


TMI, Meet IST

by Douglas L. Smith



In the old days, if your windshield wipers came on when you signaled for a left turn, it was probably a short in the steering column. But now, if your doors suddenly unlock as you punch the gas, it might be because the keyless entry system is getting cross talk from a defective accelerometer in the air bags. Today's cars have so many computer chips, says Jehoshua "Shuki" Bruck, the Moore Professor of Computational and Neural Systems and Electrical Engineering and director of Caltech's Information Science and Technology (IST) initiative, that nobody—not even their designers—has a complete understanding of them. The software in the average sedan can contain more than 35,000,000 lines of code—enough for maybe 100 copies of, say, *Grand Theft Auto*. Says Bruck, "The car industry is investing billions of dollars to figure out the interactions between the mechanical parts and the computers. Future development is actually getting stuck because they don't know how to manage the software."

But Nature controls far more complex mechanisms with ease: Consider the nematode *Caenorhabditis elegans*. A lowly roundworm about the size of this comma, it grows from a single-celled egg to an adult containing exactly 959 cells. The little fellas are clear as glass, and entire generations of grad students have spent countless hours hunched over microscopes tracking the career of each cell. The whole process takes 24 rounds of cell division—79 of the 959 cells line the guts from mouth to anus, 302 become nerve cells, and 131 die along the way. "Everything has been mapped precisely," says Bruck, who has a framed poster of this developmental tree on his wall. "But we, as engineers, don't understand how to handle all the information in that map. We don't understand what the principles are." But, somehow, the cells understand. The egg divides, and one cell has to call heads and the other, tails. The process involves the random diffusion of signaling molecules, but the result is very precise—you never end up with a

IMAGE NOT
LICENSED FOR
WEB USE

George Boole
1815–1864




Claude Shannon
1916–2001

Courtesy of The History of Computing Project, www.fhocp.net

George Boole, professor of mathematics at Queen's College, Cork, Ireland, published his masterwork, *An Investigation of the Laws of Thought, on Which Are Founded the Mathematical Theories of Logic and Probabilities*, in 1854. He pointed out the analogy between algebraic equations and logical statements—for example, if $x = \text{horned animals}$ and $y = \text{sheep}$, then $1 - x = \text{all things without horns}$ and $(1 - x)(1 - y) = \text{all things that are neither horned nor sheep}$. This means that sets of logical statements can be manipulated using algebraic operations, in what is now known as Boolean algebra.

The next breakthrough happened almost a century later. In his master's thesis in electrical engineering, written at MIT in 1938, Claude Shannon showed how to build any Boolean expression as a circuit composed of relays, thus completing the set of rules for the explicit transformation from text to math to hardware. (Shannon, who had dual careers at Bell Labs and MIT, also established the fields of information theory and communication theory.) This work by Boole and Shannon led to the field of digital logic design, which is the theoretical foundation of the microprocessor revolution.

Since any number can be rendered in binary form (ON or OFF, in electrical terms), this laid the groundwork for electronic math: The set of logical operations required to, for example, add two binary digits ($0 + 0 = 0$; $0 + 1 = 1$; $1 + 1 = 10$, send the 1 to the next adder to the left) could be encoded by a set of switches wired in the proper order. These logical operations are now commonly known as “gates”—the AND gate, the OR gate, the EXCLUSIVE OR or XOR gate (which outputs a 1 if either but not both of the two inputs are 1), and so on. The circuit for adding two one-digit binary numbers looks like this:



The green triangles are the logic gates—relays, transistors, or integrated circuitry; it doesn't matter. The + marks XOR gates, the Δ stands for AND gates, and the V is an OR gate. On the input side, x and y are the two numbers to be added, and z is the carry from the adder to the right. On the output side, S is the right-hand digit of the sum of $x + y$, and C is the carry to the next adder to the left.

The extrapolation to a Pentium is left to the reader as an exercise. □

two-headed worm. Then the other divisions have to follow in the correct order. “And even when every cell has a clock and the timetable,” Bruck points out, “they still need to coordinate their actions. It's like driving on the freeway—sometimes you need to slow down and let another car pass.” Organisms are just information made flesh.

