that drive up manufacturing costs.

Muldavin's design allows him to integrate different types of III-V components without the waste of wafer-scale bonding for die with mostly CMOS, and to space components closely without electromagnetic interference. With that new flexibility, he says, "you can just pick the best devices."

So when designing a system, an engineer could use CMOS for the digital control, and then add an RF transmitter built of gallium arsenide. Muldavin estimates that taking advantage of such efficiencies could create circuits that use 40 percent less power than those currently in use. Cell phones could run longer on the same battery. RFID could become small and cheap enough to be practical for a wider range of applications, creating a network of wireless sensors.

The design grew out of a packaging scheme for placing microelectromechanical systems (MEMS) on silicon-on-insulator substrates. That project involved etching similar cavities in a single layer of silicon on insulator to house the MEMS devices, then plating the back with gold to provide the electrical connection through the vias. This MEMS chip was essentially one half of the coaxial package Muldavin makes for the CMOS/III-V work; the new scheme adds a second etched layer on top to create the enclosed spaces for transistors. Before his work, he says, there had not been an inexpensive way to make a low-loss, hermetically sealed packaging system for MEMS. Such packaging allows MEMS devices to be designed into transmit/receive switches for radio frequency applications, such as cell phones.

Muldavin has received technology transfer funds from the Defense Advanced Research Projects Agency to transfer the work to Innovative Micro Technology, a foundry in California working on commercializing it. That same company might be an appropriate venue for the CMOS/III-V work as well, he says. He and his colleagues are finishing the fabrication of the three pieces that together become the CMOS/III-V chip package. They have to finish grinding some of the pieces to the thinness they're aiming for, carefully assemble them with gallium-arsenide transistors inside, and test the performance of the device—a set of milestones that he expects will be completed by late this year. The thinness of the silicon layers makes laboratory fabrication a delicate operation. A foundry, on the other hand, can use a different thinning process that should make the packages easier to assemble on an industrial scale—bringing the benefits of his technology to the wider world of microelectronics.

The packaging scheme makes it more practical to integrate silicon CMOS and III-V devices on the same silicon chip, potentially leading to smaller and more power-efficient devices.

ENERGY

Power to Go

Lincoln Lab–MIT collaboration challenges batteries with solar cells and a miniature burner

An army on the go needs a lot of electricity—to operate radios, run computers, power autonomous vehicles. Batteries, useful as they are, are not the ideal power source; they can pack just so much energy into a given amount of material. Besides, they can slow soldiers up, or weigh them down. “Either batteries have to be recharged, or you just have to carry a whole lot of batteries,” says Christine Wang, a materials scientist in Lincoln Laboratory’s Electrooptical Materials and Devices group.

Wang has teamed up with Klavs Jensen, the Lammot du Pont Professor of Chemical Engineering at MIT, to provide a new lightweight source of portable power. Their method: combine a tiny butane burner built at Jensen’s lab with an infrared-sensitive solar cell that Wang and her colleagues have developed. “You’re basically taking thermal energy and converting it to electrical energy,” Wang says. Wang is using her experience developing photovoltaic cells over 10 years as part of a Department of Energy project. DOE wanted to develop technology to convert thermal energy directly into electricity. To do so required development of thermophotovoltaic (TPV) cells that efficiently converted infrared radiation at longer wavelengths than those used for traditional solar cells.

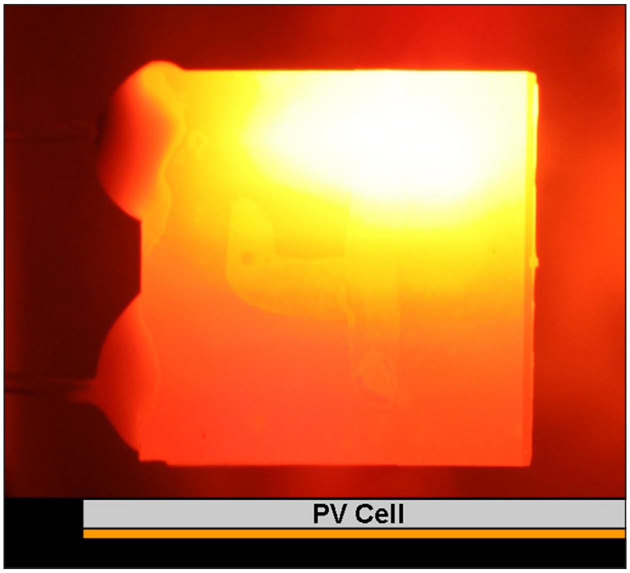

Combustion in a butane-fueled micro-reactor generates heat that a thermophotovoltaic device converts into electrical power.

This was no easy task. Standard photovoltaics that work with visible light are made of silicon. But the bandgap of silicon doesn’t match the energy levels of infrared photons. Under the DOE program, Wang’s lab developed photovoltaic cells made from the compound semiconductor gallium indium arsenic antimonide (GaInAsSb), which has the right bandgap. It’s a difficult material to work with, because with four elements involved it’s hard to produce a uniform alloy. However, Wang’s group managed to grow crystals that demonstrated near ideal performance. Wang says, though, that while gallium-antimonide-based cells are appealing, in practice indium-phosphide-based cells (also developed under the DOE program) are better candidates for practical uses. Because indium phosphide is widely used in optoelectronics applications, a manufacturing infrastructure already exists.

