ANNUM

THE EDWARD S. ROGERS SR. DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING

FINDING NEW AND SMARTER WAYS TO GENERATE AND DISTRIBUTE ENERGY IS ONE OF THE MOST PRESSING ISSUES OF OUR TIME.

In ECE, we are world leaders in energy systems research and education. We have just completed Phase 1 of a comprehensive three-stage update to our energy systems laboratories.

This renovation will open up new opportunities for both research and teaching, allowing graduate and undergraduate students to gain first-hand experience with the latest smart grid technologies, alternative energy resources and advanced methods for control of power systems.

Please help us reach this ambitious goal.

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

To learn how you can help advance our energy systems research and teaching, please contact Celeste Taylor, Director of Development, at 416 978 0291 or celeste.taylor@utoronto.ca

BOUNDLESS ENERGY

PHOTO BY RAINA+WILSON

Where research and teaching touch real-world issues

The year was 1982: the first CD players hit markets in Germany and Japan, Michael Jackson released *Thriller*, and Time Magazine declared the computer its 'Machine of the Year'. Here in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE), something prescient was developing—Professor K.C. Smith and his graduate student presented a new way to interface directly with computers by capturing signals on a flat glass surface.

It was the beginning of multi-touch, something most of us now use every day on smartphones, tablets and laptops. At the time, no one imagined how ubiquitous mobile devices would become. Today we rely on our phones to tell us where we are, remind us what we've forgotten, and inform us of everything we don't know. In return, we feed our devices a constant flow of information about ourselves. We hope our personal information is safe, but how can we be sure? This issue's cover feature, *Identity Crisis* on page 22, highlights our investigation into pressing concerns surrounding smartphone and biometric security and privacy.

One of our department's primary strengths is a close connection to industry, which influences both our teaching and research directions. I am pleased to announce a fresh collaboration with Hydro One, running through 2016. We have also founded a new ECE Industry Advisory Board composed of members from Hydro One, Qualcomm, IBM, TELUS, Altera and Huawei. Learn more about the board's role on page 39.

More than three decades since multi-touch took root in our labs, we maintain our position as the top-ranked ECE program in Canada and are proud to be among the best in the world. In these pages you'll find stories about the people that keep us there—our students, faculty, alumni and community. I view this magazine as an opportunity to invite comment about our activities and continuing evolution. I welcome your feedback and hope you will take this chance to reconnect with the department—you may reach me directly at **chair@ece.utoronto.ca**.

FARID N. NAJM,
PROFESSOR & CHAIR

ANNUM 2014

EDITOR
Marit Mitchell

ART DIRECTOR
Katina Constantinou,
Sugar Design

PHOTOGRAPHY & ILLUSTRATION
Raina+Wilson
Mark Smith

COPY EDITORS
Kelly Hayward
Avivah Wargon

EDITORIAL INQUIRIES:
ECE Communications Office
Tel: 416 978 7997
Fax: 416 978 1145
eceinquiry@utoronto.ca

Visit us online:
www.ece.utoronto.ca

Join the conversation:

All correspondence and undeliverable copies:

ANNUM Magazine
The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
Sandford Fleming Building,
10 King's College Road, Room SFB540
Toronto, ON M5S 3G4

Update or correct your contact information:
416 978 1999 or email
eceinquiry@utoronto.ca

Publication Agreement No:
42887022

ISSN 2368-7037

Printed in Canada by
Flash Reproduction.

MIX
Paper from responsible sources
FSC® C011058

FEATURES

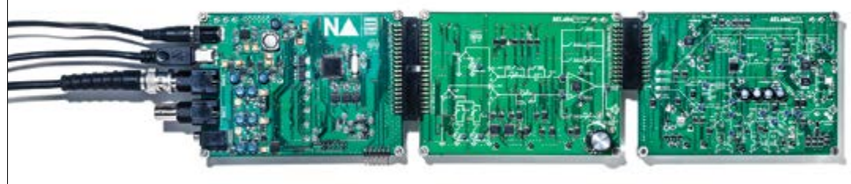

> This is the cake we want, this is how we want it to taste, you figure out the recipe and configure this thing. p.12

SECTIONS

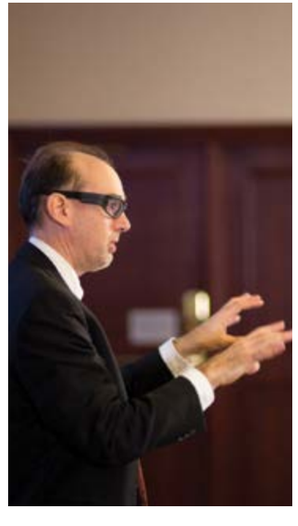

1,480
Undergraduate students
595
Graduate students
99
Professors, including Emeritus
52
Post-docs
49
Admin & tech staff
16
Research associates
6
Visiting professors

OUR STORY

Founded in 1909, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) maintains a proud history of world-leading research and innovative education. We are the top-ranked ECE department in Canada and one of the largest. With 99 professors, 595 graduate students and 1,480 undergraduates, our classrooms, halls and laboratories hum with energy and creativity.

ECE students and professors work together to solve some of the most pressing issues of our time—we are recognized leaders in the fields of smart grid technology, mobile application development, next-generation networks, and emerging areas such as biomedical engineering and robotics. Our graduates go on to prominent careers in engineering, finance, medicine, law and more—their options are truly unlimited.

QS World University Rankings 2014

#1 *in* Canada
Electrical & Electronic Engineering

#11 *in* North America
Electrical & Electronic Engineering

#23 *in* the World
Electrical & Electronic Engineering

#1 *in* Canada
Computer Science & Information Systems

#8 *in* North America
Computer Science & Information Systems

#19 *in* the World
Computer Science & Information Systems

TIMES HIGHER EDUCATION WORLD UNIVERSITY RANKINGS 2014–2015

Engineering & Technology Universities

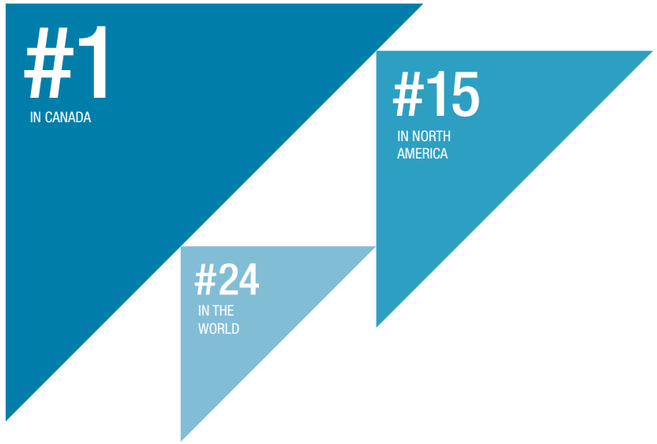

ACADEMIC RANKING OF WORLD UNIVERSITIES 2014

Engineering/Technology & Computer Sciences

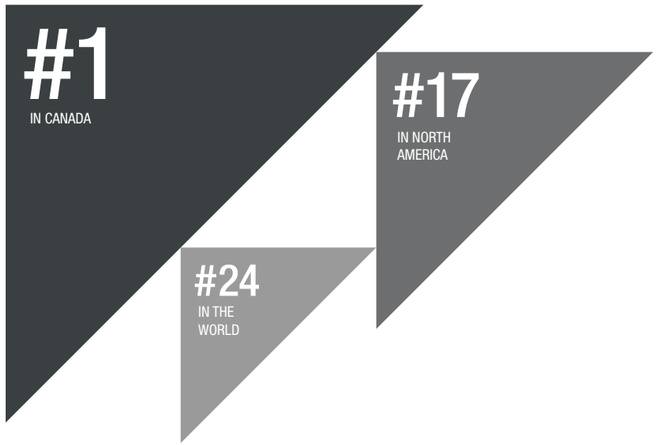

23
Chair titles held by ECE faculty in 2013–2014, including Canada Research Chairs, Endowed Research Chairs, Industrial Research Chairs and U *of* T Distinguished Professor Chairs

30
PATENT APPLICATIONS OUT OF ECE SINCE 2010

145
INVENTIONS OUT OF ECE SINCE 2010

11
STARTUPS OUT OF ECE SINCE 2010

230
UNDERGRADUATE DEGREES AWARDED FOR 2013–2014

BY MARIT MITCHELL

POWER PLAYERS

Wild weather, electricity addiction and even renewable resources put more strain on today's power grid than ever before. To make their network smarter and hardier, Ontario's largest distribution and transmission utility Hydro One turned to three ECE professors for smart grid technologies that will help keep us all plugged in

A SINGLE UNPRUNED TREE WAS ALL IT TOOK.

On August 14, 2003, one hot day in a hot summer, a power line sagged onto some branches in the small village of Walton Hills, Ohio. The resulting cascading failure throughout Ontario and the northeastern United States shut down power to more than 10 million people, becoming the second-most-widespread blackout in history and a major wake-up call for power generators and electric utilities operators around the world.

The grid has gotten a lot smarter since then, but there's still plenty of room for improvement. "The building-up of intelligence on our network will continue to go forward forever—there will never be an end," says Charles Esendal, manager of research, development and demonstration for Hydro One. "It's not a destination, it's a journey."

Helping accelerate that journey are professors Reza Iravani, Zeb Tate and Josh Taylor. Each is collaborating with Hydro One on a project to make Hydro One's network more robust for the 21st century, by improving power-system stability and monitoring, use of energy storage, and integration of distributed energy resources, including intermittent renewable generation such as wind and solar power.

The three projects launched in 2014 and will run for two years under the guidance of a steering committee composed of Esendal and Birendra Singh from Hydro One, ECE Chair Professor Farid Najm, and Professors Reza Iravani and Peter Lehn.

ISLANDS IN THE STORM

The current grid's interconnectivity, a strength when all runs smoothly, becomes a liability in adverse conditions—Americans witnessed this after Hurricane Sandy blasted the northeastern U.S. in October 2012, shutting down vast swaths of the system. As severe and unpredictable weather events increase in frequency, the ability to isolate and operate sections of the grid independently and sustainably is a high priority.

"If some event happens upstream in the system and you get disconnected, you become an electrical island," says Professor Iravani, who has had much successful collaboration with Hydro One and its predecessor, Ontario Hydro, since 1991. "We should be able to operate these islands in order to provide resiliency to the system, so it doesn't completely shut down—it should be self-sustaining." Professor Iravani is validating methods for isolating and controlling subsections of the grid within the larger system, and integrating renewable energy sources into these microgrids. Others have experimentally demonstrated this kind of microgrid islanding, but nothing has been validated and implemented on a larger scale.

We should be able to operate these islands in order to provide resiliency to the system, so it doesn't completely shut down—it should be self-sustaining.

The good news: all the essential elements of islanded microgrids—distributed resources, loads, transmission and monitoring infrastructure—are already in place on Hydro One's network. The challenge is to control them effectively, whether they're functioning as part of the whole or on their own.

"The components are there, but the point is to provide enough intelligence, decision-making and control for that mode of operation," says Professor Iravani. "All the components should be individually controlled, and all the controls should be coordinated to enable unified operation of the system. That's what we're aiming for."

THE GREAT GREEN UNKNOWN

We all know fossil fuels aren't cheap, clean or abundant—but what are we doing about it? Plenty, it turns out—the high market penetration of solar generation from individual consumers, and the uptick in electric vehicles in need of a charge are introducing unknowns into the grid on both the supply and demand sides.

"It is great to be able to say 'Let's use the sun's rays or the wind to generate electricity,' but it's very transient in nature," says Esendal. "You cannot predict the trend at all, and that is the worst thing that you can have on the grid. The system requires accurate anticipation and prediction to be able to plan and dispatch the electricity where the need is. Unfortunately that is a major challenge in integration of renewable generation."