A vast gulf yawns between our ability to describe and build complex systems and our ability to understand and manage them, says Bruck. “A Pentium chip has a hundred million transistors, but we cannot answer simple questions about *C. elegans* that has 959 cells. The bottleneck between what we see and what we understand is in our ability to abstract, and that's the power of IST.” The calculus developed by Leibniz and Newton describes the physical world, at least on the human scale; Bruck hopes IST will develop a calculus for the realm of information in all its guises. We're drowning in data, from up-to-the-nanosecond stock quotes to blogs to digital sky catalogs and protein databases, but we can't read or think any faster than we could 100 years ago. We need a new way of dealing with it all—another technological revolution, if you will.

The computer revolution happened because there are explicit ways to translate a verbal concept—“let's add two numbers”—into a mathematical expression—“ $x + y = z$ ”—that can then be turned into a series of logical operations by Boolean algebra. A mathematician and electrical engineer named Claude Shannon realized that any Boolean expression could be built as a set of wires and relays. From there to the Pentium is a bit of a technological leap, but today, with a few clicks of the mouse, you can specify what you want a chip to do and a computer will design it for you. “And that's why we can build things with a hundred million transistors,” says Bruck. “What we are missing is the ability to go backward.” Reverse-engineering things as diverse as nematodes and stock markets means bringing together people from many academic disciplines, which is a very Caltech thing to do. Bruck estimates that as many as one-quarter of the faculty will eventually participate in IST in some way.


A new building in which these folks can rub elbows will take shape soon. The international Office for Metropolitan Architecture (OMA), headed by Pritzker Prize-winner Rem Koolhaas of Seattle Public Library fame, has been chosen to design the Walter and Leonore Annenberg Center for Information Science and Technology, which will join the Gordon and Betty Moore Laboratory of Engineering on the south side of Avery Walk. Joshua Ramus, the partner in charge of the New York office, will direct the project. The building should be open for business in about three years.

To bring some structure to the initiative, it's organized into four new centers—the Center for the Mathematics of Information, the Center for the Physics of Information, the Center for Biological Circuit Design, and the Social and Information

Sciences Laboratory—and borrows from two existing ones: the Center for Neuromorphic Systems Engineering, and the Lee Center for Advanced Networking. Each new center attacks a basic question: Can we find an abstract mathematical description of information that applies across disciplines? What are the fundamental physical limits to information storage and processing? How does nature compute and communicate information? And how does information shape social systems?


The Center for the Mathematics of Information (CMI) is trying to unify three branches of engineering: computation, communications, and control. Each field deals with a scarce resource. Communications theory tells you how much information can be reliably sent through a noisy channel of limited capacity, be it a fiber-optic data line, a radio signal from a distant spacecraft, or even a CD. “Storing stuff is a sort of communication from the present to the future,” notes Leonard Schulman, associate professor of computer science and director of the center. The scarce resource here is bandwidth, or in the CD example, disk space. In control theory, the resource is real time—if your F-117 goes nose-down, a fly-by-wire system that takes five seconds to respond is going to leave a nice crater in the desert floor. And in computation, the resource is processing time: nobody likes to watch the waving Windows banner while a spreadsheet recalculates itself, and there are entire classes of useful problems that would take longer than the age of the universe for a computer to solve.

The CMI is charting the territory where these



fields overlap. Take control and communication, for example. Says Schulman, “Suppose we’re a couple of crazy teenagers. You’re driving blindfolded on an abandoned road, and I’m sitting next to you giving instructions—‘Less gas, turn right, turn harder.’” (Kids, don’t try this at home! Leave it to the professional idiots on *Jackass*.) At five miles per hour, this works. But as the driver speeds up, “there’s some maximum number of bits per second that we as humans are able to speak, and some minimum delay for us to comprehend what we’ve been told.” The communication delay makes the control system unstable, crashing it literally as well as figuratively.


