Columbia Engineering can claim that the information era began when Herman Hollerith, an 1879 graduate of the Columbia School of Mines, founded the company that was to become IBM. Today, computers, microcomputers, computerized machinery, robots, fiber optics, and all manner of digital technologies provide a research area in which many of our faculty are engaged, by advancing digital frontiers and cybersecurity to keep our information safe.

INFORMATION

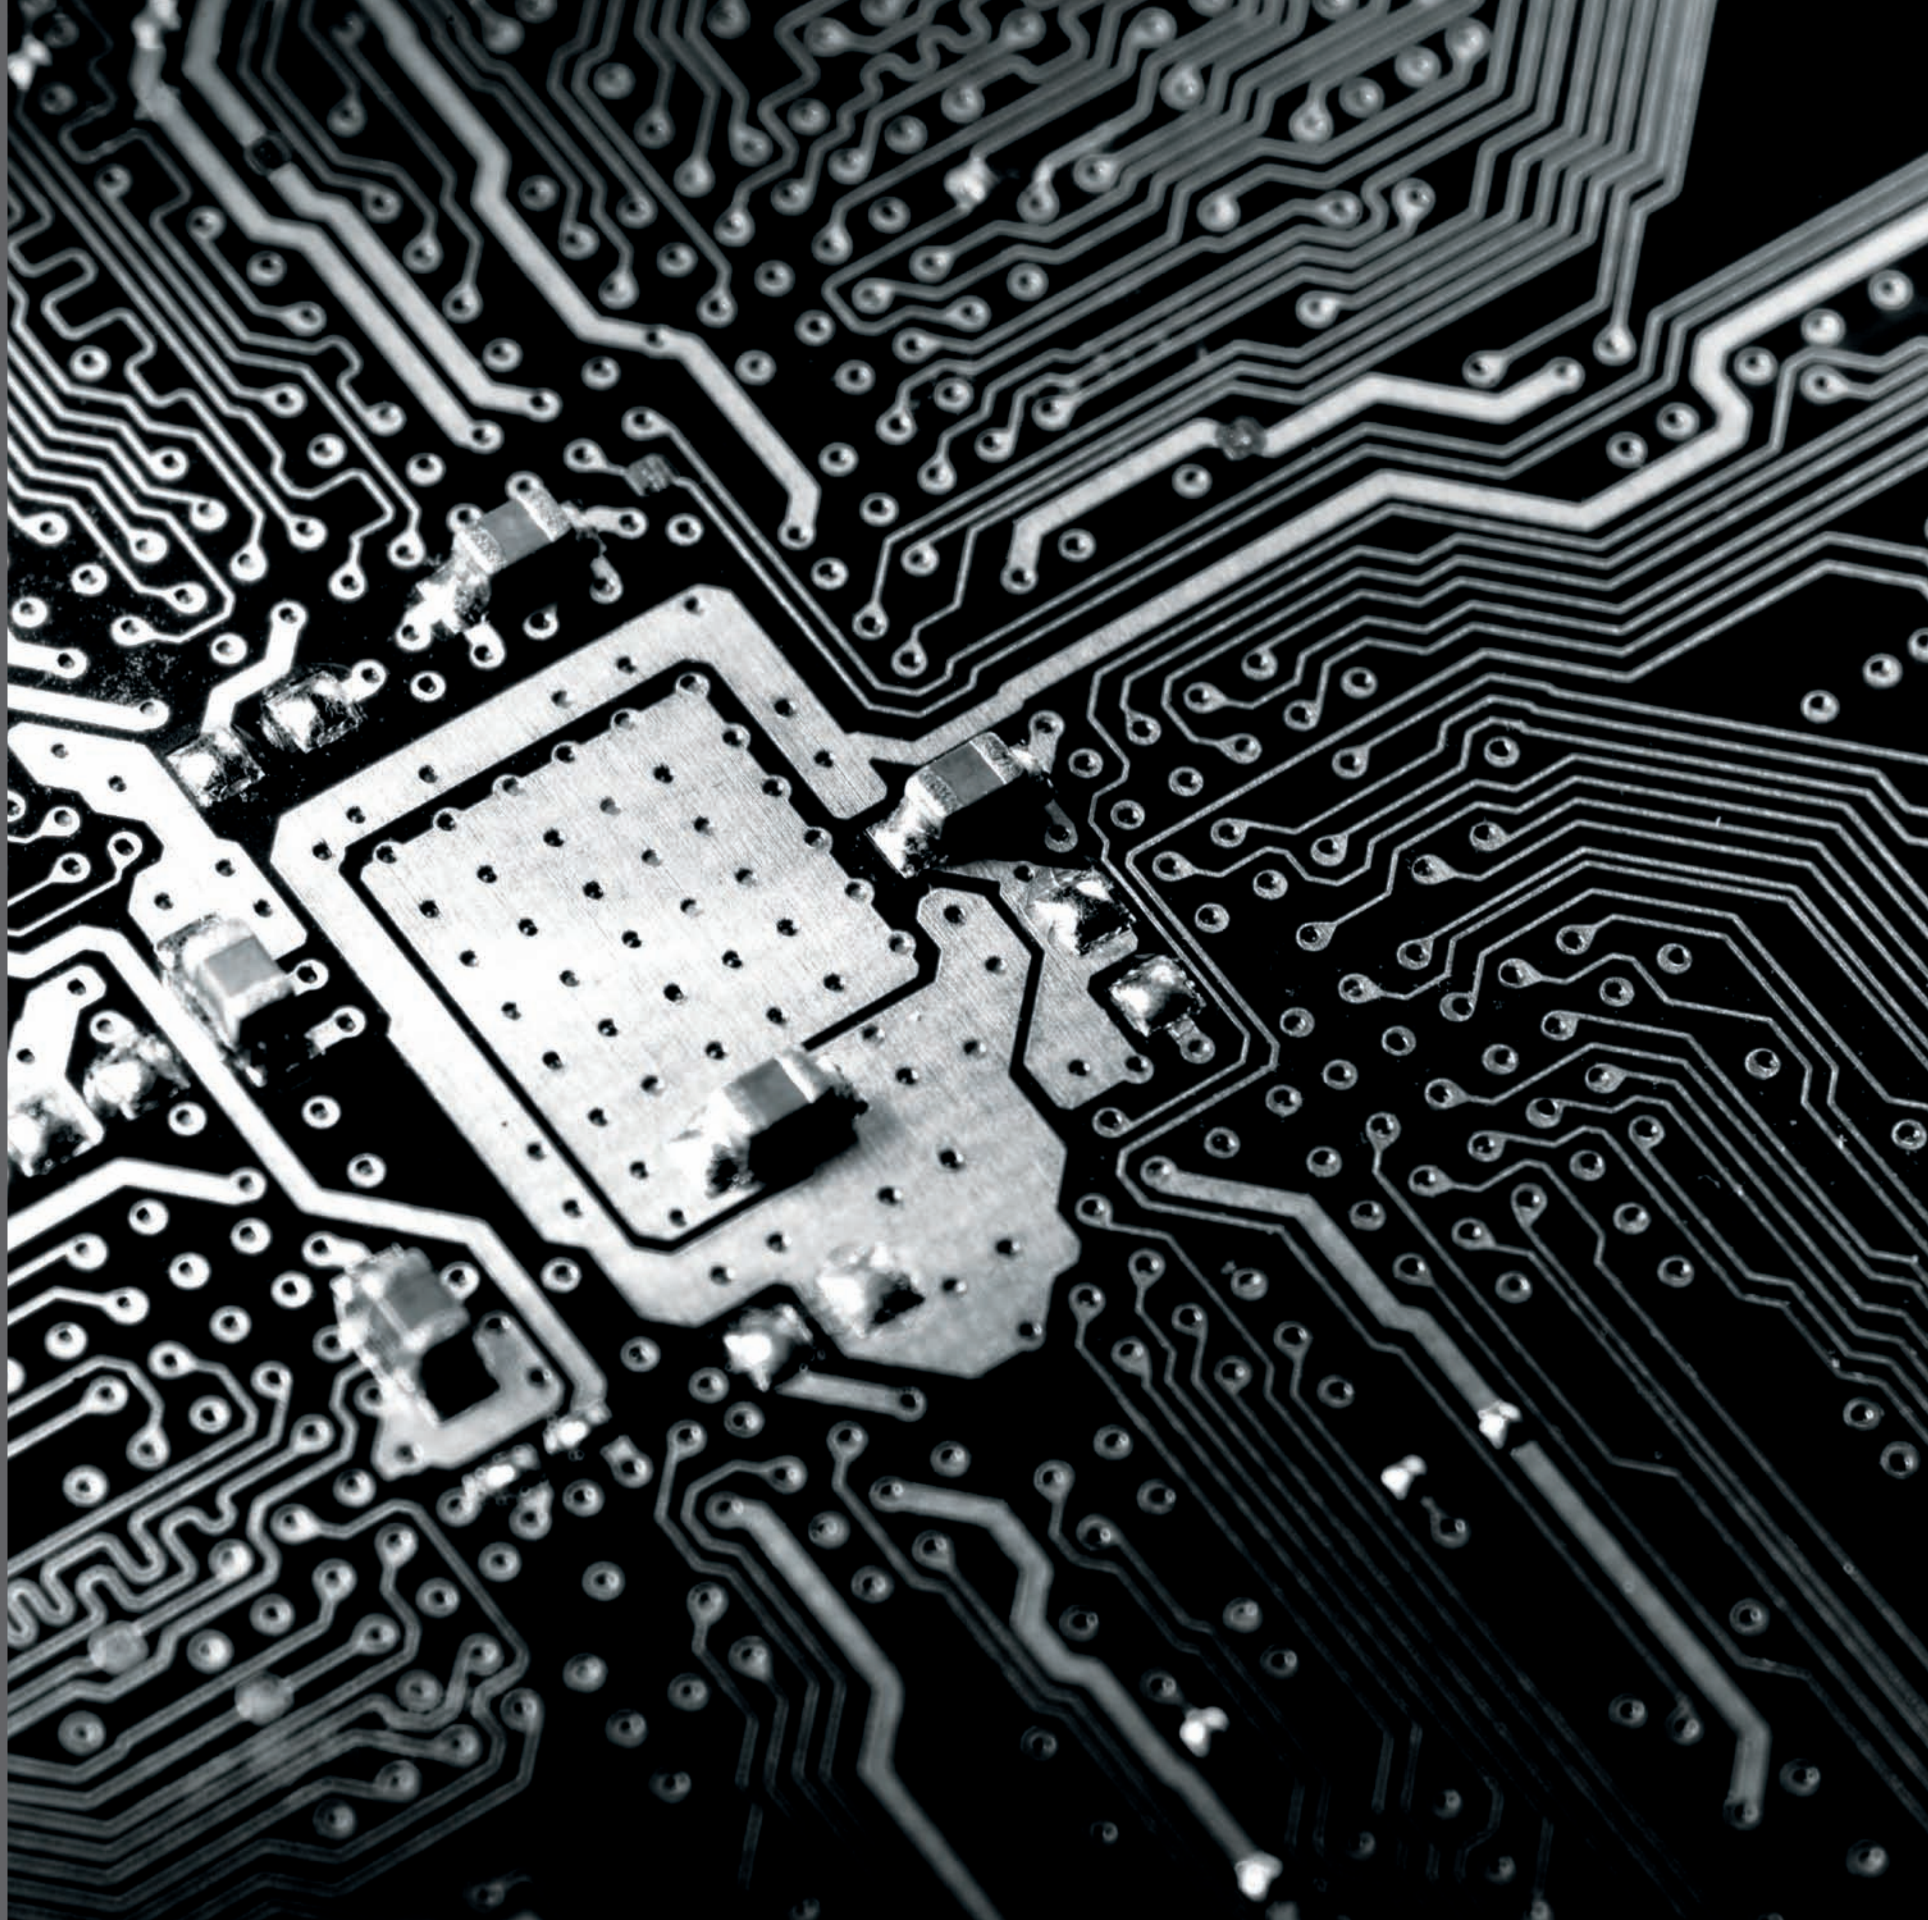

Creating Reliable Programs from Unreliable Programmers

ALFRED V. AHO

Lawrence Gussman Professor of Computer Science

For Alfred Aho, the question is simple: "How can we get reliable software from unreliable programmers?" The issue is more than academic. Aho can point to such high-profile fiascos as a $1 billion write-off for failed flight control software and hundreds of millions of dollars spent fixing an airport's automated baggage handling system.

In fact, a 2002 National Institute of Standards and Technology (NIST) study found that software defects cost the economy $60 billion annually and account for 80 percent of software development costs. Even then, Aho estimates that most commercial software has 1,000 to 10,000 defects per million lines of code.

"If you're developing a computer game, that doesn't matter much. But if you're programming a pacemaker, it's a matter of life and death," he said.

Aho's goal is to create a system that automatically tags potential problems. He hopes to do this by using the technology behind compilers, programs that translate easy-to-use programming languages like C into instructions a computer processor can understand.

When a compiler translates a program, it captures details about how it was built. Aho wants to compare this actual implementation with the program's technical specifications, which define such things as naming conventions, allowable operations, associations, data sets, and order of functions. This is similar to inspecting a building's structure, wiring, and plumbing against schematics and code.

Software, however, is more complex. "Let's say the source program has one million lines of code, and you want to look for all examples of addition," Aho explained. "It's written by several different people. Some may not have consulted the specification. They might use their own names for variables. Instead of writing 'add,' they might write it as 'plus.'"

Those subtle changes make it incredibly difficult to track errors. A plus function might use different data types than an add function, and produce unequal results. Or a programmer may discover a problem involving add functions, but fail to look for plus functions to see if the same problem exists.

"All large programs have a specification document that itemizes how the program should be written. I would like to specify a property from this document and test for its properties in the software. We already know how to create tools that do some of this in compilers. Now we want to extend these tools to software development," Aho said. "This is a long-term project, but if we can make a small dent in software development and maintenance costs, we can save billions of dollars."

B.A.Sc., University of Toronto (Canada), 1963; M.A., Princeton, 1965; Ph.D., Princeton, 1967

Helping to Rapidly Transform Materials for Engineered Systems

KATAYUN BARMAK

Philips Electronics Professor of Applied Physics and Applied Mathematics

Metallic films are critical to many modern technologies such as integrated circuits, information storage systems, displays, sensors, and coatings. In semiconductor chips, these metallic films interconnect the transistors that amplify and switch electronic signals. As the dimensions of metallic films shrink into the nanoscale regime, their structure, morphology, and arrangement of the boundaries between the grains that the material is made of change. When these changes happen, there is a profound impact on their properties and on the performance and reliability of the engineered systems they are made for.

When a complete understanding is gained about how these metallic materials form, evolve, and change, new or improved materials for engineered systems like computer hardware and advanced permanent magnets that underlie the operation of generators, alternators, and motors can be developed.

Katayun Barmak works to discover, characterize, and develop materials for engineered systems; to develop theories and models for phase transitions, structure and morphology evolution in metallic materials; and to understand the relationship between structure and property. Her aim is to quantify and to understand the differences in materials structure at the macro-, micro-, and nano-scales and to investigate the impact of these differences on the properties exhibited by the material. Her studies of materials structure immerse her in the exhilarating and powerful world of electron microscopy.

Her research interests include thin-film phase transformations and microstructures, high throughput electron diffraction-based metrology of nanocrystalline materials, identification of a next generation metal to replace copper in semiconductor interconnects, the discovery and development of rare-earth-free advanced permanent magnets, and quantitative kinetic experiments and models of alloys for extremely high-density magnetic recording media. She is also working collaboratively with colleagues in applied mathematics on the development of theories for evolution of materials structure and morphology.

Barmak is a member of the Institute of Electrical and Electronics Engineers; Materials Research Society; American Physical Society; The Minerals, Metals & Materials Society; ASM International; Microscopy Society of America; and Microbeam Analysis Society.

B.A., University of Cambridge (England), 1983; M.A., University of Cambridge, 1987; S.M., Massachusetts Institute of Technology, 1985; Ph.D., MIT, 1989

INFORMATION

Turning a New Leaf on Face Recognition

PETER N. BELHUMEUR

Professor of Computer Science

Could centuries-old techniques used to classify species hold the key to computerized face recognition? Peter Belhumeur certainly thinks so. Face recognition has many potential uses, from verifying financial transactions to recognizing criminals. Today's systems work by superimposing a subject's face over images in a database. If they align, the computer samples pixels from each image to see if they match.

The process is not very reliable. "Recognition algorithms make mistakes that they should never make, like confusing men with women, or one ethnicity with another," Belhumeur said. Belhumeur was working on improving those algorithms when Smithsonian Institution taxonomists asked for help developing software to classify plant species from photos of their leaves.

Instead of superimposing images or matching pixels, Belhumeur drew on the wisdom of taxonomists dating back centuries. They classified plants by asking a series of questions whose yes-or-no answers narrowed the choices until they came to the right plant.

To this end, Belhumeur has developed LeafSnap, a new mobile application available on the iPhone and iPad. The free app allows users to photograph a leaf, upload it, and see a list of possible matches within seconds. LeafSnap's database covers New York City's Central Park trees and the 160 species in Washington, D.C.'s Rock Creek. Belhumeur, who co-developed the software with colleagues at the University of Maryland and the Smithsonian, hopes to eventually map species across the United States and give users the ability to add their own images to the database.

The way this technology works "is exactly the opposite of how computerized object recognition is done," said Belhumeur. "Instead of pixels, we are comparing visual attributes."

Belhumeur wondered if he could use a similar strategy to recognize faces. "Could we develop software that made qualitative decisions about each image? Is it a male or female? Young or old? Broad or pointy nose? … If we could build reliable classifiers to answer these questions, " he said, "we could search for pictures based on their attributes."

Belhumeur's system uses roughly 100 labels, ranging from eye and nose shape to hair color and gender. In tests that compare a photo to a known image, like an identity card, it outperforms pixel-based technologies.

It also makes it possible to search for pictures with words that describe visual attributes. "We could search through a database based on a victim's description of an assailant, or use it to search one's seemingly endless collection of digital photos," he concluded.

B.S., Brown, 1985; M.S., Harvard, 1991; Ph.D., Harvard, 1993

Protecting Privacy in Complex Systems

STEVEN BELLOVIN

Professor of Computer Science

Ask Steven Bellovin about computer privacy and he might start by discussing aviation. "The technology is so good, there are no single causes of airplane crashes any more. But when complicated systems interact in complicated ways, you have unexpected failures," he said.

Bellovin has seen that complexity emerge on the Internet. Thirty years ago, he helped create USENET, a precursor of today's Internet forums. He wrote the first book on Internet security, and is now creating software to simplify network security. He remains an important voice in public discussions about privacy.

"Computers interact with the world around them," Bellovin said. "We cannot be only scientists or engineers. We have to bring our knowledge to the debate. We have no more right to a policy opinion than anyone else, but no less right either."

He sees the Internet's interconnected technologies eroding personal privacy. For example, nearly all commercial websites collect information about users. While some keep that information private, others do not. Anyone can cross-check for-sale databases to unearth personal information.

"In 1994, Congress mandated that telecommunications switches include technology to make it easier to tap phones," said Bellovin. "We could tell this would be abused. Sure enough, someone tapped 100 people in Greece, including the prime minister. When we see proposals like this, it is our obligation as specialists to say something."

Some privacy mechanisms fail because large Websites actually consist of many different services. Not all of them share the same privacy policies. Facebook, for example, stored pictures on servers that did not enforce privacy rules. Hackers could scrape supposedly private data by entering through those servers.

A third area of concern is anonymization, a process that wipes identifying data from database records. Yet many companies can use anonymized data to build detailed records of individuals. Google, for example, captures queries, offers check-out services that record purchases, and owns Double-Click, which tracks clicks for advertisers.

This could enable it to create detailed profiles. "Some people want to see ads about things they like. Others find it creepy that somewhere there's a repository of all your information," Bellovin said.

"Part of the solution is educational," he continued. "We can teach people to protect their privacy. But it's also a technology issue."

His group is looking at better ways to preserve privacy. This includes creating unlinkable aliases, improving the privacy of database searches, and encrypting advertising clicks so merchants cannot access private information.

B.A., Columbia, 1972; M.S., University of North Carolina (Chapel Hill), 1977; Ph.D., University of North Carolina, 1982

INFORMATION

Battling Internet Gridlock with Light

KEREN BERGMAN

Charles Batchelor Professor of Electrical Engineering

Gridlock doesn't just happen on highways. Interlocking congestion that prevents movement is also a threat to the Internet. As more people exchange more information on a more frequent basis, the Internet's traffic management system (routers) is forced to use more energy to forward and receive data between computer networks. As routers lose ground against traffic demands, performance bottlenecks occur.

Photonics, the science and technology of generating and controlling photons, could ease up electronic traffic jams by providing the solution to Internet gridlock. Through photonics, the potential exists to achieve advanced information traffic management performance along with energy efficiency by symbiotically merging the computation-communications infrastructure. Optical routers would transmit data as light, avoiding unnecessary electronic processing. In addition, they would use less power consumption while manipulating gargantuan amounts of data with complete format transparency in a smaller device footprint.

Keren Bergman leads the Lightwave Research Laboratory at Columbia University. She investigates the realization of dynamic optical data routing in transparent optical interconnection networks. Through this work, she is developing potentially disruptive technology solutions with ultra-high throughput, minimal access latencies, and low-power dissipation that remain independent of data capacity. These solutions will ultimately capitalize on the enormous bandwidth advantage enabled by dense wavelength division multiplexing.

Her work on large-scale optical networks focuses on embedding real-time substrate measurements for cross-layer communications. As envisioned by the community, this suite will support a wide range of network science and engineering experiments such as new protocols and data dissemination techniques running over a substantial fiber optic infrastructure with next-generation optical switches, novel high-speed routers, city-wide experimental urban radio networks, high-end computational clusters, and sensor grids.

Bergman's research in large-scale optical switching fabrics includes cross-layer optimized optical substrate and embedded real-time measurements. Her work in optical interconnection networks for high-performance computing systems includes data vortex optical packet switching fabric, optical network interface card and scalable optical packet buffers. Her work in integrable interconnection network systems and subsystems includes parametric optical processes and systems and nanophotonic optical broadband switches. Her work in inter- and intra-chip multi-processor interconnection networks includes on- and off-chip photonics communications for multi-processor systems and silicon photonic devices for networks-on-chip.

B.S., Bucknell University, 1988; M.S., Massachusetts Institute of Technology, 1991; Ph.D., MIT, 1994

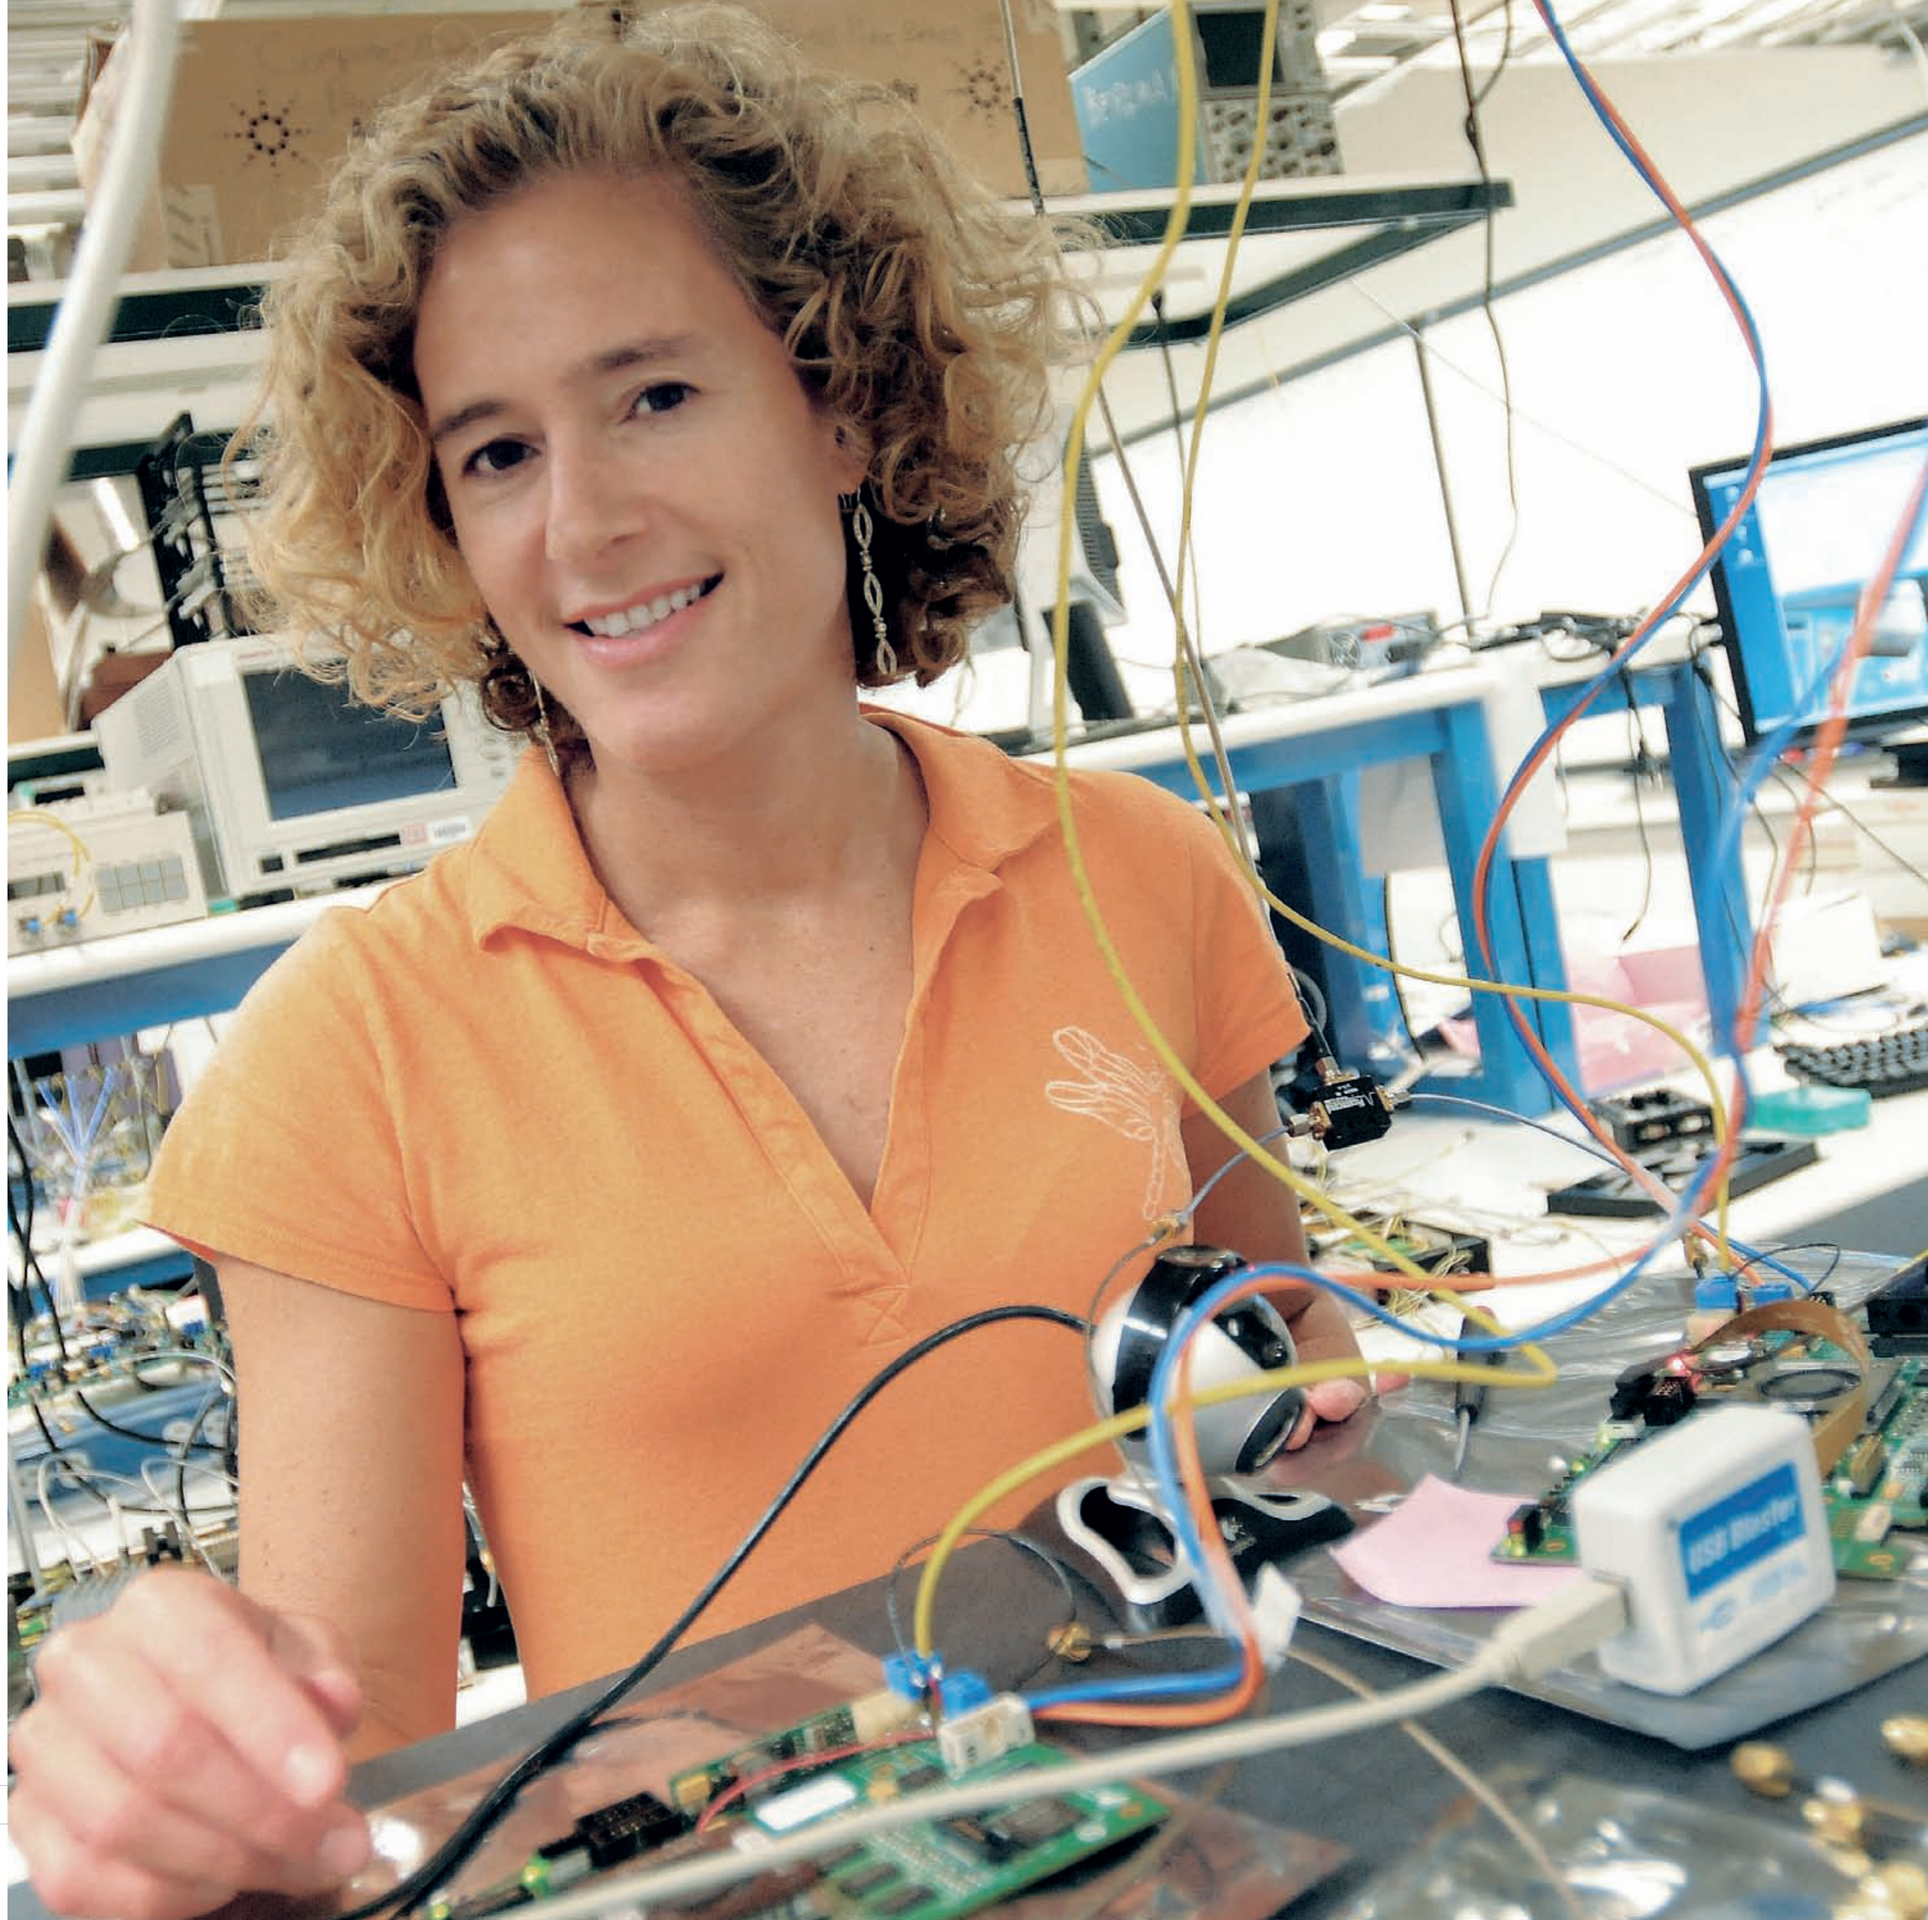

Understanding How Black Swans Evolve

JOSE BLANCHET

Assistant Professor of Industrial Engineering and Operations Research

INFORMATION

In early 2008, few investors saw the whirlwind coming. The financial crisis was what economists call a black swan, an earthshaking event so unlikely, no one anticipates or plans for it. Jose Blanchet would like to rectify this situation.