Simply placing the TPVs right near a large generator, however, wouldn’t work. “If you just took a solar cell and stuck it next to a burner, most of what would happen is you’d heat up the solar cell,” says Jensen. And heat causes a solar cell to break down. Wang and Jensen are now applying the TPV principle on a much smaller scale, where such thermal degradation isn’t so much of a threat. The researchers are collaborating to demonstrate an integrated micro-burner/TPV power generation device. Still, wresting the most electrical power possible from the micro-burner requires attention to the infrared spectrum that the burner emits. Wang and Jensen package the TPV with commercially available thin-film optical filters that allow short infrared wavelengths (less than 2.5 μm) to illuminate the cell while reflecting the longer-wave radiation back to the heat source. The reflected energy keeps the burner hot and helps it to use its fuel more efficiently. This technology is well suited for centimeter-scale power generation devices.

The device the researchers built is a sandwich—a square about 2 cm on a side and roughly 1 cm thick—of TPV cells and the micro-burner. The heart of the device is the burner, with two pipes leading into it to carry butane in and waste products out. On either side of the burner is a filter reflecting most of the heat and permitting the desired wavelengths (the ideal is about 2 μm) to pass through to the TPV cells. This electricity-generating cell, shielded by the filters and a thin vacuum, gets no hotter than 60°C.

The burner, designed by Brandon Blackwell, a graduate student in the Jensen laboratory, is a catalytic combustor. The inside is lined with alumina coated with platinum nanoparticles. In the lab, the burner is ignited by passing oxygen and hydrogen through it. The platinum catalyzes the reaction of the two gases, heating up the device until it reaches a temperature at which the butane will combust. A commercial device would probably heat the burner to its ignition temperature electrically, although the catalyst would still be needed to make the butane reaction efficient.

Jensen's laboratory at MIT is still working on perfecting the design of the combustor. The biggest challenge is controlling where the heat goes. "It's a complicated problem," he says, and tackling it requires computer modeling of heat flow. For instance, it's important for device efficiency that the burner heat up evenly, to make maximum use of the energy produced by combustion. He wants to add cooling fins to disperse excess heat, so that the outside of the device doesn't get too hot to handle. At the same time, he's trying to make the device hotter on the inside. Right now it gets to about 800°C; he'd like it to reach 900°C.

This increase has two advantages. First, as the device gets hotter, most of its radiation moves toward the shorter wavelengths that the TPV cells can convert to electricity. Moreover, the radiation output increases as the fourth power of temperature, so a hotter burner is more efficient at converting fuel to electricity. Unfortunately, the more intense heat is also more apt to destroy the packaging materials holding the various pieces of the device together. Jensen's group is trying to redesign the package to circumvent that problem. He is also experimenting with catalysts; the more efficient the catalyst, the more fuel is burned. A commercial device would probably use a material less expensive and more stable than platinum.

How efficient does the device need to be to make it competitive with existing micropower technology? "We would consider a 5 percent overall efficiency to be a success," Jensen says, "because at that point we would have beaten everything out there." Butane, he explains, provides about 13,000 watt-hours per kilogram. Five percent of that would be 650 W-hr/kg —more than triple the 200 W-hr/kg that the best battery provides. At 5% efficiency, then, 300 milliliters of butane (about the amount in a typical lighter) would provide enough energy for 50 hours of talk time on a BlackBerry—ten times longer than what a high-performance lithium-ion battery supplies.

Moreover, butane should prove no burden to soldiers on the move. "It would be a lot easier for them to be carrying a can of fuel than it would be to carry the equivalent amount of power in batteries," Wang says. And the military is used to carrying fuel.

Even at 5% efficiency, the thermophotovoltaic device could provide 50 hours of BlackBerry talk time on 300 milliliters of butane—about the amount contained in a typical lighter.

The device would be ideal for autonomous systems because it is lighter weight and has a longer lifetime than batteries. Unmanned aerial vehicles could fly farther and carry more power with less weight. Sensor networks could run longer on their own with a supply of butane than they could on batteries, which need to be changed sooner. The micro-burner/TPV device faces competition from micro fuel cells. But Jensen says that technology is still under development as well, and is complicated by the extra step of converting hydrocarbons to the hydrogen that a fuel cell consumes to generate power.

Wang and Jensen say it's less obvious whether their device would be useful in consumer products, mostly because it will probably always be too expensive to easily compete with batteries, except in niche applications. Moreover, says Jensen, the burner is still too hot to put in a laptop computer. But it might be useful in automotive applications as well as for other consumer uses, such as a portable battery charger.

The researchers expect that they'll be able to demonstrate their device by the end of the year. After that, they will look for a sponsor who wants them to develop the technology for specific applications in an increasingly mobile and power-hungry world.

PHOTONICS

A Little Light Work

Getting silicon to respond to light could usher in a new era of integrated optical devices

Michael Geis stands in a cramped lab inside the Lincoln Laboratory complex, having borrowed half an hour on another researcher's laser. Geis, a physicist in the Laboratory's Submicrometer Technology group, needs this special laser to put one of his devices through its paces—devices that will perform optical functions in silicon that have long been limited to compound semiconductors such as gallium arsenide.

