Lab Notes

NEWS FROM AROUND LINCOLN LABORATORY

ROBOTICS

Let a Robot Do the Work

They're not automatons, but these robots map, track, and locate things for you

In the popular imagination, robots not only have personalities of their own, from the fearsome Terminator to the cuddly Wall-E, they also are capable of dealing with unpredictable situations in unfamiliar environments. In the real world, such autonomy is elusive; robots either follow preprogrammed instructions and sets of rules, or they're controlled remotely by humans.

A new robotics program at Lincoln Laboratory aims to change that by developing systems that can carry out tasks with less human supervision. Through 2009 and 2010, the program was charged with seeing how existing robotics research might be applied to military needs.

"It seems like there's a lot of solid research going on in academia," says Eliahu Niewood, head of the Laboratory's Engineering Division, who is spearheading the robotics effort. "Very little of that is making it into Department of Defense applications. The research doesn't seem to be really coupled to what the warfighters need."

The issue, Niewood says, is that academic researchers tend to explore theories of how robotics should work, rather than building systems robust enough for actual use. "They take it as a solved problem, and from an academic perspective it might be, but too often it doesn't get used in the real world." Niewood hopes Lincoln Laboratory can help close the loop by taking algorithms that have already been developed and building them into usable systems.

The first stage of the project incorporated a mapping capability into a Packbot, the ruggedized machine from iRobot Corporation that the military uses for bomb detection and other tasks. The mapping subsystem, developed by Michael Boulet, relies on a ladar, an active optical sensor, to send out and measure the return of light pulses. With that technology, the robot is able to map out where walls are and where one corridor opens onto another. So far, the system maps in only two dimensions, but the researchers are trying to give it

Lincoln Laboratory staff pose with two generations of robots. From left to right are Byron Stanley, Michael Boulet, Brian Julian, and Dr. Peggy Boning.

A map of the corridors of a building from a vantage point approximately one foot above the floor will include hall widths, turns, doorways, file cabinets, recycle bins (on left), and cardboard boxes. The Packbot was programmed not to enter any room, although it looked inside each doorway and mapped the offices up to the points where its view was obstructed by chairs, desks, or other objects.

three-dimensional capability, allowing it, for example, to measure the heights of doorways.

"A big challenge is having the robot know where it is," Niewood says. Because it is indoors, it can't use a Global Positioning System (GPS) signal, so it relies on an odometer to figure out how far it has traveled and when it has returned to its starting point. However, the odometer is not accurate enough on its own, particularly for longer routes.

Boulet and Byron Stanley are also heading the "Patrol Leader" segment of the project, which focuses on in-theater robot uses. Patrols in Iraq often send a sensor-laden manned armored vehicle in front of the main patrol convoy to look for hidden bombs. As the sensors are only a few meters away from the "on-point" vehicle's cabin, it is quite possible that a triggered improvised explosive device (IED) would injure the driver. Remotely controlling an unmanned IED-detecting lead vehicle from an operator control station located in a following vehicle decreases the danger to the soldier. However, watching a moving scene on a screen in a vehicle that's moving differently can lead to motion sickness for the operators, as well as depriving them of the ability to keep an eye on the rest of their surroundings. By adding limited autonomy to the robot, both motion sickness and operator overload are reduced.

In response to the military's need for safer bomb detection, the team is designing a robot that can autonomously stay on the road 20 to 30 meters ahead of the manned convoy while moving at 10 to 15 miles per hour, instead of the current 2 mph. They need a robot that is not overly complex, with a robust, dependable behavior. "We want something that's intuitive and that can be used by the troops without their having to worry that it might not react reliably in every case," Stanley says. One potential future capability is to detect obstacles, whether *positive*, such as debris in the road, or *negative*, such as potholes.

The robot would be equipped with a combination of GPS and an inertial navigation system to have a more accurate sense of its location. Most of the computer processing capability would be on the following vehicle, where the more expensive equipment would be safer. The team is also working on communications between the robot and the vehicle. "Communication is relatively difficult in certain environments, especially if you're in an environment with jamming," Stanley says.

> They take [robotics] as a solved problem, and from an academic perspective it might be, but too often it doesn't get used in the real world.

The mapping project was completed last year; the Patrol Leader project is currently in its early stages. Niewood would like to see an actual product by the end of next year.

Further out, and more investigative in nature, is the third arm of the project, cognitive robotics. Headed by Jerome Braun, this effort aims to give a robot the ability to enter an unfamiliar situation with an open-ended task and make decisions about what actions to take. The hope is that a robot could, for instance, be sent to the site of a collapsed building, navigate and operate on its own within the rubble, while looking for survivors.

"Whatever technology we have today is not at the level of perception and understanding for the robot to make good decisions by itself," Braun says. "They typically require a large amount of human involvement."

Sending a robot into an unfamiliar or unpredictable environment—an urban battle zone, for instance—requires the robot to be able to perceive and understand the situation, not just see its surroundings, but comprehend what it's seeing. Sparking wires or a burst water main might obstruct the robot's intended path, but there might be another path or a better way to achieve its goal. The robot needs to be able to discover those possibilities, incorporate them into its decision-making process, and act accordingly. Because every situation may be different, the robot has to be flexible enough to deal with many variables. The idea, Braun says, is to mimic aspects of human cognition to allow the robot to cope with tasks that are not predefined in situations that do not allow direct human involvement.