"I study black swan events by using probabilistic methods. That doesn't mean I predict them. Instead, I use computers to understand how they evolve," Blanchet said. His goal is to help investors see the warning signs of extreme events before they occur, while they still have time to respond.

Blanchet does this by building realistic computer models of portfolios. As they evolve, he shocks them with random events, such as bond defaults and bankruptcies. Ordinarily, the portfolio absorbs the hits. Rarely, very rarely, a combination of random shocks sends values crashing, just as cascading events did to real portfolios in 2008.

The shocks, Blanchet explained, must be truly random. "If you try to model a crisis that simulates these events, you could get it wrong and it will not reflect reality," he said.

"For example, suppose this is 2006 and you want to see what happens if lots of people default on their mortgages. Rather than start with a bankruptcy, we want to start with the events that cascade to create the bankruptcy. We let the probability model capture the events that occur naturally, even if they are rare."

"We look at extreme events in such contexts as queueing networks and risk management of financial and insurance portfolios. We want to understand what happens when there are huge backlogs or when companies post enormous losses. What are the consequences of that? What is the likelihood?"

Ordinarily, it would take a week or two to run enough simulations to generate a single black swan. That is far too slow to build a large enough database to study these events for similarities and differences.

To get around this problem, Blanchet devised algorithms that generate black swans rapidly. He then runs hundreds of simulations using a variety of portfolio models to see how they behave.

"We have a family of models that capture the features we want to study, and a computational tool that lets us observe these events as they unfold," said Blanchet. "It's like watching a crack in a dam. Most of the time, nothing happens. But sometimes it propagates and then the dam goes."

B.S., Instituto Tecnológico Autónomo de México, 2000; Ph.D., Stanford, 2004

INFORMATION

Networking Chips

LUCA CARLONI

Associate Professor of Computer Science

As microprocessors grow more powerful and complex, engineers dream of putting an entire computer system on a single chip. Such chips would be smaller, faster, and more energy efficient than today's designs. To get there, though, we will need to reinvent how we design chips, Luca Carloni argues.

Today, engineers create microprocessors using tools that help them build circuits from libraries of proven designs. Yet new technologies pose many problems for traditional tools. In the past, for example, chips synchronized all operations with a single clock. "Compared to the times needed for computation, on-chip communication was basically instantaneous," Carloni explained. "Today, local calculations run so fast, it takes several clock cycles for remote signals to arrive. This is a physical issue we need to address."

Carloni ticks off other problems. New chips have multiple processors, or cores, whose parallel operations create new challenges in programmability. Billions of transistors create new levels of complexity and generate lots of heat that is hard to remove. Resolving these issues has extended the amount of time and design iterations needed to create new chips.

Those same emerging technologies also offer new opportunities. Instead of trying to develop a system-on-chip with old tools, Carloni proposes reinventing chip architectures and the tools used to design them. "We need to create communication infrastructures that make it easier to integrate new components into our designs," he said.

His solution is a network on a chip. "Our vision is to create an on-chip communication and control infrastructure," he added. "When we have a network that touches every component on a chip, we can dynamically configure the processor to optimize speed or efficiency. We don't have the solution yet, but we're working on it."

He envisions a collection of communications elements—nanoscale wires, switches, and routers—to move data around the chip. The cores would have standard interfaces to plug into the network. A new generation of tools would support component selection and network optimization.

"Instead of designing links between each circuit, you would plug components into a standardized backbone," Carloni said. "This makes it much easier to design processors. Engineers could continuously upgrade and test new components, then plug them in and know they would work on the chip."

Networked chips would also support multiple cores running at different clock rates. Chips could assign tasks to different cores to optimize speed or reduce energy use. "The path towards green computing systems starts with more efficient communication infrastructures," Carloni said.

B.S., University of Bologna (Italy), 1995; M.S., University of California-Berkeley, 1997; Ph.D., UC Berkeley, 2004

Scaling Up the Mobile Social Internet

AUGUSTIN CHAINTREAU

Assistant Professor of Computer Science

INFORMATION

Consider the power of global interconnectedness: One person's tweet about a product can influence a purchase half a world away, another person's email to a group about dissatisfaction with rules can lead to a public protest, and someone else's real-time video during a natural disaster can result in an outpouring of aid. All are astounding outcomes from social networking.

Although social networking is flourishing on today's Internet, it does not make the most of our everyday interaction, Augustin Chaintreau argues. "This is because the technology that personalizes the web to your need does not mirror how you make and keep social connections," said Chaintreau. "When you ask a friend their opinion about a good movie, you do not first have to tell her about your most recent purchases, places you have been and the websites you have visited. But that's more or less what today's computers require you to do." Today's social networking software also requires that you are connected at all times to a server on the Internet, even to interact with nearby people (or objects).

"In real life, we collect and communicate useful data sparingly, and we interact much more with our immediate environment," he said. "Why can't we do that to use the Internet socially—to update a Facebook status, email a friend, collaborate on plans?" One challenge is that humans are incredibly efficient at social interaction. By better understanding natural social networking, which predates its online counterpart, we can then mathematically model that to enhance computer networking performance and outcomes.

Chaintreau works on building algorithms that connect online social networkers more efficiently and more intuitively. What makes these algorithms unique is that they use only local information and exploit mathematical models describing users' behaviors and interactions in groups and organizations. It shows in particular that users should not surrender their privacy. "Many believe that handing out your data is necessary to connect efficiently with your friends. We want to give the users a choice," he said.

These techniques could also allow us to interact socially with many more people and objects, reaching new applications. "When you look at most urgent environmental issues, to save water or organize electricity distribution from renewable sources, many of them could greatly benefit from involving people through a fast, mobile, social Internet," he concluded.

Ancien élève, École Normale Supérieure (Paris), 2002; Ph.D., École Normale Supérieure, 2006

INFORMATION

Developing Next-Generation Visual Search Engines

SHIH-FU CHANG

Professor of Electrical Engineering and of Computer Science

Accurately searching through the glut of visual data available today—digital images produced daily in the thousands or millions—depends upon how closely your verbal description matches the words used to classify the image. It is a frustrating, time-consuming exercise that affects the general user as well as news, media, government, and biomedicine specialists who crave a richer search and browsing experience. An automated visual matching and search technology would not only enhance classification and searching activities, but could also facilitate media forensics, helping to explain if an image has been manipulated, or if it is a natural photograph or computer graphic.

In order for a search engine to visually classify and find images, or determine if images have been tampered with, computers would need to perceive the abundant visual information provided by each individual image. All that data is much like DNA: thousands of genetic concepts of objects, people, scenes, events, and domain-related syntax that make up the individual image.

Shih-Fu Chang, director of the Digital Video and Multimedia (DVMM) lab at Columbia Engineering, targets his research on next-generation search engines for digital images and videos, and has been influential in shaping the vibrant field of content-based multimedia retrieval. His group leads the Columbia ADVENT University-Industry Research Consortium, promoting industrial collaborations with the University's research teams in the media technology area. In addition, his group has actively participated in the development of MPEG-7 and MPEG-21 international standards.

Ranked by Microsoft Academic Search as the most influential researcher in the field of multimedia, Chang's research includes multimedia search, pattern recognition, media analytics, video adaptation, and media forensics. Results include a groundbreaking search paradigm and prototype tools that allow users to find content of similar visual attributes, search videos by a very large pool of visual concept classifiers, and summarize the event patterns and anomalies found in a large collection of video content.

In 1998, he developed one of the first video object search systems, VideoQ, which supported automated spatio-temporal indexing at the object region level. His work has been broadly funded by government and industry and many video indexing technologies developed by his group have been licensed to companies.

Chang is a fellow of the Institute of Electrical and Electronics Engineers and received the IEEE Kiyo Tomiyasu technical field award in 2009.

B.S., National Taiwan University, 1985; M.S., University of California-Berkeley, 1991; Ph.D., UC Berkeley, 1993

Using Game Theory to Study the Internet and E-Commerce

XI CHEN

Assistant Professor of Computer Science

The underlying impetus to every human interaction is competition, which, in turn, affects our decision making. As individuals, we decide how fast we drive in a specific lane of a highway, alongside other drivers, in order to get to a desired destination at a particular time. In business, a board of directors undertakes merger negotiations with another corporate entity in order to achieve an outcome that is profitable for their shareholders. Governments negotiate diplomatically in order to achieve economic and political benefits. In all these interactions, the need to make good decisions is important. To make good decisions, we need to understand how outside influences impact the process.

Game theorists have traditionally applied mathematics to help understand the competitive behavior of rational agents and the decision-making process in the context of economic systems. It's an area of study that is highly valued: eight game theorists have won Nobel Prizes in economics.

In the past decade, computer scientists have witnessed numerous applications of game theoretic approaches and concepts in the study of the Internet and e-commerce, where an absence of central authority opens a new frontier in understanding decision making. Much interest has centered around the new and rapidly growing field called algorithmic game theory. This theory lies at the intersection of computer science, mathematical economics, game theory, and operations research, and examines new and classic game theoretic models through the lens of computation. The goals of algorithmic game theory are to understand and even predict the behavior of selfish agents in order to make Internet-based applications more successful.

Xi Chen studies algorithmic game theory and theoretical computer science with an emphasis on natural and fundamental computational problems that arise from the game-theoretic study of Internet, e-commerce, and other decentralized systems. His current research examines algorithmic issues related to some of the most classic and fundamental models in game theory and economics, and seeks to understand and characterize the intrinsic difficulties in the computation of classical solution concepts in game theory and economics. He is especially interested in how social influence can change the computational landscape of market equilibrium problems.

Chen has won awards for his work on the computation of Nash equilibria and on the computation of market equilibria.

B.S., Tsinghua University (P.R. China), 2003; Ph.D., Tsinghua University, 2007

Exploring the Structure of Abstract Graphs

MARIA CHUDNOVSKY

Associate Professor of Industrial Engineering and Operations Research

In many ways, a good theory behaves like a rock thrown in a pond: It makes a splash and then its ripples spread. Maria Chudnovsky's work in graph theory is like that.

"A graph is a good model for many practical problems," she said. "You can think of the Internet as a graph and the computers on it as vertices; some are connected and some are not. Graph theory can tell us about its structure."

Graph theory does not involve what we normally think of as graphs. Instead, it involves groups of points, or vertices. Sometimes they form geometric objects like squares and pentagons. Other times, they are distributed as randomly as cities or cell phone towers on a map.

Graphs are characterized by the properties of their vertices and the lines, or edges, between them. They can be used to answer problems, from finding the best route for a delivery truck to routing Internet traffic to calculating the shortest itinerary on a GPS.

Chudnovsky works at understanding these attributes. In 2002, her team proved a conjecture about perfect graphs, which are graphs roughly defined as being easy to color. They showed that only two types of defects keep a graph from being perfect, and that all perfect graphs fall into a handful of different categories.

Chudnovsky's proof makes it possible to determine if a graph is perfect without coloring all its vertices. While this may sound like a strictly cerebral exercise, perfect graphs were originally conceived in order to solve a problem in communications theory.

Her work is relevant in other fields as well. Engineers could use her proof to locate wireless towers so their frequencies do not interfere with one another. Knowing whether a graph is perfect or not also helps computer scientists choose efficient algorithms to solve certain problems.

Chudnovsky continues to explore the structure of graphs. Her recent work looked at graphs that did not contain a claw. This structure occurs where three lines, or edges, emanate from a common vertex to form a three-fingered claw.

"We've explicitly described all graphs that do not contain a claw," she said. "Now that our characterization is in place, many problems that seemed to be out of reach can be solved relatively easily."

While her work is highly abstract, her results promise to solve some of the most practical of problems.

B.A., Technion Israel Institute of Technology, 1996; M.Sc., Technion Israel Institute of Technology, 1999; M.A., Princeton, 2002; Ph.D., Princeton, 2003

A Man of Many Words

MICHAEL COLLINS

Vikram S. Pandit Professor
of Computer Science

With people becoming ever more connected around the globe, statistical natural language processing (NLP), a sub-field of artificial intelligence, has become a critical area of research—as the amount of electronic data increases exponentially, so does the need for translating, analyzing, and managing the flood of words and text on the web. NLP deals with the interactions between computers and human languages, often using machine learning to approach problems in processing text or speech. One of the world's leading NLP researchers has just come to Columbia Engineering from MIT: Michael J. Collins, recently named the Vikram S. Pandit Professor of Computer Science, whose work in machine learning and computational linguistics has been extraordinarily influential.

Collins' research focuses on algorithms that process text to make sense of the vast amount of text available in electronic form on the web. The overarching thrust of his work has been the use of machine learning along with linguistic methods to handle difficult problems in language processing. His research falls into three main areas: parsing, machine learning methods, and applications.

Collins has built a parser that can obtain such unprecedented accuracy levels that it has revolutionized the field of NLP: for the first time, a system was able to accurately handle enormous quantities of text in electronic form. His parser is now one of the most widely used software tools in the NLP field.

His development of new learning algorithms has enabled him to make significant advances in several language-processing applications, greatly impacting speech recognition, information extraction, and machine translation. "A major focus of my work is on statistical models of complex linguistic structures," said Collins. "The challenge is to combine sophisticated machine learning methods with these complex structures."

Collins' research also focuses on "efficient search" in statistical models of language, an important challenge in many NLP applications. For example, in parsing, you have to search through a vast set of "possible" structures for a given sentence, in order to find the most probable structure. In translation, you need to search through a vast number of possible translations for the most plausible structure; in speech recognition, you must search through a vast number of possible sentences for the most likely sentence that was spoken.

"I find linguistics fascinating," said Collins. "I really enjoy developing mathematical models for languages. And the algorithms we're developing to process text in intelligent ways have all kinds of intriguing applications."

B.A., University of Cambridge (England), 1992; M.Phil., University of Cambridge, 1993; Ph.D., University of Pennsylvania, 1999

INFORMATION

Modeling Systemic Risk in Financial Networks

RAMA CONT

Associate Professor of Industrial Engineering and Operations Research

In 1987, automated trading programs shoved the market off a precipice. In 2008, a liquidity crisis brought the global financial system to its knees. Rama Cont, who uses probabilistic methods to model financial markets, has studied such system-wide discontinuities for more than a decade. His research on market discontinuities and systemic risk has made him a valued contributor in redesigning financial markets to reduce the impact of major shocks.

"When an epidemic spreads by contact and you cannot vaccinate the whole population, you have to prioritize vaccination resources to prevent further spread," said Cont. "We ask similar questions about market mechanisms that could lead to a financial meltdown."

Cont takes a system-wide view of financial markets. "We cannot understand why several banks failed simultaneously in 2008 by looking at individual bank portfolios," he said. "Instead, we must look at the flow of funds and assets in a network of interlinked portfolios." A theoretical physicist by training, he uses the mathematical language of science to analyze financial networks and identify where they are prone to breakdowns.

In the past, Cont said, regulators promulgated rules that restricted the behavior of individual institutions. "Now, they are trying to look at the market as a whole and assess risks in the entire system. Most markets evolved spontaneously from traders' needs. Some degree of intervention that strengthens their weakest links can make them less vulnerable to disruption," he said.

Cont believes clearinghouses can strengthen the system by acting as intermediaries for trades. They would require trading parties to register their transactions. This would increase market transparency about the price—and risk—of derivatives and other instruments that traded at wildly varying prices in the past.

Clearinghouses would also require deposits on all trades. The amounts would rise as institutions take on more risk. The deposits would act as brakes on risk and help compensate for losses if a party defaulted.

Cont is applying his systemic approach to risk management to the design of new derivatives clearinghouses. He is one of the two academics collaborating with the Market Transparency Working Group, a panel of industry officials and regulators charged with redesigning over-the-counter derivatives markets.

"Some people thought that after the market crash, financial engineering was finished. Instead, it raised awareness about the need for rigorous methods for managing risk," Cont said. "More than ever before, quantitative modeling is in demand now."

Diplôme, École Polytechnique (France), 1994; D.E.A., École Normale Supérieure (France), 1995; Doctorat, Université de Paris XI, 1998

Understanding When Models Behave Badly

EMANUEL DERMAN

Professor of Professional Practice
in the Department of Industrial Engineering
and Operations Research

Emanuel Derman knows something about models. He practiced physics after receiving his Ph.D., but moved to Wall Street in 1985. At Goldman Sachs, he co-developed one of the earliest interest rate models, and later headed their quantitative strategies group. *Business Week* chose his memoir, *My Life as a Quant: Reflections on Physics and Finance*, as one of the top 10 books of 2004.

Models Behaving Badly is Derman's tentative title for his next book. "It's about the different approaches people use to understand the behavior of the world," he said. In it, he distinguishes how theories differ from models, and explains how the unwarranted assumptions of models can lead to incorrect conclusions.

"Theories," Derman explained, "are attempts to grasp the way the world actually is, even if we don't know why. Take Newton's laws. You can't ask why they are correct. That's the way the world is. These are regularities that are always true."

Models are different. "In my view, they are metaphors or analogies," Derman continued. "We say, 'The brain is like a computer,' or 'Stock prices change the way smoke diffuses through a room.' Models are attempts to describe something by using theories that already work in a different field.

"When I first came to finance, I used the principles of physics to try to build something just as truthful. I discovered that although the techniques appear similar, the resemblance is deceptive. When we make analogies, we simplify things," he said.

Many on Wall Street believed their models represented reality. They were disabused of that notion in 2008.

"In physics there may one day be a theory of everything," he said. "In finance and the social sciences, you're lucky if there is a usable theory of anything."

Yet models still have a role to play. "I'm a bit of a Platonist," Derman added. "I think there is some truth out there. I'm trying to distinguish between finding the truth, which is rare, and building models while understanding their inherent limitations.

"Maxwell once remarked that Ampere's experiments could not have led to his results. His experiments seemed to confirm his intuition rather than point to it.

"I believe in intuitive knowledge, but you don't just wake up with it. It comes after a lot of hard work. Models are a step on that road," he concluded.

B.Sc., University of Cape Town, 1965; M.A., Columbia, 1968; Ph.D., Columbia, 1973

INFORMATION

Testing and Correcting Embedded Processors

STEPHEN A. EDWARDS

Associate Professor of Computer Science

A car takes a curve too fast. Before it spins out, its stability control system kicks in. Its microprocessors calculate and recalculate the right amount of force to apply to each wheel, adjusting the brakes many times per second until the car comes under control. Such critical systems often juggle several events at once. A car's stability system must calculate speed, momentum, spin, and dozens of other variables before each application of the brakes.

Unfortunately, its embedded processors can only perform one task at a time. To get around the problem, Stephen Edwards said, programmers slice tasks into many little pieces and have the processor hop between slices at such blindingly fast speeds, it presents the illusion of simultaneity.

This illusion—called concurrency—comes at a cost. Most programmers use the C language to code embedded processors. This involves lots of repetitive programming, and errors can creep in. "It's like writing a phone book by hand. People could do it, but there would be lots of mistakes," Edwards said.

Second, concurrent C programs are hard to test. C programs must be translated, or compiled, into the ones-and-zeros language of processors. Their sliced-up nature makes them hard to translate, model, and test. "The only way to tell if they will run fast enough to handle critical calculations is to test and re-test programs until they appear to work," he said.

Edwards has solutions for both issues. First, he has developed a language, Software-Hardware Integration Medium (SHIM), which simplifies programming concurrent events. SHIM reduces the errors that creep into repetitive programs. "We developed algorithms that automate all the bookkeeping necessary to manage simultaneous events," he explained.

He also created a customized compiler that generates testable code. It takes C-like programs and translates the concurrent parts of the programs into a (very long) series of sequential commands. It then re-compiles them back into C.

"This lets you test your program in a model to check its speed and reliability and make improvements," Edwards said. "This will lead to more reliable behavior and maybe fewer huge recalls when embedded processors fail.

"Embedded processors hide in the environment. As hardware plummeted in price, it became possible to put them everywhere. I ask students how many processors they own. They may count their computers or smart phones, but miss their coffee makers, air conditioners, and cars. We made 10 billion embedded processors in 2008."

B.S., California Institute of Technology, 1992; M.S., University of California-Berkeley, 1994; Ph.D., UC Berkeley, 1997

Delving into the Science of Listening

DANIEL P. ELLIS

Associate Professor of Electrical Engineering

There is a big difference between hearing and listening. Listening requires complex auditory processing, which facilitates learning. It's a skill humans use automatically in order to filter out background noise to understand someone's speech; remember a previously heard tune and hum along; or recognize the difference between a ringing phone and ringing alarm and understand what an appropriate response to those sounds would be.

Human listeners are able to handle such mixed signals, but machines—such as automatic speech recognizers—are vulnerable to added interference, even at levels that listeners barely notice. Consider the implications of machines that could respond when called, technology that could classify and retrieve videos by their sound tracks, or applications that could automatically search for audio data the same way we do now for text data.

To make these advances possible, it is important to understand how perceptual systems manage to make precise judgments in noisy and uncertain circumstances. This understanding can then be applied to extracting information from sound commonly encountered in daily life, identifying characteristics of the sounds, classifying them, and matching the sounds to appropriate responses.

Daniel P. Ellis is working on such advances. He is the founder and principal investigator at the Laboratory for Recognition and Organization of Speech and Audio (LabROSA) at Columbia Engineering. This lab is the only one in the nation to combine research in speech recognition, music processing, signal separation, and content-based retrieval in order to implement sound processing in machines.

His chief focus is to develop and apply signal processing and machine learning techniques to extract high-level, perceptually relevant information from sound. His intention is to test theories about how human auditory perception works and enable the creation of machines that can make use of sound information in the same way humans do.

Ellis' work in soundtrack classification pioneered the idea of using statistical classification of audio data for general classification of videos by their soundtracks. Current projects in the research group include speech processing and recognition; source separation and organization; music audio information extraction; personal audio organization; and marine mammal sound recognition.

He is a member of the Audio Engineering Society, International Speech Communications Association, Institute of Electrical and Electronics Engineers, and the Acoustical Society of America.

B.A., University of Cambridge, 1987; M.S., Massachusetts Institute of Technology, 1992; Ph.D., MIT, 1996

Transmitting Information Securely

DIRK ENGLUND

Assistant Professor of Electrical Engineering and of Applied Physics and Applied Mathematics

The assurance of confidentiality is required in all aspects of transmitting information, from the exchange of banking information and health records to military tactics and trade secrets. The problem is there is no foolproof method to ensure that confidentiality. No matter how encrypted the information is that is transmitted, as long as there is a key to decrypt it, there is a weak security link in the chain of communication.

The solution may lie in quantum photonics, the sending and receiving of data in the form of photons—the tiniest particles that make up light. By sending data encoded in photons, the data stream becomes a single-use, self-destructing key. If the message is intercepted, the stream would change, immediately alerting the receivers to the breach. In addition, by intercepting the stream, the disturbance would automatically scramble the message, making it indecipherable.

Quantum photonic networks could decrypt classically-encoded messages in a matter of minutes, rather than months or years as per today's networks. Such networks would allow for absolute security; even another quantum computer would not be able to secretly crack a coded message sent via a quantum network.

Working in the quantum world, addressing present day problems, requires an in-bred curiosity about the nature and behavior of matter and energy on the atomic and subatomic level, and a desire to develop revolutionary applications. Those are the talents of Dirk Englund, who leads the Quantum Photonics Group at Columbia Engineering. He concentrates on quantum optics in photonic nanostructures, with primary applications in communications, computation, sensing, and energy. His research focuses on implementations consisting of quantum bits (qubits) that are encoded in photons and in spins of electrons and nuclei in semiconductors.

Englund's work includes chip-based quantum networks that promise exponential speedups in computational algorithms and unconditionally secure cryptography as well as highly sensitive quantum-limited sensors. Recent works include time-resolved lasing action from single and coupled photonic crystal nanocavity array lasers, and optical modulation based on a single strongly-coupled quantum dot.

His group is also developing spin-off applications that rely on phenomena from cavity quantum electrodynamics (QED) to substantially lower the power consumption of optoelectronic systems for high-speed, low-power devices. These applications have potential for adaptation in high-performance computing. Related projects include radiation-hard electronics and radiation detectors and thin-film solar cells.

B.S., California Institute of Technology, 2002; M.S., Stanford, 2008; Ph.D., Stanford, 2008

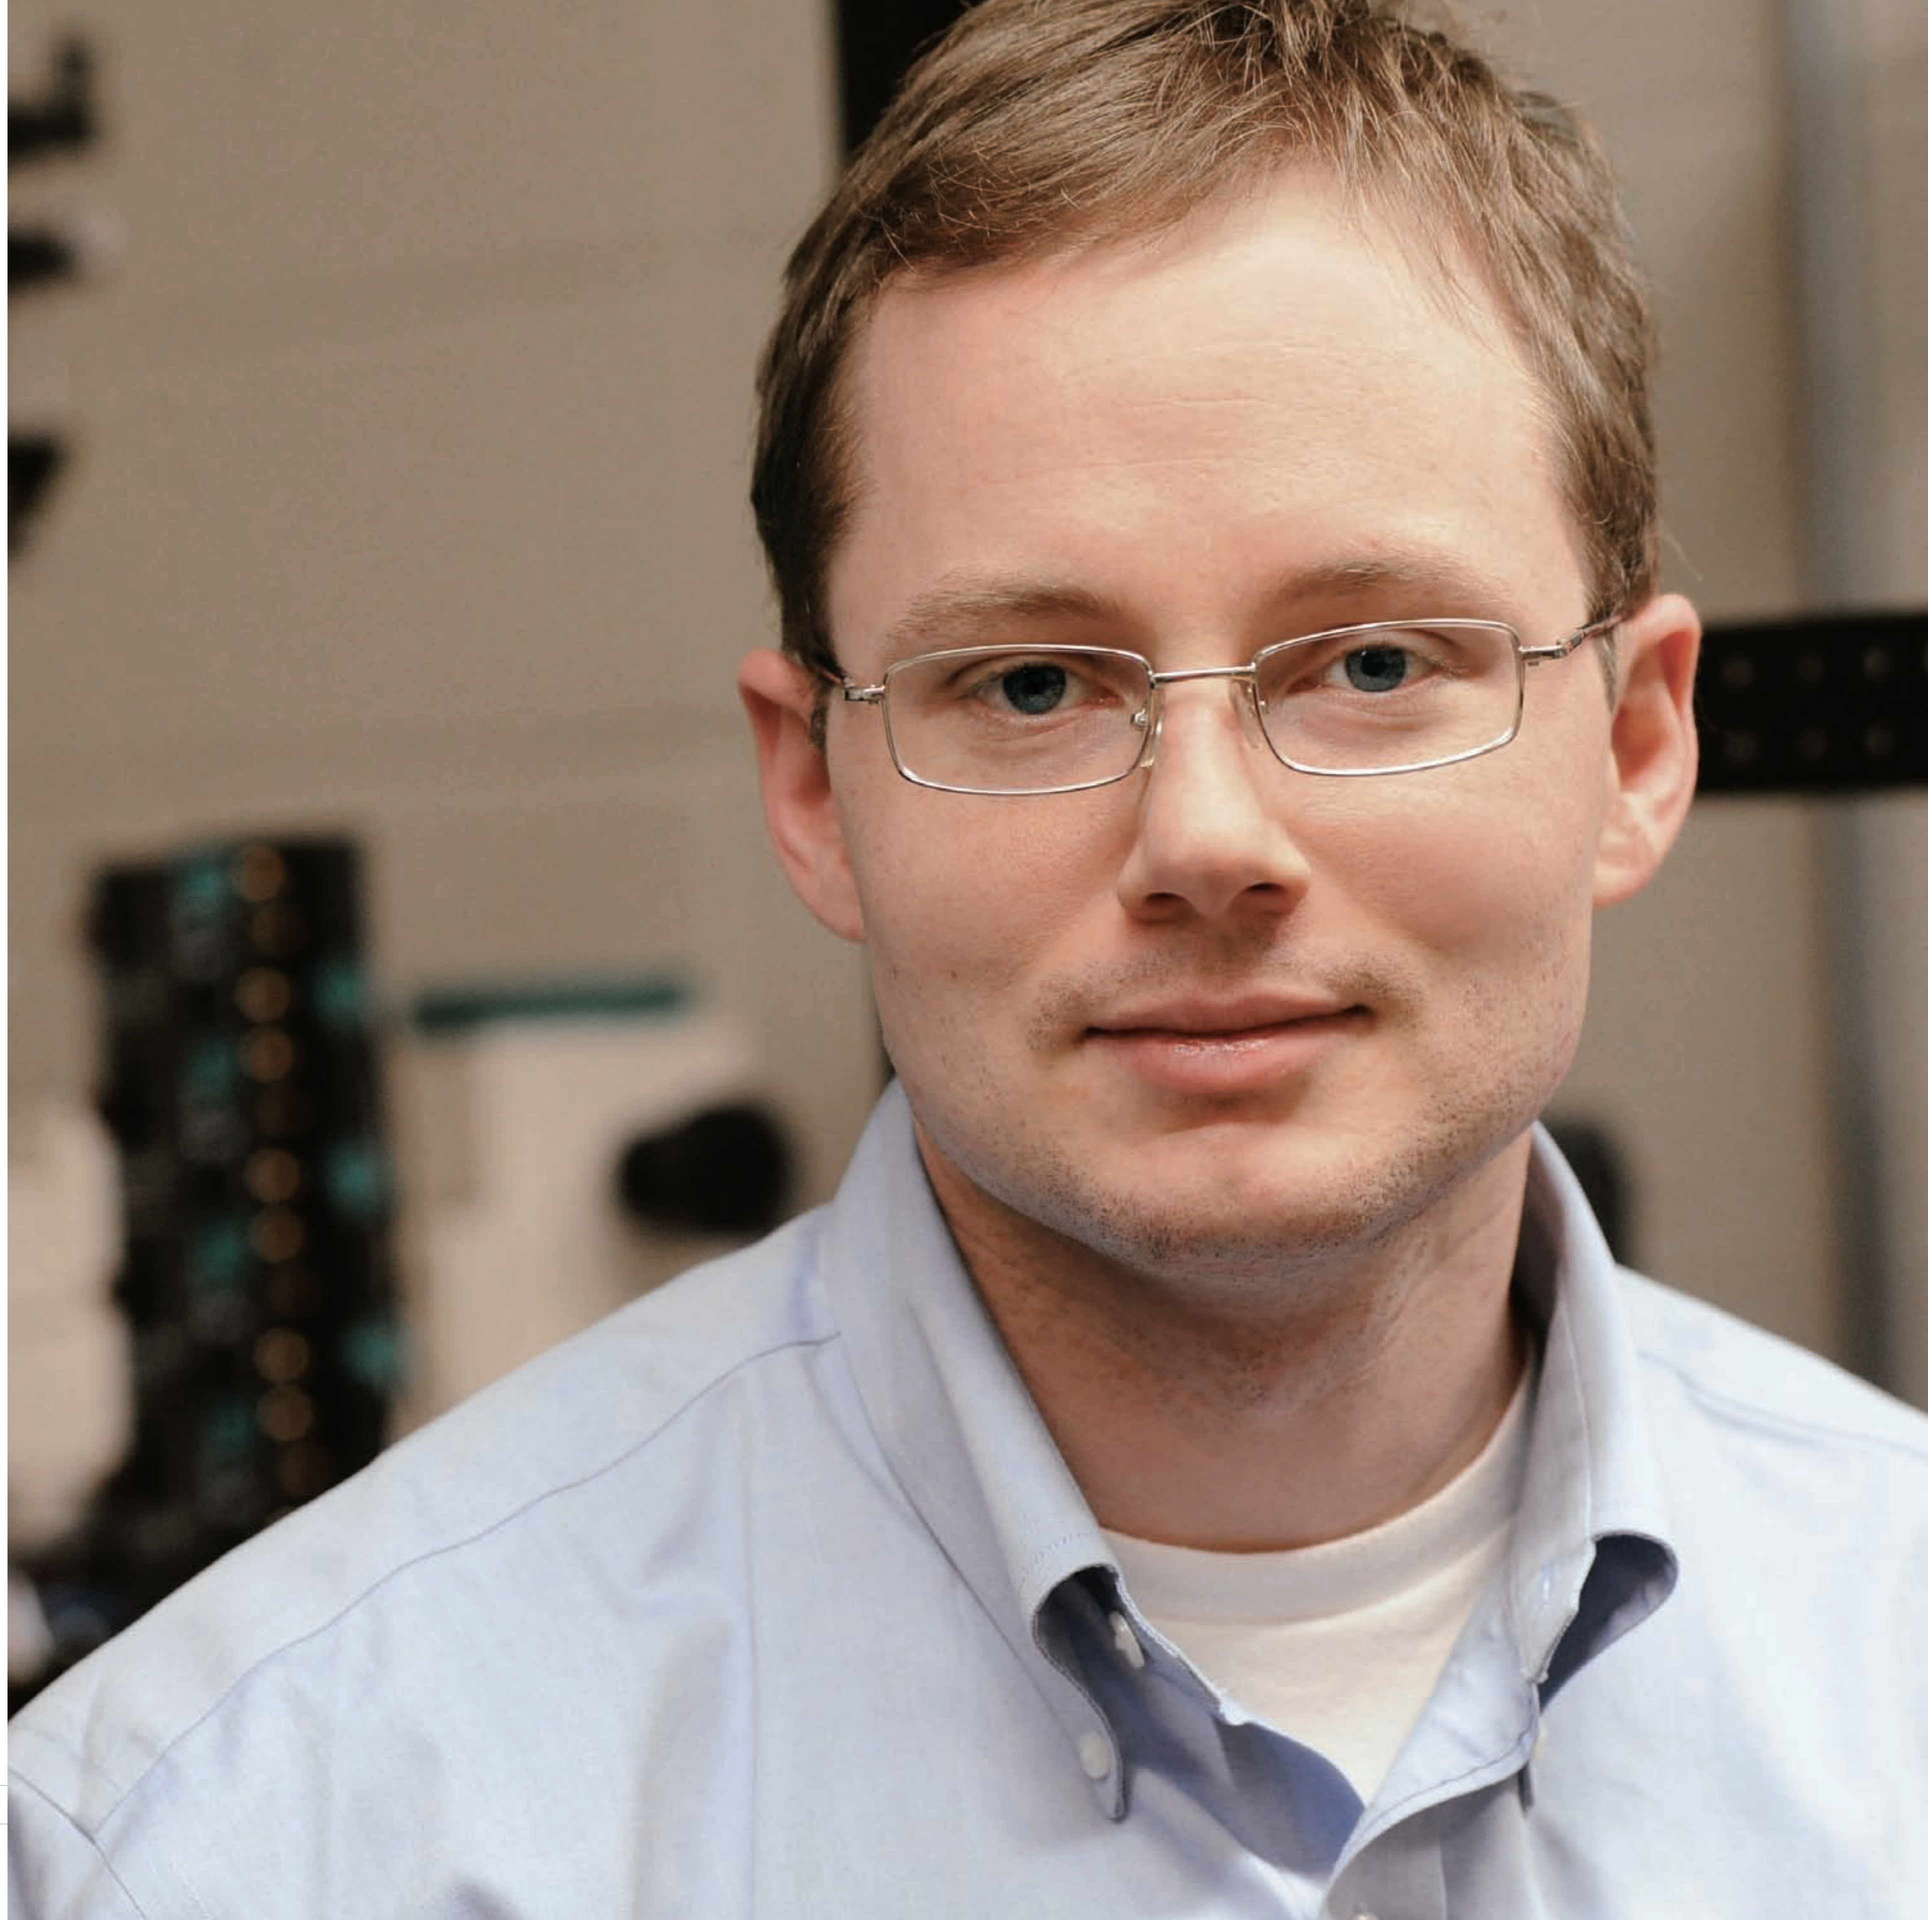

Augmenting Reality

STEVEN K. FEINER

Professor of Computer Science

What happens when a mechanic must work on an unfamiliar piece of equipment? He or she will pull out a manual and keep referring to it while making repairs. Steven Feiner has a better alternative, one that changes how we see the world around us.