Already Geis and his colleagues have made an all-silicon photodetector that's better at capturing light than any detector anyone else has built. But with researchers at labs from Intel to MIT to Cornell working on creating photonics components in silicon, it's a bit of a horse race, the researchers say. The Lincoln Laboratory group is one of many research groups worldwide building highly integrated optical devices, analogous to integrated electronic devices made in silicon. Geis, Steven Spector, and Jung Yoon, physicists working on the project, are funded by the Defense Advanced Research Projects Agency (DARPA), which is promoting the development of such integrated silicon circuits.

Much digital data travels through fiber optics at the speed of light. Electronic data processing, however, is still stuck in the sand—or, more precisely, in silicon, the material in which cheap and powerful transistors are made. Scientists around the world are working to bring together the power of light—its speed, its lack of electrical interference—with the inexpensive and reliable processes that have made computer chips so successful. Such a technological blend could lead to electrooptic chips with the high performance and low cost of today's complementary metal oxide semiconductor (CMOS) chips.

If silicon photodetectors prove to be as fast as they are responsive, reliable, and uniform, they could very well replace InGaAs photodiodes.

Geis, Spector, Yoon, and their colleagues in the Laboratory's Submicrometer Technology group, along with Matthew Grein and Robert Schulein in the Optical Communications group, have made great strides in combining the two. In particular, they have built devices out of silicon that detect, filter, and modulate light waves—devices that previously could be made only out of other materials. The Lincoln Laboratory groups are collaborating with MIT electrical engineering professor Franz Kaertner's group to show the Department of Defense they can build integrated circuits that can process optical signals. If they can make optical components out of silicon, instead of out of the compound semiconductors used in today's lasers and photodetectors, they can easily integrate the components with the digital electronic devices built on silicon chips. Such a development would allow data to be ferried across chips not by metal conductors but by beams of light.

The trouble is that silicon and optics don't mix well. At standard telecommunication wavelengths, from 1310 to 1550 nm, silicon is transparent. This quality can lead to good waveguides and filters but makes building devices such as detectors more difficult. In long-distance communications, compound semiconductors can bridge the gap between electronics and photonics, but devices made from these materials are too bulky and expensive to work on the scale of a chip.

The Lincoln Laboratory group

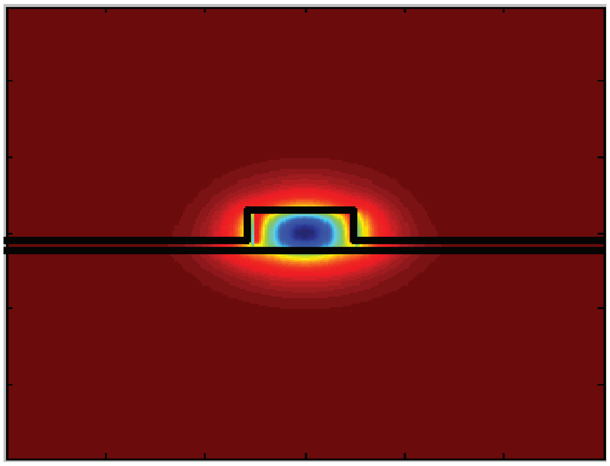

Most of the light in the silicon optical modulator passes through the center of the waveguide.

has recently made a major step toward the fabrication of practical silicon photonics: they built a silicon-based photodiode sensitive at telecom wavelengths. To achieve this goal, they took advantage of the fact that defects in the atomic structure of silicon change how it interacts with light. “You go from something that is very transparent to something that is absorbing,” says Spector. Although this phenomenon has been known for 50 years, it has been mostly ignored for silicon because it was widely regarded as useless.

Their photodiode consists of a p-i-n waveguide (a strip of undoped, intrinsic silicon sandwiched between regions of n-type and p-type material) that is connected to aluminum conductors at its edges. Because light traveling along the waveguide changes how current flows through it, measuring current provides an indication of how much light is passing. The team discovered they could control the level of optical absorption by controlling the number of defects in the silicon. Their fabrication process uses standard doping techniques, implanting the silicon crystal with silicon ions. The interaction of silicon with silicon creates two kinds of defects—gaps, or vacancies, where an incoming atom knocks an already present atom out of place, and clusters where new atoms and dislocated atoms clump together. They researchers found that annealing the crystal resulted in cluster formation, which causes even greater sensitivity to light.

“As you heat it up, these defects reassemble themselves,” Geis says. Above about 350°C, the silicon atoms move around and fill up the vacancies, so the gaps go away. But the other defects (the ones consisting of clusters of extra atoms) remain even when the material is heated to 475°C. Use of higher temperatures in fabricating the CMOS chips may eliminate these defects. However, Geis says that there is a dearth of information as to how well the defects would affect the absorption—or whether the changes they induced would stand up to the silicon manufacturing process. “We’re the first to see this defect because nobody ever did the experiment before,” Geis says. “Nobody thought it would work.”