What if you could bottle up some of that wind or sunlight for later, and redistribute it when need is high? "The most basic constraint on the power system is that supply must equal demand," says Professor Taylor. "Renewables make the supply random, but storage can make it controllable—the question is how best to do that."

Conventional storage meant pumping water up a hill and releasing it, converting potential to kinetic energy to make electricity—though it sounds old-fashioned, the method is quite efficient, but expensive and geographically restrictive. Professor Taylor is looking at both batteries and flywheels for more flexible, distributed storage and is developing an algorithm to make smart decisions about optimal locations and conditions to either store or release energy.

READY FOR ANYTHING

When a line overheats due to excessive current, as in the great blackout of August 2003, the power transferred by that line needs to be redirected, and utilities operators need to "reposition" the system in order to maintain stability. "In 2003 they didn't know the lines were out, so they didn't know they needed to reposition the system," says Professor Tate. Operators learned the hard way that accurate information is everything, and afterward deployed much more sophisticated monitoring methods.

Professor Tate is working on a project to use these smarter monitors, called synchrophasors, to detect faults in the grid and automate rapid responses. Synchrophasors sample and relay magnitude and phase data between 30 and 60 times per second, a vast improvement on old sensors that are polled only once every two to three seconds. "This gives you a much better idea of how power is flowing immediately, so you can tell whether a line is disconnected or a generator is down," he says. "Improving detection fosters efficiency—you don't have to run the system as conservatively if you know you can react instantly."

"There are a lot of challenges today—these are only a few issues that we have," says Esendal.

"It will take some time to validate all of these technologies, but we needed to start somewhere and I think University of Toronto is in a great position, not only from the expertise perspective, but also to be located just in our backyard—it's a great advantage."

If all goes according to plan, you'll never know how these projects turn out—you'll just be cozy during the next deluge of freezing rain, cool throughout the sweltering summer, and connected in the next flood, hurricane or hail storm. And soon August 2003 will be just another unbelievable anecdote of times long past. A

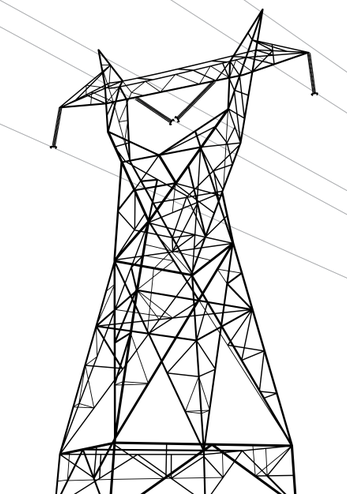

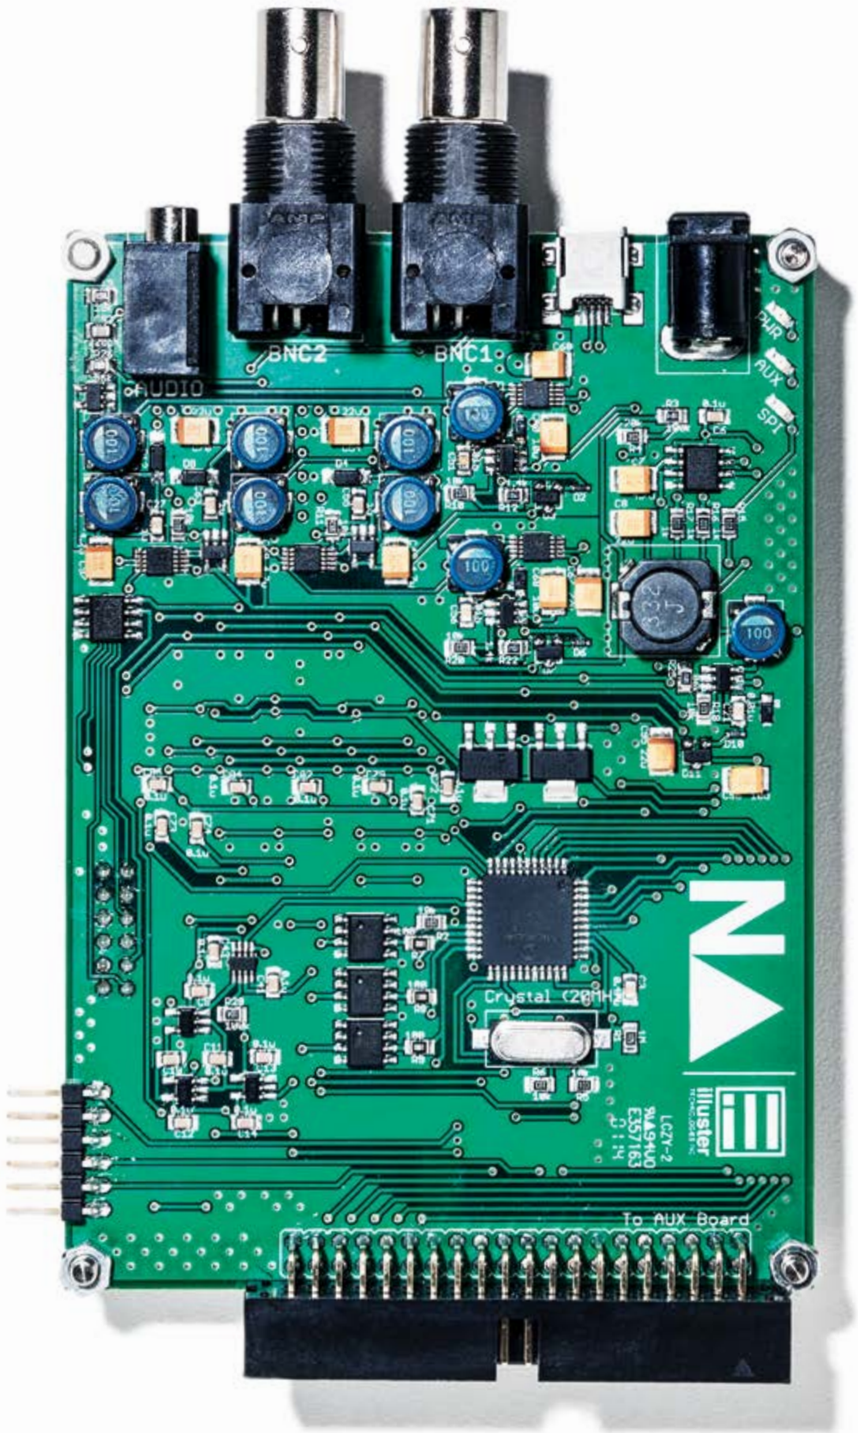

BY MARIT MITCHELL
PHOTOGRAPHY BY RAINA+WILSON

BAKING AN IMPROVED BREADBOARD

Three fresh ECE graduates are bringing sexy back to analog electronics labs, giving current undergraduates a leg up at the same time

This is a story about three intrepid alumni, two famous names in micro-electronics, one professor with a desire to run better labs, and thousands of undergraduate students who will benefit from their collaboration.

Do you remember your first-ever analog electronics lab? Whether it was 50 years ago, 20 years ago or last year, you worked with a plastic cribbage board bristling with a rainbow of tangled wires. For decades, these breadboards, or protoboards, have been the main tool for teaching electronics fundamentals to students worldwide. There's nothing wrong with that—but it no longer screams 'advanced technology'.

"When you walk into the lab as an undergrad and you see breadboards, you think 'Really? This cannot possibly be relevant,'" says Professor Olivier Trescases, who teaches ECE231 – Introductory Electronics. "Only our digital courses were using printed circuit boards. So there was this really false message that digital is sexy and modern and relevant, and analog is stuck in the ways of the '70s—in the real world, analog circuit design is a rapidly evolving and tremendously exciting field. And I think for many students, they got turned off."

It was a concern Professor Trescases had been marinating for a couple of years when he recruited second-year student Miad Fard as a summer researcher in 2010. He and Fard hatched plans to revolutionize ECE's electronics labs by replacing breadboards with digitally reconfigurable analog circuits in the form of interchangeable printed circuit boards (PCBs), with a unique board for each lesson. It was a huge undertaking, but the project got under Fard's skin—after returning from his Professional Experience Year, he convinced two classmates, Mehrad Mashayekhi and Richard Medal, to join and bring the idea to fruition as their fourth-year design project.

"PROTOBOARDS ARE EFFECTIVE AND VERSATILE, BUT THEY'RE NOT REALLY BEING USED IN INDUSTRY ANYMORE. YOU COULD GO ALL FOUR YEARS AS AN UNDERGRAD AND NOT HAVE THE EXPERIENCE OF USING PCBS," SAYS FARD. "WE WANTED TO INTRODUCE SOMETHING THAT'S MORE INDUSTRY-STANDARD AND HELP STUDENTS BECOME BETTER ENGINEERS ONCE THEY LEAVE SCHOOL."

Their hardware-software system is straightforward and elegant: a central board, called NileDelta (or NΔ), connects to a computer running software with a graphical user interface that allows students to tinker easily. It updates the setup without compromising any of the tangible equipment that brings circuitry to life—students stimulate circuits with real sources, including music fed from their cell phones, a microphone, or function generators, and observe outputs using loads such as LEDs, headphones, motors or oscilloscopes. The team concurrently developed a set of auxiliary boards, called AELabs, which plug into NΔ and let students configure and reconfigure circuits specially designed for each lesson with a different mix of transistors, resistors, capacitors and diodes. Every AELabs kit comes with a detailed lab manual with instructions and background material for each experiment, matched to its corresponding auxiliary board.

"This takes the manual component out of building and debugging circuits, which students still learn in a previous course," says Medal. "They still have the hands-on experience, which is important, but now they can step up to modern technology and dive deeper into the concepts." Adds Fard: "It also allows them more time for design-oriented tasks."

FROM L TO R: RICHARD MEDAL, MEHRAD MASHAYEKHI, MIAD FARD

The NileDelta board interfaces electric power, user activity, and external stimuli which are then used to configure the analog circuits on the auxiliary AELabs boards.

Still undergrads themselves, Fard, Mashayekhi and Medal refined and printed the system over the course of their fourth year, and in spring 2013 tested it on a subset of 20 ECE231 students in the lab. Despite expected glitches, their product was well received, and earned the 2013 Gordon R. Slemon Design Award for exhibiting the best combination of imagination, design and execution of more than 100 projects.

But that was just the beginning. In the intervening 18 months, "the three graduated from the Electrical Engineering program and founded Illuster, their startup company, to commercialize the system. University of Toronto became their first customer, purchasing 45 kits to expand the project to all second-year students in ECE231. They spent much of spring 2014 helping students navigate the lab, taking in feedback about what was working and what wasn't, and calibrating accordingly.

The Illuster system has drastically extended the scope of Professor Trescases' syllabus. "Before we said, 'Here's the precise recipe, you bake the cake and if you have time left over, tell us how it tastes.' Now we say, 'This is the cake we want, this is how we want it to taste, you figure out the recipe and configure this thing. Now with this knowledge, show us how you can make it taste even better.' There was not nearly enough time for that process in the past," he says. "The labs are totally different now—we expanded our experiments, and we hope that students have learned a lot more."

ECE *Emeritus Professors* Adel Sedra and K.C. Smith, authors of the ubiquitous textbook *Microelectronic Circuits*, have been instrumental in the project's evolution since its early days—Professor Trescases, who himself worked on many of the simulations for the book's accompanying CD as a student, brought the idea to Professor Smith right away, and later to Professor Sedra. Both are very enthusiastic about the platform and plan to mention it in their next edition. Microelectronic Circuits is an international bestseller and the endorsement could be a major break for Illuster—an opportunity to launch the next phase of their plan.

"What we've built is modular and expandable, and we intend to create additional kits with new auxiliary boards to teach other courses," says Medal. "These could be used for digital electronics, power electronics, control systems, mechatronics. That's our end goal—to reach out to schools and modernize and update the way that electronics is taught around the world." A

OPTICAL INCLUSION

PhD student Arash Joushaghani's work on a largely forgotten material brings optics one switch closer to revolutionizing our electronic status quo

BY MARIT MITCHELL
PHOTOGRAPHY BY RAINA+WILSON

Arash Joushaghani sells things people don't know what to do with.