His approach to this problem involves augmented reality (AR). Unlike virtual reality, which creates an artificial world, AR adds virtual information to the real world.

AR can guide people through complex tasks. “Instead of looking at a separate manual while disassembling a PC, imagine putting on lightweight eyewear containing a see-through display that graphically highlights the screws in the order you need to remove them,” Feiner said.

Feiner has been developing experimental AR maintenance applications for 20 years. This involves delivering information about a system, quickly and naturally, as workers move around a workpiece. He does this by tracking the position and orientation of their eyewear, then aligning information with their perspective.

In recent studies with U.S. Marines at Aberdeen Proving Ground, Feiner’s lab found that AR helped professional mechanics find the location of parts they needed to repair faster than using manuals. “With manuals, the documentation is separate from the task. Workers are always going back and forth. AR keeps them focused on the work by integrating the documentation with the task,” Feiner explained.

Feiner is also working on better ways to display AR information for people interacting with their surroundings. He has come a long way since 1996, when his lab created the world’s first outdoor mobile AR system. Before the era of smartphones, ubiquitous GPS, and Wi-Fi, it consisted of head-worn and hand-held displays—plus a 45-pound backpack stuffed with electronics.

That system let users tour Columbia’s campus, overlaying the names and websites of academic departments on their buildings. Within a few years, Feiner’s lab had added multimedia news stories and created AR restaurant guides.

Today’s smartphones are far smaller and more powerful than those early AR systems. Feiner and his students are harnessing their power, both alone and with other computers and displays ranging from wearable to wall-sized.

AR displays can create a compelling experience. That is why Feiner wants to ensure that every AR system respects the physical environment and the user’s relationship to it. “We don’t want users losing awareness of the world around them while trying to cross a busy street,” he said.

B.A., Brown, 1973; Ph.D., Brown, 1985

Pushing the Limits of Multiscale Science and Engineering

JACOB FISH

Robert A. W. and Christine S. Carleton Professor of Civil Engineering

"Imagine a world free of traditional scale-related barriers between physics, chemistry, biology, and various engineering disciplines; a world where products and processes are designed based on nature's building blocks, a world in which multiscale science and engineering will revolutionize the way engineering design and scientific discovery are conducted in the 21st century," said Jacob Fish, recently appointed as the Robert A.W. and Christine S. Carleton Professor of Civil Engineering.

Considered by many to be a pioneer in multiscale computational science and engineering, Fish has spent much of his career, first at Rensselaer Polytechnic Institute and now at Columbia Engineering, working at the forefront of this emerging discipline that bridges the gap between modeling, simulation, and design of products based on multiscale principles. His research encompasses a wide variety of science and engineering disciplines, from investigating the structural integrity of mechanical, aerospace, and civil systems, to electronic packaging, nanostructured material systems, biological systems, and energy absorption systems. He has an accomplished track record of technology transfer to industry and has worked with such companies as GE, Rolls-Royce, Lockheed Martin, Sikorsky, Ford, General Motors, Chrysler, Boeing, and Northrop Grumman.

Fish, whose research emphasizes the abundance in nature of systems that encompass interacting behaviors occurring across a range of spatial and temporal scales, believes strongly that "tomorrow's technological advances in science and engineering, including materials, nanosciences, biosciences, electronics, energy, and homeland security, cannot tolerate a partitioned view of nature." Together with his University colleagues, and in collaboration with the City College of New York and New York University, he is forming a new interdisciplinary center, Multiscale Science and Engineering Center (MSEC). MSEC, whose mission is to develop the basic science needed to revolutionize engineering practice and scientific discovery based on multiscale principles, will bring together universities in New York City, drawing upon their strengths in modeling, simulation, and experimentation across multiple spatial and temporal scales. As director of MSEC, Fish hopes to promote an ongoing research in multiscale science and engineering, develop new synergies, and pursue new funding opportunities.

"I am passionate about multiscale science and engineering," said Fish, "because I honestly believe that this field is the next frontier that will transform scientific discovery and engineering design. And I'm very excited to be able to do this at Columbia Engineering."

Fish earned his B.S in structural engineering, his M.S in structural mechanics, and his Ph.D. in theoretical and applied mechanics.

B.S., Technion (Israel), 1982; M.S., Technion, 1985; Ph.D., Northwestern University, 1989

Engineering Services (Before They Perish)

GUILLERMO M. GALLEGO

Professor of Industrial Engineering and Operations Research

INFORMATION

Fruits and vegetables are perishable inventory because they spoil if grocers cannot sell them. The same is true of hotel rooms, rental cars, and airplane seats. Unless they are filled by a certain time, these services cannot produce revenue.

Corporations have become adept at selling perishable inventory by varying prices and running sales. These adjustments are called dynamic pricing, and Guillermo Gallego is one of the field's pioneers. He originally explored how customers value such attributes as a flight's departure time, stopovers, seats, and luggage policies. His work is embedded in many of the models used to price perishable services.

Today, Gallego is working on "service engineering," a concept similar to financial engineering. "It is similar to selling options on a stock," he explained. "We take a basic service and create derivative services from it. This can be a win-win for buyers and sellers, and could dramatically change how certain services are sold."

An example is a fulfillment option. Ordinarily, airline customers buy seats on specific flights. Gallego proposes that airlines offer a discount to customers willing to fly within a certain time period, say 9:00 a.m. to 3:00 p.m., and allow the airline to pick the flight. "The buyer gets a discount. The company buys flexibility, so they can accommodate business customers who often book late and must pay a premium price for a ticket," he said.

Callable products are another possibility. Here, the seller discounts a service in return for the right to buy it back at a premium. A concert promoter, for example, might do that if a band becomes wildly popular during a tour. "In exchange for a discount, the provider can take advantage of a rise in prices," Gallego said.

Gallego is also assessing options where consumers pay an up-front fee that gives them the right to buy a service, such as a hotel room, at a discount in the future. "Companies can afford to do this because not all consumers will exercise their options," he said.

"There's an art and science to engineering and pricing services, but it is always easier when they are win-win for buyers and sellers," Gallego said. "I'm using service engineering concepts to help Hewlett-Packard redesign its warranties. We can offer annual warranties or month-to-month warranties that customers can drop at any time. Priced right, monthly warranties offer value to customers who replace products frequently and are profitable at the same time."

B.S., University of California-San Diego, 1980; M.S., Cornell, 1987; Ph.D., Cornell, 1988

INFORMATION

Increasing Control over Cloud and Mobile Data

ROXANA GEAMBASU

Assistant Professor
of Computer Science

Modern technology is both a blessing and a curse. While mobile devices and web services can deliver quick access to information and even quicker connection to other people, there is a downside: the loss of control over our data.

Consider the trouble likely to occur if your laptop is stolen: you have no way to erase the sensitive data stored on it, you cannot prevent a thief from accessing that data, and you cannot identify potentially compromised data. Or, consider that you cannot be totally certain that photos, email, or documents you try to erase from online services —like Facebook, Hotmail, or Google documents—are not maintained by these web services long after you have requested they be deleted.

To regain confidence in the privacy and security of personal data, new technology applications need to manage sensitive data rigorously and provide users with strong controls over its ownership, distribution, and properties.

Roxana Geambasu works to identify the security risks inherent in today's mobile and web technology and designs, and she constructs and evaluates systems that address those problems. She designed Keypad, a system that guarantees remote control and auditability for data stored on a stolen device; Vanish, a self-destructing data system that provides control over the lifetime of data stored in untrusted web services; Comet, a system that lets users customize the way data is managed in a storage cloud; and Menagerie, a system that offers a uniform view of a user's scattered web data.

Her interests span broad areas of systems research, including cloud and mobile computing, operating systems, and databases, all with a focus on security and privacy. She integrates cryptography, distributed systems, database principles, and operating systems techniques and advocates a collaborative approach to developing cross-field ideas in order to solve today's data privacy issues.

In 2009, she was the recipient of the first Google Fellowship in Cloud Computing. Her current research focuses on an operating system redesign for mobile devices. She has identified that the principle mechanisms, assumptions, and interfaces of mobile device operating systems have not evolved to match the unique characteristics and workloads they are meant to handle.

B.S., Polytechnic University of Bucharest (Romania), 2005; M.S., University of Washington, 2007; Ph.D., University of Washington, 2011

Photo: Bruce Hemingway

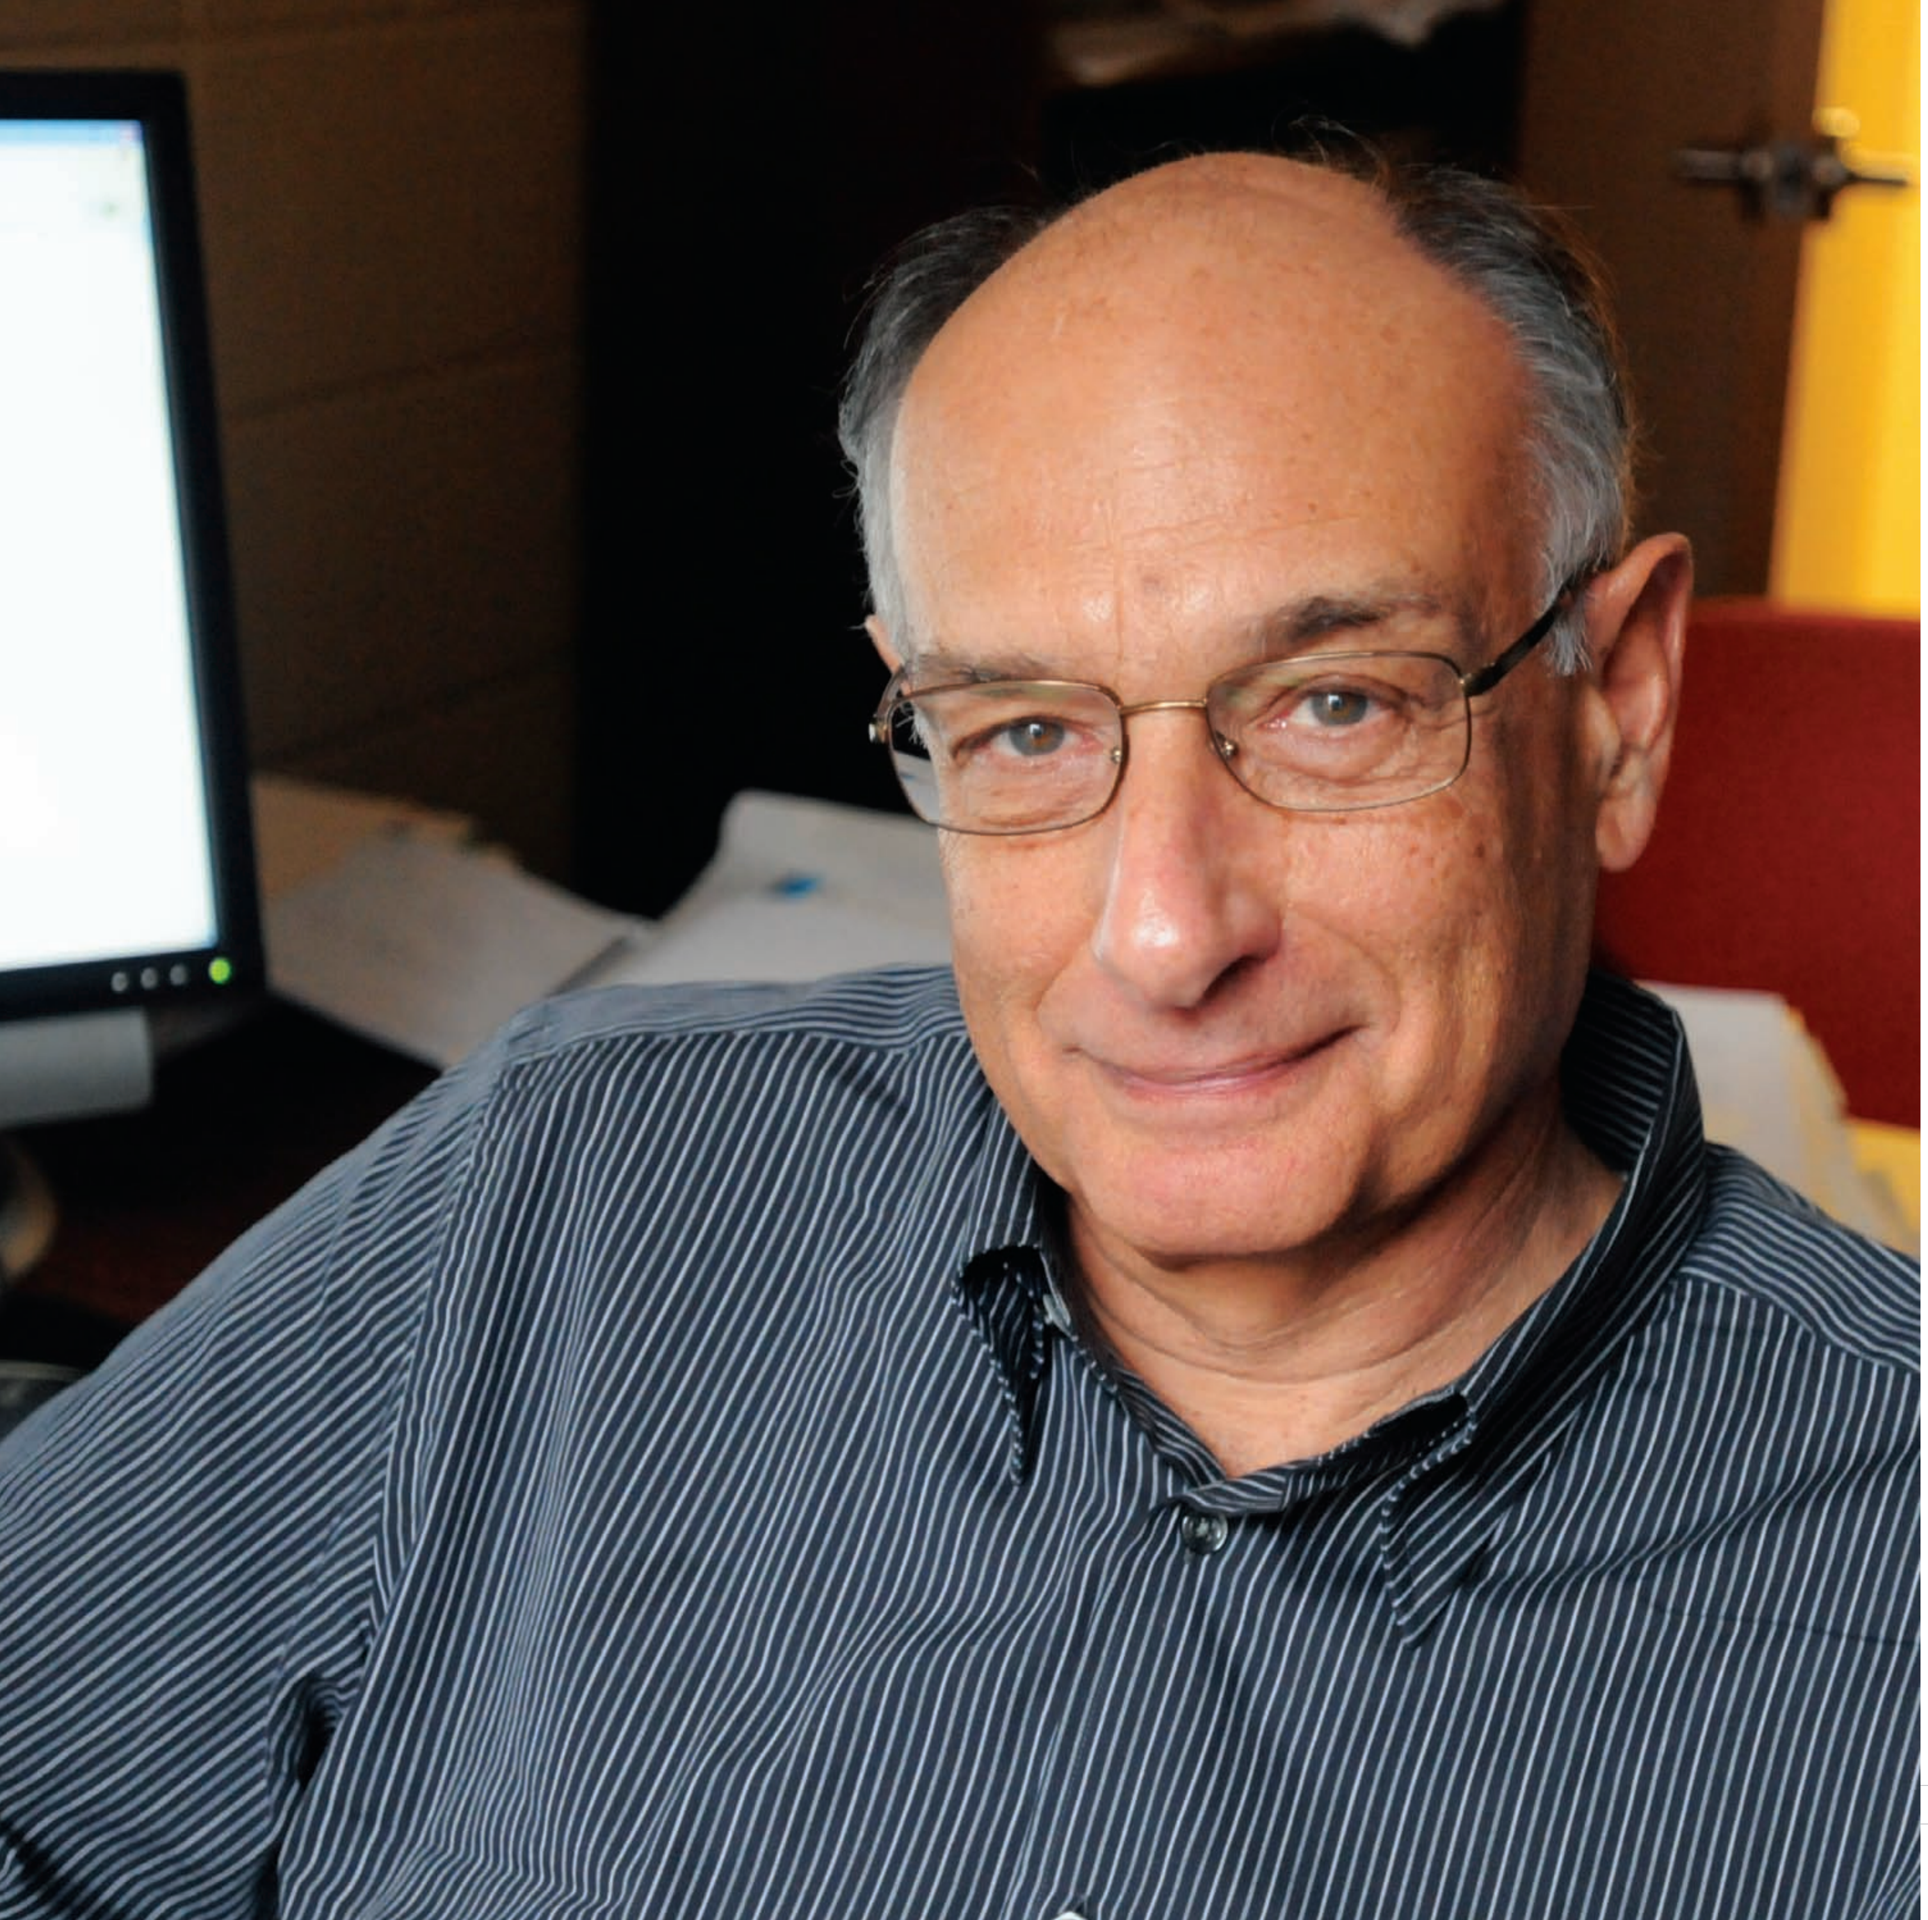

Finding a Way Around Too Much Data

DONALD GOLDFARB

Alexander and Hermine Avanessians Professor of Industrial Engineering and Operations Research

Computing power has grown rapidly, but not as fast as the problems researchers aspire to solve. “We’re dealing with enormous problems, problems so large we can’t even store all the numbers in computer memory at the same time. We cannot rely on the same methods we used for smaller problems and expect to solve them,” said Donald Goldfarb.

His work on extracting movement from surveillance videos provides an example. On surveillance videos, the background never changes. One frame looks very much like the next, except for the people moving through the space. Each frame has roughly 20,000 pixels.

“To extract moving images from a couple of minutes of video, we need to process 50 million variables and 25 million linear equations,” Goldfarb said. Doing it by brute force—one computation after the other—would take days on powerful computers. Instead, he developed a systematic optimization procedure, or algorithm, that lets a simple workstation remove the background in under an hour.

Goldfarb has a long history of developing powerful optimization algorithms. Some of his early algorithms are used in commercial software to optimize complex systems. They make it possible, for example, to adjust refinery operations on the fly instead of spending weeks plotting a production schedule.

Goldfarb’s work goes beyond just finding fast ways to solve difficult problems. “I try to prove that the algorithms I develop are not just fast for a specific problem, but will work well for any similar problem,” he said. “It’s like providing a certificate guaranteeing the algorithm’s performance.”

He also tries to discover properties about different classes of algorithms. Recently, he has focused on convex functions. Like many algorithms, they recast algebraic problems in geometric terms in order to estimate answers more rapidly. Goldfarb likens a convex function to a bowl with the minimum, or optimal, value at the bottom. Constraints usually push the answers to any given problem somewhere along the sides of the bowls.

“If you’re sitting on the side, you can see every other point inside the bowl. If you look around and every other point is higher, then you are at the optimal point,” Goldfarb explained.

Recently, Goldfarb used convex functions to optimize a method to produce MRI and CT scan images using only one-fifth the radiation. “The algorithm enables us to get an appropriate image with fewer measurements, so patients only have to spend one-fifth as much time in these machines,” Goldfarb said.

B.Ch.E., Cornell, 1963; M.A., Princeton, 1965; Ph.D., Princeton, 1966

Studying Decision Making in the Face of Uncertainty

VINEET GOYAL

Assistant Professor of Industrial Engineering and Operations Research

In almost every field, decision makers are often required to make important choices in the face of uncertainty. For instance, a financial portfolio manager must make investment decisions without being certain of asset returns. A medical doctor must prescribe a treatment without being totally certain about a patient's response to that particular treatment. An electricity grid system operator must select a set of generators to be switched on without fully knowing what the consumer demand will be and what the state of the transmission lines are. An airline makes pricing and scheduling decisions while facing an uncertain demand. Today's world of free markets, fast communication links, and availability of vast amounts of data makes the study of decision making under uncertainty or dynamic optimization extremely important.

Uncertainty in a decision-making problem is usually modeled in one of two ways: either by a probability distribution in a stochastic model or by an uncertainty set in a robust model. A stochastic model estimates probability distributions of potential outcomes by allowing for random variation in one or more inputs over time. While a stochastic model might be a good approximation of reality, the resulting optimization problem is often very difficult to solve even approximately. On the other hand, a robust model can be solved efficiently in most cases but is considered too conservative to be useful in practice as it optimizes over the worst case.

Vineet Goyal works to address a fundamental question about the relationship between these two diametrically opposite paradigms. His research focuses on providing justification for robust and other tractable approaches as a practical method to solve dynamic optimization problems. His work shows that under fairly general assumptions, the robust optimization approach provides a good approximation of the stochastic problem in many cases.

Goyal's current work focuses on analyzing the performance of various tractable approaches such as affine policies (also referred to as linear decision rules) and piecewise affine policies for dynamic optimization problems. His goal is to better understand the trade-off between tractability and performance of various approaches. This is a fundamental question and has potential for significant impact given the wide applicability of dynamic optimization.

Goyal is especially interested in applications in electricity markets where dynamic optimization is very applicable with an increasing concentration of renewable sources of generation that have a highly uncertain generation capacity. He also applies this research to problems associated with revenue management and inventory management.

B.Tech., Indian Institute of Technology (India), 2003; M.S., Carnegie Mellon University, 2005; Ph.D., Carnegie Mellon University, 2008

Supercharging Search Engines

LUIS GRAVANO

Associate Professor of Computer Science

Imagine searching for a concert and pulling up the usual web pages, plus untagged Flickr photographs, Twitter remarks, YouTube videos, and Facebook comments. Or, asking when the band will perform again and getting back a table of dates and locations.

Luis Gravano is supercharging search engines to conduct exactly those types of searches. Often, that means tapping the chaos of social media. "It's not so much about just returning a list of individual web pages as it is about combining and making sense of all information on the web to increase the effectiveness of a search," he said.

For example, many online photos are tagged to refer to specific events. Others have time and GPS data that coincide with the time and location of an event. Sometimes photos are forwarded or linked to other people who have commented about an event.

"We analyze these tags, comments, and links, and automatically cluster them to correspond to real-world events," Gravano said. His team has already shown that it can aggregate such information. It is now probing how to fit the data together to develop more powerful searches.

"If there is a concert or political demonstration, people take pictures, tweet, and form groups around these activities," he said. "We want to capture and associate this content with real-world events automatically. We'll return results that correspond to a specific event at a certain time on a particular street in New York City."

Gravano also wants to improve our ability to extract structured information, such as tables from the Internet. Today, he explained, anyone who wants to analyze the characteristics of past infectious disease outbreaks would have to sift through hundreds or thousands of search engine results.

Gravano's extraction technology searches for pages that are likely to contain the desired structured information, which is often embedded in natural language text. It then extracts, analyzes, and puts the information into a table automatically. Unfortunately, the process is prone to errors. Information is sometimes out of date or wrong. Writing is often ambiguous.

Gravano hopes to reduce errors by using such trusted sources as government documents, university archives, newspapers, and specialized websites, as well as by analyzing the frequency and context of the extracted information.

He also taps crowd wisdom to assess the reliability of popular sources. "Popularity is a step in the right direction—if you trust people to go to trustworthy sources," Gravano said.

B.S., Escuela Superior Latinoamericana de Informática (Argentina), 1991; M.S., Stanford, 1994; Ph.D., Stanford, 1997

INFORMATION

Predicting the Motion of Materials

EITAN GRINSPUN

Associate Professor of Computer Science

What do dresses, medical instruments, and the bristles on a paintbrush have in common? Their motion can all be predicted with unparalleled accuracy by techniques developed by Eitan Grinspun.

Grinspun's techniques have broad application. In the movies, they produce stunningly realistic animations of gowns swirling on dancers and animal manes billowing in the wind. "If you can compute motions that obey physical laws, you can make artistic choices about what laws you want to disobey and produce things you would never see in real life," he said.

His work is equally applicable to physics. "Think about how honey behaves when you pour it on a scone," he said. "It is a liquid, but it loops around like a rope. If we can understand how honey moves, we can understand how lava flows or the best way to bottle shampoo."

Bottling shampoo is not a trivial problem. Shampoo entrains air, which reduces its density and increases its volume. "If you can understand how shampoos move, you can reduce entrainment and pack them in smaller containers to reduce costs," Grinspun said.

Physicians have used Grinspun's techniques to test how to steer surgical needles through human tissue. Adobe has leveraged them to simulate each individual paintbrush bristle in its popular Photoshop and Illustrator programs. "Those bristles are really bending, and you get all the effects you would get with a real paintbrush," Grinspun said.

What makes Grinspun's work unique is his deep understanding of the geometry underlying physics. For example, when he looks at a long, thin surgical needle, he sees a flexible curve that bends and twists. "Computers, geometry, and physics are my ingredients. I mix them up in a bowl and what I get is a computer's ability to predict the motion of materials.

"We can visualize the problem by thinking of the boundary of North America on planet Earth," Grinspun explained. "The energy stored in bending is like the continent's perimeter, while the energy stored in twisting is its area. We have a competition between bending, which wants to keep length as short as possible, and twisting, which wants to deform the length to enclose more area."

Understanding the geometry of those forces produces fast and accurate predictions of movement. The results are readily visible in movie special effects and in basic science as well.

B.A.Sc., University of Toronto (Canada), 1997; M.S., California Institute of Technology, 2000; Ph.D., California Institute of Technology, 2003

Untying Knots with Mathematics

JONATHAN GROSS

Professor of Computer Science and of Statistics

Jonathan Gross knew little about Celtic knots before he started studying them. "I knew one when I saw it. They are characterized artistically by repetitive patterns and symmetries," he said. "Then, while browsing the Internet, I found a graphic artist's description of them so precise, I could turn it into math."

Gross uses computers to explore algebraic topology, the mathematics of translating geometric forms into algebraic expressions. "We calculate a polynomial from a picture of the knot. Once we represent the shapes with algebra, we can manipulate the math to learn fundamental truths about the shapes," he explained.

"For example, Reidemeister proved that if you make new crossings in a knot without cutting the string, the resulting figure has the exact same polynomial as before. If you hand me a knot, I can either fumble for hours trying to untie it or I can calculate a certain polynomial and quickly know that the string is really knotted," Gross said.

Gross is quick to point out that his research is theoretical. Yet some of his insights have worked their way into practical applications. "Some of my work is related to practical technology," he said. "But what motivates me are mathematical problems that involve spatial visualization and deriving algebraic formulas to count mathematical objects far too numerous and/or too intricate to count by any elementary methods."

Last year, for example, he collaborated with two colleagues in Texas to develop a computer graphics program to create designs in woven textiles. "We designed software whose mathematical models embody key principles of algebraic topology," said Gross. "A graphic artist doesn't have to know any of this to use the software to create a complicated woven pattern very quickly."

Gross has also applied mathematical modeling to social anthropology. Anthropologists used to live with a people and describe what they saw. Their descriptions were typically highly subjective. Gross worked with a team that developed an objective way to measure and compare behavior.

They started with food systems. "There are differing levels of randomness in the way people eat," he said. "When some people eat scrambled eggs, you know for sure it's breakfast. Not quite so for others. To differing extents, meal content reflects the time of day, time of year, and festivities. By measuring the information content in these patterns, we could make comparisons between different peoples."

To Gross, it was just another knot untangled by mathematics.

B.S., Massachusetts Institute of Technology, 1964; M.A., Dartmouth, 1966; Ph.D., Dartmouth, 1968

INFORMATION

Modeling the Irrational

XUEDONG HE

Assistant Professor of Industrial Engineering and Operations Research

Fill a glass half way and some people will call it half filled and others half empty. Either way, the amount of liquid in the glass is the same. Our frame of mind—optimistic or pessimistic—imposes meaning on what we see.

Investors in financial markets are not any different, said Xuedong He. They all view the same financial data, yet they draw different conclusions from what they see. Irrational biases often play a role.

"Classical economic theory assumes that investors evaluate information correctly and make decisions rationally," he said. "In reality, though, they have biases. They may overemphasize or overlook certain types of information, and this affects how they manage their portfolio."

For example, investors often miscalculate the odds of an event because they put too much weight on recent data. One common bias is to go with someone on a winning streak. "Gamblers who win two or three times in a row think they have a hot hand and are more likely to win the next time. The odds are still against them, but they over emphasize their recent success," He said.

Other investors may assess the odds correctly, but hidden biases guide their actions. "Look at people who buy lottery tickets and insurance," he said. "A lottery ticket usually has an expected value lower than the selling price. Buyers know the probability of winning is very low, but take the risk for the reward. On the other hand, people know the probability of their house burning down is low, but they buy insurance because they are risk averse."

He builds mathematical models that show how these twin engines—hope and fear—drive investment strategies. "Hope and fear coexist in investors' minds. When stock prices surge, hope takes control, so investors are more likely to invest in stocks and gamble more. When the market turns down, fear dominates and investors quickly liquidate their portfolios," he said.

"In financial engineering, not much work has been done on irrational biases. We have developed a concrete model of these irrationalities based on extensive research. We want to understand how these biases affect investor behavior and strategies," he said.

He is one of the few financial engineers researching irrational motivations. By taking biases into account, he hopes to create models that better predict market behavior and perhaps even warn when investors are being carried away by irrational exuberance.

B.S., Peking University (China), 2005; D.Phil., University of Oxford (England), 2009

Recognizing the Melody of Speech

JULIA B. HIRSCHBERG

Professor of Computer Science

INFORMATION

Anyone who has ever navigated an interactive voice-recognition system to make a reservation or review a charge knows that anger and sarcasm change nothing. But one day they might, thanks to research by Julia Hirschberg.

Hirschberg studies prosody, the intonation and melody of speech. Often, it conveys subtle differences in meaning. For example, "I like cats" may sound like a statement, but raising the pitch at the end turns it into a question.

"During deceptive speech, you experience emotions like fear if you think you'll be detected or elation if you're getting away with it. This shows up in the prosody of your speech. The best people at judging liars are criminals. Police were worse than average, and parole officers the worst of all, because they assume people are always lying," she said.

Hirschberg's goal is to teach computers to understand such subtle variations and reproduce them in natural sounding speech. This involves understanding how prosody changes under different circumstances.

"When I was at Bell Labs, we did lots of experiments that looked at people's speech, and tried to predict what words he or she would emphasize," Hirschberg related. "We looked at syntax, context, the part of speech being uttered—you use whatever information you have, and usually that's not a whole lot."