The Lincoln Laboratory group is not the first to try doping silicon to get an optical response at telecom wavelengths. A team at McMaster University built a similar diode in 2003, but that device had much lower responsivity—too low to work as a practical photodetector. Geis and Spector improved on the earlier device by shrinking the waveguide tenfold (from roughly 5 μm across to about 0.5 μm) and by optimizing the defect implantation.

At 1550 nm, their device produces 0.8 amps of current per watt of optical input—roughly the same responsivity as the best commercially available high-frequency indium-gallium-arsenide photodiodes provide. The silicon devices also have lower leakage current, a measure of the electronic noise that reduces the amount of signal a device can handle: 0.2 nanoamps as compared to 1 to 50 nanoamps for commercial devices. The silicon device has a bandwidth of 10 to 20 GHz, which is considerable, but slower than the InGaAs device’s 50 GHz. If silicon photodetectors can be shown to be faster, as well as reliable and uniform, they may replace InGaAs photodiodes.

On the way to an all-silicon system, the Lincoln Laboratory team has been exploring devices that mix silicon and other materials, which they’ll use to build an integrated silicon photonic circuit. They have collaborated with MIT electrical engineering professor Judy Hoyt on a photodetector that combines silicon and germanium. Germanium is much better than silicon at detecting light in telecommunications systems because its bandgap corresponds perfectly to the 1550 nm wavelength light that optical fiber transmits best. But integrating germanium with silicon is difficult. For one thing, the crystal lattices of the two materials don’t match. Moreover, germanium deposition requires high temperatures during manufacturing, and the current devices produce more electrical noise than the InGaAs photodiodes. In the long run, fabrication simplicity favors an all-silicon process, but optimized germanium devices may offer better performance. Ultimately, the choice may depend on the application.

One device that will be essential for an integrated photonic circuit will be a silicon modulator to imprint a signal onto a light beam. The Lincoln Laboratory researchers, in collaboration with Kaertner's group at MIT, designed a silicon modulator that's very similar to their photodetector. The modulator is based on silicon waveguides in a standard Mach-Zehnder interferometer design. The waveguides split the beam in two, and an electrical charge alters the refractive index of one arm, so that when the beams merge again on the other side, they interfere to produce the signal. This design improves on a silicon modulator developed by Intel because it operates at only 0.02 volts per centimeter of the device's length, whereas Intel's uses about 16 volts per centimeter.

Achieving sufficient voltage in as short a length as possible is essential in shrinking the modulator onto a chip. The Lincoln Lab device got its better numbers by working in forward electrical bias, as opposed to Intel's reverse bias. However, this approach causes the device to operate more slowly, a problem the researchers say they are working on. (Researchers at Cornell University have also demonstrated a design for a forward-bias device, offering different advantages and disadvantages from those of the Lincoln Laboratory approach.) In a related project, the Lincoln Lab team is developing tunable silicon filters to split different wavelengths of light into separate channels, allowing multiple signals to be sent simultaneously through a single fiber. Though such filters can be made with other materials—MIT electrical engineering professor Hank Smith, for example, is working on silicon-nitride filters—all-silicon filters offer a more direct path to a single-chip silicon system. The DARPA project requires the team to demonstrate a modulator, photodetector, and filter in a system by the end of the year.

That system won't use all-silicon devices. It will, however, include the silicon-germanium detector and the silicon-nitride filters, which are further along in development. Still, the ultimate aim is to create an all-silicon circuit. Grein is leading the Laboratory's effort to demonstrate an optically sampled analog-to-digital converter that uses these integrated optical components along with commercial devices.

The researchers are still trying to figure out the best ways to build the modulators and detectors, and are working through the trade-offs that optimizing any device entails. For instance, the photodiode is more sensitive when it's bigger, absorbing about 99 percent of light when it's 3 mm long. Increases in size produce more noise, of course, as well as lowering the frequency response. The researchers would like the photodiode to be only 0.25 mm long to fit on a chip, but such a small device length means less light absorption.

Spector says such device optimization will improve the performance of the modulator and detector by two- or threefold. The filters, he says, still need to be demonstrated. Although some engineering hiccups remain to be fixed in device fabrication, he's optimistic that the problems will be worked out and that he and his colleagues will build an integrated device for processing optical signals. "The hope is to do it all in one chip," he says, "and if the filters are there, we'll be able to do it."

SPEECH PROCESSING

No More Babel?

Exploiting word commonalities could accelerate development of automatic translators

Despite tens of millions of dollars invested by the Department of Defense on automatic translation systems, Americans stationed in Iraq still rely almost entirely on human translators to communicate with the Arab-speaking population. The slow deployment of translation devices stands as evidence of just how difficult and time consuming it is to develop a technology that can reliably convert speech from one language to another. Wade Shen and Brian Delaney, electrical engineers in Lincoln Laboratory's Information Systems Technology group,