Take his PhD research: working with a bizarre Cold War-era material used in missile-guidance systems, he designed a miniscule optical switch, invisible to the naked eye.

Who wants that? The hundred-billion-dollar global electronics industry just might.

REALLY BAD OPTICS

Your smartphone, tablet, GPS and all things computerized run on semiconductor chips built out of silicon. These chips have gotten smaller, and smaller, and smaller—now the gates routing electrons through them are about 14 nanometres in size.

Light is faster and cheaper than electricity, but optical modulators, for turning light on and off, have traditionally been a few hundred micrometres—several orders of magnitude larger—making them too bulky to incorporate into electronic devices.

"We said, instead of looking at conventional materials, like silicon, or new materials, like graphene, let's go a bit crazy—let's find something that really lets us go smaller than a micron," says Joushaghani.

The PhD student, co-supervised by photonics heavyweights professors Joyce Poon and Stewart Aitchison, obtained samples of vanadium dioxide, or VO2, from Defence Research and Development Canada's primary research centre in Valcartier, Quebec.

At first glance, it wasn't a very attractive option. "It goes from a bad optical material, to a really bad optical material," says Joushaghani. But at least it flips from bad to worse quickly—a crucial property in a switch. After fabricating several designs in the Toronto Nanofabrication Centre located at University of Toronto, his final product blinks light at relatively fast speeds using electronic signals—and it's half a micron in size.

"Engineers would say, 'Neat stuff, but that's science fiction,'" says Joushaghani. "But when you actually bring it to their own platform, with their own specs and everything, then I think you can actually start to capture people's attention."

Joushaghani didn't set out to capture Intel's attention—he just wanted a coffee.

When Dr. Tuyen Tran, engineering technology development manager at Intel's Portland Technology Development Division, visited ECE in November 2013 as part of the department's Research Partnership Lectures series, Joushaghani and a friend arrived early to secure some caffeine. He struck up a casual conversation with Tran, who invited him to Portland, Oregon for an interview and tour. He visited in February, graduated in September, and started in October with Intel's Optical & Metrology R&D division working on the next generation of 10-nanometre chips. His work will be in your hands in less than two years.

MASTER OF THE FLIP

But before Joushaghani flipped his unlikely switch into a job for himself, he was doing the same for others—primarily the researchers he worked with.

Frustrated by the length of time it takes to reproduce a colleague's research in order to build on it, he founded Sciventions, a company designed to sell the processes and products perfected by academics to their colleagues around the world, making money for some, and saving years of time and effort for others. His entrepreneurial inspiration arrived while he was still an undergraduate student in Engineering Science, and working with Professor Geoff Ozin in the Department of Chemistry. "They had these really cool optical materials, but they were sending them to Spain to get the data," he recalls. "And I was puzzled—why can't you just send it somewhere at UofT? Your neighbour is doing this, and you don't know about it."

His initial concept, to close the gap between researchers, evolved into a roundabout marketing platform: Sciventions now sells the products of academic research to other investigators, and provides the seller with more information about the commercialization potential of their work. Launched two years ago, the company now sells about 30 products online.

"Our business model is 'You never know until you try,'" says Joushaghani. "I guess that's my mentality too." A

Before Joushaghani flipped his unlikely switch into a job for himself, he was doing the same for others—primarily the researchers he worked with.

BY MARIT MITCHELL
PHOTOGRAPHY BY RAINA+WILSON

Great Expectations

How this electrical engineer became an unexpected champion for reading in the 21st century

IF Charles Dickens were alive today, he'd be an avid Wattpad user—Allen Lau is sure of it.

The ECE alumnus is co-founder and CEO of the fast-growing company that lets readers and writers share their love of literature in new and immediate ways—writers release their work chapter by chapter, and readers around the globe instantly devour new stories on their phones, tablets and laptops.

At its core it's the same format that made Dickens's work so addictive to Victorian England, when chapters of his novels were published in serial, making them much more affordable and accessible to the public. But in an era where print materials are dismissed as 'dead trees' and attention spans have shrunk to the length of a GIF, what does the future of reading look like?

Lau and his co-founder Ivan Yuen have that answer. "We believe we can build a billion-user company that entirely focuses on reading and writing," says Lau.

And with more than 32 million unique monthly visitors to Wattpad—roughly the population of Canada—the formula is working. A funding round earlier this year brought a $46-million investment from a group led by OMERS Ventures, money that will help grow the company. Its global impact is already evident: Wattpad is the number one app in the Philippines, and more than 10 per cent of its traffic comes from Spanish-speaking countries. If *Fifty Shades of Grey* showed us anything, let it be that people's voracious appetite for romance, fan fiction and sci-fi is big business.

The company has grown from two members—Lau and Yuen, hanging out in their basements—to almost 90, and just moved to a multi-floor loft-style headquarters in the King East neighbourhood of Toronto. The open-concept office is bathed in natural light from huge warehouse windows and its finished in warm wood and bright colours. Developers, marketers and business strategists work around tables together, huddle in meeting rooms or chat over snacks

and coffee in the expansive kitchen. It's the kind of space you see in recruitment materials for Silicon Valley darlings such as Twitter or Airbnb, but this isn't California—Wattpad is deliberately Toronto-based, and proud of it.

"To build a global internet company, understanding different languages is the first step, but understanding the nuances of different cultures—that's super important," says Lau. "If you talk to some of the workers here, for a majority of them, they're bilingual, have a different ethnic background, or lived or worked outside of Canada. Why? It's not because we chose to do that, it's because that's the nature of Toronto."

Lau's family arrived in Toronto in 1987, the last of his mother's side to move to Canada from Hong Kong. When it came time for university, he just asked his relatives which school they thought was best and off he went to University of Toronto. He graduated from Electrical Engineering in 1991, just before the Computer Engineering program was created, and finished his master's degree under the supervision of Professor Elvino Sousa in 1992. "I took a lot of the computer engineering, software engineering courses, and there were two areas that I focused on: computer and communications," says Lau. "My master's degree specifically focused on mobile communications. In a way it's a full circle."

Lau has always strived to combine his hobbies and passions with his day job, but didn't set out to become the three-time entrepreneur he is today. "I wanted to work for NASA—I wanted to be a scientist," he remembers. "When I watched the scientists launch the shuttle, all that equipment, it was fascinating—so inspiring. I just wanted to be part of it." His parents encouraged him to broaden his rather specific goal and play to his strengths in math and science. It's a skill set that's served him well— Lau founded his first company, a mobile gaming and advertising company called Tira Wireless in 2001. (Yuen was his first hire.) In 2002, he started developing a reading app for his Nokia Candybar before deciding the screen was impossibly small and setting the project aside. Fast forward to 2006, when he resurrected the idea for the Motorola Razr, the best phone on the market at the time.

"So I was busy in my basement coding, and one day Ivan instant-messaged me: 'Hey Allen, I'm working on a new prototype, can you give me some feedback?' and I clicked on that link, what I saw was mobile reading on a cell phone. I thought, 'No way!'" he recalls. "Ivan is the smartest person I've ever worked with. And if Ivan and I both independently came up with the same idea, it must be a good idea." Within two days he was on a plane to see Yuen, who was living in Vancouver at the time. "We were literally writing our business plan on a piece of napkin—like a Hollywood movie, but it's real life."

Will Wattpad's story have a Hollywood ending? Lau's not interested in that kind of drama just yet.

"I think in a way, startups and entrepreneurship have been heavily romanticized. I believe people hear about YouTube and Instagram, and they suddenly want a YouTube or they want an Instagram. They want WhatsApp. Those are once in a decade or twice in a decade kind of startups—most companies are not like that. And even some of the more successful ones, in the early years they may struggle a lot. It just takes time to build a great company," says Lau. "We don't think about exit, or going public—I don't even waste time thinking about this. I believe we should think about the type of company that we want to build, and that type of success or outcome will come automatically."

We'll just have to wait for the next chapter.

PHOTO BY MICHAEL TENAGLIA

IN MEMORIAM:
John Gregory Steffan

It is with deep regret that we mourn the death on July 24, 2014 of John Gregory Steffan.

John Gregory Steffan—Greg to all who knew him—was born on August 27, 1972.

Greg attended the University of Toronto and received his Bachelor of Applied Science degree in Computer Engineering in 1995, and a Master's of Applied Science in Computer Engineering in 1997. He earned a PhD in Computer Science from Carnegie Mellon University in 2003.

He rejoined the University of Toronto as an assistant professor in 2004 and launched a career characterized by outstanding research, teaching and citizenship. His research dealt with various aspects of improving computer performance through parallelism. Building on his PhD, he continued to find innovative and clever ways to enhance parallelism in single and multi-core computers, and branched out into innovative uses of Field-Programmable Gate Array implementations of processors. This work included applications to network packet processing, vector processing and an award-winning paper on novel ways to implement multi-ported memories—this paper was cited as one of the top 25 papers in the first 20 years of the principal FPGA conference.

His colleagues remember his research intellect and achievements as both logically sound and profoundly innovative. Greg's reputation and respect went beyond the boundaries of the University, where he was a highly regarded investigator whose work crossed several research communities. He was promoted to associate professor with tenure in 2009, and continued to look for new ways to make computing and computers more efficient, a vital area of inquiry.

As a teacher, his students deeply admired his efforts and ideas. He taught a wide variety of courses, ranging from first-year computer programming to a final-year complex course on optimizing compilers. He was one of the first teachers of a new course that dealt with high-performance software development, adding a key missing element to his department's software program.

His sense of humour and eye for justice came together when the Mattel Company announced a contest to determine which of several new possible careers the next Barbie Doll should portray. Greg campaigned to have his colleagues and students vote online for Computer Engineer, and he was delighted when this came to be. He bought his own Computer Engineer Barbie and kept it in his office.

Greg was a lifelong guitar player and an avid ball hockey player, specializing in goaltending in a weekly ball hockey league. He enjoyed playing golf and brought his careful analytical skills to the analysis of the universal failure at that sport. At an alumni golf tournament, he helped his group by reciting the Engineer's Chant when challenged at one of the holes.

Greg's colleagues and students were heartbroken at his sudden passing and miss his wonderful spirit of collaboration and fun. Flags across the University of Toronto's three campuses were lowered to half-mast on July 28, 2014 in his honour.

This is an abridged version of a tribute prepared by Professors Paul Chow, Andreas Moshovos and Jonathan Rose, read aloud at the October 8, 2014 meeting of Faculty Council.

BY MARIT MITCHELL
ILLUSTRATION BY MARK SMITH

IDENTITY CRISIS

We trust our smartphones with our lives—but how smart is that?

It wakes up next to you, sits by to you at lunch, hits the gym with you after work. Face it—your smartphone is your best friend.

But how good is it at keeping your secrets?

Almost two billion people have a computer in their pockets right now. And we use these smartphones for everything: not just talking and texting, but handling our finances, booking travel, mapping our next run and tracking our medications. We've taken our lives in our hands—literally.

All this personal information goes pouring into our devices through the apps we install, and floats magically into some hazily defined realm called 'the cloud'. But what if a little intel gets stolen along the way? Are we sure our apps aren't spying on us?

Anecdotal evidence, such as the flap about the Facebook Messenger app collecting audio and video recordings without the user's permission, indicates our phones may be used against us more often than we realise. "For the average user, the main threat to them is going to be getting fooled into installing some application thinking it's useful, but it's doing stuff that they didn't anticipate," says Professor David Lie, Canada Research Chair in Secure and Reliable Computer Systems. "A lot of people are worried about this, but no one can answer concretely whether it's happening."