In fact, Braun is focusing on developing biomimetic machine-intelligence approaches, inspired by biological cognition models. Such approaches, he says, can bring artificial intelligence and machine cognition to the level needed for a truly autonomous robot operation. A novel cognitive machine-intelligence architecture—the main focus of the cognitive robotics effort—draws on concepts from cognitive science and neuroscience, Braun says. As the components of the architecture are developed, they will be tested on example tasks. An example task, he says, can at first be fairly simple as long as it poses a sufficient challenge to demonstrate the cognitive operation and its strengths.

Cognitive processing by machine is obviously an ambitious goal. The cognitive robotics effort, Braun says, seeks to develop the technology underpinnings of machine cognition, embarking on the road towards intelligent and truly autonomous robots.

Niewood says the overall robotics project involves seven or eight Lincoln Laboratory researchers, with some collaboration with the Computer Science and Artificial Intelligence Laboratory (CSAIL) on the MIT campus, as well as researchers at Carnegie Mellon University. The Laboratory is funding a related project at Olin College, a small engineering school in Needham, Massachusetts, and will work with some students from Worcester Polytechnic Institute starting in the fall of 2010.

Niewood hopes to build up a community of people interested in bringing more sophisticated robotics into practical use. "This is a very attractive niche that the Laboratory could fit into."

BIOCHEMISTRY

Uncoiling DNA Analysis

Rapid DNA-fingerprinting method would identify biological organisms in the field

DNA sequencing has proven to be a useful tool for identifying organisms, but is not currently able to perform rapid identification of potential bioagents from mixed samples in the field. Now Lincoln Laboratory researchers Lalitha Parameswaran, an electrical engineer, and Eric Schwoebel, a biologist, are working to build a portable, rapid DNA-fingerprinting device that could examine samples quickly and recognize a wide variety of organisms.

The core of the technique is restriction enzymes, which are produced naturally in bacteria and protect their hosts by binding to specific known sequences on foreign DNA strands and cut-

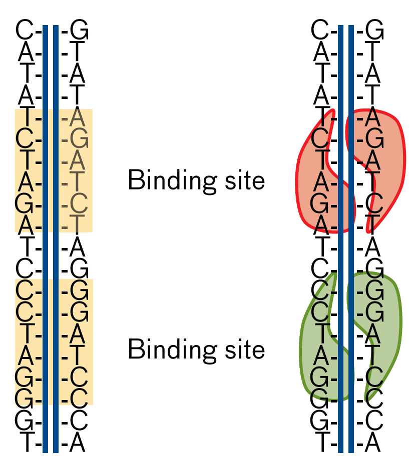

Restriction enzymes attach at specific locations on DNA. They can be illuminated with optical tags.

ting the strands into fragments. Schwoebel, a member of the Laboratory's Chemical, Biological, and Nanoscale Technologies Group, recognized that if they could get the restriction enzymes to bind but not cut the DNA and attach fluorescent markers to the enzymes, they should be able to create optically detectable patterns of binding, or "fingerprints," specific to a particular organism. Essentially they would get a string of glowing markers, arranged one way for, say, anthrax and another for *E. coli*.

"We want the restriction enzymes to bind to the DNA but retain their sequence specificity," says Parameswaran. "We already have sequence information for many organisms of interest, so we know the binding pattern that will be produced by each restriction enzyme. With this technique, we may not need to completely sequence unknown target DNA in order to identify its host organism."

There are, of course, existing approaches to DNA identification. Researchers can use the polymerase chain reaction (PCR) to create multiple copies of a strand of DNA and real-time PCR to produce a fluorescent signal indicating the amount of target-specific DNA produced. The DNA can also be sequenced by using one of several standard methods. But these approaches have many time- and labor-intensive steps, and can also require a very clean sample, necessitating stringent sample-preparation procedures to remove inhibitory components.

The team's proposal, in contrast, would use a much simpler assay. It could use a less pure sample, say from the soil, a drop of blood, or even a suspicious film scraped off a building's window if someone suspects that a terrorist has fogged the building with an airborne pathogen. The advantage of the method is that it doesn't require isolating the specific target DNA to be identified or even having foreknowledge of what target is present in the sample. One could simply extract and fingerprint all of the DNA in a sample, and the fingerprints could then be matched using already existing sequence databases.

The trick is to coax the restriction enzymes to bind to DNA sequences without actually cutting them. The pair says they have done that by altering the salts in a solution containing the DNA. The enzymes rely on magnesium ions to enable the cutting reaction, but replacing the magnesium with calcium stops them from snipping the strand. Eventually, they'd like to use restriction enzymes that have the cutting action permanently suppressed. These would be produced from genetically altered bacteria, something researchers elsewhere are working on. One issue the team is studying is whether they can interfere with the cutting and yet keep the enzymes binding to the right sites.

Another key requirement of the detection system is to create tiny channels, only a few nanometers wide, on a sample plate. An applied electric field would draw the DNA strands with bound restriction enzymes through the channels while an optical system watched for the signal from the fluorescent tags. The channels keep the DNA strand stretched out and localized; if it were to fold back on itself or ball up, it would be impossible to see the linear pattern of binding sites that identified a particular organism. Creating channels only 10 nanometers wide, or even narrower, presents an engineering challenge, but not an insurmountable one, the team believes.

The trick is to coax the restriction enzymes to bind to DNA sequences without actually cutting them.