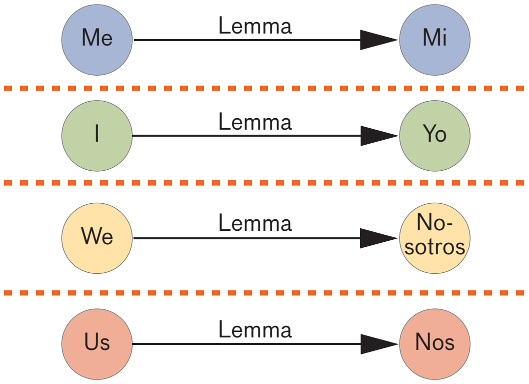

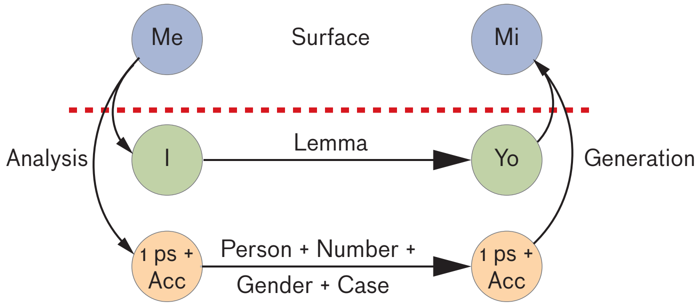

Morphology-based speech analysis enables machine translation with a smaller database of sentences. Without using morphology, me translates into the Spanish mi, but there is no built-in correlation to match I with yo. Using Shen's methods, the base word I translates to yo, and all other correlated words (me, my, we, us, our) would be translated through their person (ps) and accusative cases.

are leading work that could significantly shorten the time needed to get automatic translation capabilities into the field. Eventually, efforts such as Shen's could make possible a machine that emulates the Babel fish—the fanciful creature, described in Douglas Adams's *Hitchhiker's Guide to the Galaxy* series, that translates in real time between any two spoken languages.

Shen says that the machine translators available today have limited utility. They are adequate, he says, for enabling monolingual individuals to carry out simple exchanges in a foreign tongue—for example, asking for directions or ordering from a menu. Such systems, he says, "are typically most useful in communications requiring a very limited vocabulary." Today's technology relies on statistical methods to learn word-for-word and phrase-for-phrase translations; the translation system functions as an enormous automated phrase book. These systems need large quantities of training data—the set of sentence pairs in a language that is available for translation—and they lack breadth of vocabulary and grammatical sophistication. Furthermore, such systems are imperfect in producing translated text with words in the proper order, a deficiency that Shen and his colleagues plan to address.

Lincoln Laboratory's focus is defined by the following hypothetical scenario, as laid out by the Department of Defense: A conflict has erupted in a country where the U.S. government has no human or machine translators. There are, however, some number of sentences and phrases in the country's language that have been translated. Given this situation, how quickly can this linguistic information be used to create automatic translation devices?

Shen, working under an Air Force contract, is chiefly aiming to reduce the large volumes of training data needed while he still maintains an automated large-vocabulary phrase book. In the end, his task boils down to statistics. How can you verify that a word or phrase in one language has been translated correctly into another? Shen is applying sophisticated grammatical connections between words and phrases to produce translation algorithms that use a greatly reduced quantity of training data. Consider, Shen says, the concept of plural nouns: "Even if you have a limited sample of translated sentences, if you see the word mine in English and you see mina in Spanish in sufficient numbers, you can be fairly certain that mine translates to mina." In most cases, plurals are treated separately—an additional set of training phrases are required to match ours to el nuestros. "We want to eliminate this problem by using much smaller amounts of data in a smarter way," Shen says.

As part of an international effort involving sites in Italy, Britain, and the Czech Republic, as well as in the United States, Shen and his colleagues have devised models that attempt to make better use of this linguistic information. In these models, words are decomposed into morphological parts (e.g., houses = house + [plural], running = run + [progressive]). Each part is treated separately during translation. Next, a smaller database of translations collects richer statistics, since concurrent phrases or terms can be combined across different inflected forms (i.e., ran, run, and running all contribute to the count of the base form run). During translation, the input words are decomposed in the source language into parts. These parts are translated separately, and then recombined in the target language. The linguistic knowledge needed for the decomposition/combination processes can be learned by using statistical methods with much less data than would be required for all possible inflected forms.

When tested with a training set of only 40,000 sentence pairs, morphology-based translation approached the performance of standard systems trained on nearly 20 times the data.

For validation, Shen ran experiments by building morphology-based systems that translated English to Spanish that were trained with transcripts of European Parliament sessions. These transcripts contained some 40,000 sentence pairs, or only about 5% of a normal training set. He then compared the performance of standard translation methods that make no use of morphology—in one case using the same 40,000 sentence pairs, in another case using 750,000 sentence pairs. Shen found that the morphology-based translation approaches the performance of standard systems trained on nearly 20 times the data. The validation confirms that his translation algorithms work as effectively as he expected.

Many other linguistic generalizations could also be applied toward the translation problem. Shen's current effort, for example, attempts to incorporate syntactic information in addition to word morphology. In doing so, Shen and his associates hope to be able to resolve grammatical problems that arise from poor word ordering. This is a particularly tricky problem when translating from languages such as Japanese, where word order is relatively unconstrained. The team is extending its work in evaluation tools in preparation for the National Institute of Standards and Technology evaluation machine translation technologies for Less-Commonly Taught Languages (LCTL), to be held this fall.