At Columbia, she has analyzed the prosody of charismatic and deceptive speech. "Much of the perception of charisma is not about what people say, but how they say it," she explained. "In English, charismatic speakers are very expressive, vary their pitch contour a lot, and speak more rapidly."

She has also conducted extensive experiments in which people either lied or told the truth. In these experiments, the speakers told the truth about 61 percent of the time. Her automated computer system labeled identified truth tellers and liars about 70 percent of the time. Humans got it right about 58 percent of the time, worse than if they had just guessed "truth" every time, Hirschberg said.

She is also working on teaching interactive voice-response systems a technique called entrainment. This occurs when one speaker mirrors back the same vocabulary, pitch, and speed as another. "People like people who entrain to them more than those who do not," Hirschberg said. "We want to teach computers to change their pitch, intensity, speaking rate, and other factors to sound more like the user."

If that doesn't mollify the next generation of callers, at least the computer will recognize their anger when they express it.

B.A., Eckert College, 1968; Ph.D., University of Michigan, 1976; MSEE, University of Pennsylvania, 1982; Ph.D., University of Pennsylvania, 1985

INFORMATION

Finding the Fundamentals of Silicon for Advanced Electronics

JAMES IM

Professor of Materials Science and of Applied Physics and Applied Mathematics

Silicon, the second most abundant element in the Earth's crust, is the key material of the modern information age. Microelectronic chips use bulk-silicon wafers to power computers, and silicon is used for increasingly important electronic applications, such as inexpensive solar cells, high-resolution flat-panel displays, radio-frequency identification tags, and 3-D integrated chips. But manufacturers need high-quality crystalline silicon films in which atoms are nicely and periodically arranged.

While it's easy to obtain amorphous silicon films, they are not well-suited for making these electronic devices. Developing efficient ways to generate high-quality silicon films is a key to the proliferation of these micro- and macro-electronic applications. James Im's process for developing high-quality silicon film is playing a crucial role in developing the latest generation of flat-screens for a wide array of electronic devices.

Im has done extensive research that investigates how silicon, solid thin films and nanoscale structures behave when these materials are rapidly heated by laser irradiation, melted, and then subsequently solidify. While his studies look primarily at the scientific and fundamental issues involved, the findings also have led to various technical approaches for realizing high-quality silicon films on various technologically important substrate materials such as glass or plastics.

These laser-induced and melt-mediated crystallization processes, which convert initially amorphous or defective silicon films into low-defect-density silicon films, take place at temperatures above 1400 degrees C. According to Im, understanding how silicon melts and solidifies under these extreme conditions is critical for understanding how the atoms are subsequently packed and positioned.

"Knowing the fundamental details of how Si melts and solidifies makes it a rather straight-forward exercise for us to come up with efficient and effective ways to generate useful materials with periodically arranged atoms that make good electronic devices," said Im.

The fundamental findings and technical approaches generated at Columbia are powering the evolution of the field of thin Si-film based electronics. One method, called Sequential Lateral Solidification (SLS), is used to manufacture high-resolution LCDs, and has recently emerged as the leading method for the next generation of flat-panel TVs, which use organic LEDs.

Top display makers, including LG Display, Sharp, and Samsung, have already licensed this technology. The innovation is also applicable to smart cards, RFIDs, image sensors, and 3-D integrated circuit devices.

In addition to laser-based approaches, Im is also investigating other beam-induced crystallization techniques that could provide unconventional, yet effective solutions for various electronic devices and applications.

B.S., Cornell, 1984; Ph.D., Massachusetts Institute of Technology, 1989

Photo: Alan S. Orling

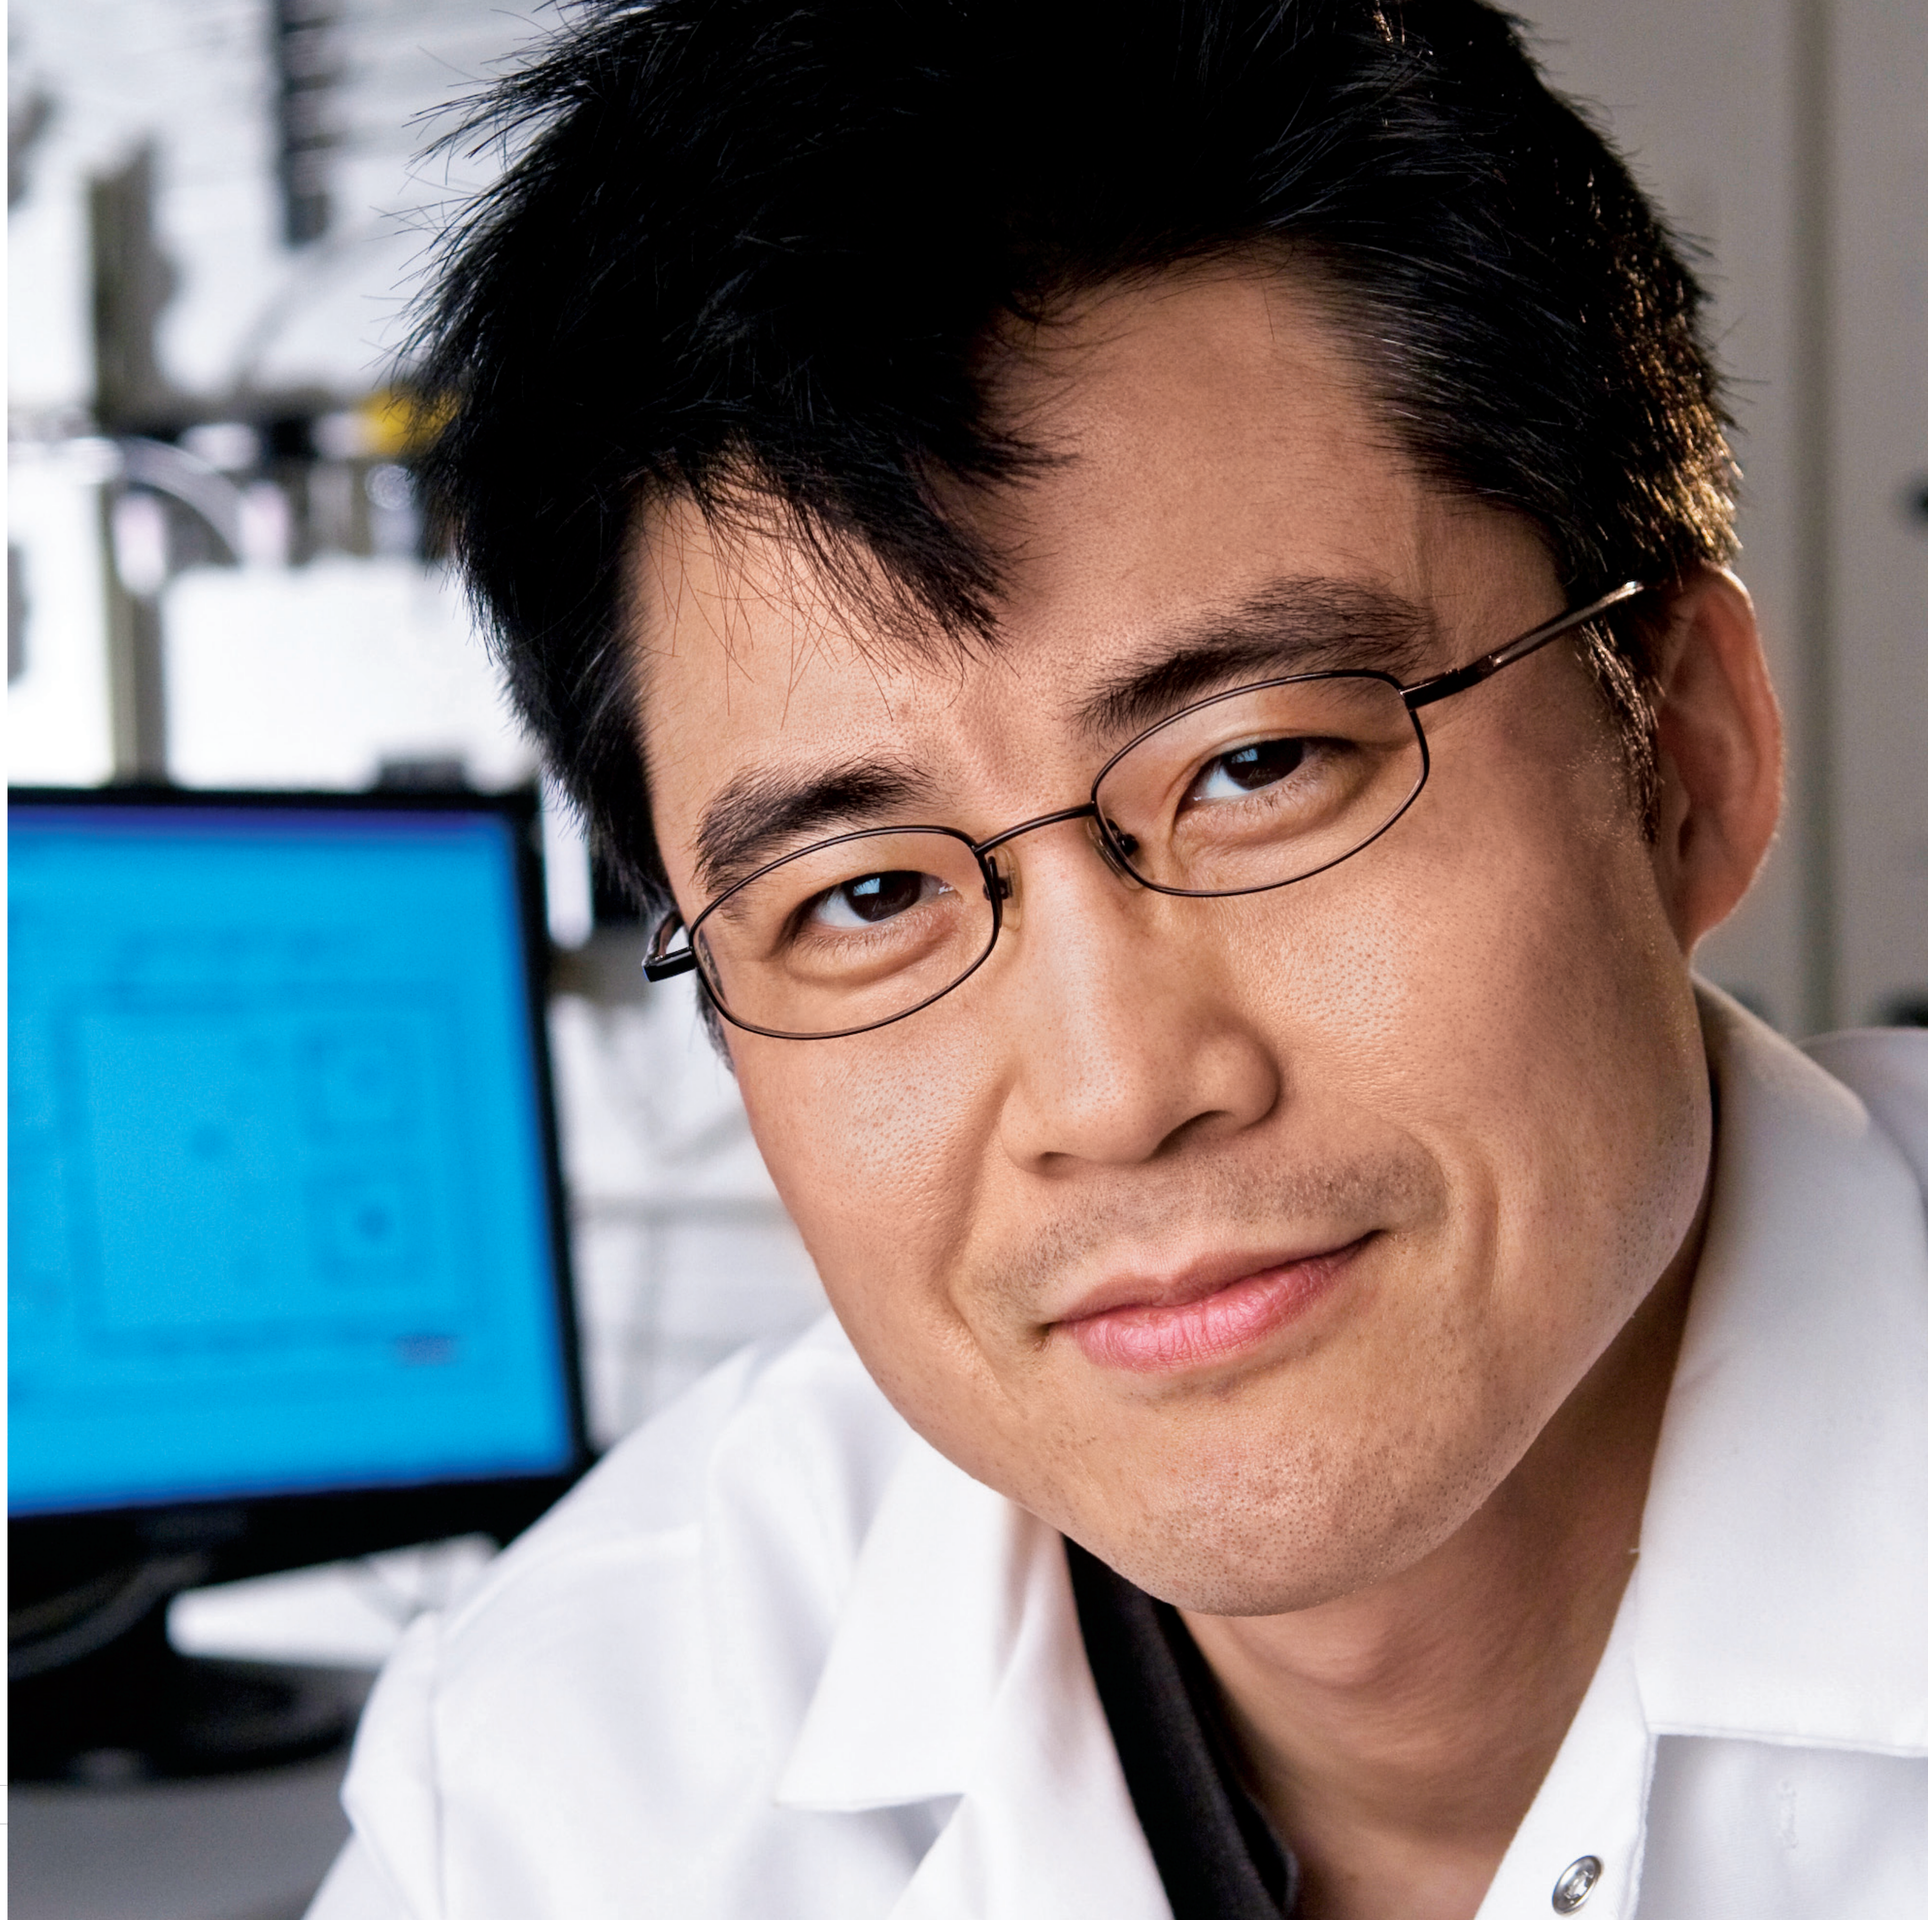

Deciphering the Mysteries of Microbial Communications

GARUD N. IYENGAR

Professor of Industrial Engineering and Operations Research

Garud Iyengar is helping to unlock the secrets of how colonies of bacteria work together, though he is not a biologist. "I'm a problem solver, rather than someone who focuses on one particular research area," said Iyengar. "My particular interest is in understanding how simple components can produce complex behavior when networked together."

Iyengar's varied background in mathematical modeling and optimization enables him to tease out insights that classically-trained biologists might miss. "My particular strength is in building mathematical models to guide experimentation by blending tools, often from different disciplines, that together work better than any single tool used independently," he said.

"Most scientists have a set of pet mathematical tools. Someone trained in statistics immediately thinks about regression to model experimental data. A computer scientist builds a combinatorial model," he explained. "An electrical engineer wants to use information theory. I've been exposed to many of these disciplines, and so my bag of tricks is bigger."

Lately, Iyengar has been trying to discover how colonies of unicellular organisms communicate in order to exploit their environment. Density sensing in Pseudomonas aeruginosa, a bacteria that inhabits the lungs of patients with cystic fibrosis, is an example. "These bacteria only turn virulent when their local density crosses a certain threshold. At lower densities, the host's immune system would overwhelm it," he explained.

It is well understood that bacteria use certain signaling molecules to sense density. A positive-feedback biochemical network triggers a switch when the signal concentration is high enough. According to classical control theory, there are many possible networks that yield the same density dependent switching behavior. Evolution, however, has selected one particular network in many different bacterial species. Iyengar is interested in understanding the reasons underlying this selection.

A more complex problem is how bacteria determine the colony's average temperature in order to optimize their metabolism. This is more difficult than it sounds. Each cell perceives only the temperature around it. Many factors, such as nearby water or chemical reactions, create microclimates that vary significantly from the average.

Iyengar speculates that bacterial colonies use a technique called belief propagation to measure spatial averages. Belief propagation is a well-known paradigm from statistical physics that describes how a particle adjusts its behavior based on the behavior of its neighbors.

"If it is used by bacteria, there are measurable consequences that logically follow," he said. "We are using our models to guide the type of experiments we need to do to quantify these consequences."

B. Tech., Indian Institute of Technology, 1993; M.S., Stanford, 1995; Ph.D., Stanford, 1998

INFORMATION

Finding Patterns in a Complex World

TONY JEBARA

Associate Professor of Computer Science

Most of the data created in human history was actually generated in the past handful of years. "Every person in the world on average generates and consumes gigabytes of text, video, Internet media, images, and music every year," Tony Jebara said.

Jebara's specialty is developing machine learning programs that sift through massive amounts of data to discover underlying patterns and make accurate predictions. "I work at the intersection between statistics and computer science, applying machine learning tools to massive data sets where the relationship between variables is often not deterministic. Our algorithms must be fast, because computer speeds are not growing as rapidly as the amount of data they must handle," he said.

Computers slice through data that would take humans years to analyze. Yet their capabilities are only as good as the underlying algorithms—the set of rules used to classify and analyze data. Computers, for example, find it hard to identify faces, a task babies master within months.

This is an area where Jebara made his start by building one of the top face recognition algorithms. His approach to face recognition used probability distributions to calculate the likelihood that two images were of the same individual. Jebara also worked on extending the standard Bayesian algorithms to minimize error rather than maximize likelihood.

Most recently, Jebara has been working on matching and graph algorithms, two promising ways of learning from massive datasets, such as those generated by social networks and the web. Viewing large amounts of data as a graph often provides a faster and powerful way to solve problems such as data labeling and partitioning.

Also, graphs allow algorithms to be implemented very efficiently by such techniques as message passing, which Jebara has worked on extensively. He has built programs that automatically visualize, label, partition, and match data in large data sets, ranging from images to social networks.

Similar algorithms also power Sense Networks, a startup Jebara founded in 2006 to analyze data from telecommunications companies. By tapping smartphone calls and GPS data, Jebara's algorithms can classify people by behavior patterns. Users can then query the network to see where people with similar tastes go to eat, drink, or shop. The phone company can use the data to filter recommendations and provide targeted advertising.

It is one more example of machine learning finding patterns in a world awash with data.

B.S., McGill University (Canada), 1996; M.S., Massachusetts of Technology, 1998; Ph.D., MIT, 2002

Unwinding Heavy Tails

PREDRAG JELENKOVIC

Professor of Electrical Engineering

INFORMATION

We are witnessing an emergence of a variety of large-scale man-made information networks, including the wireless or wired Internet, World Wide Web, social, and economic networks. Similarly, on a microscopic scale, large biological protein networks inside the cell and inter-cellular neuronal networks are just being uncovered.

While these networks operate on entirely different temporal and spatial scales, address unrelated applications, and use diverse mediums to represent information, many of them are governed by the same underlying mathematical principles. Most commonly, these networks exhibit very high variability of their parameters, either in their connectivity, the statistical properties of information they carry, or the delays for processing and transferring the information.

Predrag Jelenkovic uses the mathematical theory of heavy-tailed and power law distributions to capture the highly variable characteristics of these networks. His research focus is on mathematical modeling, analysis, and control of large-scale information networks with heavy-tailed characteristics.

His recent research resulted in a discovery of an entirely new phenomenon in communication networks that shows that retransmissions can cause long (heavy-tailed) delays and instabilities even if all messages and files in the network are relatively short (light-tailed). This finding is important in general since the retransmission-based failure recovery is at the core of the existing communication network architecture, and especially in wireless networks where communication link failures are frequent.

In addition, he focuses on developing the statistical ranking mechanisms for rapidly growing information webs (e.g., the World Wide Web, scientific data, biomolecular and neural networks, social networks, news, and e-commerce). Given that the scale and complexity of these information sets will continue to increase in the future, a new statistical approach for their ranking and understanding is needed in the same way that statistical mechanics were needed for understanding large sets of molecules. Interestingly, this research reveals that the ranks of pages on the World Wide Web, according to Google's page ranking, follow heavy-tailed power law distributions as well.

Jelenkovic is a member of the Communication & Networking and System Biology groups in the department. Within these groups, he works to advance the mathematical foundation of the underlying design principles of both man-made and biological information networks. Furthermore, his work on heavy-tailed distributions applies more broadly to insurance risk theory, financial mathematics and economy, where heavy tails are widely used.

B.S., University of Belgrade (Serbia), 1991; M.S., Columbia, 1993; M.Phil., Columbia, 1995; Ph.D., Columbia, 1996

Understanding the Dynamics Behind Pricing

SOULAYMANE KACHANI

Associate Professor of Professional Practice in the Department of Industrial Engineering and Operations Research

Some prices never sit still. Retailers discount clothing and technology products seasonally. Traders bid stocks up and down daily. Airline and hotel prices fluctuate by the hour. These are examples of dynamic pricing, where companies price goods based on cost, customer behavior, and competitive dynamics. Soulaymane Kachani's research in the field has taken him in some interesting directions.

"We are applying traffic flow theory used in transportation networks to blood rheology to prevent blood clots," he said. "Existing models are hard to calibrate for elderly patients because they require too much ultrasound data. Our models are simpler, and appear to better predict where clots will form. Our next step is to conduct clinical trials."

One recent project assessed lifecycle pricing for different generations of technology products. He found that to maximize long-term profits, companies should not discount old technology too deeply.

"These companies interact repeatedly with their customers," Kachani explained. "Once they set a price, it affects the reference price. So if they start driving down the price of older goods, they cannot go back and ask for a dramatically higher price for their next-generation product. In fact, long term, many tech companies are better off discontinuing old products than discounting to sell off inventory."

Real estate, on the other hand, could benefit from more dynamic pricing. "Imagine you're developing condominiums," Kachani said. "What price do you assign each unit?

"You don't want to sell out all the units with upper floors, good views, or two bedrooms first. If you do that, it means you did not put the right premium on the more desirable units. If the premiums are set correctly, all your different units should sell at roughly the same pace," Kachani explained.

To find the right premiums, Kachani looks at both unit sales and also what units visitors view. He uses their actions as input for a computer model that modifies prices based on real market input. This gives developers a realistic way to set prices to maximize returns.

His work also extends to fashion. Kachani compared retailers who emphasize innovation and design with those who focus on pricing. The innovators, with short product runs and high turnover, had higher profits than retailers with larger product runs who relied on periodic discounts to clear the shelves.

Yet Kachani urged the innovators to consider dynamic pricing. "They would do even better if they managed their pricing strategy better," he said.

Diplôme d'ingénieur in Applied Mathematics, École Centrale de Paris (France), 1998; M.S., Massachusetts Institute of Technology, 1999; Ph.D., MIT, 2002

Testing What Cannot Be Tested

GAIL E. KAISER

Professor of Computer Science

Software systems have a complex lifecycle, and Gail Kaiser likes to work on all aspects of it. Her research ranges from creating systems that make recommendations to finding flaws in "non-testable" programs. "I like to find solutions in one domain and then generalize them," she said.

For example, she is working on three "recommender" systems. One system monitors how biologists use tools for genomic analysis, and then gives novices recommendations based on the workflow of more experienced users. A second mines past experience to help programmers convert software to parallel code for multicore processors. The third helps computer science students solve certain errors in the code they write.

"We built all three systems independently, then noticed that they all used essentially the same architecture," Kaiser said. "This let us derive a general reference architecture that might be useful in building future recommender systems."

Kaiser is also interested in testing so-called "non-testable" programs. These include machine learning, simulation, data mining, optimization, and scientific computing systems. "I come from a software engineering background, but work at the borders of my discipline and operating systems, databases, and security," said Kaiser. "I'm concerned with how to build systems—not just coding software, but how to design and test systems over their full lifecycle.

"Ordinarily, you can look at a program's input and see if the output is correct. But what if you can't tell? After all, non-testable programs are written to answer questions whose solutions are unknown. What if the answers are wrong in some cases, but not others? These programs could have all sorts of arbitrary errors, but how would we know?"

Kaiser has developed a number of approaches to test machine learning programs. One is a technique many math students will remember. She provides a problem, then changes the order of inputs to see if the program still generates the same answer.

"We found a lot of bugs in certain packages widely used in the machine learning community by using these approaches," Kaiser said.

She has also developed methods to test for errors on computers deployed in the field. Field tests look at the widest possible range of software operating conditions. With so many variations, she can find subtle errors that elude even the most comprehensive lab testing programs.

"You can never get all the bugs out of them," she said, "but the more bugs you remove, the better."

B.S., Massachusetts Institute of Technology, 1979; M.S., MIT, 1980; Ph.D., Carnegie Mellon University, 1985

Indexing Videos Automatically

JOHN KENDER

Professor of Computer Science

Many colleges videotape classes so students can review lectures, notes, equations, pictures, presentations, computer screens, and simulations. Yet students rarely use videos to review for exams. Why? Because it takes too long to find the topics and references they need.

John Kender hopes to solve that problem with software that automatically indexes videos. Just like in a book, his index enables people to find exactly what they want in a video. "It is hard to index a video," said Kender. "Most presenters move around and change the subject. Those taping them often lack training. There are none of the classical clues, like fades or establishing shots, to indicate a change in topic. We are developing software to find those clues and create an index."

Kender has tested the software on videotaped lectures at Columbia. "We have shown that our tools helped students effectively locate the parts of lectures they wanted to study," he said. "After we gave them the tools, their grades improved between midterm and final exams. If you have a good way of reviewing educational videos, it pays off in your grades."

Indexing videos is no simple task. "A professor may start a lecture with a slide, move to a website, then stop to answer a question. That may trigger something he or she forgot to say earlier. They may start four new ideas without finishing previous ones," Kender related.

Kender's team found several ways to keep track of this convoluted discourse. One student developed software that recognized and indexed programming languages when they flashed across the screen. "Students can ask for all examples of software code and quickly page through them to find what they want," Kender said.

Another program reads words from presentations, handwriting from white boards, and captures spoken words using speech-to-text software. Although handwriting and speech-to-text identification is not highly accurate, speakers repeat key concepts often enough to locate them in the video. Another program matches the resulting index with textbook chapters and articles that cover the same material.

A current project involves gestures. Instructors typically use different gestures when reviewing old material, introducing concepts, or working through difficult problems. These gestures vary from teacher to teacher. Kender hopes to decode their meaning by correlating them with such actions as showing new slides, writing on a board, or introducing new words.

Visual indexing could change how people use videos. "It's like providing a table of contents and index for a book that didn't have them before," Kender concluded.

B.S., University of Detroit, 1970; M.S., University of Michigan, 1972; Ph.D., Carnegie Mellon University, 1980

“The barbarians are no longer at the gates,” Angelos Keromytis said about computer security. “They are inside the doors and there are not enough guards to repel them.”

Most security systems are designed to keep bad guys out, and can do little once they are inside, Keromytis explained. “We start with the proposition that attackers will compromise your system, despite your best efforts to keep them out. The only solution is to make systems that are self-healing and self-protecting,” he said.

Keromytis’ approach is to teach computers to act like the best human experts, if they had all the time in the world to react to an attack. “We want the computer to recognize an attack, see what happens, and come up with a way to modify the system so that it blocks the attack,” he said.

Most attackers take advantage of the fact that nearly all computers on a network run the same software. If an attacker finds a vulnerability in one computer, it can attack all the computers. Keromytis turns this into an asset. His software monitors each system, noting when attacks fail or succeed and looking for unusual behavior.

When the alarm sounds, his security system isolates the infected computer. Then it analyzes recent events to find the trigger—an e-mail virus, a malicious download, a tainted document—that set it off. The system automatically attempts to write software code to fix the problem, testing different approaches until it finds one that works. It then rolls back the computer to a time before the attack and inserts the fix. The entire process takes only fractions of a second.

The newly inoculated computer also passes information about the threat around the network. Each computer then builds its own fix. This build-your-own approach prevents hackers from somehow attaching viruses to fake fixes.

“What we’re trying to do is build systems where the individual computers and servers collaborate to prevent attacks, fix attacks that succeed, and then send information to other parts of the network about the vulnerability so they can fix it too,” he said.

Keromytis is currently testing the software and plans to scale up to larger systems soon. He is also looking at ways to find viruses that wait weeks or months until erupting.

The barbarians may have gotten through the gates, but in the future they will find the doors barred by a new generation of persistent guards.

B.S., University of Crete-Heraklion (Greece), 1996; M.S., University of Pennsylvania, 1997; Ph.D., University of Pennsylvania, 2001

Protecting Computers After the Barbarians are Inside the Gate

ANGELOS KEROMYTIS

Associate Professor of Computer Science

Accelerating Processing's Family Van

MARTHA KIM

Assistant Professor of Computer Science

According to Martha Kim, the typical computer processor is like the family van. It is a good all-around machine, but it achieves its flexibility by sacrificing power and performance. It will get you there, but it's not made for off-road adventure or hugging the curves at Le Mans.

Modern computers have the potential to act as vans, sport cars, and motorcycles —all on the same chip, Kim said. It is just a matter of getting under the hood and adding accelerators, small chips-within-a-chip designed to process certain types of data very efficiently.

Accelerators are possible because today's chips are so large. Many contain several separate processors and upwards of two billion transistors. "We have transistors to burn," she said. "If we could organize a few hundred thousand of these transistors into a specialized accelerator, we could handle certain types of data 100 times faster and with 100 times less power."

Those gains are possible because everything needed to process information would reside in the accelerator. "Instead of reading and decoding software instructions specifying how to manipulate the data, the accelerator could start processing immediately without waiting for software instruction," said Kim.

"Today's integrated circuits offer an embarrassment of transistors," she added. "The challenge is how to translate efficiently these raw resources into easy-to-use, high performance, low-power processors. Spending some transistors on special purpose data processors, which store and manipulate structured data types, could simultaneously boost performance and conserve power. "

Accelerators are relatively simple to define in hardware but can be very difficult to use in software. "With accelerators, the programmer not only has to write the code, but has to coordinate what parts of the program should run on which accelerators, and then reassemble the results. Also, if the number and type of accelerators differs from chip to chip, programming becomes even more complex," Kim explained.

Kim's goal is to create common interfaces and tool chains to protect the programmer from this complexity. "The programmer would write code normally. A compiler would track the libraries and data structures used by the application as well as the accelerators available on the chip" she said. "It would do the job of matching parts of the computation with the available accelerators."

In other words, Kim does not want to make her van into a Formula 1 race car. She wants to keep all of the van's flexibility and still take those turns at ridiculously high speeds.

B.A., Harvard, 2002; M.E., University of Lugano (Switzerland), 2003; Ph.D., University of Washington, 2008

Connecting Bits to Life

PETER KINGET

Professor of Electrical Engineering

We live and work in a digital society, surrounded by cell phones, laptops, cameras, iPods, and other electronic devices constantly in use. And the desire for more automation, more information, and more broadband access with better, faster, cheaper mobile infrastructure continues to increase exponentially across the globe. But the physical world around us is analog. Music, speech, images, physiological signals, radio waves, any physical signal is continuous in time and in value. As our information society transitions to more and more digital media and communications, the need for interfaces between real-world analog signals and digital signals (bits) keeps growing drastically. For instance, voices need to be converted to digitized pulses and vice versa on cell phones, music has to be translated into bits for storage and converted back to sounds we can enjoy, and images on digital cameras need to be changed to digitized pixels and then reversed.

The challenge is how to keep all our digital devices connected to the real world with better quality, more pixels, more bits, while needing less—less space, less energy, and less cost. Peter Kinget is one of the researchers leading the way in, as he puts it, "connecting bits to life."

Kinget's research is focused on designing efficient integrated circuits ("chips") that connect digital electronic circuits to the real world. The relentless scaling of semiconductor devices to nanoscale dimensions, a.k.a. Moore's law, has brought a tremendous performance improvement and cost reduction to digital electronics. But the design of interface circuits using nanoscale devices is becoming progressively harder while the performance demands keep increasing. Inventing new circuit techniques is key to keeping electronics, along with all the systems that rely on them, progressing.

These innovations are important enablers to a large variety of applications in which Kinget's group is involved. Novel wireless links using very short pulses to communicate require so little power that they can operate perpetually on energy harvested from the environment, rather than needing batteries. Such highly energy-efficient communication capabilities are key to the realization of EnHANTs (Energy Harvesting Networked Tags), a new type of tags that will enable us to connect and network everyday objects that are part of our daily lives, like wallets, keys, toys, clothing, produce, and even furniture.

But novel integrated circuits reach far beyond communications. For instance, smart power circuits used in combination with new materials and fabrication techniques to make high quality printed capacitors can also be used to convert electrical wall AC power efficiently to DC power for "greener" types of lighting, employing LEDs to replace wasteful incandescent bulbs.

M.S., Katholieke Universiteit Leuven (Belgium), 1990; Ph.D., Katholieke Universiteit Leuven, 1996

Linking Domino Theories to Real-World Pricing

STEVEN S.G. KOU

Professor of Industrial Engineering and Operations Research

INFORMATION

Many blame structured financial instruments, such as credit default swaps and collateralized debt obligations, for the 2008 recession. Yet similar products traded for decades without problems. Even now, billions of dollars in structured debt trade daily. Steven Kou has made it his mission to make these products safer.

"As an engineer, I'm interested in linking economic theory to real-world pricing of structured financial products," said Kou. "Economists understand the structure of economic forces, and statisticians understand how one event triggers another, like the aftershocks of an earthquake. We're trying to apply both to the details of financial products."