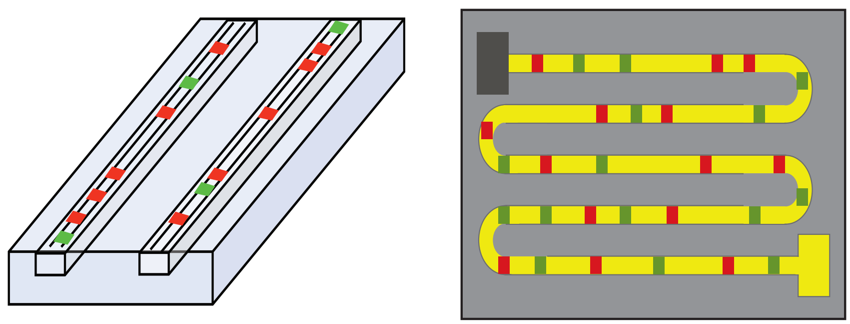

One critical part of this procedure is the straightening out of the spaghetti-like labeled DNA strand. Once it is laid out in nanochannels, shown on the left, the entire strand can be imaged, as shown on the right, to give a binding pattern signature unique to the organism from which the DNA originated.

The plan will also require a sophisticated optical system to detect small amounts of fluorescence in microseconds, as well as a computer to compare any patterns detected to known DNA sequences. They'd eventually like a system that uses a sample plate with hundreds or thousands of channels for high throughput. "It would be analogous to analyzing the amount of DNA present in the entire human genome in a matter of minutes," she says.

One plus of the method is that it could provide information about an unknown or genetically modified organism. "Even if you haven't seen this new *Bacillus* organism before, you might still be able to tell it's a *Bacillus*," Schwoebel says. It might be possible to find patterns that showed a particular strain of bacteria was resistant to a certain antibiotic so that doctors could prescribe a different one to people who had been exposed. "It could give you an idea of how to proceed without having to go through cell-culture techniques to identify a new organism," he says.

In addition to providing quick, in-the-field information in a possible biological attack, the rapid identification technique could also prove useful in medicine for bedside diagnosis and personalized medicine. It might also be helpful in food safety and forensic science.

The team's current focus is on testing their ideas about nondestructive specific binding of the enzymes, a project they hope to complete by the end of the year. If they're successful, they will seek funding to continue the development of a full system.

LASER TECHNOLOGY

A Bright Idea

Simple tweaks turn tiny diode lasers into powerhouses

Electrical engineer Kevin Creedon turns on the pump that sends cooling water flowing past the thumbnail-sized bar of laser diodes (a linear array of edge-emitting diodes, generally 1 μm high by approximately 100 μm wide), dons a pair of protective goggles, and flips a switch. Invisible to the naked eye, beams of laser light exit the diodes, pass through a lens, bounce off a grating, and move between some mirrors before emerging from a hole in a small box to produce a bright red spot on the piece of paper that Creedon's colleague, physicist Bien Chann (now at TeraDiode), is holding just outside the opening.

This setup—diodes, lens, grating, and mirrors inside a box about half the size of a shoebox—may lead to a new generation of high-powered lasers, say the researchers in Lincoln Laboratory's Laser Technology and Applications Group. It's certainly less complex than the setup on the other side of the lab bench, a couple dozen components covering the length of an optical table, which Chann says is a constant headache to try to get to work.

Chann, Creedon, and their colleagues are trying different ways to combine the beams of many tiny diode lasers into one, thus increasing the output power.

The challenge, as Chann explains, is "How do you take 10 lasers or 100 lasers or however many lasers you want and put them together so you get all the power of the group of lasers but retain the beam quality of a single element?"

Diode lasers are often used to pump other types of lasers, such as neodymium-doped YAG (yttrium aluminum garnet) crystals or erbium-doped fibers. Those larger lasers are popular in industrial applications, such as cutting and welding, in which more power makes the process faster and more efficient. But because a diode-pumped laser can't put out more energy than the diodes are putting in, it is essential to increase the power output of diode lasers.

It's common to bunch diodes together—the team uses diode bars that group together a row of 19 emitters in each bar, totaling about 100 watts. But while that increases total power, the brightness—that is, the power per unit area of solid angle that the beam spreads out into—is no greater than a single diode. Applications such as coupling the beam into a fiber require a high-quality beam that combines beams from all the diodes into one tiny spot.

To get both high power and

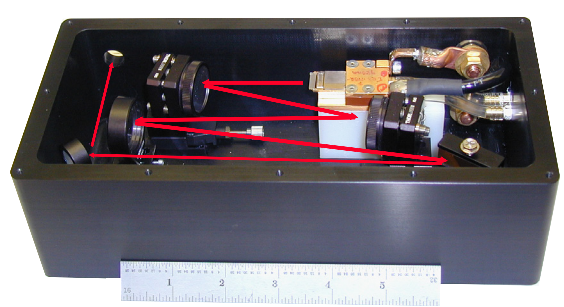

The entire unit of diodes, lens, grating, and mirror comprises a compact unit smaller than half a shoebox.

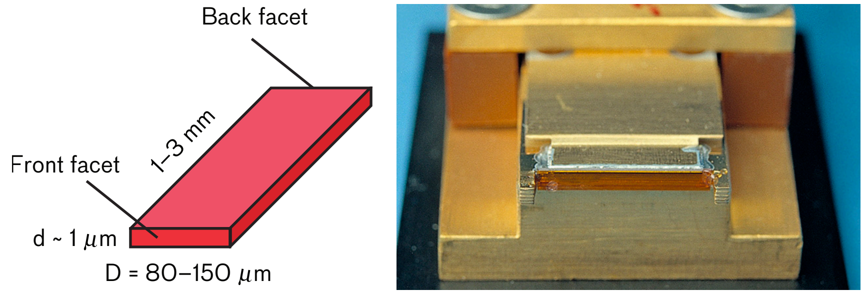

The laser diode bar, like the one shown above produced by the Electro-optical Materials and Devices Group, is an array of efficient, high-power emitters. Typical bars can emit over 100 W, but they suffer from low brightness and the "smile" (below).