Even with his advances, Shen says, we are still far from solving the machine translation problem. "We're still at the beginning stages of effectively utilizing linguistic knowledge in the translation process," he says. Asked about the possibility of a babelfish-like personal headset that performs instantaneous speech-to-speech translation, Shen is cautious. "In fifty years, we'll have something like that," he says. "Maybe."

SENSORS

Detector Avoids Speed Traps

Material manipulation yields avalanche photodiodes sensitive at 1.9 to 2.5 μm

How many layers of white paint do you need to cover a black wall? A single very thick coating will most likely do the job, but it may look streaky and uneven. A better question would be how many layers you need to change the black wall into a specific shade of gray. Now you need to proceed slowly with thin layers of white, and evaluate the results after each coating. Christine Wang, a materials scientist in Lincoln Laboratory's Electrooptical Materials and Devices group, isn't painting a black surface. Instead, she is tweaking crystalline layers to create her own version of a perfect gray—a series of lattice-mismatched buffer layers that transform a substrate lattice dimension into the base for a photodiode. And the result, if she succeeds, will be avalanche photodiode detectors (APD) for the short-wavelength infrared (SWIR) region from 1.9 to 2.5 μm. The major advantage of APDs over other infrared detectors is their greater sensitivity to light.

Wang and her colleagues, Erik Duerr and Joseph Donnelly, designed, modeled, and tested the APDs to optimize the three-active-layer structure with high

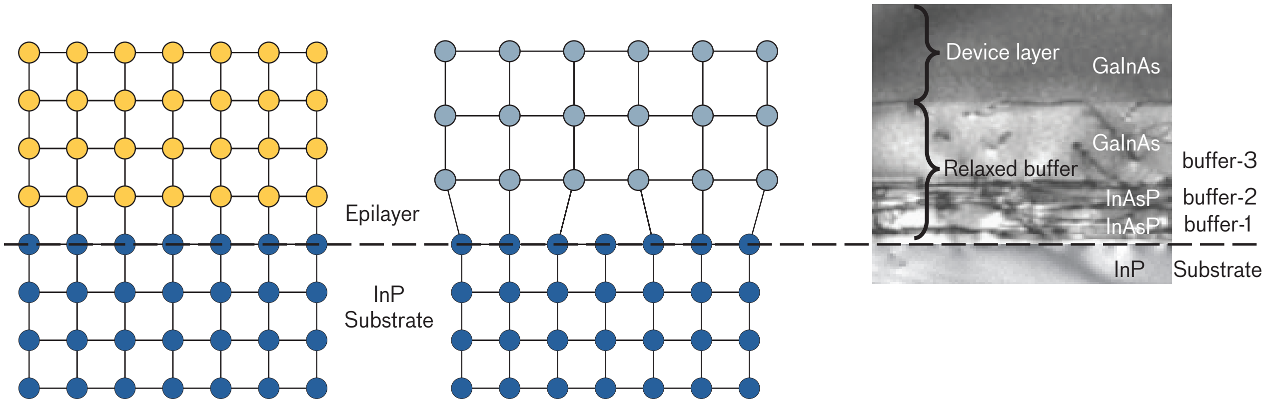

Materials whose lattices match the substrates can be deposited with no strain relief (left). Individual linear defects can accommodate the strain between slightly mismatched layers (middle). An electron micrograph shows a cross section of a series of mismatched layers (right). Note the strain relief and the absence of defects in the final buffer layer and the device layer. (The electron microscope image is from Wanlass et al., AIP Conf. Proc. No. 738, p. 427, 2004).

internal electric fields. Their aim was to ensure that the devices operate in Geiger mode—that is, past its breakdown voltage. In this mode, the detector delivers a large output signal even when struck by extremely low-intensity illumination. A photon hitting the active region of the APD generates an electron and a hole. Then the electric field drives the holes into a high-field multiplier layer, where impact ionization creates additional electron-hole pairs. The result is a cascade, or avalanche, of charge carriers. The high sensitivity of APDs makes them useful for photon counting or single-photon detection.

"The novelty of this work," Wang explains, "is that we are working on extending the detection wavelength beyond conventional InP-based APDs." Today's best high-speed APDs are sensitive only to wavelengths shorter than 1.6 μm. But there is a wide window of atmospheric transmittance between 1.9 and 2.5 μm. Photon-counting detectors operating in this spectral region would allow low-light laser radar (ladar); they may also be useful in optical communications. In addition, chemical and biological agents have signatures in the SWIR. Wang says that while existing gallium-arsenide APDs already work in this region of the SWIR spectrum, those devices "suffer from very slow reset times"—that is, the time it takes between measurable signals. Since her devices, unlike the gallium-antimonide ones, appear to be free of the material defects that trap charge carriers, they exhibit reset times of tens of microseconds—nearly a thousand times faster than the GaSb-based APDs. Defects, which act as electron/hole traps, also affect the rates of detected "signals" when there is no actually input—the dark count. Wang's materials have dark counts comparable to those of lattice-matched GaSb-based APDs.

Imaging arrays sensitive at wavelengths around 2 μm could serve in applications ranging from optical communications to land-mine detection to discrimination of bio-aerosols.