Professor Lie's research group has just launched a project with TELUS to find out for sure. He and his students work on software to improve data security and privacy, both on your device and in the cloud. There's an important distinction to be made between security and privacy, says Professor Lie. Security is like locking a safe—protecting all information from disclosure absolutely, and making sure that information can't be accessed without authorization. Privacy is a more complex problem, especially given the ubiquity of smartphones: we want to share everything, but only in specific ways, with specific people, at specific times. "The real problem is that people want to share their information, but have no way of understanding what will happen if they do," says Professor Lie.

Perhaps the worst data to give away is your biometric information—your fingerprints, retina scans or the unique signature of your heartbeat. Professor Dimitrios Hatzinakos is a leader in the field of medical biometrics and chair of the Identity, Privacy and Security Institute at the University of Toronto, a collaboration between The Edward S. Rogers Sr. Department of Electrical & Computer Engineering and the Faculty of Information. His team searches for new biometrics in the many electrical signals emitted by the human body—unique waves radiating from your brain, bouncing off your eardrum, and given off every time your heart beats.

His former PhD student Foteini Agrafioti teamed up with ECE alumnus Karl Martin to found Bionym, a startup company based on their research on biometrics and security. Bionym recently released the world's first wearable authentication system, called Nymi, to much acclaim. Nymi is a bracelet embedded with an electrocardiogram (ECG) sensor that recognizes the unique and unchanging electronic signal of your heart and uses it to identify you to all your registered devices to log you in, eliminating the need for passwords and PINs. Bionym has top-notch security, but if your unencrypted ECG were ever compromised, it's not as easy to fix as resetting your password.

"You have a finite amount of biometric data," says Professor Stark Draper. "You only have 10 fingers, as opposed to an unlimited number of passwords or credit card numbers."

That's why Professor Draper and master's student Adina Goldberg are trying to find the optimal balance between privacy and security in linked biometric systems. Imagine that your apartment complex, gym and office all use biometric data to authenticate your identity. If all three systems store the same information and an imposter hacks your gym account, they'll also be able to get into your apartment building and office, but they won't gain new information about you if they do. That's bad security, but relatively good privacy. Conversely, if each account stores a unique piece of biometric data, it's harder to gain access to all three, but each time the imposter hacks a new account they gain more of your sensitive data.

"We've seen a trade-off between the two," says Professor Draper. "Tighter security is of interest to the institution that doesn't want to get broken into, but the individual might think it's more important to keep more of their biometrics private and it's OK if someone breaks into the gym. The system designer gets to pick a point on that curve."

This tension between the individual and the institution was cast in a new light last year, when whistleblower Edward Snowden alerted the world to the U.S. National Security Agency's habit of mining personal data directly off servers run by Google, Facebook, Microsoft and others.

"I think for a lot of people in the field, the Snowden leaks came as no surprise," says Professor Lie. "When I first started doing this cloud stuff, I got a lot of pushback—the logic was, 'Why would the cloud provider attack their own customers? It doesn't make business sense.' Now we can see that they're not attacking their customers directly, but the data is still vulnerable." One of Professor Lie's projects aims to solve the problem of gaining logs from a multi-tenant server—allowing users to see exactly how their information is being accessed, without violating the privacy of all the other users on the same service.

If you don't have the savvy to interpret server logs but are keen to protect your privacy, taking a critical look at your apps is a good place to start. Professor Lie, for his part, installed a program on his smartphone to prevent applications from talking to the network unless he explicitly gives them permission—not a solution he recommends for everyone. If he does give permission, he reads the privacy policy first. "But sometimes it's just an exercise in futility," he admits. "Most of them are pretty impenetrable, or vague about what they do with the information they collect."

So sleep, eat, jog and bank on your smartphone with caution—with friends like these, who needs enemies? A

The real problem is that people want to share their information, but have no way of understanding what will happen if they do.

Undergraduate Study

TRAINING THE NEXT PROBLEM SOLVERS

Electrical and computer engineering power most of today's technological advances. An undergraduate degree from The Edward S. Rogers Sr. Department of Electrical & Computer Engineering opens doors to any career imaginable, from software development to medicine, and everything in between.

The first two years of the ECE program are common to both electrical and computer engineering students, and provide essential background in basic science and mathematics, as well as introducing important concepts in electrical and computer engineering such as circuits, digital systems, electronics and communication systems.

In third and fourth year, students choose areas of specialization depending on their individual strengths and interests. Options include photonics and semiconductor physics; electromagnetics; analog and digital electronics; communications; systems control and biomedical engineering; computer hardware; and computer software.

Our flexible curriculum is the preferred choice for students seeking the opportunity to create their own unique paths of study, and to work alongside professors who are world-renowned experts in their fields. Learn more about the flexible curriculum at **uoft.me/ececurriculum** A

15.3:1
Undergraduate student-to-faculty ratio in ECE—one of the lowest in U *of* T Engineering

STUDENTS' DESIGN PROJECTS FEATURED ON SPACE CHANNEL

Fourth-year students made their debuts on TV screens across Canada this spring, when the Space Channel show *Innerspace* dropped in to film stories about their innovative design projects.

Julie Hsiao explains her team's MySight project, which allows two viewers to watch two different programs on the same screen simultaneously.

Harsoveet Singh and Calvin Ma explain their Simultaneous Localization And Mapping (SLAM) rover to Innerspace host Morgan Hoffman. SLAM won the 2014 Gordon R. Slemon Design Award.

Team Roar, made up of Hongjing Yang, Junqian Zhang, Ming Fei Wang and Hong Shen, shows off its gesture-controlled flying robot with Innerspace host Morgan Hoffman.

Ryan Mintz demonstrates his team's mind-controlled robotic arm.

PEY Profile

Emily Miao (ElecE 1T5+PEY)
Digital Design Engineer Intern at Intel Custom Foundry, Intel of Canada, Toronto Office

Q: WHY DID YOU CHOOSE PEY?

A: I had multiple reasons to go on PEY—for students like myself who have no prior job experience, the full-year technical work experience helps build up our resumes and really makes a difference when we look for positions after graduation. Another reason is that ECE is such a broad field, and I wasn't sure what I wanted to do after graduation—whether I wanted to go to grad school or go to work in industry, and which field within ECE I wanted to pursue. I conducted summer research [with Professor Jason Anderson] after second year, and that gave me a glimpse of what grad school is like. But I wanted to use PEY as an opportunity to experience working in industry.

Q: WHAT ARE SOME OF YOUR KEY RESPONSIBILITIES ON THE JOB?

A: I work on digital verification. It consists of testing a digital design, making sure it functions the way it's supposed to. This includes running simulations to see if a design is working or not, and if not, I debug the design to see why it's not working the way it's supposed to. In addition, when we add a new feature to a design it needs to be tested, and I write the code to test the new feature.

Q: IS THERE A MAJOR PROJECT OR MILESTONE YOU'RE WORKING TOWARD? IF SO, WHAT IS IT?

A: It's company policy that I can't say what exactly I'm currently working on—but I can say that the work I'm doing is the same work that full-time employees do. I think this is a big difference between four-month internships and year-long internships—over the year, we're able to learn the technical skills needed for the job and to apply those skills to our work; we're treated like full-time employees with basically the same expectations and responsibilities. However, I'm still a student, so sometimes I need some help with work, encounter things I don't understand, or may take longer to do certain tasks than an employee who's been doing this for years. My team is very understanding and patient in teaching me new things and helping me when I need it.

Q: WHAT'S ONE THING YOU'VE LEARNED ABOUT INDUSTRY OR WORKING LIFE THAT YOU'LL BRING BACK TO SCHOOL?

A: One thing I'll take with me into fourth year is new problem solving skills—how to think outside the box to solve a problem. I think one major difference between academia and industry is that in academia, the assignments and labs we receive have a solution—they have been solved by TAs or professors before being given to the students to do. It's entirely different in industry: there is no known solution to identified problems. The problems I've encountered in this position are also much more complex and on a larger scale than the labs we see as students; most of the time in industry, there isn't an obvious solution to a problem, and we need to think creatively to invent solutions. Working at Intel has taught me to take a step back when I encounter a problem—to get the big picture first and examine the problem from all angles before digging for a solution.

PEY: MORE THAN A CO-OP

The Professional Experience Year (PEY) internship program at the University of Toronto allows students to apply their engineering education in a 12- to 16-month project-based professional internship. The duration of the work placement lets students immerse themselves in large-scale projects, build relationships with employers and accomplish significant milestones in the workplace. Running for more than 20 years, the PEY program has earned a sterling reputation in both industry and academic circles. In recent years, demand has exceeded supply: more employers have sought to hire PEY students than there were students in search of positions. In 2014–2015, 226 ECE students were working in PEY placements at 75 companies located across Canada and around the world. To learn more, visit **uoft.me/ece-pey**

ECE PEY PLACEMENTS, 2010–2011 TO 2014–2015

	CompE	ElecE	TOTAL
2010–2011	66	127	193
2011–2012	71	134	205
2012–2013	79	102	181
2013–2014	77	133	210
2014–2015	70	156	226

UNDERGRADUATE ENROLMENT, 2010–2011 TO 2014–2015

	PART TIME	FULL TIME	TOTAL
2010–2011	242	1057	1299
2011–2012	256	1096	1352
2012–2013	224	1142	1366
2013–2014	259	1133	1392
2014–2015	272	1208	1480

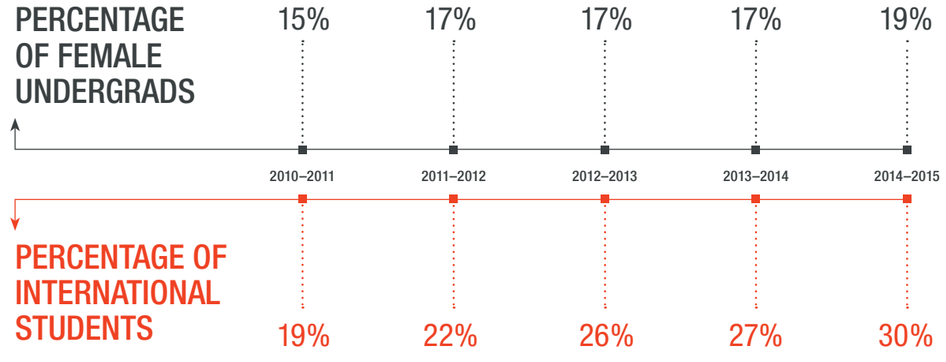

Graduate Study

AN INTERNATIONAL DESTINATION

Throughout its 105-year history, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering has borne witness to groundbreaking discoveries and developments in almost every area of electrical and computer engineering—at the core of these advances is the outstanding research conducted by graduate students and faculty members.

Graduate students in ECE choose from a wide variety of research areas including biomedical engineering, communications, computer engineering, electromagnetics, electronics, energy systems, photonics and systems control. They are supervised by our faculty of 78 professors, many of whom are internationally recognized leaders in their fields.

We offer three degrees: Master of Engineering (MEng), Master of Applied Science (MASc) and Doctor of Philosophy (PhD). For program and admissions information, visit **uoft.me/gradadmission**

OUR GRADUATE STUDENTS BY FIELD OF STUDY

BIOMEDICAL ENGINEERING // **12**
COMMUNICATIONS // **102**
COMPUTER ENGINEERING // **100**
ELECTROMAGNETICS // **31**
ELECTRONICS // **58**
ENERGY SYSTEMS // **46**
PHOTONICS // **60**
SYSTEMS CONTROL // **17**

Not Your Average Robots

We live in a thrilling time for robotics and mechatronics—researchers in both academia and industry are coming together to develop the next generation of human helpers.

At University of Toronto's Institute for Robotics & Mechatronics (IRM), we're developing intelligent solutions for society, from healthcare assistants to planetary rovers. Today's robots are more intelligent, flexible, modular and adaptive to a wide variety of applications including medicine and healthcare, rescue and exploration, military and security, education, entertainment, and manufacturing.

IRM was established by the Faculty of Applied Science & Engineering in 2010 as a central hub for multidisciplinary robotics and mechatronics innovation, uniting more than 45 faculty members across six departments and institutes, including ECE. IRM brings together researchers, students and industry partners to conduct a wide variety of exciting research, spanning fields from artificial intelligence to space robotics.