“That scenario was error-free,” Schulman continues. “We’re sitting two feet apart, and you can hear everything I say. But what if we’ve been drinking, which is why this probably seemed like a good idea, and the stereo is blasting heavy metal?” Now there’ll be transmitter and receiver errors, and a noisy—in the engineering sense as well as the auditory one—channel between. The traditional communications-theory solution uses “block codes” or “convolutional codes” in which the accumulation of successive bits builds up a picture of what the original bit was supposed to be. But you can’t retrieve that bit reliably until you’ve received a long block of code. That could take 20 or 30 rounds of communication, and by then, you’ll be upside down in a ditch. What you’d like to do is abbreviate the messages—for example, instead of saying “change heading from 263 degrees to 262 degrees,” which repeats a lot of information, just say, “-1.” But that repetition helps suppress errors, and if you take it all out, errors accumulate and eventually you’ll find that same ditch. So Schulman, Rafail Ostrovsky of UCLA, and Yuval Rabani of the Technion, the Israel Institute of Technology, devised a new class of error-correcting codes for control systems. “To do this, we needed error-correcting code theory, which everyone in electrical engineering knows, and something from combi-



When driving blindfolded with a buddy, the driver’s course (red) will not follow the navigator’s instructions (blue) with perfect accuracy, as shown above left. The trick is to keep the driver’s tracking errors as small as possible, so that the probability of the error exceeding some acceptable limit is zero, as shown above right.

Top: This satellite image was taken at 9:21 p.m. EDT on August 13, 2003, the night before the blackout.

Bottom: This one was shot at 9:03 p.m. during the outage. Local generators and other emergency systems kept the entire Northeast from plunging into total darkness, but cities including Cleveland, Detroit, New York, Ottawa, and Toronto were hard hit.



natorics called the Lovasz local lemma. It's a nice example of what can happen when you cross the lines between disciplines."


Similar gains can be made at the intersection of computation and control. The great blackout in the summer of 2003 was essentially a control breakdown, Schulman says. A relatively minor failure—one power line going out of service in Ohio—cascaded until 50,000,000 people in eight states and the province of Ontario were left in the dark. "It was a highly decentralized control system, and had they designed it properly, the outage would have been very localized. We are integrating systems that are much larger than used to be integrated, and we're pushing them much closer to their performance limits. That's what engineers do for a living, they try to get the most out of whatever hardware they've got. And in a system like that, the control mechanisms are gathering information—loads, temperatures, and such—from thousands and thousands of different sensors. Integrating all that data is a complicated computation."

And in the intersection of computing and communication, you get problems such as how to keep the Internet from clogging up as more and more people use it. The way it works now, your files get sent through any available routers to their destination. It's like flying from Los Angeles to Portland, Maine—if you have to change planes in Philadelphia and the connection is tight, some of your bags may end up on other flights. But a new process called "joint coding" promises vastly increased net-

work capacity at the expense of intensive computation at the routers. Essentially, everyone's luggage goes from curbside check-in to a wood chipper that purees it—socks, shampoo bottles, golf clubs, and all—and compacts the shredded material into space-saving bricks. Then, when the plane lands in Philly, all the baggage has to be reassembled (without musing the neatly folded clothes!) so that the items actually bound for Pennsylvania can be fished out, and the rest goes into the chipper again.

The CMI's eventual destination lies where all three fields converge and the really gnarly questions lurk, such as predicting how minor changes at individual computers will affect the global behavior of the Internet, and how to control that behavior if it's tending in the wrong direction. Says Schulman, "Engineering challenges of this magnitude can only be approached with good mathematical models. Until recently, models in computer science, electrical engineering, and control systems concentrated on the one constraint peculiar to their field. But integrating these enormous systems forces one to consider these problems as a whole. We are trying to develop the math to do that."


"Information" may be an abstract notion, says John Preskill, the MacArthur Professor of Theoretical Physics and director of the Center for the Physics of Information, "but in practice it always has some physical form. Whenever we strive to improve information technology, we are trying to find new physical processes." We'll need those processes pretty soon, because in the next few decades, our ability to miniaturize circuits in silicon will hit bottom. "Information technologies for the most part treat electrons and photons like they were basketballs," says Preskill. "You bat electrons around in a circuit, or send photons down a fiber and count them." But we're approaching the size where classical physics falters and quantum effects take over. This isn't necessarily bad—a lot of people have embraced quantum computing as the Next



The atom trap in Kimble's lab. The inset shows a close-up of the two mirrors (in the white box in the main photo), which are labeled M_1 and M_2 . The red arrows show the path of the trapping laser.