A few years ago, Wang explains, Bernadette Johnson (leader of the Laboratory's Biodefense Systems group) indicated that she was unable to locate any APDs that worked in the 2 to 2.5 μm region. Still, Johnson wasn't sure of their potential usefulness. "If you don't have it, you don't know how effective it will be," Wang says. However, she adds, "if we build it, they will use it."

Johnson is now actively interested in Wang's results. She says that analysis of existing hyperspectral imaging data as well as measurements with extended InGaAs detectors "have identified some promising applications for sensitive

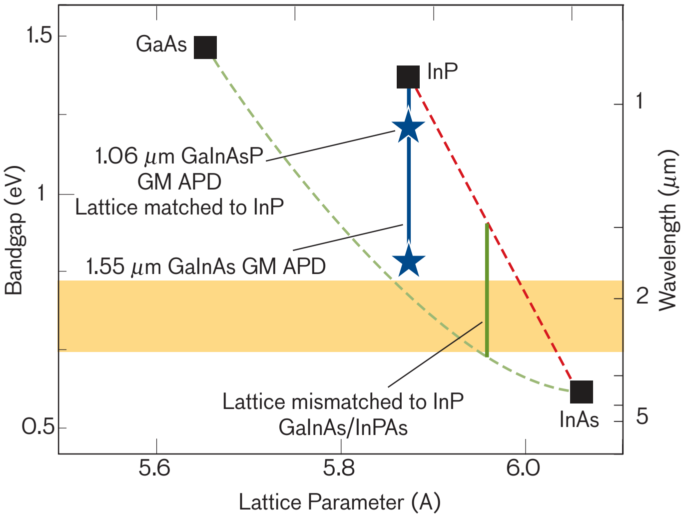

The orange band highlights the wavelengths of interest for III-V detectors. The blue vertical line indicates the series of materials that are lattice matched to an InP substrate. The green vertical line details the lattice and chemical information for Wang's work. The red line linking InP to InAs shows the effect in both lattice parameter and wavelength of a chemical shift in the In(P, As) compounds.

imaging arrays" around 2 μm. Such applications, Johnson says, include "land mine detection, especially in vegetative ground cover, and bioaerosol discrimination and classification."

To fabricate their APDs, Wang and her associates need to alter the chemistry of the active layers so that the device responds at the appropriate wavelength. For example, in the formulation $Ga_xIn_{1-x}As$ lattice-matched to a substrate of InP, the detector's wavelength of peak sensitivity changes from approximately 1.6 μm for high gallium concentrations to near 2.5 μm for high indium composition. But changing the chemical composition does more than tune the sensitivity wavelength. It also alters the distance between the atoms in the semiconductor crystal's lattice, leading to a lattice mismatch between the active layer and the InP substrate. Somehow, the black wall of the substrate lattice dimension must morph into the gray of the active layer.

Lattice-mismatched detectors that work at the longer wavelengths require manipulating lattice chemistry to create designer substrates on which InGaAs/InPAs detectors reside. Wang does this manipulating by creating a series of buffer layers of $In(P_yAs_{1-y})$ that build on an InP substrate, each with a slightly higher value of y as the phosphorous substitutes for the arsenic, until she achieves the proper lattice parameter base for the device layer, which has the formula $(Ga_xIn_{1-x})(P_yAs_{1-y})$. Wang's associates Dan Calawa and David Chapman measure both the lattice parameters and chemical doping composition of the devices to confirm that their technique is working properly. Epitaxial growth is a slow process, in which layers are deposited one atomic layer at a time with a process called organometallic vapor-phase epitaxy (OMVPE). Wang must be careful to change the chemistry of each subsequent layer slowly. Otherwise, strain between mismatched layers will produce defects. It's "painstaking work," she says. "If you try to rush it, you have an inferior product."

Wang must make one additional significant change. Her current work is based on nontraditional sources of arsenic and phosphorus sources: tertiary butyl arsine and tertiary butyl phosphine. Both are liquids. The industry standard OMVPE process, however, uses arsine and phosphine gases. "If we ever want this technology to be transferred to production," she explains, "we have to go mainstream." But following the standard process poses a potential difficulty. The liquids decompose at lower temperatures than arsine and phosphine gas. Wang is concerned that the higher temperature growth associated with arsine and phosphine gas may drive a faster deposition rate and, therefore, multiple starting points for layer growth. These islands of growth eventually connect with one another—but in doing so, create defects. Wang's process, by contrast, lays down a complete monolayer before starting the next layer. In this way, she increases the probability of relieving interlayer strain while minimizing formation of defects. "We need to confirm that our growth kinetics still work with the new sources at the higher temperatures," Wang concludes. When she makes her first full-scale 2.0 μm GaInAs/InPAs detector, using arsine/phosphine, that shows the same performance characteristics as her current material, her next question to Johnson and others might be: What wavelength do you want next?

SATELLITES

Shadowy Work

An array of telescopes could figure out what geosynch satellites are doing by studying their silhouettes

Jane Luu and Leaf Jiang are chasing shadows. Specifically, the two members of Lincoln Laboratory's Active Optical Systems group want to capture the shadows of geosynchronous satellites and examine them to see what they can learn about the satellites' functions.