To learn more about research at IRM, including industry collaboration, undergraduate and graduate studies, visit **irm.utoronto.ca**

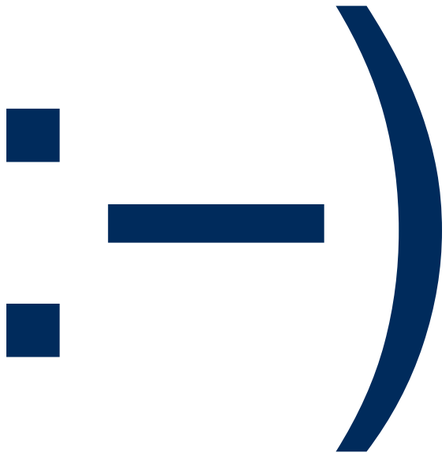

Engineering Happiness

There's no denying that engineering can be stressful—but the experience doesn't have to be all academics, anxiety and all-nighters. Four ECE graduate students banded together to create eHappy, a new student-run club to promote a simple concept: happiness and well-being are important in life.

Through talks, social events and a photo contest, founders Suthamathy Sathananthan (MASc 1T4), Mariya Yagnyukova (MASc 1T3), Amber Houle (MEng 1T3), Rebecca Vaddadi (MEng 1T3) and chair Pisek Kultavewuti (PhD 1T5) brought a few rays of sunshine to engineering students. The club is now planning to stage a Ball Pit of Fun, which is exactly what it sounds like—a big box filled with colourful plastic balls that bring out the fun-loving child in everyone who climbs in. No one leaves the box without a grin and a new friend!

$2.3 million
Graduate student scholarships awarded annually in ECE

ARTHUR MONTAZERI WINS $50,000 WESTON FELLOWSHIP

Arthur Montazeri spent years bouncing between studies in math, chemistry, programming and physics before landing in the one field he had never considered: ECE. Now a PhD candidate in the Advanced Photovoltaics and Devices Group under the supervision of Professor Nazir Kherani, Montazeri was named one of just four 2014 Weston Fellows from the University of Toronto. Montazeri was chosen for his multidisciplinary work uniting elements of photonics, electro-magnetics, nanoengineering and medicine. One of Canada's most prestigious awards, the $50,000 prize allows Weston Fellows to collaborate with scholars anywhere in the world. He will spend the next year at the University of California, Berkeley, advancing his work on converting infrared radiation to electricity and improving medical imaging.

Enrolment in Research-stream Programs 2014–2015

MASc 173 | PhD 253

TOTAL 426

TOTAL GRADUATE ENROLMENT, 2010–2011 TO 2014–2015 (HEADCOUNT)

	FALL 2010	FALL 2011	FALL 2012	FALL 2013	FALL 2014
MASC	146	153	174	167	173
MENG	67	114	155	140	169
PHD	218	212	236	249	253
TOTAL	431	479	565	556	595

GRADUATE DEGREES AWARDED, 2009–2010 TO 2013–2014

	2009–2010	2010–2011	2011–2012	2012–2013	2013–2014
MASC	79	61	57	56	70
MENG	33	43	25	56	85
PHD	32	39	34	37	27
TOTAL	144	143	116	149	182

Research

A CONCENTRATION OF POWERFUL MINDS

Researchers in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering tackle challenging fundamental and applied problems of strategic importance to global industry and society. We work across and beyond disciplinary boundaries to advance current understanding in key emerging areas such as smart grid technology, mobile and cloud computing, and biomedical engineering.

Our department alone occupies more than 4,400 square metres of research laboratory space across six buildings on the University of Toronto's downtown St. George Campus. We also house the Toronto Nanofabrication Centre, home to several state-of-the-art nanofabrication facilities that are available to both academic and industry clients for open-access prototype development and testing.

As one of the largest ECE departments in Canada, we are a research powerhouse. Our professors partner with countless industry leaders worldwide to stimulate, enhance and translate our research into application. We continue to seek opportunities to collaborate with industry, government and other academic institutions to improve quality of life in Canada and around the world.

$263K – AVERAGE ANNUAL RESEARCH OPERATING FUNDING PER ECE FACULTY MEMBER

Our Faculty by Research Group

BIOMEDICAL ENGINEERING **7**
COMMUNICATIONS **15**
COMPUTER ENGINEERING **21**
ELECTROMAGNETICS **5**
ELECTRONICS **9**
ENERGY SYSTEMS **7**
PHOTONICS **9**
SYSTEMS CONTROL **5**
EMERITUS **21**

TOTAL 99

SELECTED AWARDS WON BY ECE FACULTY, 2013–2014

THE ROYAL SOCIETY OF CANADA, FELLOW // **PROFESSOR TED SARGENT**
THE ROYAL SOCIETY OF EDINBURGH, CORRESPONDING FELLOW // **PROFESSOR J. STEWART AITCHISON**
INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, FELLOW // **PROFESSOR WEI YU**
CANADIAN ACADEMY OF ENGINEERING, FELLOW // **PROFESSOR ALBERTO LEON-GARCIA**
UNIVERSITY OF TORONTO FACULTY AWARD // **PROFESSOR JONATHAN ROSE**
FACULTY OF APPLIED SCIENCE & ENGINEERING SUSTAINED EXCELLENCE IN TEACHING AWARD // **PROFESSOR GLENN GULAK**
ENGINEERING INSTITUTE OF CANADA, FELLOW // **PROFESSORS PAUL CHOW, SHAHROKH VALAEE**
THOMSON REUTERS HIGHLY CITED RESEARCHERS LIST OF MOST INFLUENTIAL SCIENTISTS // **PROFESSOR WEI YU**

PROFESSOR NATALIE ENRIGHT JERGER WINS PEO ENGINEERING MEDAL

Already known as one of the top computer architecture researchers in the field, Professor Natalie Enright Jerger received the 2014 Professional Engineers Ontario (PEO) Engineering Medal – Young Engineer. She is known for her outstanding contributions to the critical areas of interconnection networks and parallel architectures, and is widely recognized as a leading expert in the area of on-chip networks, an active and vibrant area in computer architecture driven by the dramatic shift away from increasing processor speeds to replicating processor cores. A pressing problem in the community is how to connect an increasing number of cores in an energy-efficient way. Professor Enright Jerger directly addresses this challenge through innovations in on-chip networks to provide scalable communication and better performance. She is also a co-author of the widely used textbook On-Chip Networks. A dedicated educator, Professor Enright Jerger has redesigned the undergraduate fourth-year computer architecture course and introduced a new graduate course about on-chip networks. She is a ardent advocate for increased gender diversity in STEM disciplines of science, technology, engineering and mathematics, and dedicates a great deal of time to this cause.

Introducing: Hai-Ling Margaret Cheng

Our newest faculty member joined the Institute of Biomaterials & Biomedical Engineering (IBBME), with a cross-appointment to ECE, in July 2014. She comes to U of T Engineering from The Hospital for Sick Children where she worked as a scientist in the Research Institute. Marit Mitchell sat down with Professor Cheng to talk about her career, research and striking a balance between science and art.

Q: WELCOME! WHY DID YOU CHOOSE TO JOIN ECE AND IBBME FROM SICKKIDS?

A: Prior to coming to the Faculty, I was engaged in research 100 per cent of the time. That's one thing I'm really passionate about—I love what I do, but even as an undergraduate I always wanted to become a professor. I love teaching—I know, some people find that hard to believe—but I feel it's an obligation to pass on my knowledge to the next generation and inspire them to have curiosity about things around them, inspire them to pursue science for the love of knowledge. You only get to do that if you're a professor on campus. You wouldn't get that opportunity as a scientist in a research institution—and that was my primary motivation.

Q: YOU MADE ANOTHER CAREER MOVE EARLIER ON, TO PIVOT FROM AEROSPACE TO BIOMED. HOW DID THAT COME ABOUT?

A: After I finished my master's degree, I decided that I didn't really have a clear vision of what I wanted to do for my PhD, and at that point I needed to give myself some time to explore industry, and to see if that meshed better with my personality than academia. So I spent two years working in industry, developing software interfacing for surveillance for defence aircraft. It was challenging work, but something was missing. I got up in the morning and went to bed at night thinking, 'I'm not really contributing anything to society.'

My sister [Hai-Ying Cheng, now a professor at University of Toronto Mississauga] was doing her graduate degree at the time in medical biophysics, and she encouraged me to enrol in the same department, but in the physics stream, where all the engineering geeks and physics geeks were! It was the best choice. I could apply my engineering knowledge and apply physics to medical imaging to solve different healthcare problems—I never regret that day.

Q: WHAT ARE YOU WORKING ON NOW, AND WHERE ARE YOU GOING FROM HERE?

A: At the hospital I was working on magnetic resonance imaging technology to look at cancer and tissue regeneration, but within IBBME there's Michael Sefton, Milica Radisic, and other experts in the field of tissue engineering. My goal is to advance MRI, non-invasive imaging, to enable advances in tissue engineering. So the fact that I'm here makes a lot of sense and gives my research program, and hopefully my colleagues', a lot of synergy.

The electrical engineering is also beneficial, because now I can start to work with graduate students and colleagues on the ECE side to develop hardware—better imaging coils—to look at small samples or animals on a human scanner. It opens up new research avenues for me.

Q: WHAT DO YOU PURSUE WHEN YOU'RE NOT THINKING ABOUT RESEARCH?

A: I love art. When I was a little girl, and up until graduate school, in my spare time I would just sit and draw. And I love reading. When I read, I don't read science books. I read history, I read about religion, I read about different cultures. I find that it enriches my mind in a different sphere. I cannot have science 100 per cent of the time. It also feeds the spiritual needs—to have those human aspects in your life is really important. That's another reason why I like education—educating not just in terms of knowledge, but in terms of the whole person.

ECE INVENTION DISCLOSURES, 2010–2014

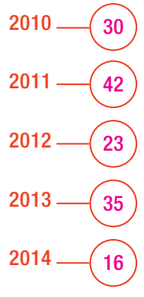

The University of Toronto's Invention Policy requires inventors to completely disclose their inventions. In most cases, rights to their inventions are then jointly owned by the University and the inventor.

ECE RESEARCH FUNDING 2008–2009 TO 2012–2013

	FEDERAL	PROVINCIAL	INDUSTRY	OTHER	TOTAL $
2008-2009	7,898,653	4,862,786	2,709,325	2,876,148	**18,346,912**
2009-2010	8,598,523	3,948,865	1,481,958	2,917,487	**16,946,833**
2010-2011	8,914,271	3,201,687	2,228,062	3,518,860	**17,862,880**
2011-2012	10,137,402	1,569,027	2,765,890	3,469,600	**17,941,919**
2012-2013	9,713,015	2,482,000	2,304,819	3,377,035	**17,876,869**

ECE is the top U *of* T department over the past five years for new invention disclosures filed, new license and option agreements executed, and new startup companies formed.

NSERC FUNDING FOR ELECTRICAL AND ELECTRONICS ENGINEERING AT LEADING CANADIAN UNIVERSITIES, 2009–2010 TO 2013–2014

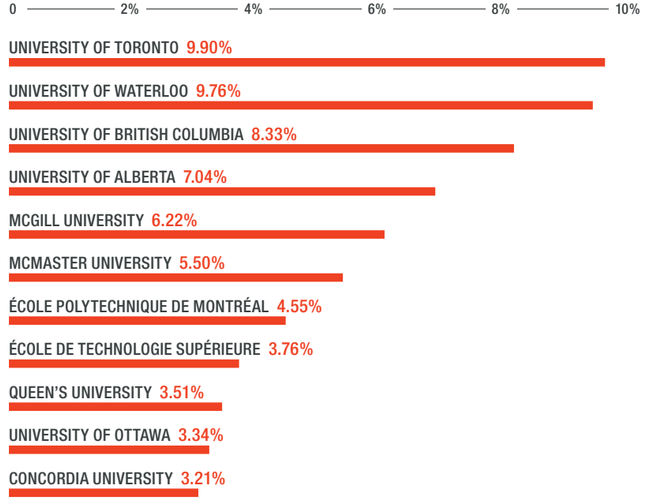

Industry Partnerships

TACKLING REAL-WORLD PROBLEMS

In The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, we pride ourselves on cultivating both longstanding and new partnerships with visionary corporations, foundations and institutes the world over. These ties keep our research relevant, our teaching fresh, and continually inspire us to find multidisciplinary solutions to technological, economic and social problems both here in Canada and abroad.