Structured instruments reduce risk, he said. He points to instruments that pool corporate bonds. Instead of buying a bond from one company, investors can buy a diverse portfolio of bonds from companies in different industries.

Financial firms typically divide this basket of bonds into risk categories, or tranches. The top tranche has the lowest risk but the lowest returns. It loses money only if 30 percent of the bonds default. This is highly unlikely, and it trades like a highly rated bond. The bottom tranche has the highest return but loses money if only a smaller percentage of bonds default.

"There's a value to this," Kou explained. "Pension funds, for example, cannot invest in bonds rated less than AAA. Many strong companies have lower credit ratings. If their bonds are included in the top tranche, a pension fund can buy them without great risk and still receive a higher return."

Many investors were lured by that combination of higher returns and lower risk. They believed that even if conditions in one industry forced a company to default, diversification would keep their investments safe.

In 2008, though, that assumption was upended. "The model we had been using was no good. During a severe crisis, we found that when one company defaults, others outside its industry are more likely to default," Kou said.

Kou calls this "default clustering." To understand how it affects risk and value, he builds models that draw on both economics and financial engineering.

Kou said the models will help set more realistic prices for structured financial instruments. Initial results are promising. Just before Lehman Brothers went bankrupt, conventional models set the cost of insuring the top tranche of corporate bonds at about $7,000. His model priced it at around $52,000.

"That's more consistent with what happened in the market," he said.

M.A., Columbia, 1992; Ph.D., Columbia, 1995

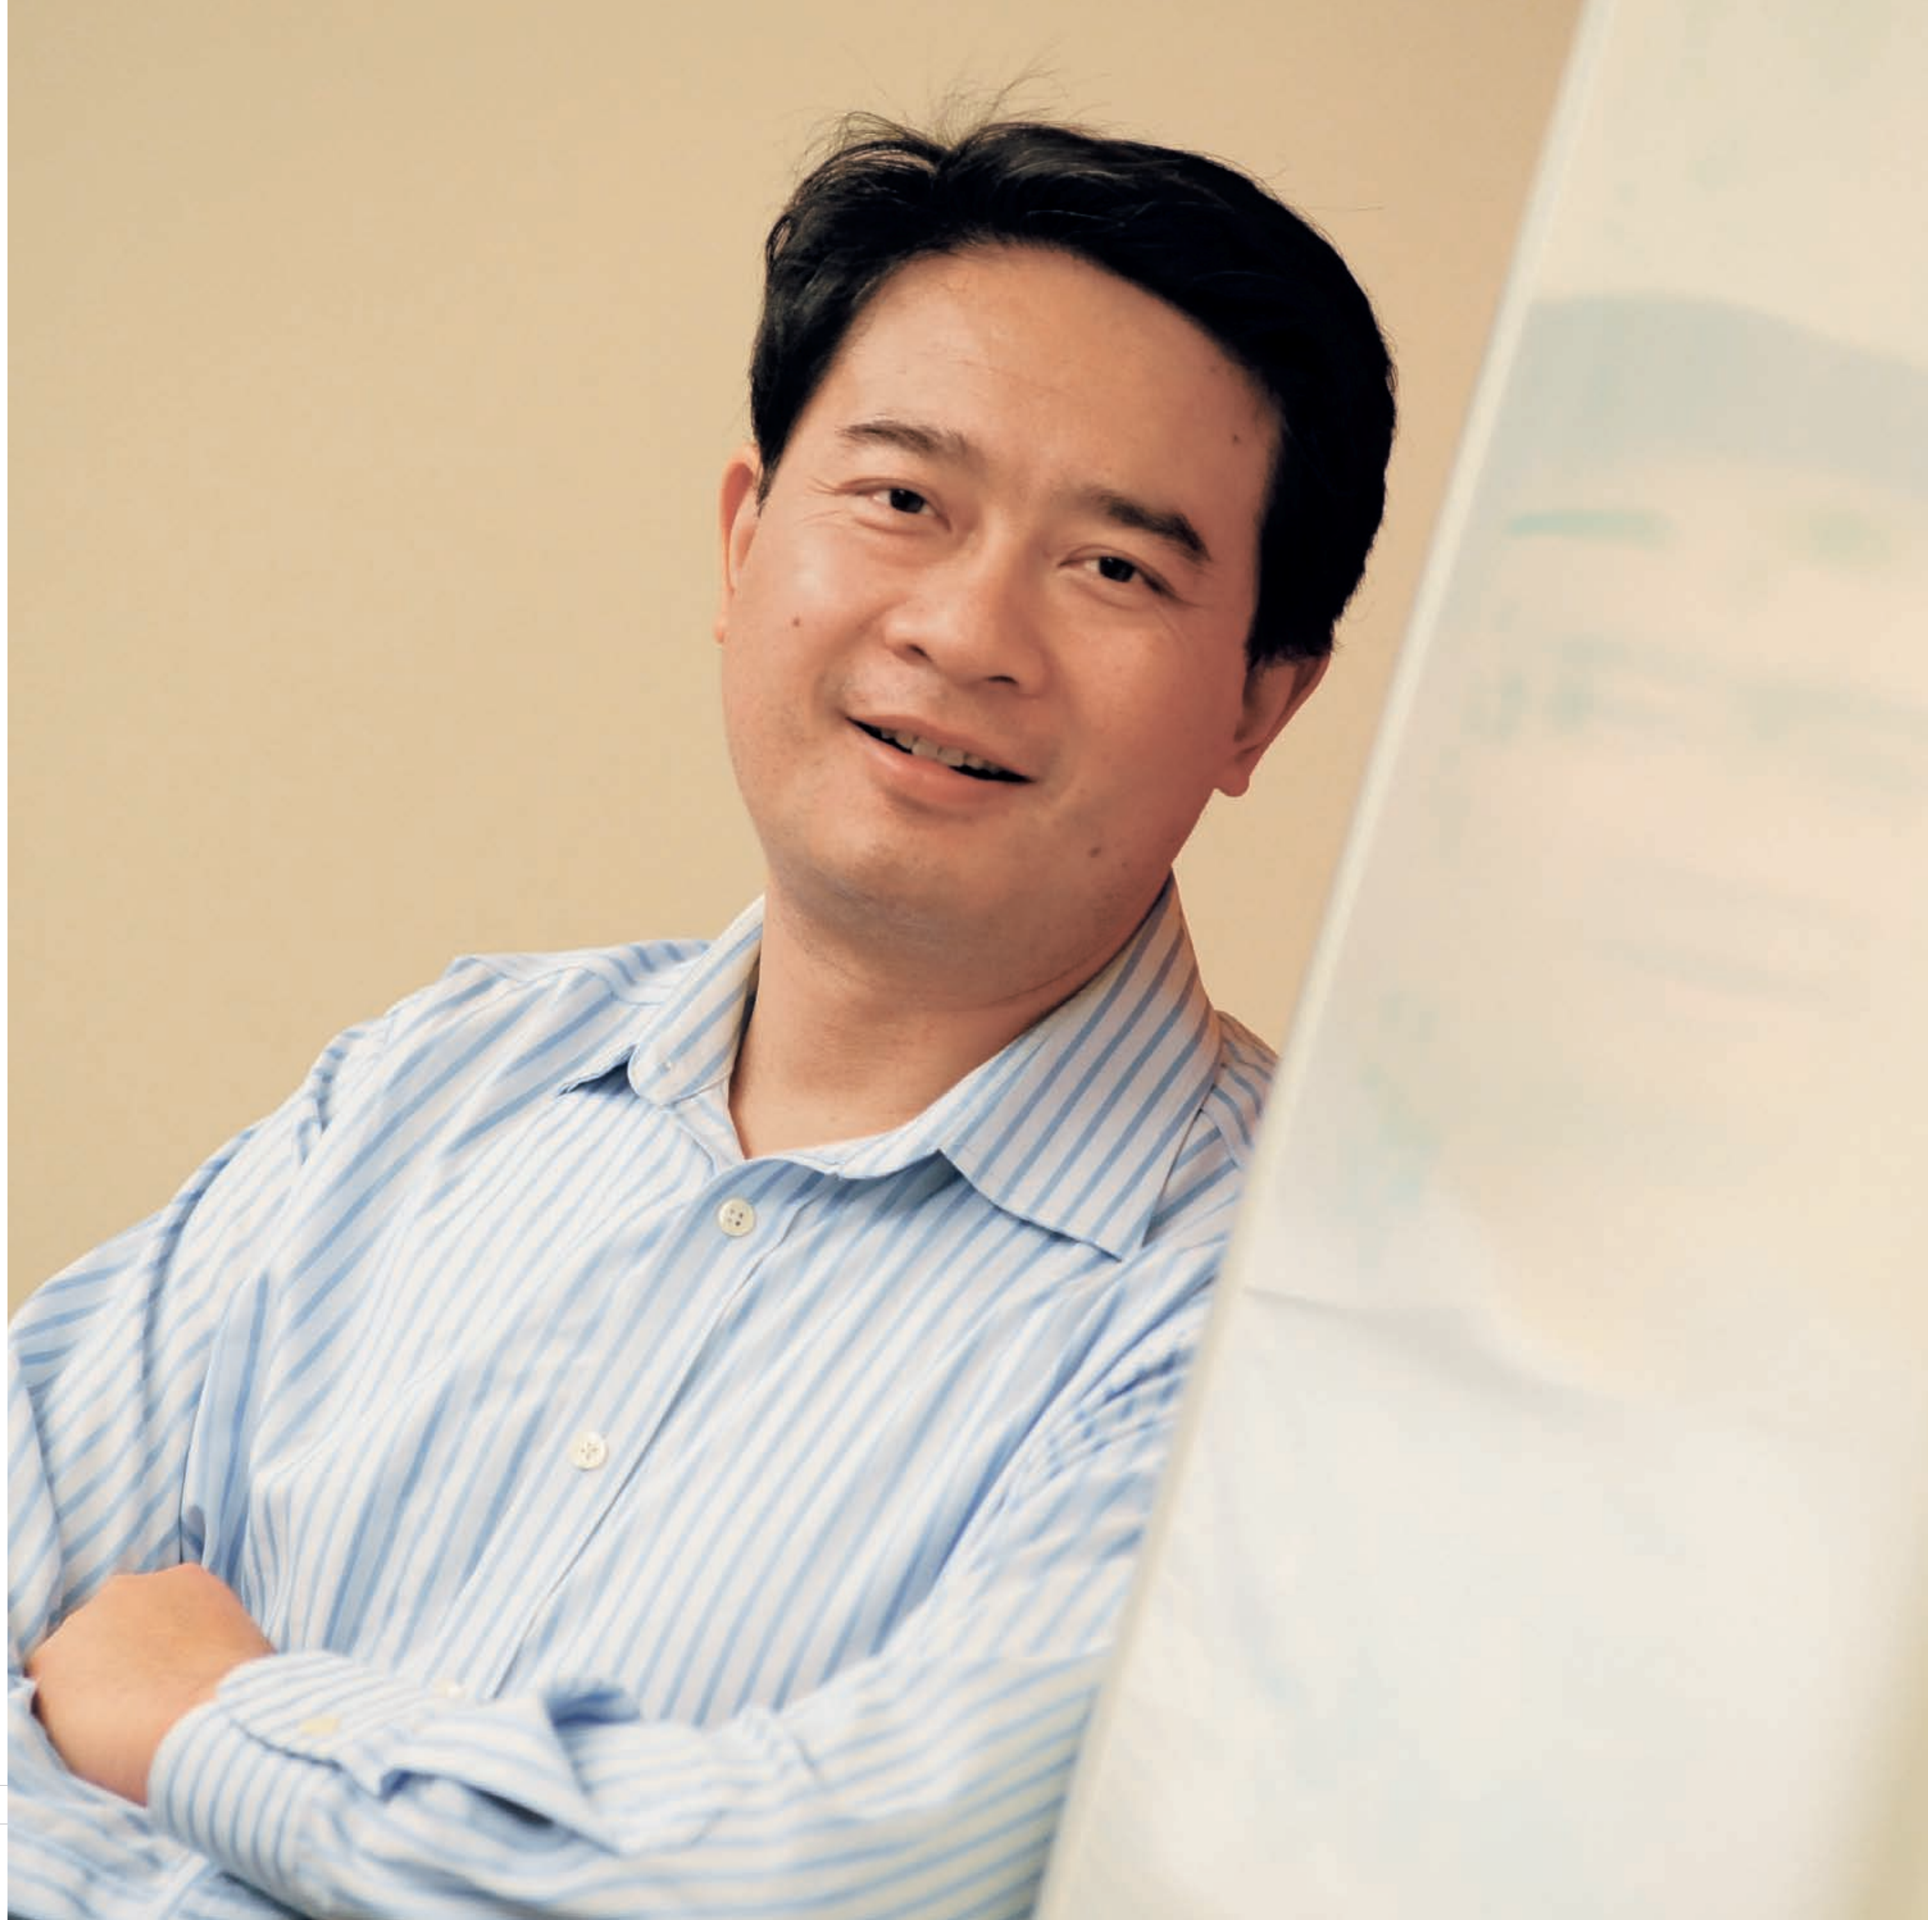

Pushing the Performance of Silicon-Based Systems

HARISH KRISHNASWAMY

Assistant Professor of Electrical Engineering

INFORMATION

Automobile collisions account for tens of thousands of fatalities in the United States annually. While the most expensive automobiles have on-board collision avoidance systems, such technology is priced out of the market for most drivers. The cost has much to do with the technologies that are currently used to implement these systems. Current sensors rely on multiple integrated-circuit chips based on compound-semiconductor technologies, resulting in systems that are large, bulky, power-inefficient, and expensive.

Silicon-based millimeter wave technology could make automobile collision avoidance systems as common as seatbelts in the cars of the future. Millimeter waves deliver good directionality, and offer a large amount of available bandwidth not currently being used, making them functionally comparable to fiber optics without the financial and logistical challenges. Silicon-based technologies offer the opportunity to integrate complex sensors onto a single chip, greatly reducing power, cost, and size. This technology's utility is wide ranging and includes collision warning systems, blind spot analysis, and pedestrian detection. It also is being explored for high-data-rate personal area networks for future "wireless homes," non-invasive medical imaging, airborne chemical sensing, and concealed-weapon detection for security systems.

Integrating extremely high frequency electronic circuits and systems into silicon-based technologies is one of the grand challenges of electronics, and where Harish Krishnaswamy is applying his research efforts. He pioneered silicon-based, nonlinear, multifunctional circuits, and systems which, when coupled with millimeter wave technology, allow multiple simultaneous functions to be performed on a single, compact, power-efficient chip. A nonlinear, multifunctional phased-array transceiver chip won the prestigious Lewis Winner Award for Outstanding Paper at the 2007 IEEE International Solid-State Circuits Conference. Krishnaswamy is also working on new Multiple-Input, Multiple-Output (MIMO) radar concepts that use multiple transmitting and receiving antennas to capture a more detailed and accurate image of the scene around the vehicle.

The Krishnaswamy Group at Columbia University analyzes, designs, and experimentally verifies novel integrated devices, circuits, and systems for a variety of radio frequency and millimeter-wave applications. His research efforts blur the boundaries between circuits, electromagnetics, device physics, and communication/signal processing theory. Results include a variable-phase ring-oscillator based architecture for radio-frequency (RF) and millimeter-wave phased arrays, architectures and circuits for single-chip MIMO radar, timed arrays for ultra-wideband beamforming, and high-performance RF and millimeter wave building blocks for wireless transceivers.

B.Tech, Indian Institute of Technology (Madras), 2001; M.S., University of Southern California, 2003; Ph.D., University of Southern California, 2009

Designing Ways to Account for Foreseeable Financial Risk

TIM SIU-TANG LEUNG

Assistant Professor of Industrial Engineering and Operations Research

In order to attract and retain top executive talent, many firms develop sophisticated compensation arrangements that include employee stock options (ESOs) and securities. In fact, almost half the compensation for corporate CEOs is usually in the form of stock options.

But what's the true value of these ESOs? It's difficult to determine because value is dependent upon fluctuations in the stock market and when an ESO owner exercises the option to cash in those compensation vehicles. Because timing is variable, valuation of ESOs can be somewhat random. Without a viable model for ESO valuation, the actual and true costs of these compensation vehicles cannot be correctly reflected in any company's financial bottom line. This puts a burden on other shareholders, the company, and the economy.

Tightening the gaps in accepted practice of valuation for compensation vehicles like ESOs can have direct impact on businesses as well as the stabilization of the economy. Key to accomplishing that is the application of mathematical acumen and practical financial knowledge—components of financial engineering.

Financial engineering is both the art and science of evaluating, structuring, and pricing financial instruments and designing strategies to reduce risk and maximize opportunities. Through innovative, analytical procedures, financial engineers help individual and institutional investors as well as regulators understand and manage financial risk.

Tim Siu-Tang Leung uses financial theory, engineering methodology, and mathematics to build reliable models that account for foreseeable financial risks. His research interests are in financial engineering, especially in the valuation of ESOs and credit derivatives. The National Science Foundation is underwriting his research in stochastic modeling of risk aversion and its implications for derivative pricing and risk management.

He has made significant contributions to the field with a revised ESO valuation model that takes into account the complex contractual features and the realistic behaviors of ESO holders (i.e., due to their heightened risk perceptions—fear of market crash or job termination—ESO holders usually exercise options to cash in early). He has also developed strategies to help employees hedge some of the risk involved with owning ESOs. His research has led to interesting mathematics including analytical and numerical studies of several combined stochastic control and optimal stopping problems. The mathematical tools from his research are also being applied to tackle other financial engineering challenges.

B.S., Cornell, 2003; Ph.D., Princeton, 2008

Making Robots Learn

RICHARD W. LONGMAN

Professor of Mechanical Engineering and Professor of Civil Engineering and Engineering Mechanics

During Richard Longman's sabbatical in 1984 he initiated research in three new fields, becoming one of the very early contributors to each. With support of ex-doctoral student Robert Lindberg at the Naval Research Laboratory, an Egleston Medal recipient, he started research on robotics in space. The shuttle arm can handle a load of mass similar to the shuttle, and this creates a question: Which end of the arm is the base and which is the load? Two of his early papers appeared in the first book on space robotics produced by the Carnegie Mellon Robotics Institute.

With German collaborators, he started research on time optimal control of robots, something that challenges numerical solution methods. One research focus was a press chain on the Mercedes production line near Stuttgart. The objective was to increase productivity by making the slowest robot get its job done faster. A series of publications progressed from idealized investigations to ones including detailed hardware constraints. Similar productivity problems appear in the production of semiconductor chips.

When a robot is commanded to follow a trajectory, it will repeatedly follow a somewhat different path. Robots often do the same operation hundreds of times a day, making the same errors each time. Longman considered this a bit stupid—can't we make a control system that learns from its experience to do what we ask? He started work on this at the University of Newcastle in Australia. Since then, this problem has developed into the fields of iterative learning control (ILC) and repetitive control (RC).

Longman has produced some 250 publications in this area, and is known for advancing the theory in a way that produces improved real-world performance. Experiments on a robot at NASA improved tracking accuracy by a factor of 1000 in just 12 iterations for learning. The methods can apply to a very large number of feedback control systems, creating high precision motion by improved algorithms instead of higher precision hardware.

At Seagate Technology, experiments reduced the repeatable error in computer disk drives by 98 percent. Similar experiments improved paper handling in copy machines at Xerox. Experiments also demonstrated improved beam focus at the 8 GeV (one thousand million electron volts) accelerator at Jefferson National Accelerator Facility. Longman is currently working on similar experiments at the Naval Postgraduates School on jitter control in laser optics on spacecraft.

ILC and RC aim for high precision motion and optimal control aims for fast motion. Longman is working to develop a marriage between these research areas to simultaneously get the benefits of both—aiming for higher quality products created with improved productivity.

B.A., University of California-Riverside, 1965; M.S., University of California-San Diego, 1967; M.A., UC San Diego, 1969; Ph.D., UC San Diego, 1969

INFORMATION

Securing the Lock after the Key is Stolen

TAL MALKIN

Associate Professor of Computer Science

From online transactions and ATM machines to databases and voting, cryptography lets us share critical information while keeping it safe. Yet cryptographic systems have a weakness. They rely on keys to code and decode messages, and keys can be cracked or stolen.

"Traditional cryptography depends on the assumption that an attacker has no access to secret keys," Tal Malkin said. "Yet sometimes an attacker can hack into a computer or tamper with your hardware. Part of my work is to maintain security even against such adversaries."

She envisions systems that respond when attacked. "We can build systems where the key evolves to protect against an adversary who reads or changes part of the key. Even if an adversary reads the entire key, we can protect future transactions," she said.

Another assumption underlying cryptographic systems is that there is some hard problem that no attacker can solve. For example, the public key software used for secure Internet transactions often relies on the assumption that it is hard to factor the product of two very large prime numbers.

"No one can prove the factoring problem is hard to solve," said Malkin. "We assume it is because people have worked on this problem for decades. They have developed sophisticated techniques that are much better than the more obvious approaches, but even those procedures require as many operations as the number of atoms in the universe. But if someone does find an efficient solution, it would break all encryption on the Internet."

As part of her research, Malkin also studies the mathematical foundations of cryptography, searching for the minimal assumptions needed to guarantee security. This starts with studying primitives, such as one-way functions, that act as cryptographic building blocks. "Primitives are small, simple to describe, easy to compute, and hard to crack," Malkin said. "Cryptographers can combine small primitives to form complex, multilayered security systems."

Malkin has also focused on general systems for secure computation among two or more parties, as well as optimizing their performance for specific purposes. One example is the no-fly list. The government wants to keep it secret, while airlines want to protect passenger privacy. Malkin has developed a fast way to exchange critical information without showing compromising data. Other applications of secure computation include online voting, sharing national intelligence, and bidding on projects.

In today's increasingly interconnected world, Malkin's work on provably secure cryptographic protocols could help protect some very important secrets.

B.S., Bar-Ilan University (Israel), 1993; M.S., Weizmann Institute of Science (Israel), 1995; Ph.D., Massachusetts Institute of Technology, 2000

Summarizing the News (Automatically)

KATHLEEN MCKEOWN

Henry and Gertrude Rothschild Professor of Computer Science

At first glance, the nine-year-old Columbia Newsblaster website (newsblaster.cs.columbia.edu) looks like Google News. Both feature the day's top stories plus sections on national, world, financial, and science/technology news.

The difference is their technologies. Google lists the first few sentences of one news article and links to similar stories. Newsblaster publishes summaries of a dozen or more articles—all written and edited by software developed by Kathleen McKeown.

"Newsblaster summarizes multiple news articles," said McKeown. "We're using similar technology to answer questions from information on the web. Today, users read the documents their search returns to see if they are relevant. Our software takes the next step. It looks into the documents, pulls out the relevant information, and summarizes it in a paragraph."

McKeown's software starts by scraping 25 different websites for news every night. It uses key words to cluster articles and categorize topics, counting the number of articles in each cluster to determine its importance.

Once classified, the software uses several approaches to generate summaries. First, it extracts sentences from important sources, such as stories from prominent newspapers and wire services.

It also pairs each sentence with every other sentence in the cluster. It analyzes their similarity and groups related themes together. "The software lines up the sentences in each group side by side and looks at where they overlap or intersect," McKeown explained. "It is looking for phrases that say the same thing, where words overlap or there is paraphrasing.

"The software parses the sentences for grammatical structure, so it knows that this phrase functioned as a noun and that phrase acts as an adjective. This helps it align similar sentences and fuse phrases to create summary sentences. It then generates the summary by ordering the sentences, using information about chronological order of the events. It also edits for coherence, substitutes proper nouns for pronouns, and adds or removes references, depending on whether a person or place is well known or not," she said.

The core technology has found other uses. A small company is using it to power smartphone applications that track and create timelines for breaking news on specific topics. Another application responds to open-ended questions, generating summaries of information about, for example, a particular event or a particular person. A third creates English summaries from news sources in other languages.

While some Newsblaster stories read like newspaper articles, others are choppier. Still, the technology could become an important tool for making sense of all the information on the web.

B.A., Brown, 1976; M.S., University of Pennsylvania, 1979; Ph.D., University of Pennsylvania, 1982

INFORMATION

Boosting Profits with Peer-to-Peer Networks

VISHAL MISRA

Associate Professor of Computer Science

Peer-to-peer (P2P) networks exploded onto the scene around 2000. That is when Napster, LimeWire, BitTorrent, and similar services made it possible for anyone to download libraries of music and movies for free over the Internet. The new technology gutted music industry profits and led to massive layoffs and downsizing.

Yet peer-to-peer networks are not inherently bad for profits, Vishal Misra argued. In fact, they may prove the most efficient and least expensive way to share media over the Internet.

In 2000, large peer-to-peer networks were something new. Instead of warehousing information on a central computer, they took advantage of files distributed on PCs throughout the network to store and send files to other users.

"Smartphones need lots of bandwidth," said Misra. "Wireless providers want us to buy femtocells, small broadcast towers to improve performance in our homes and offices. Instead, they should give us femtocells. Then they could offload traffic from their cell towers and reduce the number of new towers they need to build to support their smartphones.

"Everyone agrees P2P is a great technical solution. The more users, the more resources the network provides and the faster it responds to requests," Misra said. This is the opposite of today's centralized client-server model, which must keep investing in more servers as network demand grows larger.

"There shouldn't be this war between P2P users and people who own music and movie copyrights on the other," he continued. "We need an economic reboot so that the system works for both camps."

To understand how that might be possible, Misra used game theory to analyze the problem. Ordinarily, models that involve cooperative interactions are extremely hard to calculate, especially for millions of users. Misra simplified those calculations by applying theories based on fluid flow to the continuum of users and peers. "It's like analyzing a glass of water as a fluid instead of trillions of water molecules. By representing millions of peers as a fluid," he said, "it is easier to see their behavior and compute the right incentives."

For example, Misra estimates that providers of such content as live TV and video-on-demand could save over 90 percent of their Internet distribution costs through user-based P2P networks. "The stores could save lots of money, and people who own legal copies of media might be willing to share them if they receive part of those savings," said Misra.

"Peers, or users, can help providers reduce costs, as long as incentive structure are in place to reward them," he concluded.

B.S., Indian Institute of Technology, 1992; M.S., University of Massachusetts-Amherst, 1996; Ph.D., University of Massachusetts-Amherst, 2000

Picturing the World in New Ways

SHREE K. NAYAR

T. C. Chang Professor of Computer Science

INFORMATION

Shree K. Nayar's work is all about seeing things differently. "The basic principles of photography have remained unchanged since the earliest camera obscura," Nayar explained. "Cameras use an aperture to capture light, a lens to focus it, and some medium to capture the familiar linear perspective image. In the 1990s, I started asking whether we could use new optics and a computational processing to produce new types of images."

One of Nayar's first inventions was the Omnicam. Its combination of lenses and mirrors captures panoramic 360 degree images in a single click. "The image is distorted, since you can't map a sphere to a flat surface without distortion, but we corrected that with mapping software," he said. "In fact, a single 360 degree image could be used to generate any number of traditional views of the scene.

"Placed in the middle of a table of people, it gives the illusion of multiple cameras pointed at individuals during a video conference, although it is one camera with no moving parts," Nayar said. The camera is also used for surveillance.

Nayar's next invention, a high dynamic range camera, takes better photographs of scenes that mix dark and light areas. "Let's say you try to take a picture of a scene with shadows and a bright sky. Today's digital cameras cannot reveal details within the shadows and the sky. If the sky comes out well, the shadows do not, and vice versa," Nayar said.

Nayar's solution is to use an image sensor with a patterned optical mask on it. The mask ensures that neighboring pixels on the sensor have different sensitivities to light. His software decodes the captured photo to produce one that captures the shaded clouds in the sky and the objects in the shadows. Sony has prototyped the technology for use in its digital cameras.

A third camera enables photographers to focus on close-up details without blurring background features. Nayar does this by physically sweeping the image sensor of the camera through an entire focal range, during the exposure of a single photo. The captured photo is again processed by software to obtain one where everything appears in focus.

Nayar has also launched a project to help children around the world learn science, art, and culture by assembling and using a digital camera. His Bigshot Camera has panoramic and stereo imaging capabilities, and makes it easy to post photos on the web.

"Each picture is a window on another culture, and youngsters can learn about those cultures from their peers," he said. To Nayar, it is just another way of seeing things differently.

B.S., Birla Institute of Technology (India), 1984; M.S., North Carolina State University, 1986; Ph.D., Carnegie Mellon University, 1990

Delivering Desktop Computing from the Cloud

JASON NIEH

Associate Professor of Computer Science

INFORMATION

Cloud computing—delivering software and services from a central computer to desktop terminals—is arguably the hottest topic in computing today. The reasons are economic. PC hardware prices continue to fall, but maintenance costs continue to rise.

"If you're a large corporation with 50,000 or 100,000 desktops, you're fixing broken hardware, guarding against viruses, and patching and upgrading software for each one of them. The costs are astronomical," explained Jason Nieh.

If the PC-on-a-desktop paradigm is broken, what will replace it? "Most analysts believe we are moving to cloud computing, where corporate computers run only in secure data centers where they are protected, secure, and easier to manage and service. If a desktop fails, it doesn't matter because all the memory and files actually reside in the data center," said Nieh.

But cloud computing has a weakness: speed. Centralized applications run slower than the same program on a local PC. This is especially true for programs with graphical displays.

"A modest display has 1024 x 768 pixels, and each pixel has 32 bits of data," added Nieh. "Displays update 30 to 60 times per second or more, so you're potentially sending a gigabyte or more of data per second to each PC on the network, and that can slow response times."

Computer scientists have tried to compress data to reduce the load. This helps, but it requires additional computing power and fails to handle gracefully today's complex graphical interfaces.

Nieh uses intelligent software to reduce data flows and response times from the cloud. In Nieh's scheme, the application draws the screen on a virtual display. Then his program analyzes what is on the display, and sends commands to the desktop terminal, instructing it on how to redraw the screen. Many of the most common commands are embedded in the graphics card's hardware, so they operate very fast. The system updates the terminal by sending only those portions of the display that change, enabling very fast response times.

The big payoff comes when connecting to the Internet. Data centers almost always have the fastest Internet connections. "They update web pages much faster than local desktops, laptops, or smartphones," Nieh said. "If a carrier uses this technique, you don't have to settle for the limited functionality of smartphones that run some software but not others.

"You get improved functionality and improved performance that makes you feel like you're right there, and you get it on your smartphones, desktops, and laptops."

B.S., Massachusetts Institute of Technology, 1989; M.S., Stanford, 1990; Ph.D., Stanford, 1999

Marching Without a Beat

STEVEN NOWICK

Professor of Computer Science and of Electrical Engineering

For decades, computer processors were typically organized like marching bands: a conductor kept time and band members stepped to the beat. In processors, a clock's pulse determined when all computations and data movement occurred. Today, that paradigm is breaking down, Steven Nowick explained.

Modern processors consist of a handful of smaller processors, or cores. "When you have four separate cores, it is difficult for one clock to keep them in lockstep," Nowick said. The problem will only worsen when future processors have dozens of cores.

Today's transistors also pose problems. As they shrink to a few tens of nanometers, they become much more variable. "Their speeds vary depending on temperature, voltage, and how they are manufactured. Their unpredictability is a major design challenge," Nowick said.

Nowick and colleagues at other institutions have been pursuing an alternative approach: eliminate the clock and let digital components operate at their own speeds. "Let them communicate as conditions require, and make their own decisions with their neighbors about when they need new data and when they will output results," he stated.

"Most digital systems have clocks running at billions of cycles per second. Everything operates in lockstep with that clock," he continued. "As circuits get larger and more complex, imposing fixed timing on billions of transistors and millions of components is a huge design effort. We think we can solve these problems with asynchronous, or clockless, circuits."

It sounds chaotic, but the Internet works the same way, Nowick said. "People around the world add, update, and remove web pages individually, without any centralized control mechanism."

In addition to solving timing issues, asynchronous digital systems could provide other advantages. In synchronous chips, even idle components are activated every clock cycle, like band members marching in place. In contrast, the on-demand components in asynchronous systems respond only when necessary. This conserves energy and can prolong battery life in laptops, smartphones, and other portable devices.

Asynchronous processors are potentially easier to design, since new circuits do not have to be synchronized with the entire chip. "It's a Lego-like system, which can be snapped together," Nowick said. Hurdles remain. Engineers need new software tools to design asynchronous circuits, and face subtle issues in designing these circuits correctly.

Nowick is currently working on both challenges, including projects to design a flexible asynchronous interconnection network for future desktop parallel computers, and ultra-low energy signal processors for hearing-aids and medical implants.

B.A., Yale, 1976; M.A., Columbia, 1979; Ph.D. Stanford, 1993

Searching for a Heavy Tail

MARIANA OLVERA-CRAVIOTO

Assistant Professor of Industrial Engineering and Operations Research

Search Google and within a fraction of a second it will return a list of the most popular websites on the subject. Or will it? Mariana Olvera-Cravioto has been trying to answer that question by understanding what makes a website popular on Google.

The principles behind Google's page-ranking system are well known. It weighs links to and from a page, as well as links of other pages on the same website. Yet the details remain unclear. For example, what counts more, a few links from such important websites as Wikipedia or Technorati, or many links from less significant pages? And what happens to the rankings if you change Google's search algorithm ever so slightly?

To probe those questions, Olvera-Cravioto relies on a form of probabilistic theory called heavy tail theory. To understand it, consider a normal bell curve. Most samples are grouped close to the center, or mean, and decline rapidly towards the ends of the curve.

Heavy tails have far more outliers at the ends of their curves than normal distributions. They are surprisingly common. They show up in the distribution of wealth (few people own more assets), oil reserves (a few have the most value), insurance payouts, and the time supercomputers spend completing tasks.

"Internet video transmission is an example of a heavy tail distribution," she said. "Streaming video only transmits pixels that change. Most of the time, that's relatively few pixels. But then the camera changes angles and the whole screen is refreshed. It happens less often, but accounts for most of the transmitted data." Google search results also have a heavy tail distribution.

"The mathematical techniques needed to solve heavy tails are completely different from what we use with well-behaved distributions," Olvera-Cravioto said. Those techniques provided some deep insights into Google's page rankings.

"Before we started our analysis, it was not obvious what determined the relative rankings of websites," Olvera-Cravioto said. Heavy tail analysis, for example, shows how large numbers of links outweigh important links in popularity rankings.