Reprinted with permission from "Researchers Achieve Lasing from a Single-Trapped Atom," *Physics Today*, vol. 57, no. 1, p. 16, January 2004, © 2004 American Institute of Physics.

Right: A scanning electron micrograph of one of Vahala's photon race-tracks—the flared region around the rim of the mushroom's cap. Far right: An idealized representation of how a stored photon (red) could change the state of a passing photon (blue) in a fiber-optic line. In reality, the photon's color is one property that could NOT be changed, but it's easier to draw than, say, phase or polarization.



Reprinted with permission from Vahala et al., *Nature* vol. 421, pp. 925–928, © 2003 Nature Publishing Group.

Big Thing, because by exploiting a system that is in all possible states at once until you measure it, “you can spectacularly accelerate the solution of a big class of problems.”

But we’re a long way from a quantum Pentium. People like postdocs Warwick Bowen and Tobias Kippenberg (MS ’00, PhD ’04) are still trying to build individual logic elements in which one photon changes the state of a second one—giving it a left-hand twist instead of a right-hand one, for example. The catch is that, unlike Jedi light sabers, photons pass through each other unhindered. They do interact weakly with atoms, however, providing a potential middleman, and Jeff Kimble, the Valentine Professor and professor of physics, greatly enhances this interaction by placing a single atom in the tiny void between near-perfect mirrors. A reverberating photon within this optical resonator smacks the levitated atom a million times or so, and, like a transistor, this turns a small signal into a big one. And Kerry Vahala (BS ’80, MS ’81, PhD ’85), the Jenkins Professor of Information Science and Technology and professor of applied physics, builds ring-shaped silicon microstructures that store light—photon race-tracks some 60 microns (millionths of a meter) in diameter and six microns thick—that sit on stalks like little silicon mushrooms. The cramped dimensions intensify the photon’s electric field enormously, and Kimble’s methods can be used to trap a single cesium atom within that field. The pumped-up field distorts the atom—enough, Bowen hopes, to some day affect passing photons one by one, providing a basic building block for the quantum Internet.

So much for quantum computing—what about computing quanta? “Information science is ripe to illuminate a lot of other fields,” says Preskill. “What new insights can we get into physics? Information lost inside a black hole gets coughed up in the form of Hawking radiation, which is a quantum effect. I think the really juicy issues arise when we think about information confronting quantum physics.”

Postdocs Frank Verstraete and Guifré Vidal have invented new methods for doing quantum many-body physics on classical, i.e., ordinary computers. This has been a burgeoning field for 30-some

years as people try to simulate the behavior of materials that owe their properties to quantum effects—high-temperature superconductors, for example. Most simulations use the so-called Monte Carlo method, which generates random samples for statistical analysis. It’s very straightforward—if you can ensure that the samples include a proportional representation of all the possible states of the system. A more sophisticated method called the Density Matrix Renormalization Group (don’t ask) has been stalled since the early ’90s, says Preskill. “People have had Moore’s Law on their side, so there are bigger and bigger computers that can solve bigger and bigger problems, but the techniques have not advanced very much in 15 years. Verstraete and Vidal have made tremendous advances in six months, because they had a much deeper understanding of how information is carried by quantum systems.”

Quantum entanglements affect all parts of a system at once, making them fiendishly difficult to simulate. There’s no shorthand way to write down all the correlations and, says Verstraete, “Each particle doubles the size of the computation. So if 10 particles takes 10 minutes to run, 11 particles takes 20 minutes. The time increases exponentially.” But there are degrees of entanglement, and most of the systems of real-world interest aren’t Gordian knots. Says Verstraete, “Most of the correlations are redundant, so we found a way to compress the uninteresting ones and extract the very few numbers that tell you about the physical state of the system.” “It’s just an amazing achievement, and it’s having a really big effect,” says Preskill. Until now, people have mainly simulated ground states at zero temperature because modeling excited states—which is where all the action is—was just too difficult. But Verstraete and Vidal can track the dynamics of hundreds of atoms as an excited state is induced, peaks, and then decays.