We enjoy active partnerships with more than 85 industry funders and collaborators. In an effort to open our doors to some of these contacts, as well as to strengthen connections between ECE and researchers across the University, in summer 2014 we launched the first in a series of one-day workshops designed to unite thinkers from a diverse array of fields on a common topic. Our first workshop, 'Tools to Tackle Big Data', was hosted by ECE in July 2014 and saw more than 200 registrants —we welcomed investigators from across Engineering and Computer Science, to Public Health and Dentistry, as well as industry representatives from IBM, Public Health Ontario, Sciencescape and more. Learn more about upcoming events at **uoft.me/eceevents**

ECE typically accounts for about 50 per cent of all U*of*T Engineering startup companies and invention disclosures

WHERE THE MONEY GOES:

OUR 2014-2015 FINANCIAL PICTURE AT A GLANCE

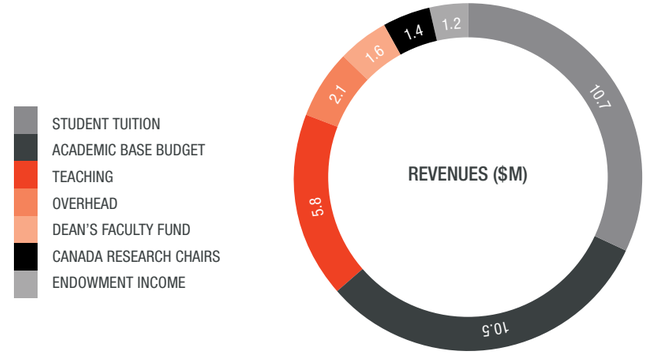

$11.5 million Corporate funding to ECE over the past five years

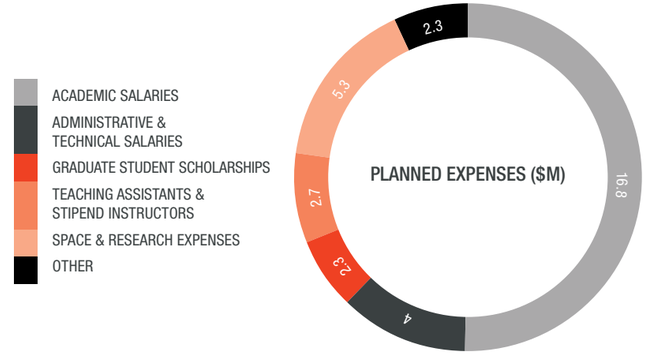

ECE FOUNDS NEW INDUSTRY ADVISORY BOARD

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering is pleased to announce the formation of its first Industry Advisory Board. The board will provide ECE with objective advice about its role in contributing to technological and economic development, and engineering practice. Convened by ECE Chair Farid Najm and Associate Chair, Research Tony Chan Carusone, its founding members include:

PAUL BROWN, DIRECTOR OF DISTRIBUTION ASSET MANAGEMENT // HYDRO ONE
SAI-KIT ENG, DIRECTOR OF PROGRAM MANAGEMENT // QUALCOMM CANADA
ALLEN LALONDE, SENIOR EXECUTIVE // IBM CANADA RESEARCH & DEVELOPMENT CENTRE
SAID MOKBEL, VICE-PRESIDENT, WIRELESS AND TELEPHONE/MESSAGING SERVICES // TELUS
DESHANAND SINGH, SOFTWARE ENGINEERING DIRECTOR // ALTERA
SONG ZHANG, DIRECTOR, TECHNOLOGY PLANNING & PARTNERSHIP // HUAWEI CANADA

226

Number of students on Professional Experience Year industry placements, 2014–2015

Electrical Engineering: 156
Computer Engineering: 70

75 EMPLOYERS HIRING ECE STUDENTS ON PEY:

407 ETR CONCESSION COMPANY LTD
ACUITY ADS
AGFA GRAPHICS NV-BELGIUM
ALTERA CORPORATION
AMD (ADVANCED MICRO DEVICES, INC.)
ANALOG DEVICES INC.
ARUP
AUTOLIV ELECTRONICS CANADA INC.
BELL
BELL MOBILITY
BLACKBERRY LIMITED (FORMERLY RIM)
BLUECAT
BMW GROUP CANADA
CAST SOFTWARE INC.
CELESTICA INC.
CGI GROUP INC.
CIRBA INC
CISCO SYSTEMS INC.
CITY OF BRAMPTON
COMMUNICATIONS AND POWER INDUSTRIES CANADA INC.
DEMONWARE
DESIRE2LEARN
ENVIRONMENT CANADA
EPSON
ESNA TECHNOLOGIES
EVENTMOBI
FU TIN BUSINESS CO. LTD.
GENERAL ELECTRIC CANADA
GENERAL MOTORS OF CANADA LTD.
GENESYS CANADA LABORATORIES INC.
HUSKY INJECTION MOLDING SYSTEMS LTD.
HYDRO ONE
IBM
INTEL CANADA
KNOWROAMING
LOBLAW COMPANIES LTD.
MARIN SOFTWARE
MARVELL TECHNOLOGY GROUP
MAXXIAN
MICROSEMI CORPORATION
MINISTRY OF CHILDREN AND YOUTH SERVICES
MINISTRY OF TRANSPORTATION
MODIFACE
MOTOROLA SOLUTIONS/PSION
NEXJ SYSTEMS INC.
NVIDIA
ONTARIO MINISTRY OF HEALTH & LONG-TERM CARE
OPG (ONTARIO POWER GENERATION)
PALANTIR TECHNOLOGIES INC.
PERSPECSYS
PHILIPS (CHINA) INVESTMENT CO., LTD.
QUALCOMM
RED HAT CANADA LTD.
ROYAL BANK OF CANADA
SAP CANADA INC.
SCOTIABANK
SEMTECH
SIEMENS
SMITH AND ANDERSEN CONSULTING ENGINEERING
SOUTHPAW TECHNOLOGY INC.
SUNCOR ENERGY
SURFEASY
TECK RESOURCES LTD.
TELUS
THALES RAIL SIGNALLING SOLUTIONS
THE HOSPITAL FOR SICK CHILDREN
THE INDEPENDENT ELECTRICITY SYSTEM OPERATOR (IESO)
TORONTO HYDRO
TRANSCANADA
TRAPEZE GROUP
TUNNELBEAR INC.
TXIO
TYCO ELECTRONICS CANADA LTD.
XAGENIC INC.

HOW TO MULTIPLY YOUR MONEY

There has never been a more exciting time to unite with ECE at the University of Toronto: our industry partners enjoy an effective multiplier of up to 5x their spending on joint research projects with our department. Cash and in-kind contributions by industry partners can be used to leverage additional cash investments from the provincial and federal governments, and from the University. Contact us to explore the boundless potential of research and technology development.

JASON CHANG, MBA
DIRECTOR OF FOUNDATION & CORPORATE PARTNERSHIPS
JASON.C.CHANG@UTORONTO.CA | 416 978 7890

MARCIUS EXTAVOUR, PHD
DIRECTOR OF GOVERNMENT & CORPORATE PARTNERSHIPS
MARCIUS.EXTAVOUR@UTORONTO.CA | 416 978 6990

WITHIN U OF T ENGINEERING, WE LEVERAGE INDUSTRY FINANCIAL SUPPORT WITH FEDERAL, PROVINCIAL AND INSTITUTIONAL MATCHING TO ENABLE $25 MILLION IN INDUSTRY-PARTNERED RESEARCH PER YEAR. A $1.40 INDUSTRY CASH OR IN-KIND CONTRIBUTION GROWS TO $7.60, MORE THAN 5X LEVERAGE ON INITIAL INVESTMENT.

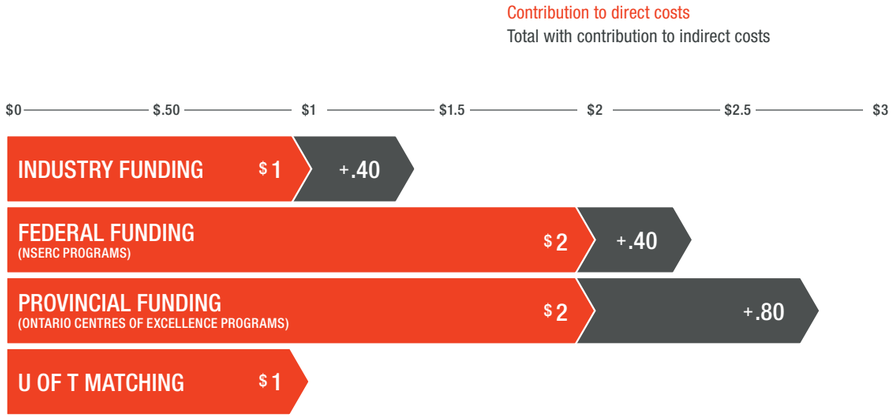

$6
Total Contribution to *direct costs*

$1.60
Total Contribution to *indirect costs*

=

$7.60
Leveraged *5x*

The outcome: a joint project with the support of world-renowned researchers, expert administration, project management and motivated graduate students.

Alumni

A STRONG INTERNATIONAL NETWORK

With more than 11,000 members, the ECE alumni community is a powerful global network. Our graduates are leaders in fields ranging from medicine to mining, entrepreneurship to entertainment. ECE alumni are extremely generous with their time, and many volunteer to help current students—whether through participation in the Engineering Alumni Mentorship Program, by meeting fourth-year students on the brink of graduation at our spring Fourth-Year & Alumni Reception, or by sharing stories of startup successes and failures as part of our Engineering Entrepreneurship Speaker Series.

We value input from our alumni on our evolving priorities and direction, and we like to know where our graduates are now! If you have a piece of news you'd like to share, thoughts on our performance or would like to hear more about how to get involved with your fellow alumni or current students, we'd love to hear from you.

Contact Senior Communications Officer Marit Mitchell at **eceinquiry@utoronto.ca** or 416 978 7997. A

Total number of active alumni worldwide:

11,165

Make a Donation to ECE

ECE alumni developed the skills and network they needed for success, as well as lifelong friendships, at the University of Toronto. Many remember this by contributing either time or money to projects meaningful to them. Our donors give back to the department by creating scholarships to allow future undergraduates to focus on their studies, establishing research chairs that free our professors to tackle important problems, or supporting vital capital projects, such as the renovation of ECE's energy systems laboratories or the new Centre for Engineering Innovation & Entrepreneurship building.

There are many ways to give to U of T Engineering's Boundless campaign. Donations can be made by way of cheque, credit card, gifts of securities or other planned gifts such as naming the University as an estate beneficiary.

For more information on how you can support The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, please visit **boundless.utoronto.ca/how-to-give** or contact Celeste Taylor at **celeste.taylor@utoronto.ca** or 416 978 0291.