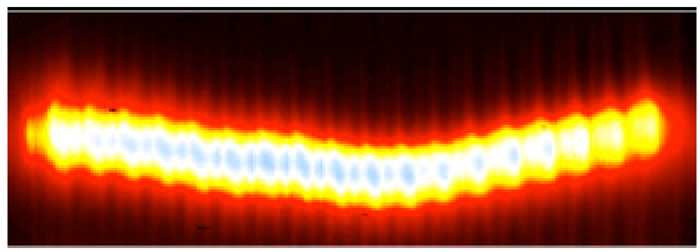

high brightness, Chann and Creedon use a setup designed a few years ago by group leader Antonio Sanchez-Rubio and associate leader Tso Yee Fan. They buy standard diode bars that have an internal cavity coating at the back that reflects 99 percent of the light, while the front coating reflects 5 percent back. The emitted light passes through a lens onto a grating and bounces from there to a mirror that has a 10 percent reflective coating. The light then follows the same path backwards, onto the grating, through the lens, and back to the diode, where it injection locks the emitters so that they lase at the wavelength they should to combine into a single beam. Essentially, they've turned the whole setup into a single laser resonator cavity.

Commercially available diodes generally have a gain of about 50 or 60 nanometers. That means the light coming out of a diode that nominally emits at 1000 nm can be "locked" at any wavelength between 970 and 1030 nm. The 5 percent reflectivity coating that is optimized for high power output reduces this range to approximately 4 nm.

But when the beam strikes the grating at a particular angle, only light at a very specific wavelength is directed along the path that will return it to the laser, so the diode will only lase at that specific wavelength. The beam from the diode next to that one will hit at a different angle on the grating, requiring a different path, and thus will lase only at a slightly different wavelength. "It's automatic. You don't have to do anything to it," Chann says. Once all the elements lase at the right wavelength, the grating acts like a prism in reverse, combining them into a single beam, with the power of each, but with vastly improved brightness.

How do you take 100 lasers and put them together so you get all the power of the group of lasers but retain the beam quality of the single element?

But there is still a problem. Even diode bars manufactured to the strictest tolerance suffer from a phenomenon called "smile." The separate emitters don't align perfectly horizontally, but tend to sag into a crooked smile of laser spots because of the dissimilar thermal expansion during soldering to the heat sink. Even in the very best laser bars, emitters can be out of place by 2 to 3 micrometers. Because the setup relies on each beam hitting at its particular angle on the grating, the fact that they're out of alignment throws a monkey wrench into the system.

Chann fixed the problem with a very simple idea. He flipped the diode bars on their sides, so the rows of emitters are vertical rather than horizontal. Because of the shape of the lens's focal plane and the design of the grating, a vertical displacement of the emitters doesn't have the same effect as the horizontal displacement, and the setup works to make a bright, powerful beam with a spectrum that fits within the 4 nm locking range, so it's compatible with standard, high-power facet coatings.

"It's so simple it makes you wonder, 'Why didn't I think of that?'" says Sanchez-Rubio.

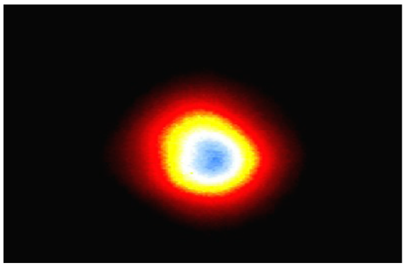

The beam-combining optics merge all the laser beams into one beam that has both high intensity and the beam quality of a single laser.

Having designed the system, the staff focused on building an actual working version. In a two-year program that began in March 2009, Sanchez-Rubio, and Steven Augst are building a system that combines lasers together to produce a coupled 100 m fiber system with an output of 2 kW and a bandwidth of 2 nm. By contrast, the best such beam commercially available is 50 watts. Much of the work necessary to realize this 40-fold increase will involve finding components able to handle such high powers. The researchers will also have to work on reforming the beam into the optimum shape to be coupled into a fiber.

The Laboratory team is also using the same setup to build a 200 W system for the Laser and Optics Research Center at the U.S. Air Force Academy. This system will act as a diode pump for an alkali-vapor laser, resulting in a tool powerful enough for industrial applications. Lasers that are based on alkali metals, such as potassium, rubidium, and cesium suspended in buffer gases, should produce beams of high power and high optical quality and not suffer from the heat-management issues that afflict other lasers, but they have been limited by a lack of suitably narrow, high-brightness pump diodes. Narrower linewidths from the pump lasers would allow the alkali lasers to operate at lower temperatures and gas pressures, slowing down chemical reactions between the alkalis and the buffer gases and making the laser more efficient.

Boris Zhdanov, a research scientist at the Academy, is waiting for the Laboratory to complete the work so Zhdanov and his colleagues can test the setup on their alkali laser where, he says, a narrow-line pump source is key. “What they are doing is very important for us, and not only for us, but others who work in this field,” Zhdanov says.

However it is applied, the beam-combining setup, Chann says, should make diode arrays into lasers much more powerful than they are now. “There’s no reason we couldn’t do 10 kilowatts.”