Other center members are trying to figure out how to integrate photons into the silicon world, which won’t fade away any time soon, and are looking at molecules, such as carbon nanotubes, that could be adapted for computing. But building complex machinery from molecule-sized parts is no cakewalk—how do you put all those tiny pieces in

the right places? Nature uses a program encoded in the genes. Inspired by this, Senior Research Fellow Paul Rothemund (BS '94) and Assistant Professor of Computer Science and Computation and Neural Systems Erik Winfree (PhD '98) are making DNA "tiles" that spontaneously assemble into complex patterns based on information contained in the DNA. This raises some interesting questions about how information can be used to direct physical processes, Winfree says. "How can self-assembly be programmed to create a desired shape or pattern—such as a circuit layout for molecular electronics—and how can mistakes in self-assembly be controlled?" Like many faculty members, Winfree thus has a foot in two centers, the other one being the Center for Biological Circuit Design.

Cells do amazing things with seemingly slapdash components. The body heals broken bones and fights off diseases, and we walk around and we do crossword puzzles, all with flimsy, floppy protein molecules packed into cells that keep dying. There's nothing magical about the stuff we're made of, so clearly the miracles are in the circuits—broadly defined—that they're organized into. How do these circuits work? And what else can be done with the same components? Can we find Bruck's "calculus" for biology, and will it ultimately lead to a software package that will accept a high-level design and spit out the genes that will automatically grow that circuit?


The goal of the Center for Biological Circuit Design (CBCD), says Paul Sternberg, Morgan Professor of Biology, investigator, Howard Hughes Medical Institute, and director of the center, "is to learn about biological circuits by trying to build them." Fortunately, a huge catalog of parts is available—every protein or regulatory network that has ever been published. There are actually three nested levels of circuitry, says Sternberg: networks of signaling molecules within a cell that handle such things as regulating metabolism or allowing an amoeba to find and engulf its prey; circuits consisting of several cells, such as the ones that coordinate our defense against infection; and the vast neuronal circuits that are responsible for, say, understanding speech. The CBCD will initially tackle the first two, leaving the brain to the ganglion of neuroscientists on campus. Says Sternberg, "The whole point of IST is to try to abstract what's general. And here, in terms of circuits, we believe that the

general principles will apply across different levels." By biological standards, the human brain with its 20 to 50 billion cerebral-cortex neurons is only middlingly complex—a protein molecule can have 10 thousand atoms, a cell can contain a billion macromolecules, and the heftier *E. coli* reader might consist of up to 100 trillion cells. That's 27 orders of magnitude of organization from an atom to a person, which is like going from the diameter of an atom to the distance to Sirius.

On the intracellular level, Assistant Professor of Biology and Applied Physics Michael Elowitz is examining "primitives"—basic functions that show up pretty much everywhere. One really basic function is gene regulation, in which turning on one gene produces a protein called a transcription factor that turns another gene on or off, stimulating or suppressing the production of *its* protein, which may in turn be another transcription factor, and so on. Elowitz, Caltech staff member Jonathan Young, Nitzan Rosenfeld and Uri Alon from the Weizmann Institute of Science in Rehovot, Israel, and Peter Swain from McGill University have been tracking the concentration of a specific transcription factor (fluorescently tagged to light up yellow) and the protein that it regulates (tagged to light up cyan) in a single *E. coli* bacterium through many cycles of cell division. The idea was to see how noise in the regulatory circuit—the randomness of biochemical reactions in the face of many competing processes, differences in the cell's environment, and the state of the cell itself—affected the circuit's performance. Elowitz calls it "popcorn biochemistry" because "we can determine how biochemical parameters vary from cell to cell, or in a single cell over time, just by watching movies of these cells." The study showed that gene regulation embodies a fundamental trade-off between speed and accuracy, Elowitz says. "If you want a cellular circuit to really accurately control the level of a transcription factor, it would take a very long time." In real life, speed is usually more important.

On the cellular level, Frances Arnold, the Dickinson Professor of Chemical Engineering and Biochemistry, grad student Cynthia Collins, and Ron Weiss, Subhayu Basu, and Yoram Gerchman at Princeton have developed circuits in which sender cells emit a tracer molecule called acyl-homoserine lactone, or AHL, which the surrounding bacteria detect. Each bacterium has been bred to respond

Below: A "movie" of one of Elowitz's fluorescent bacteria as it divides and becomes a colony. For greater contrast, the yellow-fluorescing cells have been colored red, and the cyan ones green. The insets show the original bacterium outlined in white. The numbers are elapsed time in minutes.