OTHER BOARD MEMBERS ARE:

ALAN BOYCE
ElecE 7T8

JOHN EAST
ElecE, MBA UC Berkeley

ALEX GRBIC
CompE 9T4, MASc 9T6, PhD 0T3

CATHERINE LACAVERA
CompE 9T7, JD/MBA U of T

ALEX SHUBAT
ElecE 8T3, MASc 8T5, PhD Santa Clara, MBA Stanford

ECE ALUMNI BOARD OF ADVISORS:

WELCOME NATASHA LALA

Formed in 2013, the ECE Alumni Board of Advisors meets several times a year to advise the department on its alumni engagement and advancement activities, and help inform its strategic direction. This year we welcomed Natasha Lala to the board. Ms. Lala graduated from ECE with a BASc in Computer Engineering, class of 9T8. She quickly discovered an aptitude for bridging the gap between technology and business teams. She was soon leading teams in the successful execution of complex technology programs all over the world for the CRM, telecom, and financial industries with SOMA Networks and IBM.

Ms. Lala is currently chief of staff at OANDA Corporation, a company that uses innovative computer and financial technology to provide Internet-based foreign exchange trading and currency information services. Since joining OANDA in 2003, she has held several leadership positions within the company, including vice president of engineering, where she was responsible for the program and teams that build the product for a high-frequency trading system that transacts billions of dollars a day globally. Now based in San Francisco, Ms. Lala is passionate about engaging women to pursue careers in technology.

DAVID HE NAMED TO MIT TECHNOLOGY REVIEW'S INNOVATORS UNDER 35

Alumnus David He (ElecE 0T5) found himself in good company this summer when MIT Technology Review named him to its 2014 Innovators Under 35 list. The biotech entrepreneur co-founded startup Quanttus out of his PhD research at MIT. "We are still in stealth mode," he said of the company, which aims to transform personal health by providing people with detailed data and insights about their body from a wearable device.

"This could be the first time that we have gained access to an unprecedented amount of continuous vital signs data in real-life settings," he said. "These data may lead to new discoveries about the heart and how our lifestyles affect our bodies." He credits three ECE professors with setting him on his current path. "Professors Andreas Veneris and Khoman Phang introduced me to the pursuit of academic research and the joy of creating something that has not been done before. That carries with me to this day. Professor Berj Bardakjian taught the class 'Cellular Bioelectricity', which inspired me to explore the multidisciplinary field between human physiology and electrical engineering."

FALL ALUMNI NETWORKING LECTURE & RECEPTION

ECE alumni from all years returned to Skule to hear a talk from Professor Steve Mann on the evolution of the digital eyeglass system he invented, and 35 years at the forefront of wearable computing.

PHOTOS BY WILLIAM YE

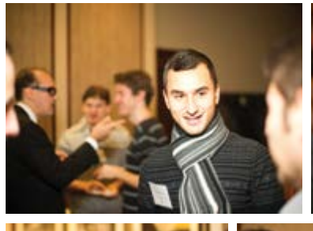

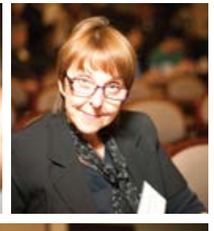

SPRING REUNION

Sixty-five years of ECE history was represented at this year's Spring Reunion, where we welcomed back graduates from the classes of 4T9 all the way to 1T4 for lunch, lab tours, and a stress-free lecture from Professor Olivier Trescases. Special thanks to alumnus Grant Allen (ElecE 4T9) for travelling all the way from Calgary to join in. Mark your calendars fo r the next Spring Reunion lunch on Saturday, May 30, 2015!

PHOTOS BY MICHAEL TENAGLIA

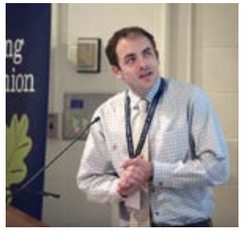

FOURTH-YEAR AND ALUMNI CELEBRATION

Each spring, on the same day fourth-year students finish their final capstone design projects, ECE hosts a reception to celebrate its imminent graduates. This year Arshia Tabrizi (CompE 9T5, Law 9T8) shared his congratulations and welcomed the fourth-years into the alumni community they would soon join.

PHOTOS BY JIMMY LU

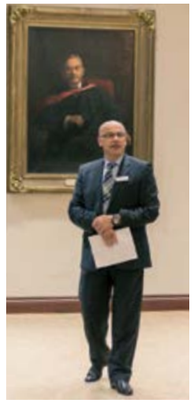

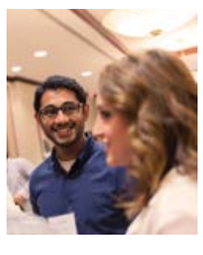

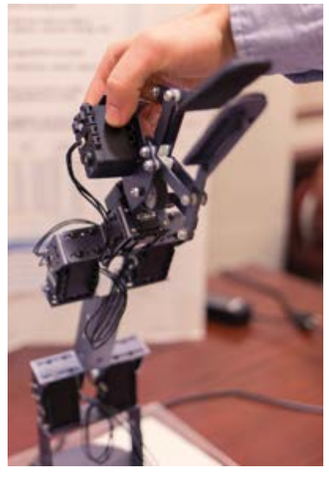

The Edward S. Rogers Sr. Department
of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

Distinguished Lectures Series
2014 - 2015

Bruno Sinopoli
Carnegie Mellon University
Sept. 11, 2014

On the Security of Cyber-Physical Systems

Shanhui Fan
Stanford University
Sept. 18, 2014

Nanophotonics for Energy and Information Applications

Garth Gibson
Carnegie Mellon University
Oct. 9, 2014

ML Systems: Experience with Machine Learning as a Distributed Systems Problem Space

Mustafa Khammash
ETH Zurich
Oct. 30, 2014

Cybergenetics: Feedback Control of Living Cells at the Gene Level

Rudolf Kalman
ETH Zurich & University of Florida, Gainesville
Nov. 27, 2014

How Old Mathematics and Old Engineering Became New Engineering and New System Theory

Gregory Wornell
Massachusetts Institute of Technology
Jan. 22, 2015

On Digitally Enhanced Antenna Array Architectures and the Possibilities of Intelligent Co-Design

Dustin Tyler
Case Western Reserve University
Feb. 26, 2015

Touching the World Again: Neural Interfaces Restore the Sense of Touch and Position Following Limb Loss

Lihong Wang
Washington University in St. Louis
March 26, 2015

Photoacoustic Tomography: Ultrasonically Beating Optical Diffusion and Diffraction

Ali Hajimiri
California Institute of Technology
April 9, 2015

1+1= 3, or How I Learned to Love Holistic Circuits

All lectures take place at **3:00 p.m.** in room 1105 of the Sandford Fleming Building.

uoft.me/ece-dls

DIRECTORY

The Directory section contains a listing of the research projects currently underway in ECE.

DIRECTORY

The Directory section contains a listing of the research projects currently underway in ECE.

The Edward S. Rogers Sr. Department
of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

ECE RESEARCH
DIRECTORY

THE EDWARD S. ROGERS SR. DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING
10 KING'S COLLEGE ROAD, TORONTO ON M5S 3G4 / PUBLICATION MAIL AGREEMENT NUMBER: 42887022

COMMUNICATIONS

Communications Equipment
Communication Systems and Services: Planning, Organization, Services
Communications Technologies: Satellites, Radar

COMPUTERS

Computer Communications
Computer Software/Hardware

ENGINEERING/SCIENCES

Aerospace
Electrical and Electronic Machinery, including Computer Hardware
Ground: Road and Rail
Instrumentation Technology and Equipment, Computer Hardware
Materials Sciences
Mathematical Sciences
Physical Sciences

INFORMATION TECHNOLOGY

Information and Communications Services
Information, Computer and Communication Technologies
Information Systems and Technology

ENERGY

Alternative Energy Resources
Electrical Energy
Energy Efficiency
Energy Resources: Production, Exploration, Processing, Distribution and Use
Energy Storage and Conversion

LIFE SCIENCES

Advancement of Knowledge
Biomedical Engineering
Human Health
Life Sciences, including Biotechnology
Medical Equipment and Apparatus

AARABI, PARHAM				•		
ABDELRAHMAN, TAREK				•		
ADVE, RAVIRAJ				•		
AITCHISON, STEWART			•			
AMZA, CRISTIANA				•		
ANDERSON, JASON		•		•		
BALMAIN, KEITH			•			
BARDAKJIAN, BERJ			•			•
BETZ, VAUGHN		•		•		
BROUCKE, MIREILLE			•			
BROWN, STEPHEN		•		•		
CHAN CARUSONE, TONY	•	•		•		
CHENG, HAI-LING MARGARET						•
CHOW, PAUL		•		•		
DAVISON, EDWARD			•	•	•	
DAWSON, FRANCIS					•	
DRAPER, STARK	•	•		•		
ELEFTHERIADES, GEORGE			•			
ENRIGHT JERGER, NATALIE		•				
FRANCIS, BRUCE				•		
FREY, BRENDAN	•					•
GENOV, ROMAN	•					•
GOEL, ASHVIN		•		•		
GULAK, GLENN	•			•		
HATZINAKOS, DIMITRIOS	•		•	•		•
HELMY, AMR S.			•	•		
HERMAN, PETER			•			
HUM, SEAN	•			•		
IIZUKA, KEIGO			•			•
IRAVANI, REZA					•	
JACOBSEN, HANS-ARNO				•		
JOHNS, DAVID				•	•	
JOY, MIKE						•
KHERANI, NAZIR			•		•	
KHISTI, ASHISH	•					
KSCHISCHANG, FRANK	•					
KUNDUR, DEEPA	•				•	
KWONG, RAYMOND			•	•		•
LEHN, PETER					•	
LEON-GARCIA, ALBERTO	•	•	•	•		
LEVI, OFER			•			•
LI, BAOCHUN	•			•		

	●	■	▲	⬢	▰	⬟
LIANG, BEN	•	•		•		
LIE, DAVID		•				
LIEBEHERR, JORG	•					
LISCIDINI, ANTONIO				•		
LO, HOI-KWONG	•					
MAGGIORE, MANFREDI	•		•			
MANN, STEVE		•	•	•		•
MOJAHEDI, MO			•			
MOSHOVOS, ANDREAS	•	•		•		
NACHMAN, ADRIAN			•			•
NAJM, FARID		•				
NG, WAI TUNG		•	•		•	
PAVEL, LACRA			•			
PHANG, KHOMAN				•		
PLATANIOTIS, KONSTANTINOS N. (KOSTAS)	•	•		•		•
POON, JOYCE	•		•	•		
PRODIC, ALEKSANDAR			•		•	
QIAN, LI	•		•	•		
ROSE, JONATHAN		•				•
SARGENT, EDWARD			•		•	•
SARRIS, COSTAS				•		
SCARDOVI, LUCA			•			•
SHEIKHOLESLAMI, ALI	•		•	•		
SMITH, PETER W. E.			•	•		
SOUSA, ELVINO				•		
STUMM, MICHAEL		•				
TATE, JOSEPH (ZEB)					•	
TAYLOR, JOSHUA					•	
TRESCASES, OLIVIER			•	•	•	
TRIVERIO, PIERO			•	•		•
TRUONG, KEVIN		•				•
VALAEE, SHAHROKH		•		•		
VENERIS, ANDREAS		•				
VENETSANOPOULOS, ANASTASIOS				•		
VOINIGESCU, SORIN		•	•	•		
WONG, WILLY						•
WONHAM, MURRAY			•			
YOO, PAUL						•
YU, WEI	•					
YUAN, DING		•				
ZHU, JIANWEN		•				