HYBRID TECHNOLOGY

A Route to Energy Savings

Predictive control of plug-in hybrid vehicles will analyze your route to work and optimize battery and gas usage

The graph of power demands during Nicholas Judson’s drive from Lincoln Laboratory to his home in Cambridge, displayed on his computer screen as a series of vertical lines, would make perfect sense to anyone familiar with the route. There’s the increase in power consumption as he gets on Route 128 and picks up speed, a brief drop as he slows down for the off-ramp to Route 2, and a precipitous decline as he goes down the steep hill from Belmont into Arlington. The graph of his morning commute from home to laboratory is pretty much a mirror image. Measurements made when he skips the highways and sticks to Massachusetts Avenue through Lexington show a trip that takes about ten minutes longer, but uses about 20 percent less energy.

Judson and his colleagues Alan Millner and group leader Bill Ross in the Advanced Electro-optical Systems Group hope that gathering data on power demands during driving will help make plug-in hybrid electric vehicles (expected to be introduced in the next few years) much more energy-efficient. They’re aiming to provide such vehicles with predictive control, a system that makes decisions on when to use the battery or gasoline to power the engine. By timing the trip so the battery reaches its minimum charge just as the driver reaches his destination, they predict they can get a 7 to 10 percent boost in engine efficiency, raising mileage to roughly 71 miles per gallon from the 66 mpg expected from certain plug-in hybrid electric vehicles.

The idea is to use the gasoline engine only at its most efficient operating point, while using as much battery power as possible. “Any battery power that isn’t used up because you’ve used gasoline instead, means higher cost because gasoline costs more than electricity

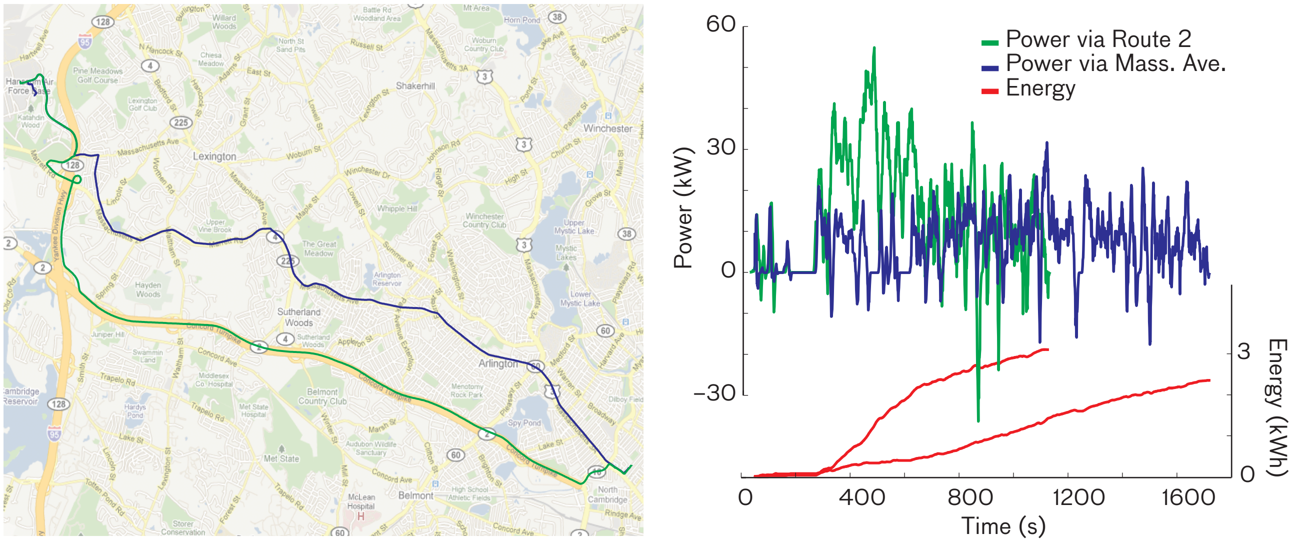

The route chosen for a trip can have a big effect on power and energy usage. Prior or assumed knowledge of the destination and the required power profiles, which depend on the route (green signifies Route 2; blue is Massachusetts Avenue), driver, average speed, changes in altitude, and many other factors, will affect how the algorithm controlling the use of stored energy in the vehicle will meet upcoming power requirements. (Image is copyrighted 2010 Google, map data copyrighted 2010 Google.)

for the same power," Judson says.

"Now you can use any kind of power source you want for transportation, rather than being stuck with gasoline or diesel fuels," says Millner. Many sources of electricity contribute far less to global warming than oil and would allow gains in efficiency with the use of central power generation. Millner says that if half the passenger cars in the United States could be converted to plug-ins that get 100 mpg, it would lead to about 500 million fewer tons of carbon dioxide being spewed into the atmosphere annually, roughly an 8 percent drop in emissions. It would also mean saving an amount of energy comparable to what the U.S. purchases from the Middle East.

In a hybrid vehicle, the gasoline engine kicks in when there's a demand for increased torque, caused by rapid acceleration or climbing a slope. But it also turns on in an effort to keep the battery charged at a certain level. In a hybrid that doesn't plug in, such as today's Toyota Prius, the system tries to maintain a battery charge of around 50 percent of capacity. In a conventional plug-in, the battery drops from nearly a full charge to roughly 20 percent, and then uses gasoline to maintain that level until it can be plugged in again. In the scheme Judson, Millner, and Ross are working on, the battery drops steadily from a full charge to the minimum—about 20 percent charge—with the aim of hitting bottom at the end of a trip when it can be plugged into the electrical grid and recharged.