A schematic of Arnold's cellular band-pass filter. The sender cell emits molecules of ALH (purple dots) that diffuse evenly out in all directions, so detector cells at greater distances get diminishing doses. At close range, the ALH receptor protein (LuxR) turns on both the $LacI_{M1}$ and the CI proteins (green arrows). The $LacI_{M1}$ protein inhibits (red T-bar) the production of Green Fluorescent Protein (GFP), trumping the CI protein that inhibits the production of another protein called LacI that in turn inhibits the production of GFP. At intermediate ALH concentrations, not enough $LacI_{M1}$ is produced to shut down GFP production, but CI still inhibits the other inhibitor. This causes the cell to light up. At still lower ALH levels, CI turns off too, freeing the LacI protein to shut down GFP production. Got all that? And this is a very simple regulatory scheme, as these things go...

to AHL at a specific concentration—the cellular equivalent of a band-pass filter—and when it does, it turns on a fluorescent gene that makes it glow. “It’s a little model of how organisms develop,” says Arnold. “The cells communicate via AHL and turn on different genes. In this case, it creates a bull’s-eye pattern in a homogeneous lawn of bacteria.” Taken to its logical conclusion, this ability to lay down a gene-expression pattern of your choosing gives you a way to grow complex structures, maybe even molecular computers, automatically. Or the bacteria could be used as sensors by adapting them to recognize other substances—a whiff of TNT in a suitcase, perhaps. And since much of biology these days has to do with tracing signals carried by very rare proteins, a sensor with a big, easy-to-read signal could be a biotech bonanza. But more importantly, says Arnold, “we demonstrated that you can cobble together all these weird pieces from various organisms to make a human-designed system that does something nature doesn’t do.”

Sternberg sees biological computation not for general-purpose processors (at least, not any time soon!), but for embedded control “chips” to manage other microbes. “Even if they’re slow, and don’t do your taxes, they could run a little ecosystem on Mars that makes sugar. That’s been in science fiction for decades.” Assuming that we can fill a spaceship with modified pond scum from the lakes that lie beneath Antarctica’s ice fields and send it to Mars, it could arrive months or years before the astronauts do, and a maintenance-free biological controller would be handy. Closer to home, one could foresee bioreactors—brewer’s vats—in which kidneys, hearts, and other transplantable organs are grown. The biosensor cells would make sure that the right growth factors kick in at the proper times to form healthy organs. Or, to really get down to earth, these supervisory cells could run insulin-adjusting implants for diabetics.

Says Sternberg, “In 10 years, I think there will be a new technology of circuit design. There will be components, and circuits, and people will be using them. We’re still in the days of making computers that fill a room and can add a couple of two-digit numbers—in fact, we’re not quite even there yet. We’re just trying to get *anything* to work.” It helps that the CBCD houses people who are building artificial circuits and people who are reverse-engineering real ones. “Now we say, ‘This cell has switchlike




Reprinted with permission from Basu et al., *Nature* vol. 434, pp. 1130–1134, © 2005 Nature Publishing Group.

behavior—what mechanism is it using?” It would be nice if you could say, ‘Well, there are four different ways that cells usually do that.’ It would be even better if you could say, ‘Well, there’s one way that they usually do it, let’s go test that one first.’”

The theoretical underpinnings will emerge naturally, Sternberg thinks. “The word on the street is that biology doesn’t have that many abstractions. We want to generalize from special cases, lumping phenomena into mechanisms, and lumping mechanisms into variations of the same mechanism. And another good thing about IST is that our nonbiologist colleagues insist on abstractions. They’re not going to listen to 20 hours of special cases. So they push us, push us, push us, and we’ll get there faster.”

Then there’s the ultimate information-processing system—humanity en masse. Each of us as individuals holds little nuggets of information—some



of it incorrect, some of it opinion—that somehow produces a computational result, be it a stock price or a new president. The Social and Information Sciences Laboratory (SISL, pronounced “Sizzle”), directed by Matthew Jackson, the Wasserman Professor of Economics, looks at existing social and economic institutions to see how they work, and attempts to apply these insights to the design of new ones.