"Once you understand how it works, you can engineer search algorithms for specific purposes," she said. "Maybe you want to rank stores by number of sales rather than links, or measure the importance of a paper by how many times it is cited by reliable websites. Adding these things to a search algorithm could make it easier to find the page you're after."

B.S., Instituto Technológico Autónomo de México, 2000; M.S., Stanford, 2004; Ph.D., Stanford, 2006

Improving Large-Scale Disaster Response

FENIOSKY PEÑA-MORA

Morris A. and Alma Schapiro Professor and Professor of Civil Engineering and Engineering Mechanics, of Earth and Environmental Engineering, and of Computer Science

INFORMATION

When disaster strikes, the interdependent complexity of the environment (utilities, transportation, communication infrastructures, homes, and office buildings) can result in a cascading effect that quickly exacerbates the crisis. Large-scale disasters, such as Hurricane Katrina and the earthquake in Haiti, have graphically demonstrated the need for reliable initial disaster preparedness, response, and recovery. In such cases, the immediate availability of critical real-time data is crucial to saving lives.

Feniosky Peña-Mora, dean of The Fu Foundation School of Engineering and Applied Science at Columbia, has developed a new disaster response framework—Collaborative Preparedness, Response, and Recovery (CP2R)—that makes a significant difference in the outcome of such disasters. As part of this framework, he and his research team have created a mobile workstation using an all-terrain, heavy-duty Segway personal transporter outfitted with a payload that can include a Tablet PC, infrared and thermal still and video cameras, Global Positioning System receivers, and other advanced data collection technology. These instruments can collect, archive, analyze, and report large quantities of data to provide better situation awareness of an emerging disaster response scenario, and automatically generate digital models that can be used for disaster response.

By deploying these modified chariots manned by civil engineers, real-time data from first responders can be transmitted to coordination centers by wireless voice and data communication infrastructures. "This new cohort of first responders will provide accurate, real-time information to support technically sound decision-making processes during both the initial disaster response and the recovery phases," said Peña-Mora. "With a legion of mobile workstation chariots, we will be able to mitigate the dynamics of the disaster by improving the dynamics of the disaster response."

Recent testing of the mobile chariot has shown the potential for its success in the field. Despite additional weight from mounted instruments, the unit retained its stability on uneven surfaces and in differing weather conditions. Using digital images collected at the disaster site, decision-makers in coordination centers can evaluate infrastructure stability, study how the first responders are reacting to changing situations, and collect data for future analysis.

Peña-Mora holds appointments as professor of civil engineering, computer science, and earth and environmental engineering, is the author or co-author of more than 150 scholarly publications, and holds five patents, one provisional patent, and one technology disclosure.

B.S., Universidad Nacional Pedro Henríquez Ureña, 1987; Post-Graduate, Universidad Nacional Pedro Henríquez Ureña, 1988; S.M., Massachusetts Institute of Technology, 1991; Sc.D., MIT, 1994

Creating Nanoscale Devices

ARON PINCZUK

Professor of Applied Physics and Applied Mathematics and Professor of Physics

Creating the next generation of electronic devices—be they computers, smartphones or displays—will depend on understanding the properties of materials on the nanoscale—one-billionth of a meter.

Aron Pinczuk's research projects employ advanced optics methods in condensed-matter science, with a focus on understanding the properties of novel materials and the physics of exotic states of matter that emerge in semiconductors at extremely low temperatures. His research findings address issues used by scientists seeking the development of quantum computing and cryptology. The research on graphene, a single atomic layer of graphite, contributes to the quest to initiate a new era in the creation of electronic components.

Pinczuk conducts his research at the Nanoscale Science and Engineering Center at Columbia, in the Department of Applied Physics and Applied Mathematics, and in the Department of Physics. His laboratory had support from the Keck Foundation and his research is funded through the National Science Foundation, the Department of Energy, and the U.S. Office of Naval Research, which support projects that span disciplines in science and engineering.

His research has explored the properties of gallium arsenide, a semiconductor, which is used in advanced optoelectronics, lasers, microwave circuits, and solar cells. To determine material properties in condensed matter systems, he subjects gallium arsenide to temperatures below 0.1 Kelvin, a temperature at which almost everything freezes. At these temperatures, the electrons cool down to make a liquid, emit light, and exhibit new, unexpected behaviors.

His research with gallium arsenide also has added to the basic science needed to develop a quantum computer, in which computational operations are executed in quantum bits. Theoretical studies show that quantum computers can solve certain problems quicker than classic, digital computer systems.

His findings also have assisted those looking to develop ways to use complex quantum states to build a key used to encrypt computer information. Such encrypted keys could be used to improve the security of computer systems.

Pinczuk's research with the carbon material, graphene, is part of the effort to develop a new generation of electronics that use carbon components. He studies the properties of carriers of an electric charge as it travels through a single layer of graphene, which is two-dimensional. Scientists are working on larger scale integration of these layers, which will create multi-layer structures with new properties.

"In the case of graphene, there are new properties that develop when you put all the layers together," said Pinczuk. "It's a field that is rapidly evolving."

Licenciado, University of Buenes Aires (Argentina), 1962; Ph.D., University of Pennsylvania, 1969

Processing Parallel Insights

KENNETH ROSS

Professor of Computer Science

The volume of data we want to analyze is growing even faster than computing power. Kenneth Ross is looking for ways to close the gap. “People are coming up with ever-more challenging database projects, like analyzing the differences in genomes, which have billions of base pairs, among thousands of patients,” Ross said.

Until now, computer scientists have relied on raw increases in computer power to crunch more data. Today, those advances have been harder to achieve. To keep moving forward, engineers reinvented the microprocessor, dividing it into two or more smaller processors, or cores.

Dividing tasks among cores works best when the answers do not depend on the previous step. Databases are like that. “The work you do on one record is pretty much what you do on another, you can process them in parallel,” Ross said.

Yet parallelism comes with its own set of problems, such as cache misses and contention.

Cache misses occur because computer processors have fast and slow memory. They waste hundreds of processing cycles retrieving data from slower memory. Those lost cycles—cache misses—waste half the time needed to perform some tasks.

Ross wants to reorganize data to take up less space in memory. The hard part, he said, is doing this without spending too much time or resources.

“I’m trying to take advantage of relatively recent changes in computer architecture to make database software more efficient,” said Ross. “Computer processors are now made up of four to eight smaller processors, or cores. We have to take advantage of those cores by developing code that runs in parallel.”

Contention occurs when several parallel jobs all need to update a single item. “Each of those jobs needs exclusive access to the item for a short time to keep them from interfering with one another. If the item is sufficiently popular, those jobs get stuck in line waiting for their turn to access the data rather than working in parallel,” Ross explained.

Ross’ recent research seeks to automatically detect contention and then create several clones of the busiest data items. “We want to distribute processes among the clones and then combine results. Again, the key is to do this without using more computer resources than we are saving by eliminating contention,” he said.

From genomics to climate, the sciences are accumulating data at a faster rate than ever before. Ross’ work will help make it possible to analyze that data and see what they really mean.

B.Sc., University of Melbourne, 1986; Ph.D., Stanford, 1991

INFORMATION

Networking Your Wallet, Credit Cards, and Keys

DAN RUBENSTEIN

Associate Professor of Computer Science

Imagine a world where library books tell you they are on the wrong shelf and fruit reports it has gone bad to grocers. It is a universe where you can always find your keys or remote control.

This world is under construction in Dan Rubenstein's laboratories. His team is working with small tag-like devices that attach securely to everything from books to baseball bats. "They will let you track all the things you want to track without being tracked by entities you don't want to track you," Rubenstein said.

The devices are called EnHANTs, which stands for energy-harvesting active network tags. "They're designed to soak up energy from the environment to form a network with the tags around them. The networked tags then keep track of one another," said Rubenstein. "Unlike similar radio frequency identification (RFID) tags, which turn on only when activated by powerful radio transmitters, EnHANTs would generate their own power by harvesting energy from ambient light, tiny vibrations, or temperature changes.

Unfortunately, this is not enough power to stay turned on all the time, communicate more than 10 feet, or send lots of information at a time. To get around those limitations, EnHANTs must network with other nearby EnHANTs and devices.

"Existing network protocols waste too much power to work with devices of EnHANTs' size. We have to be more efficient," Rubenstein said.

He imagines a room with 10 tagged possessions. The devices sleep to conserve energy, but turn on periodically to see what devices are nearby. Over time, the EnHANTs identify the other devices in the room.

A more powerful device, such as a home wireless network or smartphone, would query the EnHANTs and ask them what they see. Over a period of time, the network would build a map of the room's contents and any sensor data the EnHANTs had to communicate.

"If you start to leave your house and your wallet knows it should be with your belt, coat, and keys, it could tell the network to text a reminder to your cell phone," Rubenstein said.

Meanwhile, Rubenstein's group continues to work on shrinking prototypes to postage-stamp size. "We are really scaling back the components that go into a tag to see how small we can make it," he said. If he succeeds, we may never forget our wallet, keys, or bank cards again.

B.S., Massachusetts Institute of Technology, 1992; M.A., University of California-Los Angeles, 1994; Ph.D., University of Massachusetts-Amherst, 2000

Sensing Our Connected World

HENNING G. SCHULZRINNE

Julian Clarence Levi Professor of Mathematical Methods and Computer Science and Professor of Electrical Engineering

The latest smartphones automatically plot your location and update traffic and weather. If a Facebook friend calls, they automatically find and display his or her picture. Slowly, we are weaving together the different strands of the virtual world. Henning Schulzrinne wants to make that fabric richer by making it easier to connect those services and adding sensors to the mix.

Sensors let computers measure and interact with the physical world. "Imagine you're driving home," Schulzrinne suggested. "If the temperature is above 80 degrees F, your GPS-enabled cellphone could turn on your air conditioner. It would then turn it off when the last family member leaves home."

Working behind the scenes, Internet-enabled automation could use sensor data to tailor its response to the situation. Interconnected sensors could warn when household appliances need repairs, water the lawn only when it is dry, analyze traffic so you leave home with enough time to make your dinner reservation, and even check for signs of disease.

"Today, many of these web services are available to other applications," Schulzrinne continued. A savvy developer could query a calendar program for today's appointments or a weather program for a forecast.

"We want to leverage these services into more interesting and comprehensive systems," Schulzrinne said. "We want to program anything that can be controlled through the Internet, from your lighting and heating to your e-mail and smartphone. We want to make it easier to build smart offices and homes, and to link your calendar with your phone."

To make that happen, Schulzrinne is focusing on two first steps. One is to develop simple ways to interconnect services, sensors, and applications. "Today, you have to learn Java or other programming languages, or rely on tools from Internet companies. We want to make it easy for the nontechnical to moderately technical users to link things together in interesting ways," he said.

He is also pushing for standardized interfaces that make it easy to plug sensors into the web. "There is no reason why every sensor maker should not use the same format to convey information," Schulzrinne said. "We want to develop a standardized interface, a platform that other people can create modules that use sensor and Internet data to trigger events like services. For example, a module might trigger a stock sale depending on its performance. Another might see if it is going to rain before watering the lawn."

Ultimately, it could lead to a physical world as interactive as the virtual world that ties it together.

B.S., Technical University of Darmstadt (Germany), 1984; M.S., University of Cincinnati, 1987; Ph.D., University of Massachusetts-Amherst, 1992

Improving Human Health with Low-Power Cyber Physical Systems

MINGOO SEOK

Assistant Professor of Electrical Engineering

INFORMATION

Technological innovation on the health care front means better data management—like being able to monitor multiple vital signs of a patient on an operating table—as well as enhanced patient outcomes, such as the use of implanted pacemakers that use electrical impulses to prompt a heart to beat at a normal rate. Now, science is exploring how to take technology to the next level in order to further improve human health. That step will require the design of complex, interoperable medical devices that would be able to vary their operation to suit changing body conditions, detect minute physiological changes that signal disease, and transmit such data to medical professionals, who could take remedial action before the disease is significantly developed.

The development and use of cyber physical systems that interconnect the human body and external computers (and thus medical professionals) will be dependent upon several things: minute scalability of the system and power source, long-term operability, functional robustness regardless of environmental factors, and security of the transmitted information. Creating nearly invisible implantable medical devices is challenged by conventional circuitry and system-design techniques that fail to deliver energy efficiency to satisfy a lifetime of service.

Mingoo Seok works to combine new circuitry and architectural design elements with ultra-low-voltage systems to make the possibility of millimeter scale implantable medical devices possible. He has demonstrated a very small (1 mm^3) computer that consumes pico- to nano-watts of power—consumption that is more than 1,000 times smaller than previous state-of-the-art technology.

His research interests are in low-power digital and mixed-signal design and methodology, and he has devised approaches that deliver record-setting energy efficiency in microcontrollers, embedded memories, power conversion circuits, and DSP accelerators. As part of the technical staff at the research and development centers of Texas Instruments, he focused on developing ultra-low-power security-enhancing circuit techniques.

B.S., Seoul National (South Korea), 2005; M.S., University of Michigan, 2007; Ph.D., University of Michigan, 2011

Photo: Mingoo Seok

Playing "20 Questions" with Geometry

ROCCO A. SERVEDIO

Associate Professor of Computer Science

In "20 Questions," one player thinks of an object and the others get 20 yes-no questions to guess its identity. "That's easy, but what if you let the answerer lie three times? That makes it much more difficult," Rocco Servedio said.

That is the type of problem researchers face when false signals, or noise, corrupt data. Servedio's goal is to develop robust algorithms that learn complicated rules even in the presence of noisy data. Such algorithms could learn patterns that improve sensor performance, predict earthquakes, or forecast financial markets.

One of Servedio's most powerful tools is geometry. "When you cast a learning problem in a geometric framework, you're often on the way towards solving it," he said.

Imagine, for example, a piece of paper with red plus signs and green minus signs on opposite sides of an unknown dividing line. A few pluses are mixed with the minuses and vice versa. "In this two-dimensional example, you can eyeball the data and see which points don't belong. In higher dimensions, where each point has many coordinates, this is much more difficult, though we can sometimes pull it off with tools from high-dimensional geometry," he said.

"The way people understand something is by drawing pictures. I'm usually working in high-dimensional Euclidean spaces where it's tough to draw accurate pictures," he said, "but thinking geometrically still provides useful insights."

Servedio also takes a geometric approach to studying rules used to classify information. One popular approach is the decision tree. Like "20 Questions," it uses a sequence of yes-no questions to decide how to label data points.

"If you think of this logical representation geometrically, you can sometimes see properties that would have otherwise remained hidden. These insights can lead to better learning algorithms," Servedio said.

Servedio also uses geometry to compensate for missing data. Imagine that it takes 1,000 coordinates to describe a data point completely. What kind of learning is possible if only one of those coordinates is available?

"There are ways to compensate for massive amounts of missing data," Servedio said. "It might sound impossible, but doctors do something like this all the time. They could potentially run thousands of clinical tests on a patient to fully describe his or her condition, but a good doctor can make a useful diagnosis from just one or two tests."

A.B., Harvard, 1993; M.S., Harvard, 1997; Ph.D., Harvard, 2001

INFORMATION

Designing Secure Hardware

SIMHA SETHUMADHAVAN

Assistant Professor of Computer Science

All computer software has one thing in common: it runs on computer hardware. But what if you could not trust the hardware to securely run software? That's the question posed by Simha Sethumadhavan. "If the hardware is hacked, then it can subvert all software and software security countermeasures,"said Sethumadhavan. "Since hardware is the root of trust, attacks on hardware are potentially very dangerous."

Until recently, computer scientists never suspected that someone could tamper with hardware. Yet investigators have found unusual additions in military chips. One way to prevent hardware hijackings is by passing tokens every time data moves within hardware. Sethumadhavan likens this to sending a thank you card after a gift.

"Let's say Charlie wants to contribute $100 to Alice's charity, but has to send it to Bob first," he said. "Bob takes $10 for himself and pays the rest to Alice. One way to find out if there is a problem is for Alice to write Charlie a thank you note for the $90 donation. When Charlie sees the discrepancy, he asks accountants to trace the missing money."

Sethumadhavan proposes creating similar triangle-like structures within a computer processor. "They would monitor any irregularities. We want to create a chain of monitored data and sound the alarm if any of the links break. These lightweight monitoring additions incur very little processing overhead," he said.

"We are taking a clean slate, ground-up approach to designing secure systems," he added. "As a foundational step, we have designed methods to protect processors, the core of all computing infrastructure. Once processors are secure, we can securely build out support for protecting other hardware and software."

Sethumadhavan is also working on other techniques for securing processors. "All hardware back doors have triggers and payloads. The triggers are usually time or data input values that activate the payload," he said.

"We are working on ways to silence the trigger," he said. "For example, we might be able to reset the processor's counter so it never reaches the threshold value needed to trigger an event. Or we could use lightweight encryption to obscure data values." Only when we fully trust our processor can we fully trust other security procedures, Sethumadhavan concluded.

Sethumadhavan is leading a project on rethinking security, making it a priority instead of an afterthought, with three other Columbia Engineering professors and a team from Princeton University. The project, titled "SPARCHS: Symbiotic, Polymorphic, Autotomic, Resilient, Clean-slate, Host Security," is funded by a federal grant for more than $6 million.

B.S.E., University of Madras, 2000; M.S., University of Texas, 2005; Ph.D., University of Texas, 2007

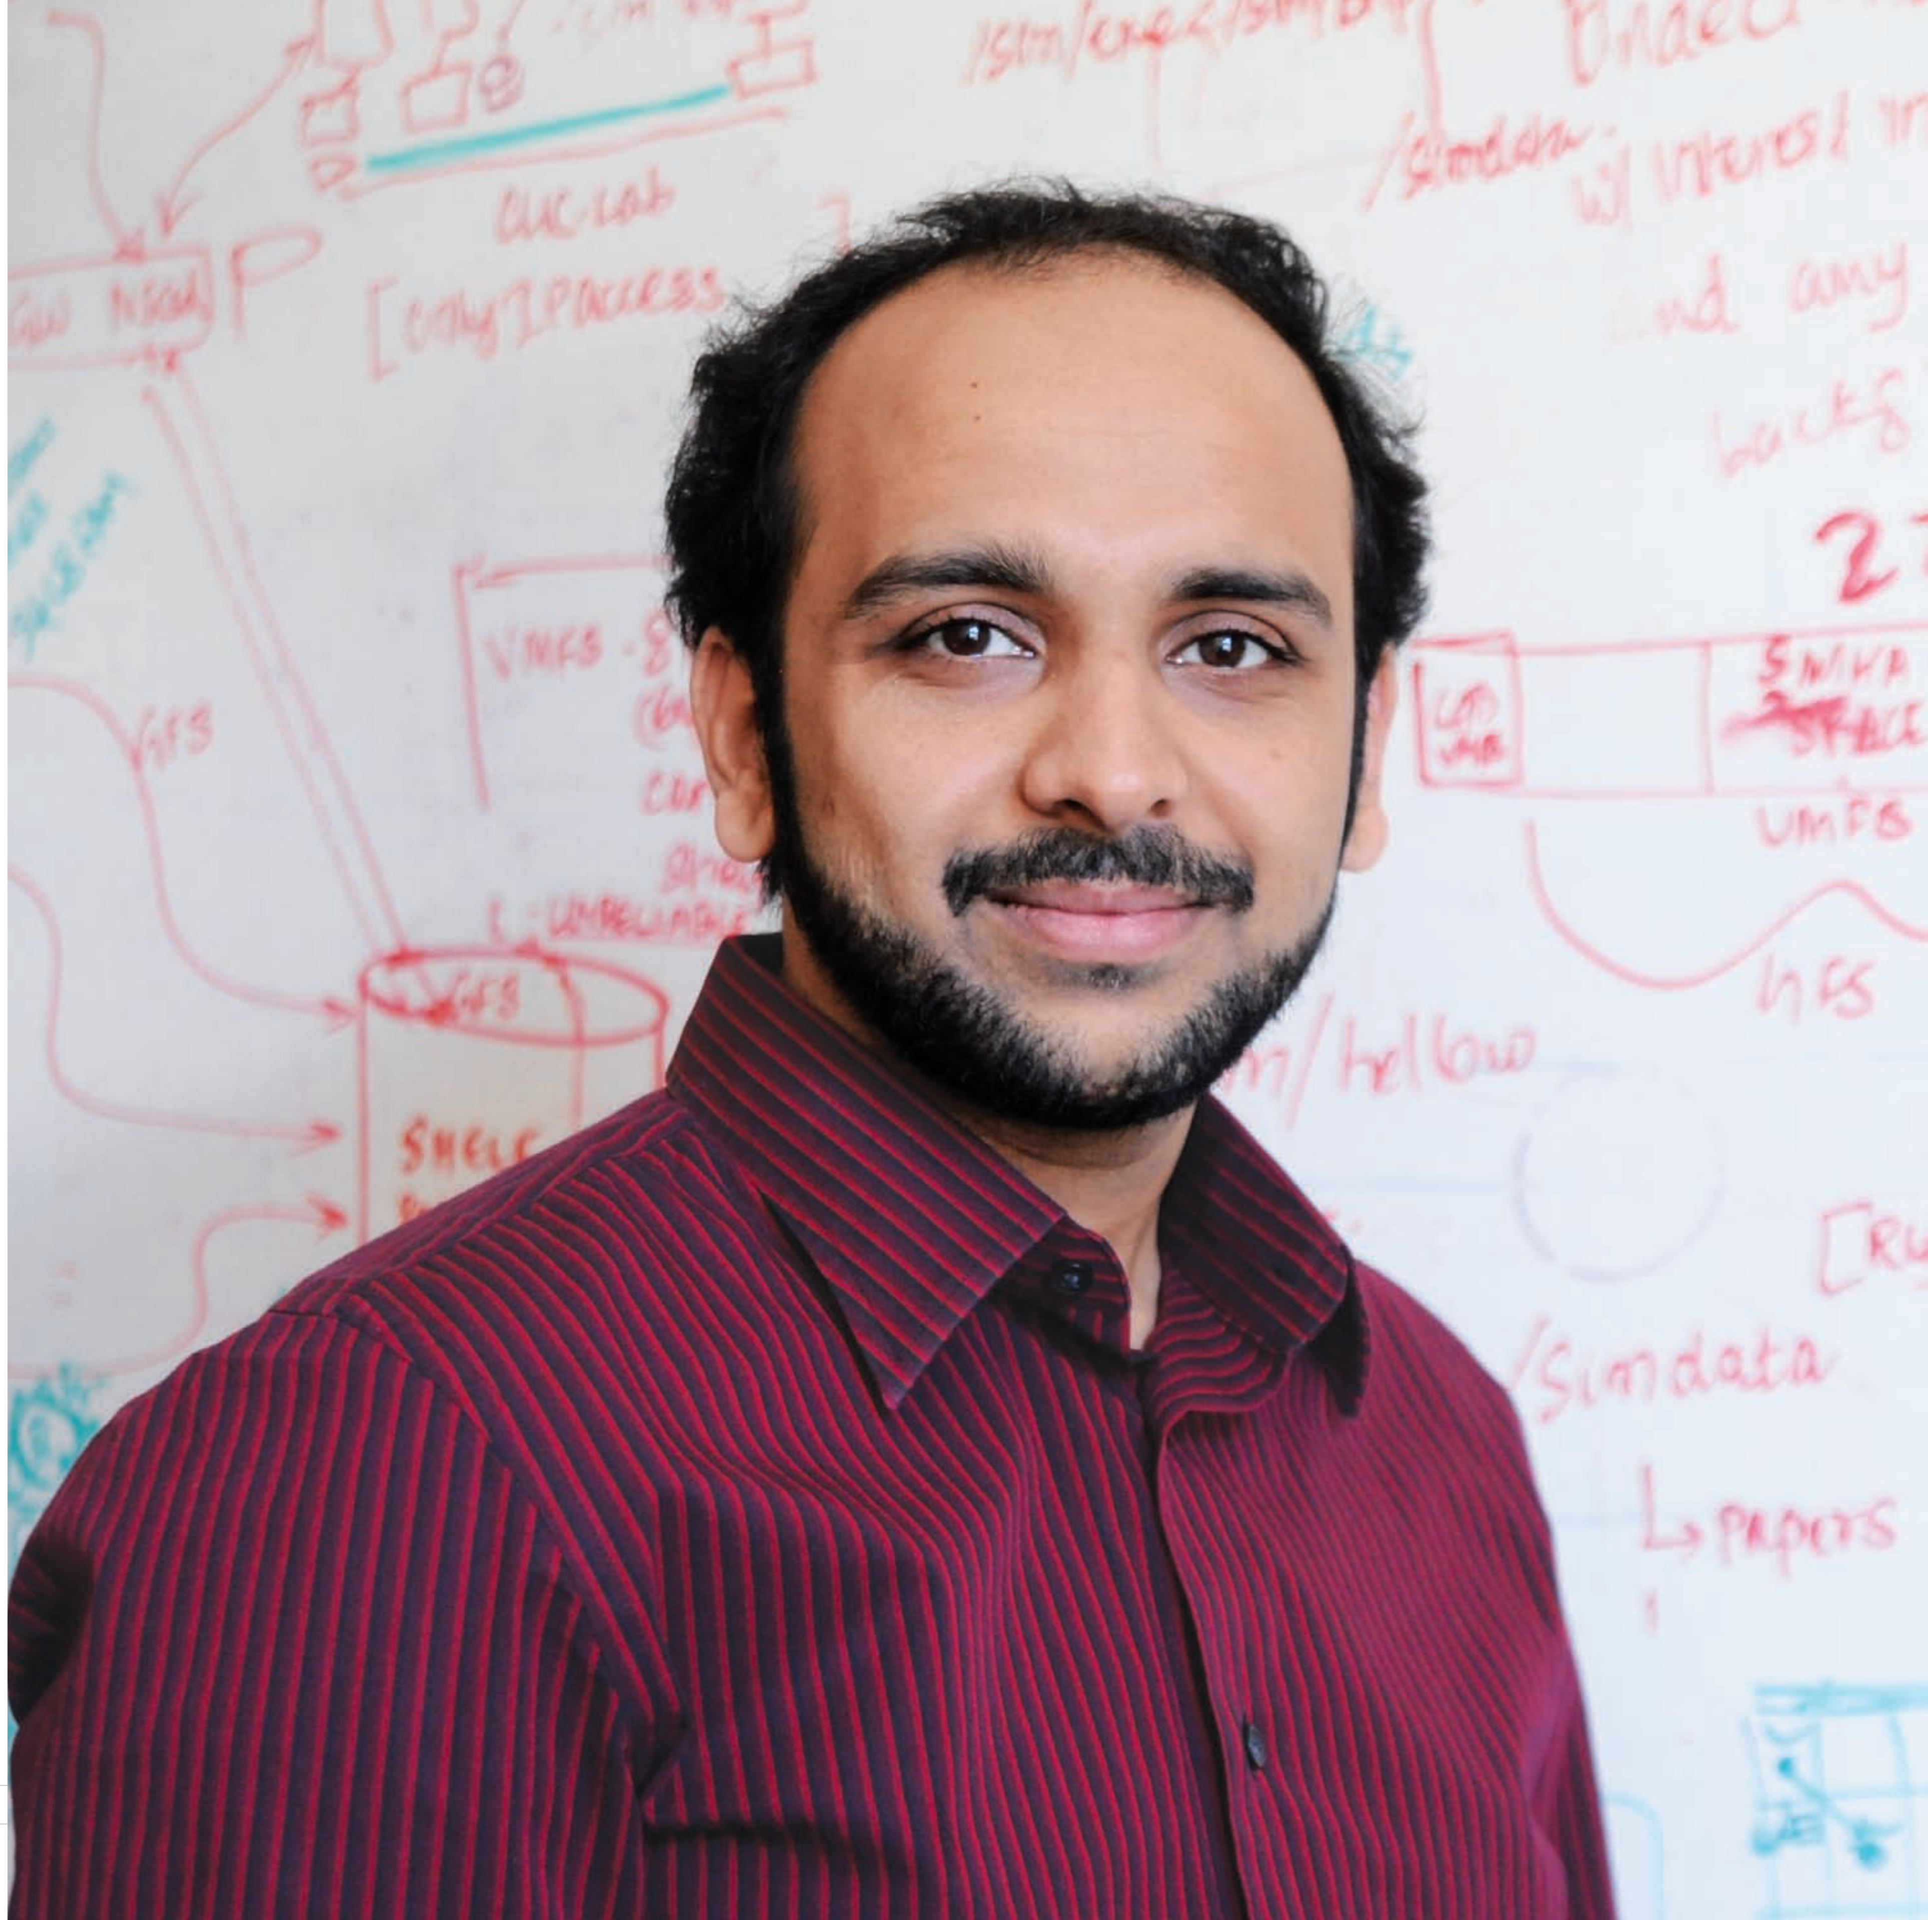

Finding a Fairer Way to Admit Students

JAY SETHURAMAN

Associate Professor of Industrial Engineering and Operations Research

Jay Sethuraman began his career by matching sets of jobs with machines to improve factory performance. More recently though, he has used operations research to find the fairest way to admit students to top public high schools.

To Sethuraman, the two problems are similar, with one major exception: "In a factory, the machines don't care what job they do. But schools do care about which students they admit," he said.

New York City's selective schools choose students based on admissions test scores. Students who do not get into their top choice can appeal. In fact, schools may set aside a certain number of seats for appeals, in addition to those seats lost when students they admitted leave for another school.

Resolving appeals fairly and efficiently discourages students from gaming the system, Sethuraman said. They may be willing to go to several schools, but list only one school if that increases their chance of placement. Or they may list schools that are unlikely to admit them if it improves their odds of getting into the school they want.

"A better system would give students an incentive to list their true preferences without penalizing them for doing it," Sethuraman said. "We want to maximize the number of students who get into their top choice, but treat all students in a fair, systematic way."

Under Sethuraman's approach, each student starts with a seat in a school that he or she wants to trade. Rather than trade individual seats, students exchange their seat for a fraction of a seat in the schools they want to attend. Those fractions, which add up to a full seat, are computed based on seat availability plus the desirability of the student's existing seat.

"At the end of this procedure, a student may have one-half a seat in school A, one-third a seat in school B, and one-sixth a seat in school C. This determines their probability of getting a seat in the lottery," Sethuraman said. Students who do not complete a trade move onto the next round of lotteries, where their odds are reset to account for the remaining available seats.

"Listing all the schools you are willing to attend increases your chances of staying in the game longer and getting into a school you really want," said Sethurman. "We give students an incentive to list all their acceptable schools without trying to game the system."

B.E., Birla Institute of Technology and Science (India), 1991; M.S., Indian Institute of Science, 1994; Ph.D., Massachusetts Institute of Technology, 1999

Predicting the Probability of Congestion

KARL SIGMAN

Professor of Industrial Engineering and Operations Research

When people refer to the World Wide Web as an information superhighway, they rarely consider traffic jams. Yet congestion slows the movement of information around the web, and appears naturally in systems as diverse as highways and hospitals.

Karl Sigman uses probability tools to build and analyze mathematical models of congestion, also known as queueing. A simple example is an ATM machine, where people arrive randomly and sometimes find themselves waiting in line to use the machine.

Successfully analyzing queueing models can help optimally route requests to a set of web servers, staff a call center, process jobs in a manufacturing plant, and schedule surgeries in a hospital.

The mathematics of probability gives Sigman many insights into a model's evolution. Still, many models remain breathtakingly complex due to the inherent randomness involved in the real world.

"Randomness, such as when the next request arrives or when something breaks, affects all these systems," said Sigman. "The further you look into the future, the more random it can become. It's like stock prices. Tomorrow's price is likely to be similar to today's, but the price next week is less certain.

"I'm interested in the relationship between what system users see and what the system actually does," he explained. "A user might click a link on a website. How long he or she waits to see the page is a measure of congestion from the user's perspective."

A system observer's viewpoint is different. "He or she looks at the web server over time and asks, 'How many users are trying to access a given page?' It does not look at the experience of any given user," Sigman said.

"This is also a measure of congestion and system performance, but from different perspectives," Sigman added. Yet the two views are interrelated. In fact, the solution to a problem from one perspective can sometimes be transformed into the desired solution from the other perspective.

Sigman has spent years teasing out those connections. "Sometimes the model looks very complicated from the perspective of a user, but it proves easier to solve from the perspective of an observer," he said.

Sigman joined Columbia Engineering in 1987. He was the recipient of the Distinguished Faculty Teaching Award both in 1998 and in 2002. He teaches courses in stochastic models, financial engineering, and queueing theory. Before joining Columbia, Sigman was a postdoctoral associate at the Mathematical Sciences Institute at Cornell University.

B.A., University of California-Santa Cruz, 1980; M.A., University of California-Berkeley, 1983; M.S., UC Berkeley, 1984; Ph.D., UC Berkeley, 1986

Monitoring Structural Health with Sensor Data Fusion

ANDREW W. SMYTH

Professor of Civil Engineering and Engineering Mechanics

Aging infrastructure is a major problem around the world and monitoring the health of structures, from bridges to dams to buildings, is critical to our modern society. Andrew Smyth specializes in structural health monitoring, using the dynamic signature of a structure to determine its condition. This can include assessing a structure's day-to-day performance, locating and quantifying potential areas of damage, or calibrating a model that can be stressed in a computer simulation for a heretofore-unseen loading event.

One of Smyth's recent projects has focused on monitoring vibrations on New York City's Manhattan Bridge. To assess the bridge's performance subsequent to a major retrofitting and strengthening program, and to calibrate a mathematical model of the bridge to predict its performance in the event of a potential seismic event, Smyth and his team placed a variety of different sensors that detected dynamic motions on the bridge over a two-month period. With the recorded data and their newly developed data fusion algorithms—a new technique that combines data from multiple sources— the team was able to identify the dynamic characteristics of Manhattan Bridge.

Smyth has also pioneered the use of differential GPS technology in conjunction with the data fusion technique to obtain highly accurate measures of low-frequency bridge deformations. He continues to develop data fusion algorithms for other civil and mechanical systems that combine information from a network of different kinds of sensors used to measure the dynamic response of a system. He says that, by taking advantage of the various levels of data redundancy, one can get high-fidelity virtual-sensing information that plays to the respective strengths of different types of sensors.

"Basically our work allows us to better understand the condition and performance of the built environment," said Smyth. "This really is our society's most valuable physical asset and the backbone of our way of life. Structural health monitoring allows us to better allocate our resources to maintain and improve our infrastructure, and keep us safe."