A predictive control system would gather route and elevation data from a Global Positioning System (GPS) navigation system in the vehicle, along with data already available in the car's computer systems, and learn about a driver's trips and driving habits. Drivers could either tell the system where they were headed, or the system could figure that out on the basis of experience—if it's 8:30 a.m. on a weekday, it would know the driver is probably headed to work, for instance. Knowing when the demands for high torque would come, it could use the battery as much as possible and level the load on the gasoline engine over the entire trip, thereby operating the engine at optimal efficiency. And instead of using preset parameters, such as battery charge levels, to switch to gasoline, the system could make decisions on the basis of the trip. For instance, if the battery were getting low but the car was two minutes from home and had a little battery energy left, it could decide not to use the gasoline engine.

Judson envisages the predictive control algorithms they are developing could be included in GPS navigation systems. The algorithms would function in response to what it learns about a particular driver's driving habits and usual routes. Standard commercial systems could also be designed to tie into map systems and other data to apply the algorithm to unfamiliar routes. For instance, a route with a lot of traffic lights could be predicted to be less efficient than a trip along a freeway because of the numerous slow-down/speed-up instances. Or the system might be linked to real-time traffic data and suggest an alternate route that's not only smoother but more energy-efficient. Judson even speculates that some people might opt for slightly longer routes to reduce their energy consumption.

There are limits to how useful data from sources outside the car would be. Individual factors such as driving habits and even tire inflation are likely more important variables, Judson says. But he's also looked at some publicly available GPS data that measured driver behavior for about 200 vehicles in St. Louis over a period of ten weeks. What the data showed was that the range of accelerations in real-world driving is far broader than that used in the government standards used to give mileage ratings for cars.

You can use any kind of power source you want for transportation, rather than being stuck with gasoline or diesel fuels.

The Laboratory's research program includes not only the predictive control but also an effort to optimize the batteries and bring their costs down. Millner says that plug-in vehicle batteries should not only draw from the electrical grid but also feed power back into the grid, thus smoothing out power demands for buildings and reducing costs, making the whole plug-in vehicle infrastructure more economical. For example, if the car "sensed" that it had more than enough reserve battery power to get itself back home (where it could recharge during off-peak hours of electrical consumption), it could feed some of that power into the grid.

The team has built a vehicle energy lab at Lincoln Laboratory and is awaiting a test vehicle they can use to gather more real-world data. They spent last year on analysis and modeling, and are seeking to build a working system in 2011.

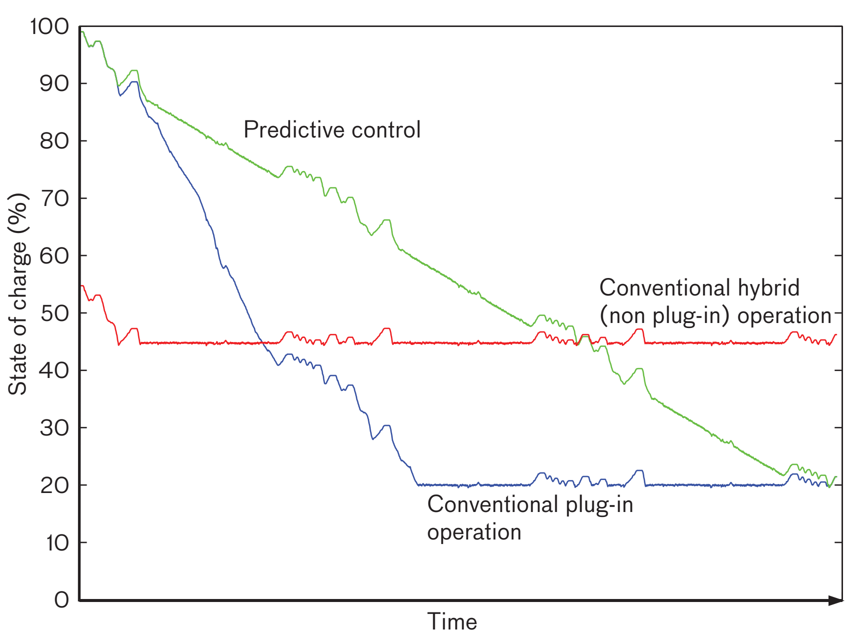

Battery power consumption can be optimized further than what current hybrids (and proposed plug-in hybrids) consume. With predictive control, the battery can be used throughout the entire trip, allowing the gasoline engine to run closer to its optimal point of efficiency, thereby saving significant amounts of fuel.

ASTRONOMY

Vision Correction

Two projects should help astronomers see further across the universe

The history of the universe is laid out across the sky, with the earliest events—the birth of stars and formation of galaxies—the furthest away, at roughly 13 billion light-years. To study the oldest, dimmest parts of the cosmos, astronomers are building larger, more sensitive telescopes. Lincoln Laboratory is developing technology that could help the planned next generation of ground-based telescopes see stars no human has seen before.

The Laboratory's Advanced Imaging Technology (AIT) Group is working on two instruments to aid the Thirty-Meter Telescope (TMT), an instrument that should be built on Mauna Kea, Hawaii, within the next decade. With a 30-meter-diameter mirror, made up of 492 individual segments, the TMT vastly increases the light-collection area, and therefore the sensitivity, over that of the 8- and 10-meter telescopes in use today. The AIT Group, in collaboration with Sean Adkins of the W.M. Keck Observatory in Hawaii, is working on a charge-coupled-device (CCD) sensor that will let this telescope see more clearly.