Some kinds of information flow are quite subtle. “Statistically, over a broad range of professions, more than 50 percent of people find jobs through social contacts,” says Jackson. So forget the want ads and Monster.com—the more friends you have, and the better placed they are, the better your access to jobs. Conversely, if all your friends are unemployed, you’re in a classic negative-feedback loop and you might as well stay in bed. “While labor economists have worked for a long time to explain why there are pockets of unemployment, there’s a lot we don’t know. Now we can begin to try to model these geographic patterns, and other socioeconomic patterns. Different social networks have different properties, and networks differ across societies and ethnic groups.” Ultimately, Jackson hopes to be able to figure out what kinds of policies would help people trapped in the wrong sorts of networks.

SISL melds engineering analyses and studies of human behavior, says Jackson. “For instance, in economics, we’ve always assumed that people can handle an auction protocol where you might have to bid on a large number of items at once. Say you’re bidding on broadcast-frequency licenses for cell phones from the FCC, and you’re thinking, ‘Well, I really want the license in Los Angeles only if I can also get the license in Riverside, so if my Riverside bid isn’t going well I want to drop out of


the L.A. auction, but if I drop out there, do I want to get the San Francisco license instead?’ Computing your optimum bid is a very complicated problem.” So postdoc Ron Lavi has been using techniques from his computer-science background to develop multi-object auctions that people can actually use without their brains exploding. And economists traditionally deal with equilibria, that is, the final prices of things, says Jackson. “With all the information we have about markets, we still don’t understand price formation. We know what equilibria look like, but how you get there, and *when* you get there, or *if* you get there, remains a mystery.” But engineers are used to systems in motion, so postdocs Sean Crockett, an economist, and Tudor Stoiculescu, an electrical engineer, are trying to apply engineering methods to track the forces at work in the marketplace.

In a similar vein, John Ledyard, the Davis Professor of Economics and Social Sciences; Richard Murray (BS ’85), professor of mechanical engineering; and Mani Chanday, the Ramo Professor and professor of computer science are looking at electricity markets. Part of the project involves experiments in Caltech’s Social Sciences Experimental Laboratory, in which subjects play the parts of the various utilities, consumers, network operators, and so on. The idea is to blend economics and engineering to design better distributed control systems without having to run a full-scale experiment on the state of California, as we did a few years ago. Says Ledyard, “Most analyses of power grids—both economic and engineering studies—rely on equilibria, which do not provide much insight into robust control.”

These new centers join the Center for Neuro-morphic Systems Engineering (CNSE) and the Lee Center for Advanced Networking, which served as a model for them. For years these two centers have been drawing faculty from across campus to work on problems that lie in the cracks between disciplines, and supporting studies that are hard to get funded through traditional means.

“Everything we do in CNSE is IST-related,” says director Pietro Perona, professor of electrical engineering. “We take neurobiological principles and use them in engineered systems, and use engineering expertise to try to understand the brain.” The center hopes to one day build autonomous intelligent machines. This may summon up visions of heroic robots rescuing little girls from burning buildings (or evil robots for global domination, depending on your predilections), but the reality is much more mundane. “Right now you have lots of machines around you—your car, your washing machine, your telephone. Many of them have microprocessors, and memory, and sensors, so they *could* figure stuff out about the world, but we don’t know how to do it,” Perona says. A high-end digital camera could learn to locate all the human faces in the viewfinder, for instance, and meter off of them instead of the bookshelves that hap-

Reprinted with permission from Basu et al., *Nature* vol. 434, pp. 1130–1134, © 2005 Nature Publishing Group.



Clockwise, from above:

1) A petri dish with the sender colony visible in the middle. 2) An experiment with two strains of sensor cells that have been tuned to different concentrations and fluoresce in different colors. 3) Additional sender colonies make more complex patterns possible.