In 2008, Smyth was awarded the prestigious Walter L. Huber Civil Engineering Research Prize of the American Society of Civil Engineers. The award recognizes notable achievements by younger faculty members in research related to civil engineering. Smyth was recognized "for fundamental contributions in the highly efficient identification and modeling of nonlinear deteriorating structural dynamics." The selection committee commented that his research is characterized by "thoroughness, novelty, relevance, and intelligent breakthroughs."

B.A./B.Sc., Brown, 1992; M.S., Rice, 1994; Ph.D., University of Southern California, 1998

Estimating Solutions to Difficult Problems

CLIFFORD S. STEIN

Professor of Industrial Engineering and Operations Research and of Computer Science

Cliff Stein has built a career on finding algorithms to solve difficult problems—but not precisely. Stein specializes in algorithms that estimate the answer to problems that are difficult to solve. In operations research and computer science, these are problems that grow exponentially more complex as the number of inputs grows.

This contrasts with simple problems, like alphabetizing words. Double the number of inputs—words—and it takes only about twice as long to accomplish the task.

A well-known difficult problem is calculating the most efficient route for a salesman to visit different cities. To find the most efficient route, a computer must calculate all possible outcomes. For five cities, there are 120 potential paths. For 10 cities, 3.6 million. “For 80 cities, there are roughly as many possible answers as there are atoms in the universe,” Stein said.

“No conceivable advance in computing power would enable us to solve that problem precisely,” Stein noted. “But if you’re willing to solve it approximately, you can do so more easily and efficiently.”

Many algorithms already exist for estimating the solution for the traveling salesman and other difficult problems. Stein prefers to break new ground, studying the fundamental structure of problems to develop new algorithms.

“There’s a collection of algorithmic tools that are commonly used to solve many problems,” he said. “But often there are problems that are important to solve. It is worth investing the time to study their mathematical or combinatorial structure to come up with a solution specific to that problem.”

Much of his work deals with scheduling everything from computer systems to factories. Scheduling starts with jobs and the machines needed to complete them. Constraints—jobs take different amounts of time, some are more important than others, some tasks depend on others—add to the difficulty. So do different objectives, like fast completion, minimal resources, and rapid response.

Stein is looking at ways to apply scheduling to computer processors in order to save energy. “Most chips can run at four or five different speeds,” he said. “If you run at half speed, you decrease energy use by roughly a factor of four. But no one has figured out how to give chips the intelligence to know when to slow down, so they typically run at top speed all the time.”

By estimating a chip’s workload, constraints, and performance goals, Stein believes he can achieve significant energy savings. Even if his estimates are not precise.

B.S.E., Princeton, 1987; M.S., Massachusetts Institute of Technology, 1989; Ph.D., MIT, 1992

Using Anomalies to Defend Against Insiders

SALVATORE J. STOLFO

Professor of Computer Science

INFORMATION

If your credit card company ever called to confirm a purchase, you have entered Salvatore Stolfo's world. Stolfo specializes in detecting anomalies, events that stray too far from expected patterns. In addition to fraud, anomaly detection can be used to monitor engineered systems, sensor networks, ecosystems, and computer security.

Stolfo entered the field after inventing an algorithm that let marketers merge lists of consumers and purge bad records. "I realized I was aiding and abetting people who pierced personal privacy. It was an ethical dilemma," he recalled.

His interest in privacy led to cybersecurity and eventually to the study of insider attacks. "Most security breaches are the fault of the humans. Someone didn't implement something, or stole an identity, or had a grudge against an organization," Stolfo said.

This differs from most security research, which aims to keep out hackers. University researchers are more ambitious, developing inherently secure programming languages and self-repairing systems. "These are important aspects of security, but they don't matter if your adversary is already inside," Stolfo said.

"There are many different types of insiders, and they all do things in different ways," he added. "We think of it as a chess game. What if insiders can control system access? If they can blind the system to their actions, they can get away with anything. We want to stop them."

The most common type of insider threats is unintentional users. They may disable security measures to do their job more easily, or inadvertently push two buttons and erase a day's work. "These are the most prevalent and least dangerous insiders," Stolfo said.

Masqueraders include credit card thieves with stolen credentials. "The credentials make them insiders," Stolfo said. He works with banks to model consumer transactions. "We're always looking for ways to use more data to find problems sooner," he said.

Maliciously intentful insiders use their own credentials to copy secret government or corporate documents, steal money, and even sabotage the system. Highly privileged insiders have a similar agenda, but they are the ones responsible for detecting other intruders.

To foil these intruders, Stolfo looks at how their behaviors vary from company norms. By plotting how users interact with software and documents, he hopes to find patterns that suggest malicious intent. He has also developed decoys to ensnare bad guys.

"Ultimately, we want to define metrics for what it means to be secure," he said. "Then we can start to build a science of security."

B.S., Brooklyn College, 1974; M.S., New York University, 1976; Ph.D., NYU, 1979

Computing Quantum Potential

JOSEPH F. TRAUB

Edwin Howard Armstrong Professor of Computer Science

Joseph Traub is best known as a pioneer in the computational complexity of continuous problems. This involves understanding the least amount of resources—time, memory, communications—needed to solve a computing problem.

"My strategy is to start a new area of research or get into something fairly early," he said. "Then I can just walk along and pick up diamonds of knowledge and insight. I never have to strip mine for them."

It is probably as good a background as any for his investigations into the potential of quantum computing. It is a quest at the intersection of physics, mathematics, and computer science.

Quantum computing stands conventional computing on its head. For example, bits are the basic unit of information in today's computers. They can have one of two values, either zero or one, which microprocessor transistors represent as on or off.

Quantum computers are built around qubits, which have a property called superposition. This means they can be in many quantum states between zero and one, all at the same time. The more qubits a processor has, the more potential states it allows.

Qubits also have a property called entanglement. For reasons not yet understood, changing the quantum state in one of two entangled particles instantaneously changes it in the other. "That enables qubits to work together without wires," Traub said.

Because quantum computers are not limited to on-off states, they can calculate many possible answers at once. This could make it possible to calculate very complex problems rapidly.

"What I'm trying to do is ask, 'Where are the big wins?' In particular, what kind of problems could a quantum computer solve that physicists and chemists are really interested in solving," Traub said.

One of those problems is calculating the lowest energy state, or ground state energy, of a large number of particles. "This is a central problem in computational chemistry, and it would allow us to predict chemical reactions better," Traub explained.

The problem, he explained, is that ground state energy calculations are difficult and soak up computer resources. A quantum computer's ability to make multiple calculations simultaneously could give chemists the tool they need to predict particle interactions in large systems.

"We're theoreticians, trying to understand the type of problems quantum computers might be able to solve. Physicists may never succeed in building one, but if they do, we want to be ready," Traub concluded.

B.S., College of the City of New York, 1954; M.S., Columbia, 1955; Ph.D., Columbia, 1959

Creating New Circuits for Interfacing the Computer to the Physical World

YANNIS TSIVIDIS

Charles Batchelor Professor of Electrical Engineering

INFORMATION

While the transistor revolutionized the field of electronics and paved the way for personal computers, it made way for several perpetual challenges: deliver more power in smaller sizes, enable real-time interaction with the real world, and constantly adapt to technological change. Solutions to those challenges can make possible, for example, biomedical ingestible pills, containing chips that aid, or give information about, the body; sensor networks that provide information about the environment or physical infrastructure; or wireless communication technology that uses less battery power but provides more range.

One of the challenges in making this new era a reality lies in advancing the development of single silicon chips that perform both analog and digital signal processing. Analog and digital signal domains have significant technical differences, yet new technology demands more and more complex mixed-signal design. The development pace is relentless, driven by demands for increased performance. New techniques need to be invented, and fundamental limitations must be better understood, to make such analog/digital circuits with improved performance possible.

Yannis P. Tsividis has been an important contributor to the field of silicon chips that mix analog and digital circuits. He and his students have done extensive research in this field at the device, circuit, system, and computer simulation level.

In 1976, Tsividis designed and built the first fully integrated MOS operational amplifier and demonstrated its use in a coder-decoder for digital telephony. These results were widely adopted by the industry in the first massively produced mixed-signal MOS integrated circuits, which incorporate both analog and digital functions on the same silicon chip.

Tsividis and his students have since been responsible for several important contributions, ranging from precision device modeling and novel circuit building blocks to new techniques for analog and mixed-signal processing, self-correcting chips, switched-capacitor network theory, RF integrated circuits, mixed analog-digital Very Large Scale Integrated (VLSI) computation and the creation of computer simulation programs. This work has resulted in several patents in several countries, and has been incorporated by the industry into products we use every day.

Tsividis is a fellow of the Institute of Electrical and Electronics Engineers.

B.E., University of Minnesota, 1972; M.S., University of California-Berkeley, 1973; Ph.D., UC Berkeley, 1976

Unlocking a Complex World Mathematically

VLADIMIR VAPNIK

Professor of Computer Science

"When the solution is simple, God is answering," Albert Einstein once commented. He believed we could discover nature's laws only when they connected a few variables, like the relationship between temperature and pressure or energy and mass. "When the number of factors coming into play is too large, scientific methods in most cases fail," Einstein said. Of course, Einstein did not have computers. Vladimir Vapnik does.

Vapnik works in machine learning, a discipline that uses algorithms to detect automatically those laws of nature that depend on hundreds or even thousands of parameters. This enables computers to make better predictions, and also provides insights into the elusive nature of human learning.

Today's machine learning technology requires many examples to generate accurate rules. Yet humans clearly learn to understand their complex world from far fewer examples. This led Vapnik to consider how teachers provide students with what he calls "privileged information," holistic knowledge often delivered as metaphors and comparisons.

Master classes for musicians are an example. "The teachers cannot show students how to play an instrument because their technique is not as good," he said. "Instead, teachers may use metaphors or comparisons to show students how to understand a piece. This may sound like nonsense in terms of musical technique, but it helps them play better."

Vapnik has shown mathematically that privileged information could slash the samples needed for machine learning by the square root of the original number. "Instead of 10,000 examples, we would need only 100," he said.

He demonstrated this using privileged information to help a computer identify handwritten numbers. He asked Professor of Russian Poetry Natalia Pavlovitch to write a short verse describing her feelings about each number sample. The information was subjective and not available by analyzing only the numbers. Including it during training yielded more accurate results than training with the numbers alone.

Vapnik also used surgeons' descriptions of biopsy pictures—from "quiet" to "wide aggressive proliferation"—to improve the classification of tumors. The notes were impressionistic, but improved the computer's ability to identify cancerous cells.

Humans frequently use such holistic privileged information to make sense of complex phenomena. Providing it to machines could open a new door onto a complex universe.

"For 2,000 years, we believed logic was the only instrument for solving intellectual problems. Now, our analysis of machine learning is showing us that to address truly complex problems, we need images, poetry, and metaphors as well," Vapnik concluded.

M.S., Uzbek State University, Samarkand, 1958; Ph.D., Institute of Control Sciences, Moscow, 1964

Devising a Design Framework for Next-Generation Wireless Technology

XIAODONG WANG

Professor of Electrical Engineering

INFORMATION

Advancing wireless communication technology to a new generation of application and service is one of today's prime research disciplines. Demands for higher capacity drive the need to create novel signal transmission techniques and advanced receiver signal processing methods. Challenging design requirements are compounded by the complexity of the nature of the transmitter and receiver: a complicated system consisting of radio frequency, analog and mixed-signal components. Plus, heated competition in the development arena forces tight time-to-market deliverables.

To develop effective next-generation wireless technology under the constraint of thousands of variables, it is important to use mathematical modeling and analysis, computer simulations, fast calculations, and data summaries to thoroughly account for manufacturing process variations before build-out. Using these tools to analyze production provides a comprehensive transistor-level statistical design and verification framework. With it, designers can troubleshoot and devise design enhancements to solve the issues of fading, impulsive noise, and co-channel interference in the concept phase.

Xiaodong Wang is a leading researcher in signal processing, computing, and communications. His broader research interests include information theory, algebraic coding theory, wireless communications, optical communications, communication networks, statistical signal processing, and genomic signal processing. Results of his research have included extensive publication in these areas, most recently in the areas of chip-level asynchronism on a Code Division Multiple Access (CDMA)-based overlay system for optical network management; modulation classification via Kolmogorov-Smirnov test; Generalized Likelihood Ratio Test (GLRT)-based spectrum sensing for cognitive radio with prior information; and blind frequency-dependent I/Q imbalance compensation for direct-conversion receivers.

Wang also has become active in the emerging field of genomic signal processing (GSP). The aim of GSP is to integrate the theory and methods of signal processing with the global understanding of functional genomics, with special emphasis on genomic regulation. He took part in a National Science Foundation-funded multidisciplinary collaborative project to develop a structural health monitoring (SHM) system using a wireless piezoelectric sensor network.

Wang is a fellow of the Institute of Electrical and Electronics Engineers (IEEE). He received the 1999 NSF CAREER Award and the 2001 IEEE Communications Society and Information Theory Society Joint Paper Award. He has served as an associate editor for the *IEEE Transactions on Signal Processing*, the *IEEE Transactions on Communications*, the *IEEE Transactions on Wireless Communications*, and *IEEE Transactions on Information Theory*. He is listed as an ISI-Highly-Cited researcher.

B.S., Shanghai Jiao Tong University, 1992; M.S., Purdue, 1995; Ph.D., Princeton, 1998

Predicting Waves Mathematically

MICHAEL WEINSTEIN

Professor of Applied Physics and Applied Mathematics

Understanding the behavior of waves in complex environments holds the key to advances in a wide range of applications—from optical communications and computer technology to the prediction and detection of seismic, atmospheric, and oceanic phenomena. Wave phenomena are described using partial differential equations, which are a mathematical encoding of physical laws.

But significant challenges arise because phenomena are both multiscale—they derive from activity and interactions among very small spatial scales all the way up to very large scales—and nonlinear, which leads to waves that distort dramatically and "scatter" differently as their size is changed. These general features limit the solvability of problems on even the fastest computers.

Michael Weinstein develops hybrid analytical/computational approaches, which combine asymptotic mathematical analysis with computer simulation. Asymptotic analysis yields approximate, but fairly explicit and detailed information, on the very small-scale phenomena. With these degrees of freedom "solved for," the computer can then focus on the larger scales and efficiently give approximate, yet very accurate predictions.

Applying these approaches to the partial differential equations of optics, Weinstein has discerned how "soliton" light-pulses travel and interact within communications lines. This work has a wide range of practical applications: from determining the stability of optical pulses to ideas on how to robustly encode information in streams of optical pulses.

He has proposed designs of novel optical media to slow or even stop light pulses in micro-structured waveguides, and has proposed their application to optical buffering of information. A recent project exploits parallels between the equations of electromagnetics with those arising in the theory of shock waves in supersonic flight, to understand the generation of broadband, multi-colored light from laser light of a single color. Broadband light sources have applications ranging from communications to imaging science.

Other recent work he is addressing concerns metamaterials: specially engineered microstructures, which act as a macroscopic device, and achieve properties not possible using naturally occurring materials. One application studied by Weinstein is the attainability of the cloaking effect. Cloaking involves surrounding a region of space by an appropriate metamaterial. Anything in the surrounded region is undetectable by exterior sensors, and anything within the shielded region is isolated from the exterior world. Other application areas of metamaterials envisioned include improved solar energy cells, secure communications and sensors.

Weinstein is a fellow of the Society for Industrial and Applied Mathematics (SIAM), elected for his "contributions to the analysis and applications of nonlinear waves".

B.S., Union College, 1977; M.S., Courant Institute of Mathematical Sciences at New York University, 1979; Ph.D., Courant Institute, 1982

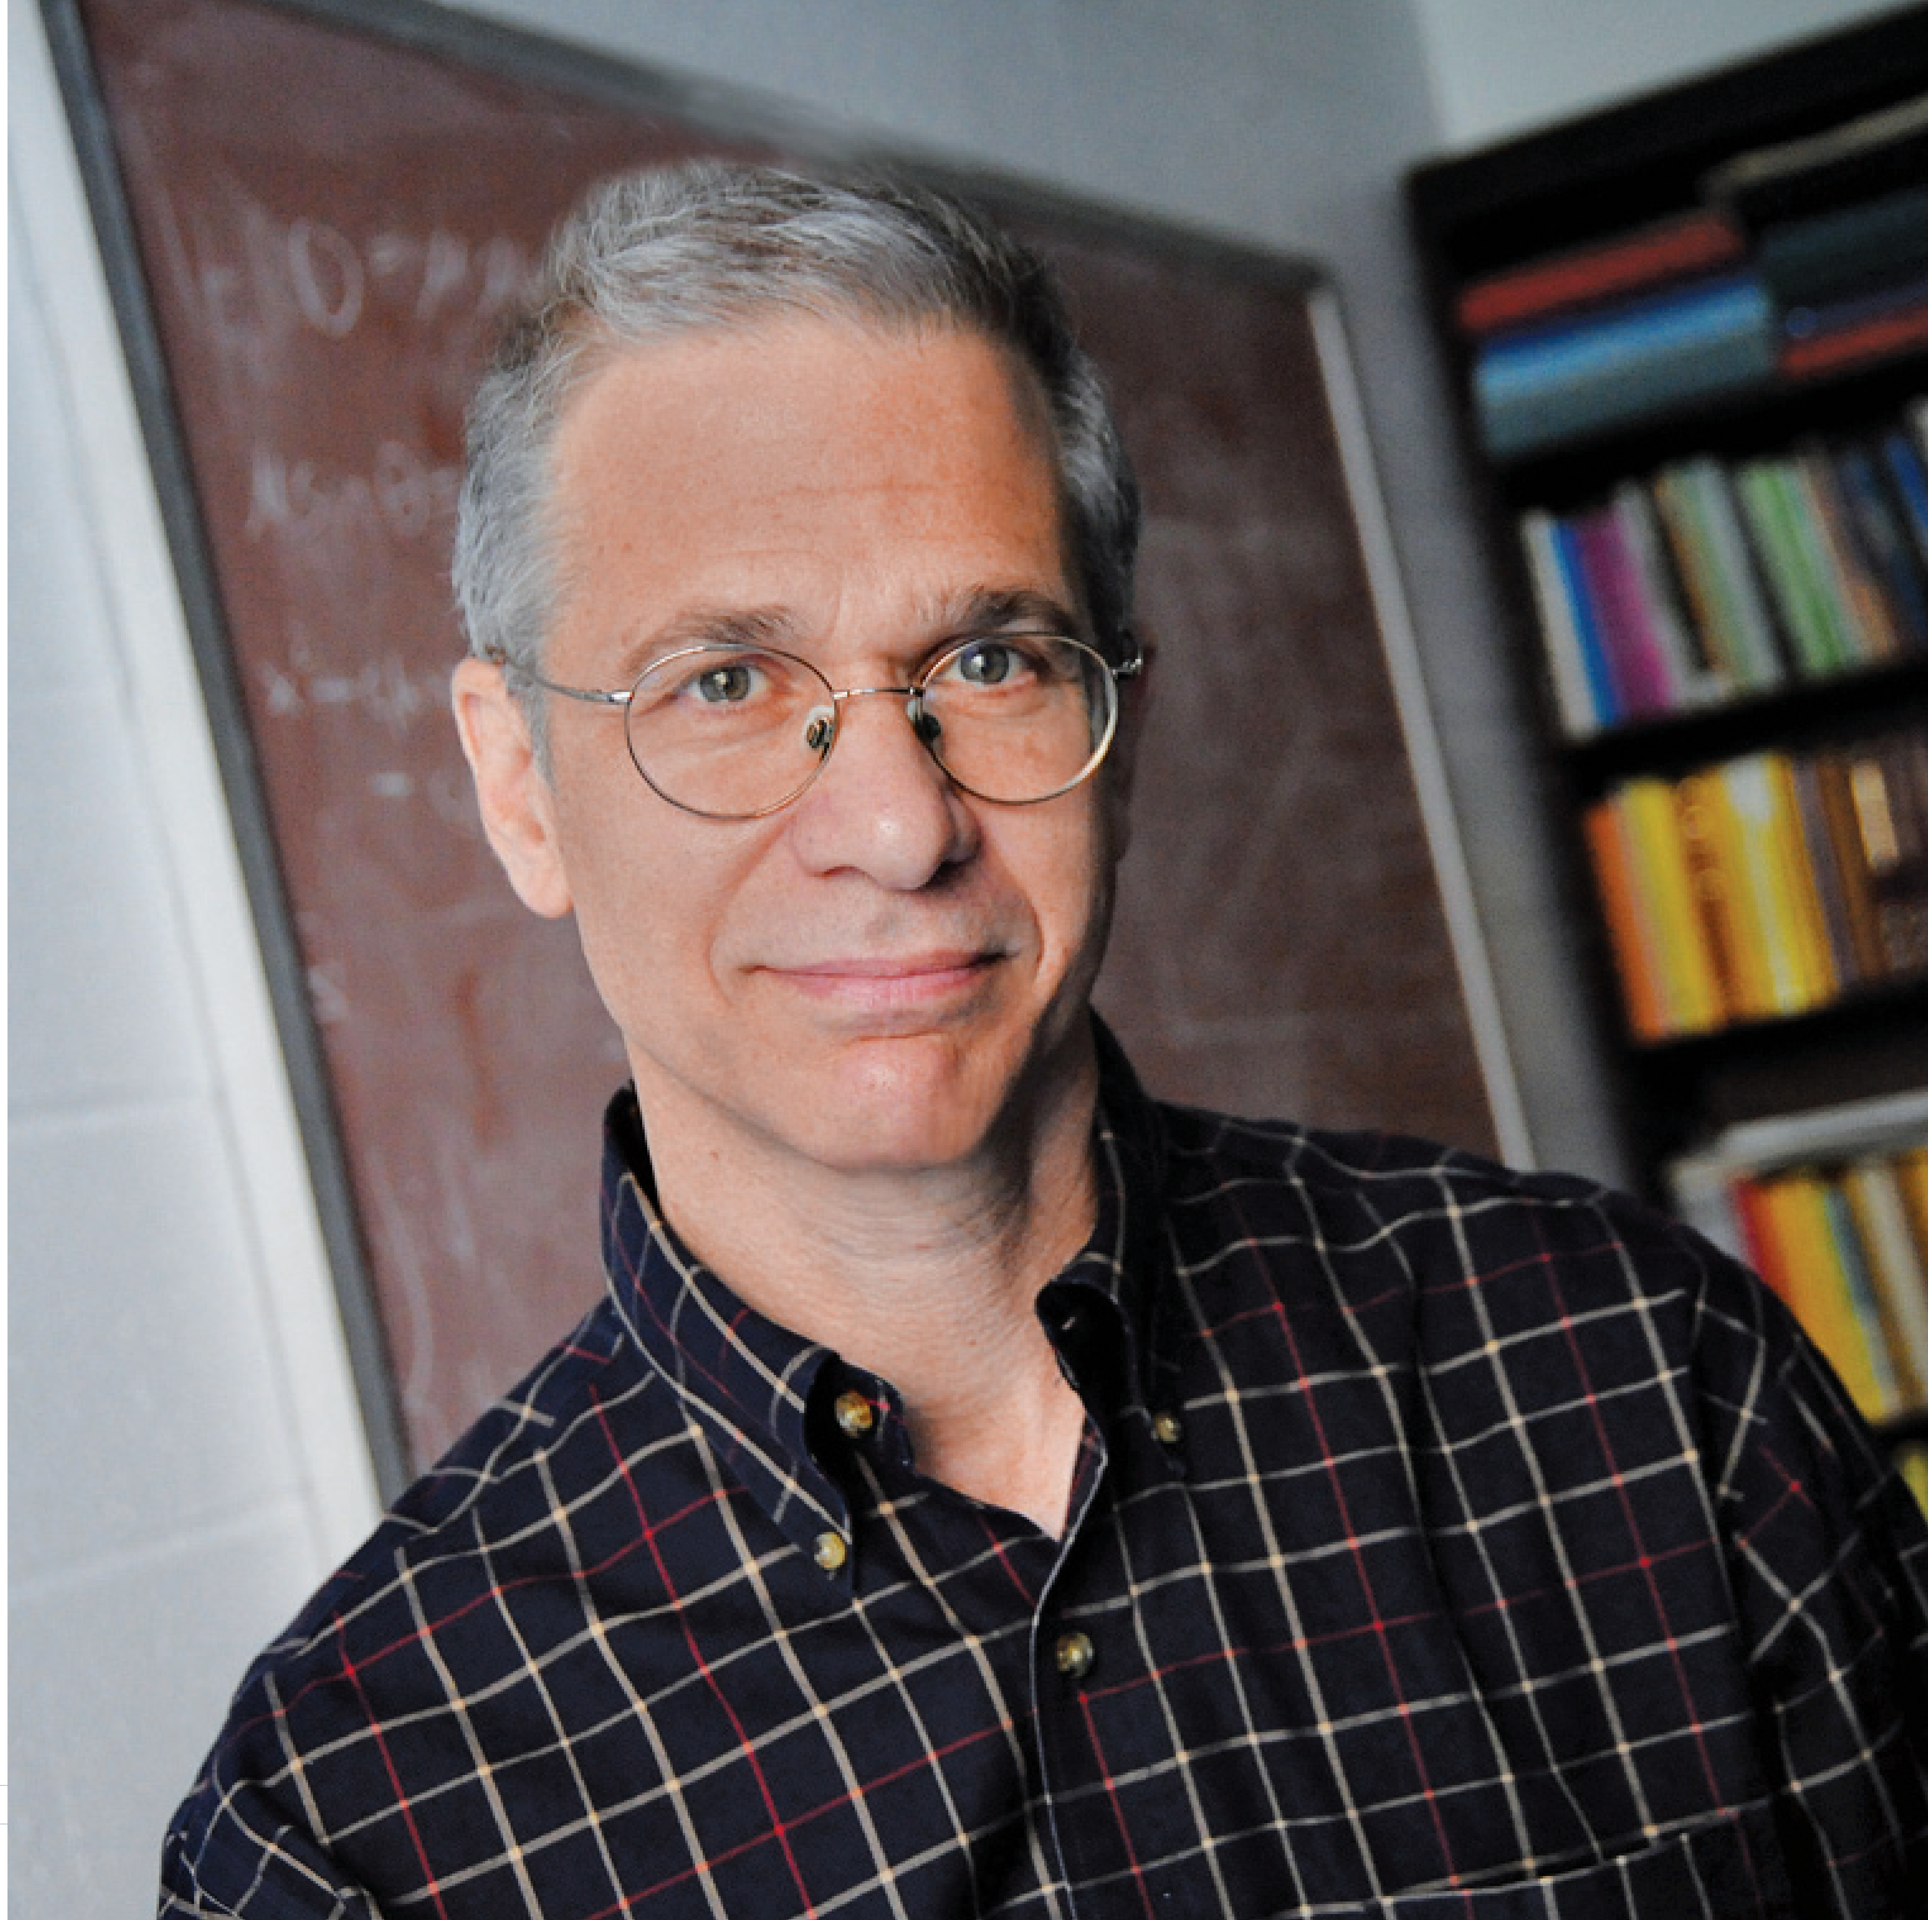

Unraveling the Mysteries of Congestion

WARD WHITT

Wai T. Chang Professor of Industrial Engineering and Operations Research

Anyone who has driven on highways understands that random events affect congestion. Even in relatively light traffic, with no accidents or obstructions, cars will suddenly bunch up, slow, and then speed up again.

Ward Whitt studies the enigma at the heart of this process. His discipline—queueing theory—examines how random fluctuations in flow, waiting, and processing cause congestion in complex systems.

Examples are everywhere. "We all spend too much time waiting on lines, from physical lines in a supermarket or bank to invisible lines on hold for a call center or waiting for a web page to load," Whitt said. Queues are equally present in the waiting times of a computer processor or the movement of parts through a factory.

One major goal of queueing theory is to reduce waiting. Understanding congestion helps engineers specify the right number of telephone switches, Internet servers, and even call center personnel.

Despite their wide use, queuing models have a significant weakness. "The standard queueing models assume random flow, but the rate of that random flow is assumed constant. In reality, the arrivals to a system occur randomly, but the rate of that random flow is not constant," Whitt said.

Whitt tries to capture that systematic variation in the flow rate together with the uncertainty about that flow rate. He builds and analyzes models that reflect both these features of everyday queueing phenomena. "This produces high fidelity descriptions of congestion that go far beyond standard textbook queueing models," he said.

Whitt is also applying these insights to complex networks. "Queues do not appear in isolation, but appear in networked systems with multiple flow paths and queues," he said.

One way to tackle complex, networked systems is to see how they would behave as they scale up. "Sometimes," Whitt said, "a larger model tells a clearer story. Toss a coin 20 times and you expect to average 10 heads and 10 tails, but you may see from seven to 13 heads. But toss the coin one million times, you are likely to get closer to a 50-50 split."

He has developed mathematical techniques that show how congested systems behave at larger scales. He then compares the model with computer simulations of the system or data from that system.

"When you do this, you can end up with a fairly simple story that tells you a lot about your system," Whitt said.

A.B., Dartmouth, 1964; Ph.D., Cornell, 1969

Bringing Order to High-Dimensional Datasets

JOHN WRIGHT

Assistant Professor of Electrical Engineering

It's a data-driven world out there. Every day, streams of data in the form of images, videos, biomedical observations, Internet links, and more are fed to scientific organizations, businesses, and governments worldwide. And while management and warehousing of these prodigious amounts of data are important, equally important is developing the technological capability to understand the structure of the datasets.

Much of the data collected today is in digital imagery, each made up of several million pixels. With millions upon millions upon millions of pixels residing in any given dataset, finding order within those datasets is critical to being able to efficiently search and find specific data. Add to that the challenge when data in any set is unreliable (e.g., "dead" pixel(s), a disguised face, shadows, or occlusions), and the classical algorithms used to search and find specific data break down.

John Wright considers the area of high-dimensional data analysis a gold mine for great mathematical and algorithmic problems, with the potential for profound impact on applications that can deal intelligently with imagery data.

His research has developed new theory and algorithms for uncovering several important types of low-dimensional structure in high-dimensional datasets, even in the presence of gross observation errors. This combination of efficient algorithm and good theoretical understanding has led to new, highly accurate algorithms for recognizing human faces, even with occlusion or disguise; for recovering the shape of three-dimensional objects from two-dimensional images; and for building three-dimensional models of urban environments.

He is currently working on new techniques for finding good representations of data—searching for a "dictionary" that can most compactly represent a given set of data samples. Recent results have shown that if it is possible to find efficient data representations, those representations can be used to acquire signals and images more accurately, and using fewer resources. Through collaborations at Microsoft Research, he is investigating the use of these techniques to efficiently acquire images for cultural heritage preservation.

Wright is a member of the Association for Computing Machinery, the Institute for Electrical and Electronic Engineers, and the Society for Industrial and Applied Mathematics.