The TMT will rely on adaptive optics to overcome distortions caused by turbulence in the Earth's atmosphere. Adaptive-optics systems, like the one in use at the 10-meter Keck telescope, fire a laser beam into the thin layer of sodium located about 90 kilometers above the ground, causing the sodium to fluoresce and produce a glow that serves as a "guide star." The specialized CCD measures wavefront distortions in the light coming back from the sodium layer and uses these measurements to tell a series of actuators how to push and pull on the telescope's mirror to compensate for the atmospheric distortions. Fluctuations in the atmosphere happen on a time scale of about 10 milliseconds, so all the measurement and correction has to be done within that time.

"It's not very poetic," says materials scientist James Gregory, assistant leader of the AIT Group, "but we're taking the twinkle out of the stars."

To measure the shape of the wavefront across the area of the telescope aperture, the CCD is placed behind an array of microlenses. These microlenses divide the sensing device into "subapertures," each comprising multiple pixels. Each microlens focuses the light from the guide star into a spot, and the position of the spot on the CCD pixels is used to compute the tilt of the wavefront at the corresponding point in the telescope aperture.

The trouble with using a laser guide star with the TMT is that the sodium layer is generally 15 to 20 km thick. If the laser beam is sent straight up from the center of the telescope, then someone standing right next to the laser could look up and see a spot of light. But an observer standing further out toward the edge of the telescope would see the side of a column of light going through the sodium layer—that is, the spot now appears as a streak. While a spot would illuminate only a handful of pixels on a CCD sensor, the streak will illuminate many of them. Because each pixel will inevitably produce a certain amount of electronic noise in the act of reading out the data, a higher number of pixels leads to a lower signal-to-noise ratio.

"You don't have a lot of photons coming back to work with in adaptive optics, so signal-to-noise is everything," explains AIT Group leader Robert Reich.

To limit the number of pixels involved, Brian Aull, Barry Burke, Bradley Felton, and Kay Johnson of the AIT Group are designing a new type of CCD sensor, conceived in collaboration with Sean Adkins of Keck Observatory, Jim Beletic of Teledyne Imaging Systems, and Jerry Nelson of the Center for Adaptive Optics at the University of California, Santa Cruz. Instead of the standard contiguous rectilinear array of pixels, they are making what they call a polar coordinate detector. Each subaperture has its own miniature CCD, with just enough pixels to image the streak

The CCDs under development at Lincoln Laboratory will be used in the Thirty-Meter Telescope, shown in an artist's conception. (Image courtesy of TMT Observatory Corporation. Rendering by Todd Mason.)

of light. Furthermore, each CCD is oriented to form columns of pixels radiating out from the center of the sensor to the edge, like lines of longitude radiating out from the North Pole in a map looking down on the top of the Earth. Because the CCD columns are aligned along the streak of laser light, the CCD can be narrow in the direction perpendicular to the streak. This technique minimizes the number of pixels needed behind each microlens; as a result, the sensor has fewer pixels to read. Because the signal-to-noise ratio varies by the square root of the number of pixels read out, cutting the pixel count in half is equivalent to making the signal 40 percent stronger.

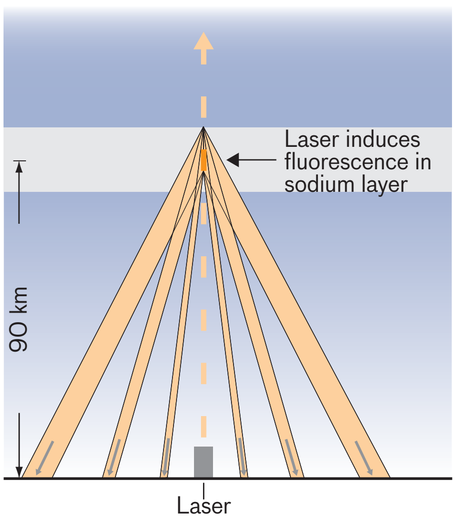

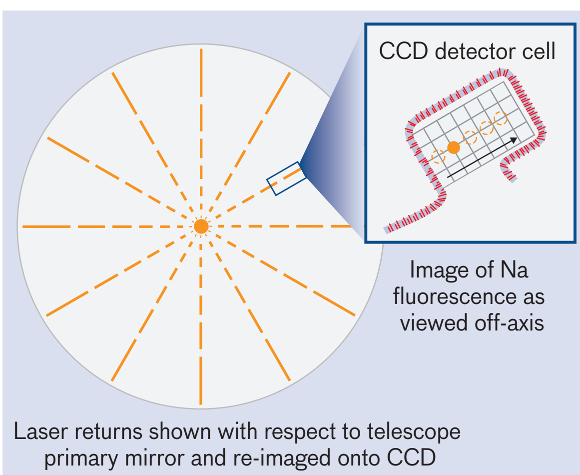

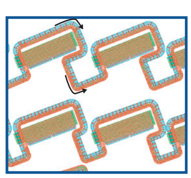

The CAD layout of CCD cells at the bottom shows interconnecting readout registers, which transport the charge to the periphery of the CCD for conversion to video signals.

Because the light streaks get longer toward the edge of the telescope aperture, the column lengths of the CCDs must also lengthen. Whereas there might be an array of 6 6 pixels behind the center microlens, out near the edge the pixels would be stretched out into a 6 15 formation.