The U.S. Air Force tracks these satellites, which orbit at an altitude of 36,000 kilometers, but their telescopes see, essentially, blobs. It's possible to tell where the satellites are but not what they're doing. "Even the biggest telescope on Earth can only get a resolution of three meters, which is very poor by defense standards," says Luu. If the satellite being observed is 10 m long, 3 m resolution will produce an image based on only three points.

Radar does a good job of imaging satellites that have discernible motion. But because geosynchronous satellites don't move with respect to the Earth's surface, and the newer ones don't spin, there is no movement to produce the Doppler signals that radar needs to make out a shape.

Luu, a former astronomer, and Jiang, an electrical engineer, think the solution is to capture the shadows that satellites cast on the ground as stars pass behind them. Using a line of 150 separate 30 cm telescopes, they can create the equivalent of a single telescope with a 45 m aperture. The resulting half-meter resolution produces images sharp enough to get a clear silhouette of a satellite that will show its shape, where it's pointing, and whether it's carrying a telescope or other device—the sort of information that will give clues to the satellite's function.

As a star passes behind a satellite, the light is diffracted around the edges, casting an array of ripples as the light waves interfere constructively and destructively. To reconstruct the image, the researchers measure the variations in intensity of the ripples, then run the data through an inverse Fourier transform, a standard processing algorithm. That's the easy part. The hard part is catching the shadow in the first place.

Depending on the size of the satellite, a star passes behind it in a few milliseconds. To capture such a fleeting signal, the team uses avalanche photodiodes, which can detect single photons. The observations are designed to be made in blue wavelengths; this spectral selectivity screens out reflected sunlight from the satellites, which tends to comprise mostly yellow and red.

Luu and Jiang need to know with a fair degree of precision when a star will pass behind a satellite. Because a 10 m object in orbit casts a shadow that is only 15 m long, there's not much margin for error. The researchers can use Air Force data to calculate the orbital position of a satellite fairly closely, but can wait months for a satellite and a star to line up. To solve this problem, they put the telescopes on a track, so that they can move them back and forth and line them up where they need to be. Luu and Jiang estimate that a 10-km-long

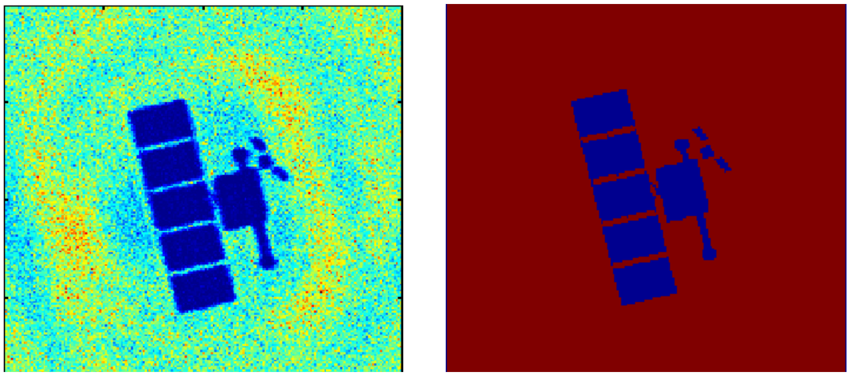

At left, a simulation made by 150 telescopes, each with a 30 cm aperture, in a 45-meter-long array, portrays shadow imaging of two geosynch satellites as produced by a star of magnitude 7.5. Compare the simulated image to the actual shapes of the satellites at right.

track should serve most purposes, leading to one occultation per satellite per night.

“If you could build a 45-meter telescope on the ground, the image quality would be just as good as with our method—but no such telescope exists,” says Luu. Nor is it likely to. Such a telescope would be vastly more expensive than an array of 30 cm telescopes, which cost about $6000 apiece. And it would have to overcome the distortions caused by the weight of that large a mirror.

Scaling up the system should be a simple matter, Jiang says. “Just buy more telescopes and add them on.” He thinks it should be possible to get resolutions in the tens of centimeters by adding more telescopes, although that would also mean dealing with minor sources of distortion, such as chromatic dispersion, in which the atmosphere at higher latitudes spreads out different wavelengths of light.

While their aim is to track satellites for national defense purposes, the same set-up could be used to look for small objects, such as asteroids, within the solar system, as well as Earth-scale planets in other solar systems, where so far only Jupiter-sized bodies have been found. “Anything that goes in front of a bright source, you can find,” Luu says.

The Lincoln Laboratory team performed simulations showing that the idea should work, and is now aiming to demonstrate experimentally, possibly as early as September, using two telescopes. One or two telescopes will not produce an ideal image, but Luu and Jiang hope it will be enough to encourage a sponsor to endorse a full-scale project. Part of their work involves writing software to determine the satellite’s position and to automate the alignment of the telescopes. Once they get sponsorship, deploying an actual system would take about one year, according to Luu. “The concept is simple,” she says. “The actual implementation is hard.”

“Although still easier than building a 45-meter telescope,” Jiang puts in.

Using a line of 150 separate 30 cm telescopes can create the equivalent of a single telescope with a 45 m aperture.