B.S., University of Illinois at Urbana-Champaign, 2004; M.S., University of Illinois, 2007; University of Illinois, Ph.D., 2009

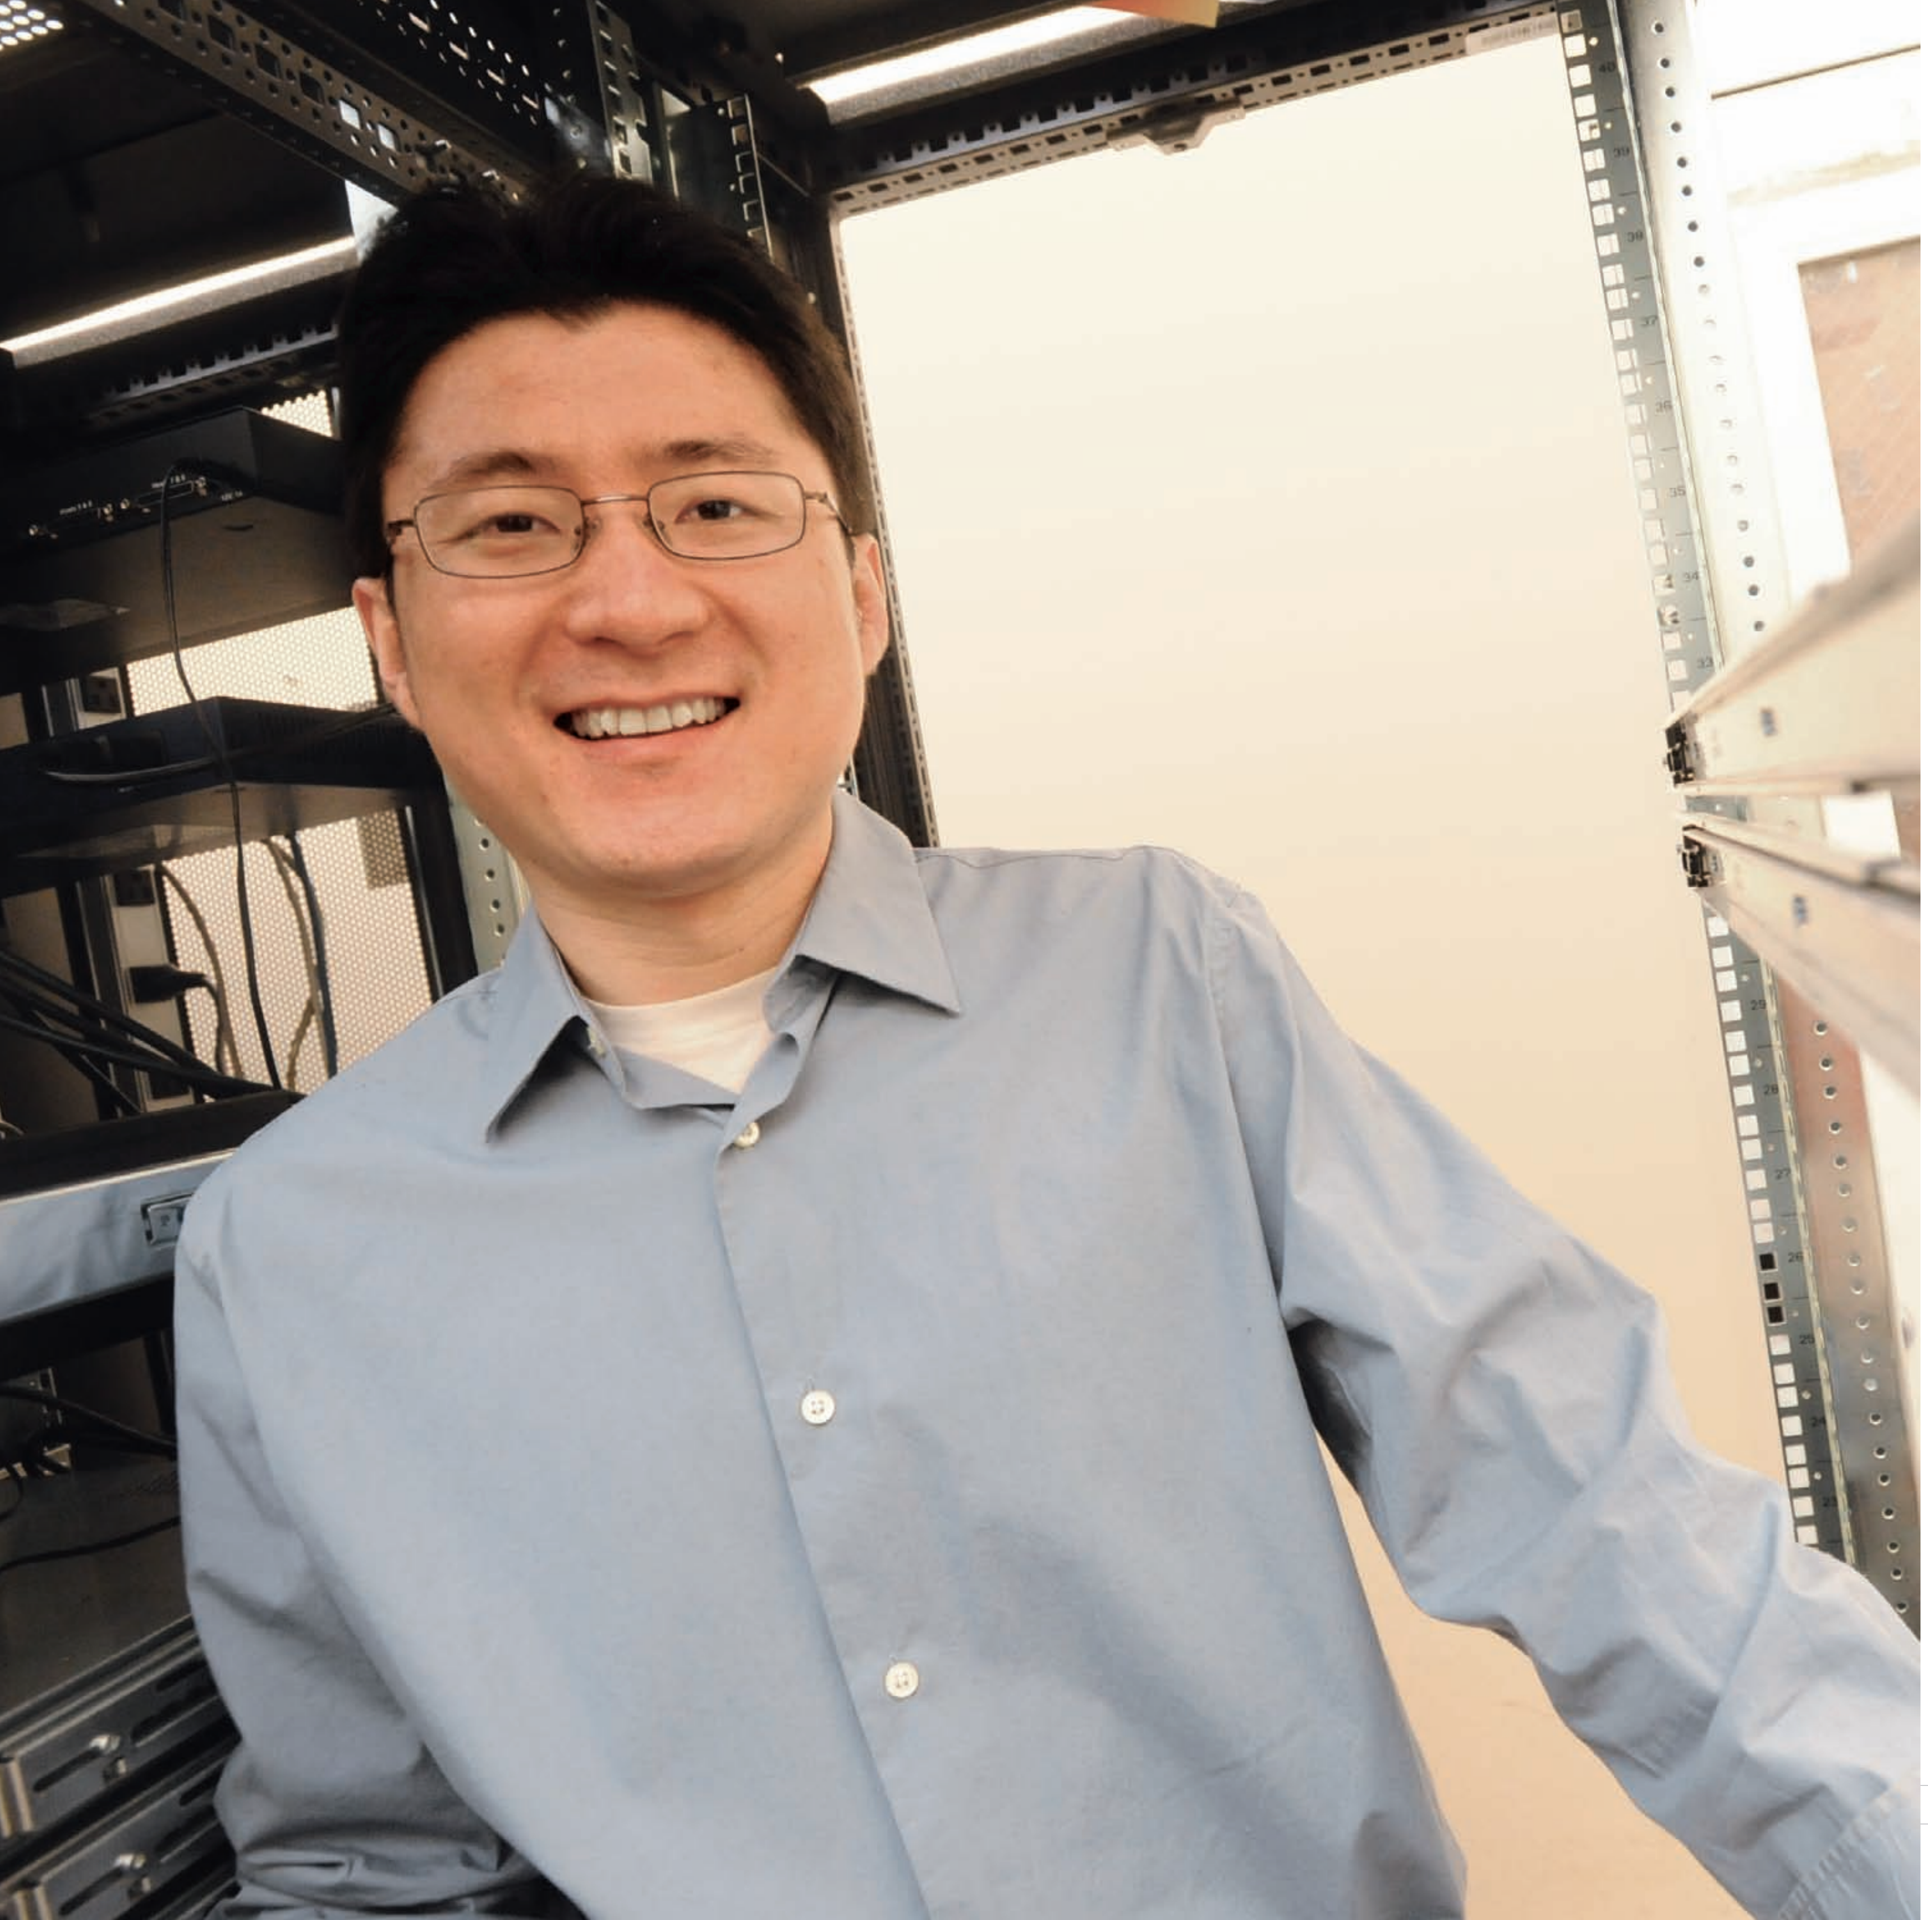

Weaving More Reliable Software

JUNFENG YANG

Assistant Professor of Computer Science

Even the best written software contains errors. Junfeng Yang wants to unmask and correct those often subtle defects. The software bugs are costly. In 2002, the National Institute of Science and Technology put their cost at $60 billion annually. Bugs do more than crash computers. They contributed to the northeast power blackout in 2003, and delivered lethal doses of radiation to hospital patients.

"My research involves finding ways to make software more reliable," Yang said. In graduate school, he developed an automated method to detect storage system errors. "Past tests were like throwing darts and hoping to hit a problem area. We developed systematic ways to test all possible storage states," he said.

After joining Microsoft, he extended his work to distributed storage systems on large networks. "People knew they were losing data, but not why. Our tool helped them find those bugs," Yang said. His work led to numerous patches for Microsoft's production systems and the Linux Operating System.

Now Yang is focusing on the reliability of multithreaded programs. Unlike programs that run all their instructions sequentially, multithreaded programs consist of segments, or threads, that run concurrently. Multithreaded programs are significantly faster than sequential code.

They are also more difficult to write, test, and debug. "This is because they are not deterministic," he explained. In other words, a multithreaded program may behave somewhat differently each time it runs. "It may act correctly or buggy, depending on such variables as processor speed, operating system scheduling, and what data arrives when during operations," Yang said.

Lack of determinism makes it difficult to reproduce errors, much less fix them. Yang's research makes multithreaded programs execute deterministically, so programmers can isolate problems.

Explaining his approach, Yang likens threads to cars driving down a four-lane highway. "The cars drive in parallel lanes. During nondeterministic execution, they can change lanes whenever they want. When they do, sometimes they collide and cause the program to crash.

"To make threads execute deterministically, we've placed barriers between the lanes. We only allow threads to change lanes at fixed locations, following a fixed order. This prevents random car collisions," he said.

Yang records this path and makes every subsequent group of cars follow it. "Because we know the path causes no collisions, there should be no collisions when another group of cars use it," he said. By attacking multithreading, Yang hopes to weave more reliable software.

B.S., Tsinghua University (Beijing), 2000; M.S., Stanford, 2002; Ph.D., Stanford, 2008

Calculating What Is Possible

MIHALIS YANNAKAKIS

Percy K. and Vida L. W. Hudson Professor of Computer Science and Professor of Industrial Engineering and Operations Research

Computers are solving ever more complex problems, yet some problems have resisted intense efforts for many decades. How can we tell which problems can be solved efficiently and which cannot? How do we find the most efficient algorithms? And for intractable problems, how do we find the best solutions possible in reasonable amounts of time? These are some of the challenges taken on by Mihalis Yannakakis.

One line of his research seeks to understand the inherent computational complexity of problems. "It turns out that many computational problems from diverse fields are intimately related to one another," he said. For example, optimizing network designs, scheduling jobs, and folding proteins all exhibit essentially the same type of computational difficulties. Yannakakis seeks to find the underlying features that characterize the complexity of different problems and identify their unifying principles.

Many optimization problems are computationally hard, in the sense that we cannot compute efficiently the optimal solution. For these cases, Yannakakis has been working on algorithms that compute near-optimal solutions. His goal is to design efficient approximation algorithms with provable performance guarantees.

Yannakakis' third research thrust involves trade-offs when making decisions. "We care about a design's quality and also its cost, or a health treatment's benefits and risks. Typically, there is no one solution that is optimal for all criteria, but rather many incomparable solutions that encapsulate the trade-offs between different criteria," he explained.

For two criteria, Yannakakis visualizes these trade-offs as a curve on the plane. As the number of criteria rise, the trade-offs form a surface in a higher dimensional mathematical space. "It is generally impossible to generate all the points on the trade-off surface because there are usually an exponential or even infinite number of them," he said. "But we wish to generate enough points to represent the whole design space, so decision makers have an accurate enough view of the trade-offs to make an informed choice."

Yannakakis' approach is to design algorithms with guaranteed succinctness and accuracy to compute a carefully selected small set of solutions that offer the best possible representation of that space.

"Computers could work forever on a task," he summarized. "We try to characterize what we can actually compute efficiently for a specific task and in general. Like trying to understand the physical world, we're trying to find the laws that govern the computational world. We want to determine the powers and limitations of computation."

Dipl., National Technical University of Athens (Greece), 1975; M.S., Princeton, 1979; Ph.D., Princeton, 1979

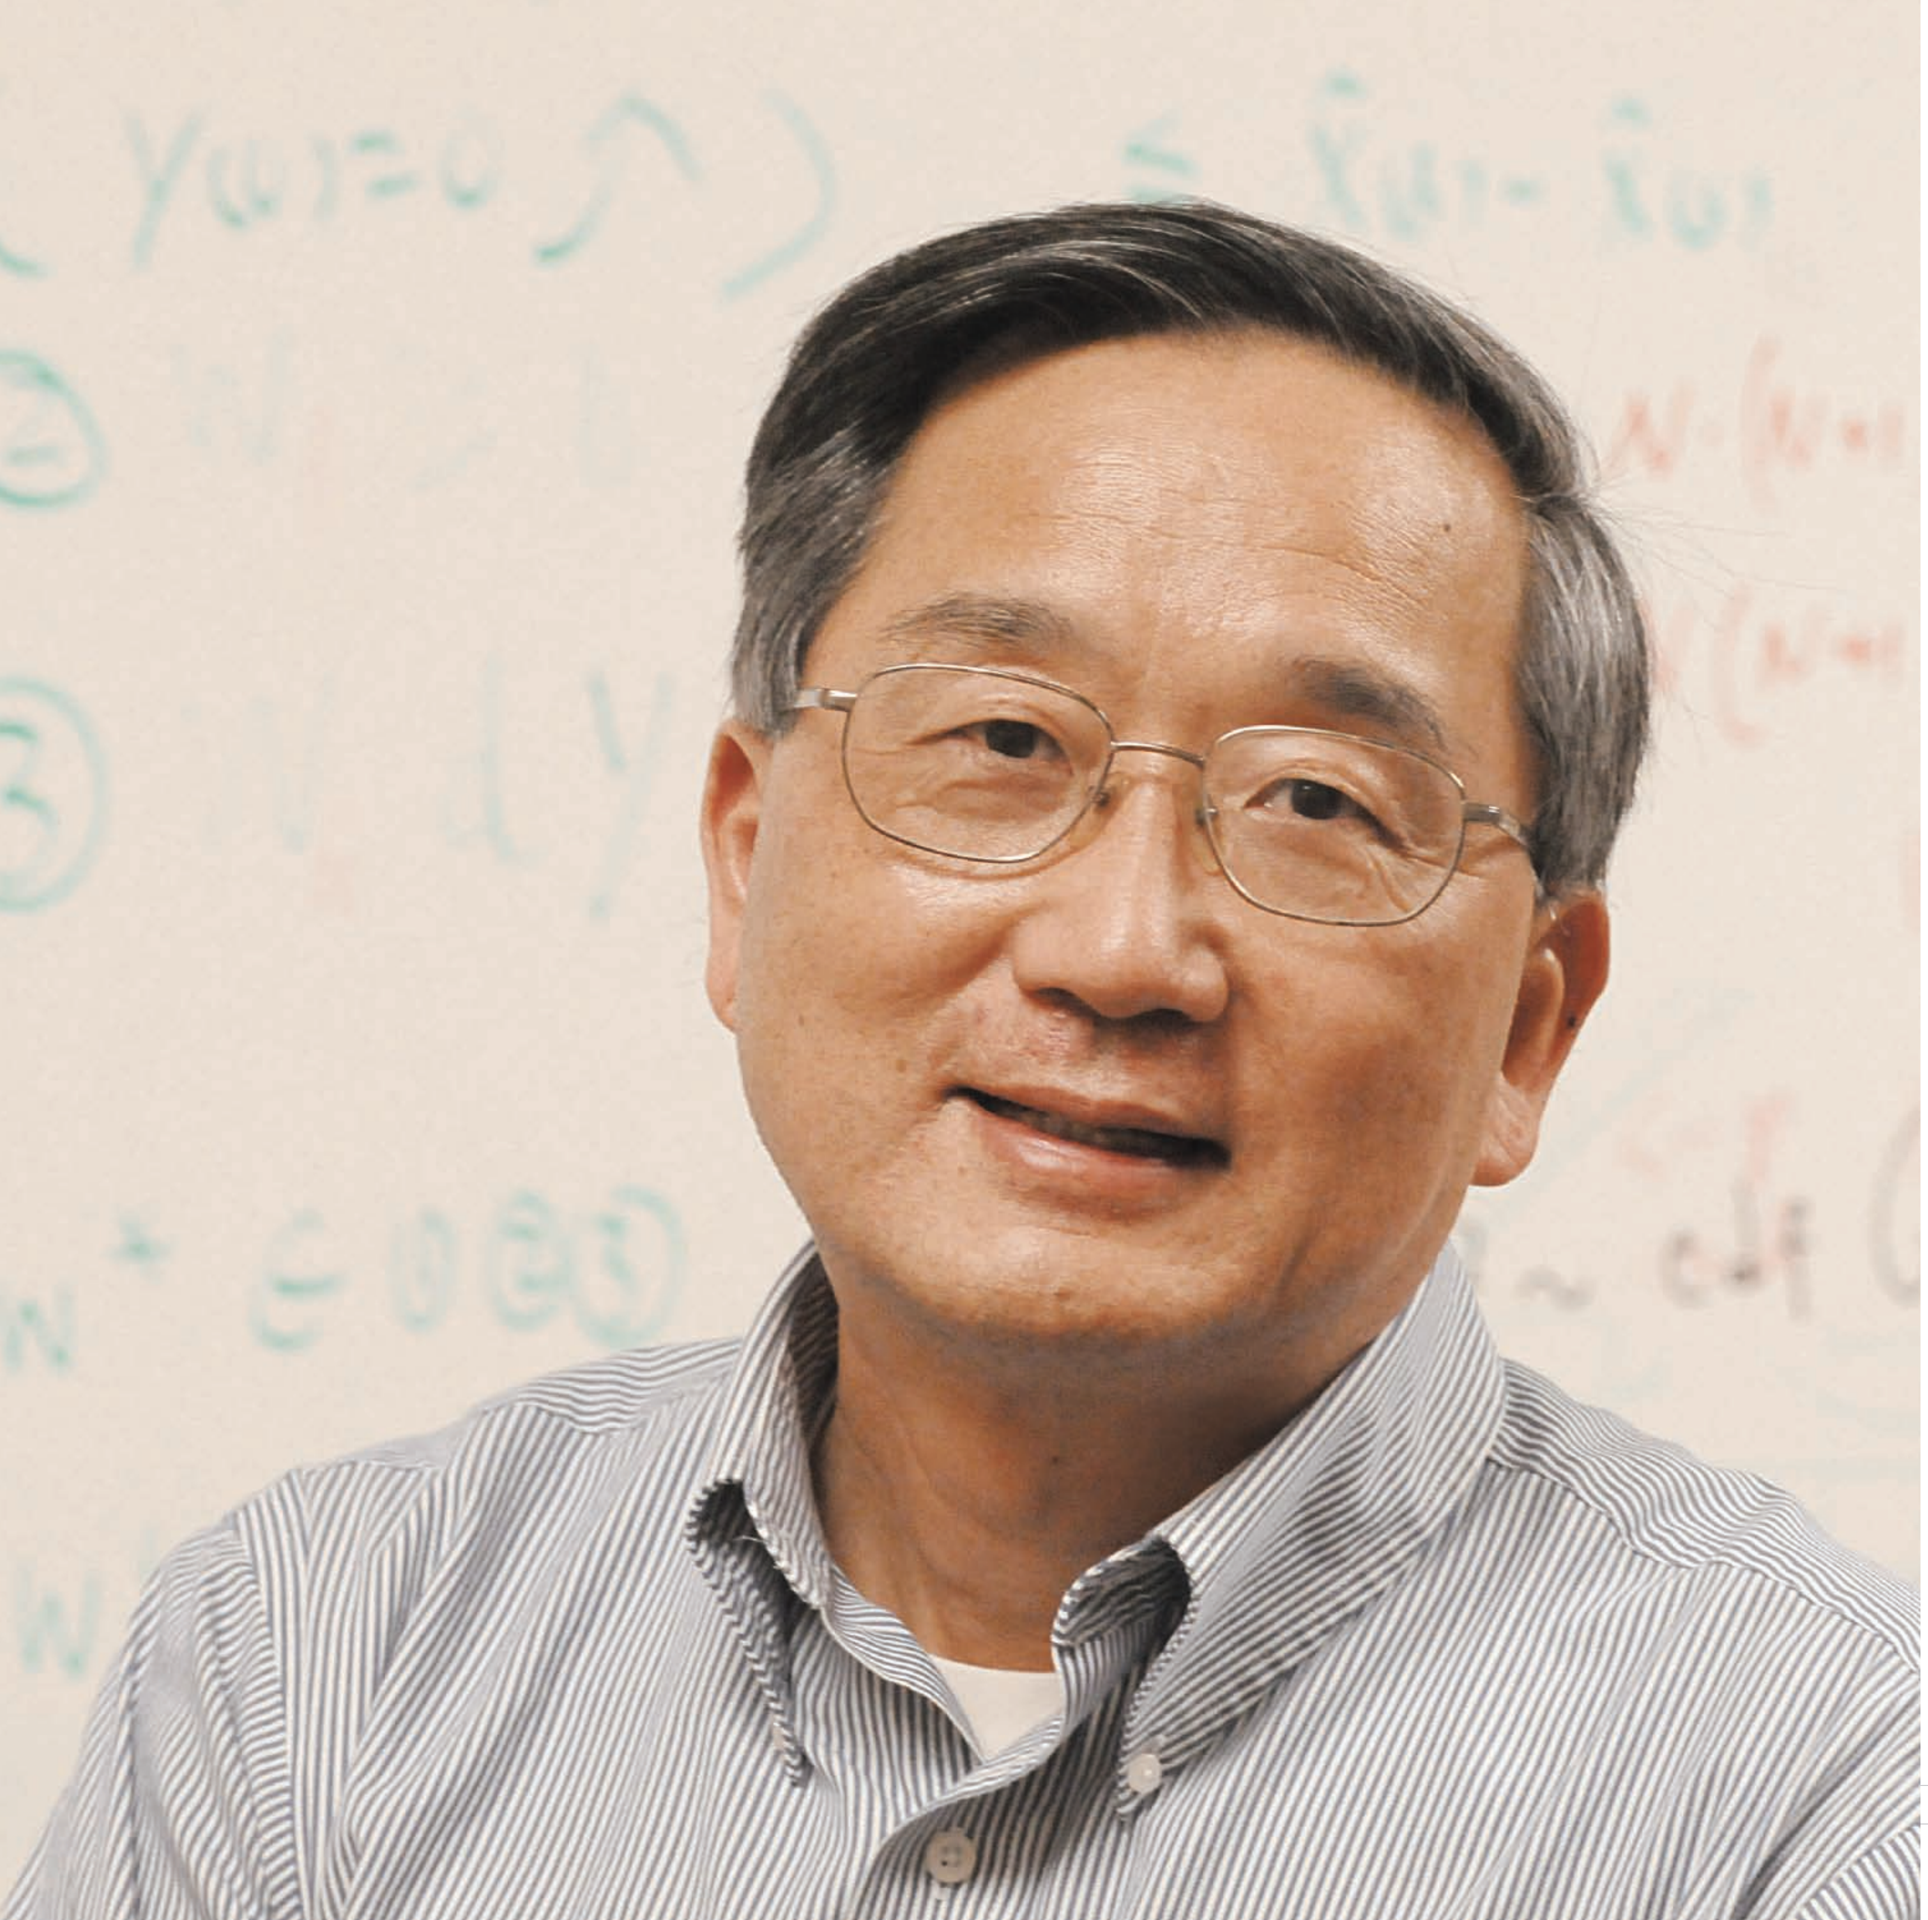

Optimizing Networked Resources

DAVID YAO

Professor of Industrial Engineering and Operations Research

INFORMATION

What do hospitals, airlines, supply chains, and the Internet all have in common? According to David Yao, they are all complex networks that must bring together multiple services and assets to accomplish any task. They must also share these same resources among different classes of customers, who pay different amounts for service.

Organizations want to manage their resources efficiently to maximize profits. But if they are too efficient—Yao likens it to filling a highway with cars so traffic slows to a crawl—they sacrifice quality of service. Balancing efficiency and service across complex networked resources is an exercise in extreme juggling. Yao wants to help by giving organizations the tools to do it in real time.

He points to airlines as an example. They must divide a limited number of seats among first, business, and several types of economy classes. Each class sells for a different price.

Airlines maximize revenue when they fill every seat. They can do this by discounting and by overbooking flights, since they know there will always be some no-shows. They also reserve some tickets to sell at higher last-minute prices.

That leads to problems. "The price they pay for overbooking is that they may have to ask people to get off the airplane. They also don't want to hold too many last-minute tickets, or they will have unfilled seats," Yao said.

Airlines estimate how many seats to sell and reserve by looking at past data. "That does not capture the real-time dynamics of the network," Yao stated. "On a particular day, a plane might be delayed and those passengers will need new connecting flights. Now their planes must carry their own customers plus passengers from the delayed flight."

Yao's models capture that type of real-time information and use it to optimize the entire system rather than a specific resource, like a single flight, a bank of servers, or a hospital bed. On airlines, his models assign all seats a shadow price, the revenue they could potentially earn if they sold a reserved ticket, and compare it with the probability of delays and other events as they evolve. It shows them the most profitable way to reroute passengers and flights.

"We look at the probability of events, but also at how we can hedge our bets if that probability is wrong," said Yao. "We want to create models that are predictive but robust, so if you're off, you won't walk away from money on table."

M.A.Sc., University of Toronto (Canada), 1981; Ph.D., University of Toronto, 1983

Turning Students into Entrepreneurs

YECHIAM YEMINI

Professor of Computer Science

Dell, Yahoo, Google, and Facebook were founded by college students, Yechiam Yemini tells his Principles of Innovation and Entrepreneurship class. He wants to teach students how to create innovative technologies and transform them into successful startups.

Yemini has combined academia with serial entrepreneurship. His first company, Comverse Technology, co-founded in 1984, revolutionized voice messaging technologies. Ten years later, System Management Arts created the first products to diagnose network failures automatically.

Startups, he explained, are another way of disseminating basic knowledge. "High-tech startups distill the value in raw, basic technologies by creating innovative products and introducing them to the market," he said.

Yemini's course rests on three legs. The first is understanding how to identify opportunities. The most fertile areas are those where new ideas disrupt established ways of doing things, such as integrated circuits, the Internet, and wireless networks.

"Today, the biggest transition is from cellular phones to mobile computing. Now your phone is a tool to go shopping, access content, play video and read books. It's a wonderful opportunity to launch companies that exploit this," Yemini said.

The course's second leg involves startup mechanics. "We look at the engines of value creation," Yemini said. "Different engines make products, exchange information with the market, and manipulate the flow of financial resources. We look at how to design these engines to optimize the value they create while minimizing risks and errors."

Yemini's third leg is product development. "Many companies fail because they spend all their time creating a product and then look for a market," he said. "They didn't manage the risk that customers wouldn't like their implementation, or that market needs might change." He advises students to begin talking with customers from day one, and to keep improving products incrementally until they are happy.

Yemini is focusing on managing mobile services. "Mobility presents a disruptive change in delivering network services. It presents research opportunities to create new technologies, which may one day lead to new startups," he said.

"A startup company is a bunch of engines that express the value of a technology," said Yemeni. "Think of it as a mechanism, a black box. There are ways to build a better box, ways to engineer it to better distill the value of the underlying technology. My course on innovation and entrepreneurship tries to teach how to engineer a technology company, much as one engineers other innovative constructs."

B.Sc., Hebrew University of Jerusalem (Israel), 1972; M.Sc., Hebrew University of Jerusalem, 1974; Ph.D., University of California-Los Angeles, 1978

Plugging the Leak in Circuit Efficiency

CHARLES ZUKOWSKI

Professor of Electrical Engineering

Computer chips are the building blocks that allow billions of transistors to fit in a small area. These chips have enhanced everyday life, and enable the design of electronics of increasing functionality and lower cost, making most modern-day technology possible. But as transistors continue to become smaller and faster, new challenges for circuit designers constantly arise. The research field of Very Large Scale Integration (VLSI) addresses these challenges.

One of those challenges is transistor current leakage, which is becoming a bigger problem as transistors in computer chips continue to shrink, leading to problems with power and reliability. While current leakage and power dissipation in each transistor remain quite small, they can add up to a significant amount over billions of transistors, potentially limiting function and performance. Solving this problem could have a big impact on industry, and the feasibility of critical future applications of electronics.

Charles Zukowski, past chairman and current vice chairman of the Department of Electrical Engineering, has worked in the area of VLSI throughout his career and has contributed to the progress of integrated circuit technology in a number of areas. His chief focus now is twofold: circuit techniques such as monotonic logic to reduce the impact of current leakage in future integrated circuit technologies; and special-purpose hardware prototypes for the simulation of gene regulatory networks. Through this work, his intention is to further the capability of integrated circuit technology and to explore new applications.

His research has covered both circuit design and circuit analysis, results of which include a patented circuit technique for generating high data-rate serial data from a number of lower data-rate channels, and an approach for mixing digital and large-signal analog computation for simulation. He derived a number of results for bounding the behavior of digital integrated circuits that were compiled into a research monograph, and based on this work, he received a National Science Foundation Presidential Young Investigator Award. He later developed a technique for measuring the convergence of waveform relaxation algorithms for simulating digital circuits. He also proposed a technique for significantly reducing the power consumption in certain content-addressable memories and investigated the use of various memories and circuit techniques in internet routing hardware. Throughout, he has consulted for industry in the field of Complementary Metal Oxide Semiconductor Integrated Circuit (CMOS IC) design.

B.S., Massachusetts Institute of Technology, 1982; M.S., MIT, 1982; Ph.D., MIT, 1985

Improving the Efficiency and Resiliency of Wireless Networks

GIL ZUSSMAN

Assistant Professor of Electrical Engineering

The design and deployment of mobile and wireless networks has undergone an extraordinary transformation. While this technology already forms the backbone of crucial systems such as health care, disaster recovery, public safety, manufacturing, and citywide broadband access, it has even greater potential. The flexibility inherent in cellular, sensor, mobile ad hoc, mesh, and wireless local area network technologies delivers an almost endless range of applications, including mobile banking, inter-vehicle communication, space exploration, and climate-change tracking.

Despite their promise, efficiently controlling wireless networks is a challenging task, due to interference between simultaneous transmissions, mobility of the nodes, limited capacity of the wireless channel, energy limitations of the devices, and lack of central control. Such distinct characteristics set wireless networks apart from other networking technologies and pose numerous challenging theoretical and practical problems.

To tackle those problems, Gil Zussman focuses on designing new wireless networking architectures and on improving the performance and resilience of existing networks. Due to the special characteristics of these networks, Zussman designs architectures and algorithms that are optimized across multiple layers of the networking protocol stack. For example, he has been working on energy-aware protocols that take into account energy consumption and battery status while making joint decisions regarding routing and scheduling. Zussman has been recently focusing on developing algorithms and prototypes for Energy Harvesting Active Networked Tags (EnHANTs). These tags harvest their energy from the environment and can be used in various tracking applications, and particularly, in disaster recovery applications.

Moreover, in order to enable the efficient operation of distributed algorithms which usually have inferior performance to centralized algorithms, Zussman has been working on identifying topologies in which distributed algorithms obtain maximum throughput. His results in this area enable the partitioning of networks to subnetworks in which distributed algorithms operate very well, thereby improving the overall network performance.

Other research projects of Zussman's group focus on controlled mobility of wireless nodes, dynamic spectrum allocation and cognitive radio, interfaces between wireless and optical networks, and resilience of networks to geographically correlated failures. Results regarding the latter include identifying vulnerabilities of networks to large-scale attacks, such as Electromagnetic Pulse (EMP) attacks, and mechanisms to mitigate the effects of such attacks.

Zussman was a postdoctoral associate with the Massachusetts Institute of Technology as a Fulbright Fellow and Marie Curie Fellow. He is a senior member of the Institute of Electrical and Electronics Engineers.

B.Sc., B.A., Technion-Israel Institute of Technology, 1995; M.Sc., Tel Aviv University, 1999; Ph.D., Technion-Israel Institute of Technology, 2004

Engineering as the Newest Liberal Art

As a faculty member at Columbia Engineering for more than a half century, I would like to share with you my perspective on the School's history and future. Since 1956, Columbia Engineering has been my academic and professional home.

As a newly appointed professor in the Department of Civil Engineering and Engineering Mechanics, I joined a faculty that boasted many of the top names in their fields—Rudolf Kalman, known for the Kalman filter; Lotfi Zadeh, who invented fuzzy logic; Ferdinand Freudenstein, father of modern kinematics; Raymond Mindlin, known for his work in the theory of elasticity; Cyrus Derman, known for optimization of stochastic systems; and Elmer Gaden, the father of biochemical engineering.

As you read about our current faculty, you can see that today we have similarly outstanding and innovative researchers who themselves are, or are becoming, the equals of the legendary faculty giants who were here when I arrived more than 55 years ago.

What has changed is the nature and practice of engineering, which has now become central to almost all human intellectual activities, ranging from pure science to business and economics. In fact, engineering is now sometimes called the newest liberal art.

This change has served to encourage, and even demand, that a great university such as Columbia have an engineering school ever stronger in its engineering-based programs. This impetus has spurred the creation of additional departments—Applied Physics and Applied Mathematics, Computer Science, Biomedical Engineering, and Earth and Environmental Engineering—and many new programs, such as financial engineering, fusion energy, stem cell research, biological systems research, materials and process research at the atomic and molecular levels.

Engineering and applied science research now plays a greatly expanded role in the rapidly advancing biological, physical, chemical, and mathematical sciences, and as such, in the intellectual life of the University.

I am privileged to be a member of this vibrant Columbia Engineering faculty and I know you are as proud as I am of its accomplishments throughout its history and of the bright future that lies ahead.

Morton B. Friedman

Morton B. Friedman
Senior Vice Dean and Professor of Civil Engineering and Engineering Mechanics