Transferring the signal from the polar imager off chip proved to be challenging. A rectilinear imager design would have simple wiring and selection logic circuitry at each pixel array associated with a microlens. The addressing and multiplexing of the signal off chip would be straightforward. But for the polar-imager design, in which sizes and directions vary, such a method would have been a nightmare of complexity, Reich says. Instead, his team, following a suggestion by Adkins, figured out how to pipeline the readout registers of the pixels in series and load all the data from behind each microlens into the register so there were no gaps.

The design work is well under way, Reich says, and the team should be able to have a one-quarter chip fabricated early next year. If that device works, they'll build a full chip. Once they're satisfied, they can pass the CCD on to the team building the TMT.

"We are definitely planning on using it and factoring it into our design," says Brent Ellerbroek, group leader for adaptive optics on the TMT project, which is a partnership of the University of California, California Institute of Technology, and the Association of Canadian Universities for Research in Astronomy. "It's a very important innovation that simplifies the system and improves the performance in a lot of ways."

The second instrument the team is developing is not part of the adaptive-optics system, but will take advantage of the tiny number of photons collected by the telescope to make better pictures of distant galaxies. Aull and team member Daniel Schuette, along with Felton, and in collaboration with Donald Figer of the Rochester Institute of Technology, are designing a photon-counting imager based on Geiger-mode avalanche photodiodes. This imager will improve the sensitivity of the telescope so it can see dimmer objects at a given distance or find bright objects further away.

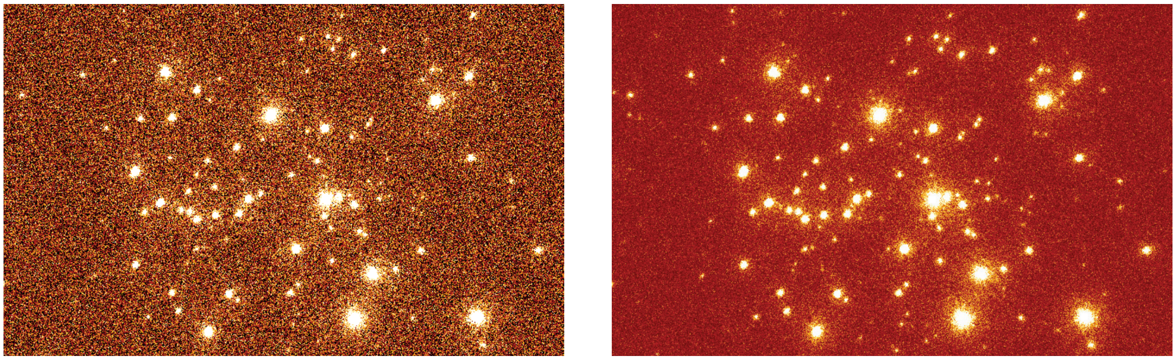

These simulated images of the Arches Cluster near the Galactic Center are based on Keck/Laser Guide Star Adaptive Optics (LGSAO) data and are courtesy of the Rochester Institute of Technology. The image on the left has read noise comparable to that in detectors currently planned for TMT. The image on the right has zero read noise.

The imager grows out of the team's work using similar photodetectors in laser detection and ranging (ladar) systems. A Geiger-mode avalanche photodiode is put into an unstable state by the application of a high voltage, typically about 30 volts. A photon that strikes the device creates a pair of charge carriers—an electron and a hole. When one of these carriers is drawn into and accelerated by the diode's high-electric-field region, it gains enough energy to create new electron-hole pairs each time it collides with atoms in the semiconductor lattice. The electrons and holes released by these collisions are then themselves accelerated by the high field, creating a cascade of electric current. The net result: capture of a single photon triggers an easy-to-detect burst of electrons. Ladar systems use these photodiodes to detect how long it takes for a photon leaving the laser to bounce off an object and return, thereby allowing calculation of the distance to objects. Compiling a large number of such range readings yields a three-dimensional picture of a scene.

It's not very poetic, but we're taking the twinkle out of the stars.

Astronomers are not as interested in when photons reach the telescope as they are in accurately measuring the total number of photons that the telescope sees. Analog detectors generate some noise every time a pixel is read out, making it difficult to see faint objects spread over many pixels. In a photon-counting imager, however, observation of each photon is recorded as a bit in custom all-digital circuitry that is directly linked to the photodiode array. Consequently, there is no readout noise—a signal is produced for each photon that is captured, and no signal is produced if there's no photon. "You're essentially at the limit of how good you can be," Reich says. "It's as quiet as you can get. You're counting every photon."

Reich says he expects within the next two years to build an array of avalanche photodiodes for the TMT by using silicon, but that the project may ultimately move to a different Lincoln Laboratory group that works with compound semiconductors rather than silicon. The reason for this shift is that astronomers need detectors to work at wavelengths between 1 and 2.4 micrometers, where materials such as indium-gallium-arsenide are more sensitive (light that originates from stars a vast distance away is red-shifted from the visible into the infrared by the expansion of the universe). Ellerbroek points out that, because longer wavelengths are less disturbed as they pass through the atmosphere, adaptive optics work better in the infrared, where astronomers can get near-diffraction-limited performance.

But compound semiconductor materials present a number of challenges that make them, currently, poorly suited for an astronomy imager. A team lead by Simon Verghese, assistant leader of the Laboratory's Electro-optical Materials and Devices Group, is working to address these issues. For now, silicon is better understood and better developed for this sort of detector. Says Reich, "We'll be the pathfinders."