Your cell phone is part of an electronic network, and the people in your cell phone's phone book are part of a social network. Are they governed by the same mathematics?



pen to be in the center of the frame. Then, if the camera were told by your computer that you tend to brighten your pictures in Photoshop, it could even learn your preferences. “Machines now are sitting lumps of matter, and we have to turn knobs, or read manuals that train us how to use menus. We work for the machine in some way, which is paradoxical since the machine should help us,” he added as he fiddled with a webcam trained out his office window. Despite his ministrations, the cam resolutely adjusted its exposure to the shadows of the foreground arcade, washing out the vista beyond. “There is no reason why we cannot design docile behavior in machines.”

And finally, the Lee Center was founded in 1999 by Caltech trustee David Lee (PhD '74) to create the technology needed for a global wireless and fiber-optic communication system that would be as ubiquitous and reliable as indoor plumbing. It was the first big center at Caltech to be privately funded, says director David Rutledge, the Tomiyasu Professor of Electrical Engineering and associate director of IST, and “it opened our eyes to a different kind of flexibility. David Lee wanted us to start a lot of small projects, so we fund 13 faculty members, and they decide what to do. When people follow what they are interested in, it often leads to quite new things.” Indeed, the Lee Center has been a fruitful source of start-ups and spin-offs, which “suggested that we think about a bigger, much more ambitious project, which is IST.” Lee also had the radical notion of funding the center for 10 years, period, on the logic that by then we'd either have solved the networking problems it was set up to address or we'd quit throwing good money after bad.

IST is taking a leaf from Lee's book—its four founding centers expire a decade from inception and new ones will take their places, ensuring a steady supply of fresh ideas. For the same reason, most of IST's seed money from the Gordon and Betty Moore Foundation is going into graduate

students and postdocs. Says Preskill, “We're trying to attract exceptionally bright people at the peak of inventiveness in their careers, and see if something exciting will happen.” Sternberg agrees, saying, “The postdocs are running around campus, coming up with ideas, and instigating things, and that's the glue that holds us together. Someone says, how about building this, and someone else says, you know, I've always wanted to try that. Now a lot of those projects will actually get implemented, which is the leverage that we really want.” IST hired 23 postdocs last fall, and Bruck notes that a couple of them deferred faculty positions for a year in order to come. The initiative is also hiring several junior faculty members, the first of whom, Assistant Professor of Electrical Engineering and Computer Science Tracey Ho, will arrive this fall to work on joint network codes with Professor of Electrical Engineering Michelle Effros.

Setting up multidisciplinary research programs is the easy part. IST should also define a curriculum for this emerging discipline, says Bruck. Some of the core courses already exist—Boolean algebra, probability theory, and the like—but they haven't coalesced into a logical sequence, and new classes will be needed to fill in the gaps. “What should we teach? How do we integrate research into basic classes at the freshman level? That's still not clear. I wish we could have a Feynman's *Lectures on Physics* on information. Physics was the way to educate the generalists of the Industrial Age, and it was extremely successful. Electrical engineering and computer science emerged out of physics. But now we need to educate the generalists for the Information Age.”

It's starting to happen. In 2002, Assistant Professor of Computer Science Andre DeHon and Winfree launched the Computing Beyond Silicon Summer School, which exposes a select group of undergrads from across the country to the emerging fields of bio-, molecular, and quantum computing. Last year Murray, Elowitz, and Assistant Professor of Chemical Engineering Christina Smolke did a SURF summer school on synthetic biology, which is what the art of growing logic elements and circuits in bacteria is called. And Chandy and Ledyard are teaching an upper-level undergrad course at the intersection of economics, game theory, and computer science. The class looks at “networks of systems that integrate markets with physical constraints,” says Ledyard, who goes on to note that this includes health-care systems as well as power grids.

Says Bruck, “In time, I think ‘information’ will be a first-order concept. So in 20 years, if a high-school student asks her friend, ‘Do you like information?’ like, ‘Do you like algebra?’ the other girl will say ‘Yes,’ or ‘No,’ or ‘Yes, but I hate the teacher.’ But the other day I asked my daughter, a high-school junior, ‘Do you like information?’ and she said, ‘*What?!?*’” □

PICTURE CREDITS:

6-7 – Erik Jorgensen; 9, 15 – Doug Cummings; 10 – NOAA/DMSF, Jeff Kimble; 11 – Deniz Armani; 12-13 – Jonathan Young & Michael Elowitz