Aarabi, Parham

WWW.APL.UTORONTO.CA

⬢ Internet Video, Audio and Image Processing

In the past few years, we have seen an exponential increase in the number of videos and images that have been recorded and placed on the internet. Smart mobile phones (Blackberry, iPhone, etc.) now enable seamless recording, transmission and sharing of videos in near–real time. Whereas just two decades ago there were a few video broadcasters and publishers, today there are millions, if not billions, of video broadcasters and online publishers. With all this visual content, how do we find what we want? How do we categorize the content? How do we develop search engines that bring order to visual content just as text-based search engines (Google, Bing, etc.) brought order to the textual web? With videos and images, there are of course certain tags manually entered by users that define and categorize the video. However, the 10–15 words that usually accompany a video or image can hardly describe the entire content and at best help to generally categorize the video or add a specific note regarding it. In fact, the vast amount of visual information online is untagged and inadequately described and as a result is difficult if not impossible to find. Finding all images and videos is but one problem. The appropriate categorization of visual content can result in more appropriate contextual advertisements (leading to better monetization of visual web/mobile sites). It can help in finding duplicate versions of the same video or image (which is useful for detecting copyright violations, among other applications). It can also help identify the important parts of a video segment or the most content-filled section of an image (which is useful for compressing videos and intelligently resizing images for mobile devices). The more information we can extract automatically from images and videos, the more we can address a range of practical problems including better search, better (and more contextual) monetization, copyright detection, video compression, and intelligent image resizing, as well as a broad range of other important applications. Although different approaches for visual information understanding have been explored in the past, one of the most promising directions is that of utilizing Extremely Large Datasets (ELDs). ELDs allow for greater accuracy in extracting information from images and videos, but in return require a substantially greater number of computations for each image processed. Examples of successful ELD systems include the Tiny Images image categorization system, which used a database of 80 million tiny images for image classification, or the Tiny Videos framework, which utilized a large library of videos for video classification. In both cases, the images and videos were resized to a “tiny” representation in order to minimize the rather large computational load. We aim to extend these research directions by exploring new hardware and software solutions that enable real time image and video searching using large databases. Our goal can be described in the following two phases: (1) algorithm research and development (including finding ways to utilize ELDs for better image and video understanding and improved visual classification accuracy); and (2) hardware acceleration of the developed algorithms in order to enable accurate real time searching of images and videos using ELDs. To summarize, using currently available images and videos that are either tagged or partially tagged, it is possible to develop highly accurate (and computationally demanding) systems that use this information for understanding and classifying vast amounts of untagged images and videos. In turn, the computational load can be addressed through FPGA-based hardware acceleration, which would enable the classification of an image or video to be performed in real time.

Abdelrahman, Tarek

WWW.EECG.UTORONTO.CA/~TSA

⬢ Architectural Support for Parallel Programming

The last few years have witnessed the introduction and then dominance of multicore processors as the mainstream approach to improving performance. Multicores have more than one processor on a single chip. Today’s multicores have four processors on a single chip and industry experts predict that by the year 2020 the number of cores on a chip will exceed 100. However, to exploit the performance potential of multicores, software must be developed to execute on multiple processors. This process, referred to as parallel programming, is a difficult one and is currently one of the major challenges in the field. This problem of parallel programming is exacerbated by the fact that multicores are designed by replicating the hardware of single-core processors. We believe that novel multicore architectures are needed to ease the tasks of parallel programming. The goal of this project is to explore architectural support for parallel programming. One example of such support is our recent work on the memory versioning scheme. In this scheme, hardware is added to each processor to monitor the accesses the processor makes to shared data. The hardware is designed in such a way that processors access shared data in exactly the same order as in a single processor execution. The result is a parallel programming model that requires little effort on the part of the programmer and delivers good performance. This is achieved at the expense of hardware that is dedicated to versioning. We continue to explore efficient implementation of this hardware as well as other novel approaches to supporting parallel programming at the architecture level.

⬢ Compiler Support for GPU Programming

Graphics Processing Units (GPUs) have emerged in the last few years as a promising platform for cost-effective and energy-efficient computing. These highly parallel systems are the core that processes data for graphics-intensive applications, such as computer games on almost every desktop computer. Recently, vendors such as AMD and NVIDIA have provided standard high-level programming interfaces that allow the use of GPUs for more general-purpose non-graphics computation. However, one of the major obstacles still facing GPUs today is the considerable effort needed to program them. This effort is needed to extract and package computations in the form of parallel kernels, to partition computations among threads, to explicitly manage the complex memory hierarchy that exists within the GPU and finally to deal with the various resource constraints that exist inside a GPU. The goal of this project is to develop compiler support that eases the programming difficulty of GPUs. We have defined a directive-based language that provides users with directives allowing them to carry out common tasks such as kernel definition, computation partitioning, data movement, local memory allocation, etc., in the sequential code base. Such an approach

will make GPU programs easier to develop, debug and maintain. We are exploring common optimization patterns performed by GPU programmers and are developing compiler optimization to automate the applications of these patterns.

Dynamic Acceleration of Soft Processors

Soft processors have gained popularity as a means of implementing general-purpose computations on Field-Programmable Gate Arrays (FPGAs) using the familiar sequential programming model. However, soft processors are slow compared to dedicated FPGA designs. In this project, we explore the dynamic acceleration of soft processors using traces. Our approach is to execute a program on a soft processor, detect at run-time hot paths of execution (i.e., traces) in the program and then dynamically synthesize circuits on an FPGA to speed up the execution of these traces. This approach has the advantage of leveraging the considerable resources on an FPGA to match the dynamic characteristics of an application in a transparent way. However, it faces several challenges. They include (1) the efficient detection of traces; (2) the quick synthesis of a trace at (the use of traditional CAD tools is prohibitive); and (3) the dynamic reconfiguration of the FPGA to realize the synthesized trace circuit at run time. Our goal is to address these challenges through a novel overlay architecture that we refer to as the Virtual Dynamically Reconfigurable FPGA (VDR-FPGA). We further use a trace synthesis approach that exploits the fact that traces are short straight-line segments of code, which makes them more amenable to analysis and optimization at run time.

Adve, Raviraj

WWW.COMM.UTORONTO.CA/~RSADVE

Adaptive Signal Processing for Wireless Communications and Radar Systems

Our research takes two directions: (1) adaptive processing in the physical layer in communication systems: exploiting the spatial and temporal dimensions to improve the quality, capacity and reliability of wireless communication systems; enabling co-operation for energy savings in wireless sensor and data reliability and/or access-point networks. (2) signal processing for radar systems: the detection, identification and tracking of tactical targets in stressful interference environments using advanced signal processing techniques.

Aitchison, Stewart

PHOTONICS.LIGHT.UTORONTO.CA/AITCHISON

Nanophotonics for Optical Signal Processing and Sensing

Our research falls within three areas: (1) electron beam lithography and process development; (2) photonic wires for wavelength conversion applications; and (3) photonic wires for optical sensing applications. In 2009, we officially opened our new electron beam lithography system, which allows features down to 10 nm to define across large areas. The high beam current and low stitching errors possible with this tool allow a wide range of structures to be patterned, including nanostructured surfaces for biology, sensing and photonics. Optical frequency conversion, based on second- or third-order nonlinearities, provides a mechanism for generating new wavelengths and has applications in telecommunications for agile channel allocation in a wavelength division multiplexed system and for the generation of mid-IR wavelengths for optical sensing. The use of high-refractive-index contrast waveguides to implement wavelength conversion has many advantages. The small core size increases the local intensity, the waveguide structure can be used to dispersion engineer the waveguide to enable phase matching and resonators can be used to further enhance the conversion efficiency. Under this theme we will use the almost ideal nonlinear properties of the III-V semiconductor AlGaAs to develop efficient wavelength conversion devices based on second-order nonlinearities (difference frequency generation) and third-order effects (four-wave mixing). The ability to engineer the dispersion and field profile in a nanowire waveguide has applications in optical sensing. By narrowing the waveguide and incorporating a photonic crystal or defect state, it is possible to control the overlap of the optical field with the sensing material. Typically sensing can be done through a change in refractive index, or absorption of an intermediate material which is sensitive to the substance to be measured, for example, platinum for hydrogen detection. Using this approach it is possible to develop a single chip, with multiple sensors which could detect multiple gases, temperature, humidity and pressure.

Amza, Cristiana

WWW.EECG.TORONTO.EDU/~AMZA

Automated Self-Management in Cloud Environments

The economy has been transformed by the delivery of web services over the Internet in the past three decades. Behind the scenes of web-based service delivery technologies are large-scale, complex information systems. These information systems consist of thousands of servers, which store client data in a consistent and persistent manner and are multiplexed among several applications. Many aspects of content delivery still currently depend on manual fine-tuning and troubleshooting by humans. As a result, the service provider is strapped, with huge human administrator costs for performance debugging, resource allocation to applications and infrastructure configuration adjustments. This is a major impediment to the efficiency of this industry, limiting reinvestment, research and development. To address this problem, we investigate innovative self-diagnosis and adaptive reconfiguration techniques for scalable and available information systems. We have designed and developed novel online performance modelling and anomaly detection algorithms and tools that form the basis for self-configuring, self-tuning and self-healing servers. We use these techniques in our data centre laboratory towards automatically providing quality of service for a range of dynamic content services such as e-commerce, online bidding and massively multi-player games.

System Support for Parallel and Distributed Software Transactional Memory

Because of the increase in complexity and ubiquity of large-scale parallel and distributed hardware environments, simpler parallel programming paradigms become key. Transactional Memory is an emerging parallel programming paradigm for generic applications that promises to facilitate more efficient, programmer-friendly use of the plentiful parallelism available in chip multiprocessors and on cluster farms. We developed and optimized libTM, a Transactional Memory library that can be used in connection with C or C++ programs. libTM implements Transactional Memory (TM) for generic applications, and it allows transactions on different processors (or machines) to manipulate shared in-memory data structures concurrently in an atomic and serializable (i.e., correct) manner. There is no need for the application to do explicit fine-grained locking by acquiring and releasing specific locks on data items. Instead, a cluster-based run-time system automatically detects data races and ensures correct parallel execution for generic parallel programs. Any detected incorrect execution resulting from a data race is rolled back and restarted. In this project we have focused on reducing the software overhead of run-time memory access tracking and consistency maintenance for Transactional Memory support. We currently support applications with highly dynamic access patterns, such as massively multiplayer games. We have shown that Transactional Memory not only simplifies the programming of these applications, but can also improve performance and scaling relative to that obtained by using traditional locking techniques for code parallelization for the same application.

Anderson, Jason

JANDERS.EECG.TORONTO.EDU

A Self-Profiling Adaptive Processor: High-Level Hardware Synthesis

High current density in sub–100 nm ICs has created a power wall, limiting the rate of clock speed scaling in general purpose microprocessors. Attaining higher speed performance and improved energy efficiency motivates the need to develop processors that are customized to specific applications. Performing computations in custom hardware can deliver orders-of-magnitude improvement in energy efficiency and throughput. However, custom processor design, as with any hardware design, is difficult and time-consuming compared to software design. Moreover, hardware design in VHDL or Verilog lies outside the skill set of most software programmers. A further challenge is identifying what custom hardware would be of benefit to the speed and/or power of an application, as this depends on the run-time and power profile of the application under specific datasets. There is a need then for the rapid and automated design of processors tailored for particular applications. We are building a self-accelerating adaptive processor by modifying the architecture of a standard processor to create the ability to profile the execution of its own code. Using this profiling ability, our unique processor will be able to identify sections of its code that require optimization. Specifically, the profiling results will drive the selection of program code segments to be retargeted to custom hardware from their original high-level language implementation. C-to-RTL synthesis will be used, with the RTL subsequently compiled by standard backend tools. Once the hardware “compute accelerators” are available, the program binary will be modified to access the accelerators accordingly. Programmable logic devices, such as field-programmable gate arrays (FPGAs), are an ideal implementation platform for such adaptive processors, as FPGAs can be configured in milliseconds to implement any digital circuit. The reconfigurability of FPGAs also permits functionality to evolve over time, based on application needs.

Circuit and Architecture Techniques to Improve FPGA Speed, Power, Area and Ease of Use

Field-programmable gate arrays (FPGAs) are computer chips that can be programmed by the end user to implement any digital circuit. FPGAs can be thought of as “configurable” computer hardware, making them an ideal platform to realize application-specific hardware accelerators that are used in tandem with standard processors to improve computational throughput and energy efficiency. However, as programmable chips, FPGAs naturally consume more power, are slower, and use more area than fixed-function chips. In this research thrust, we are undertaking several circuits/architecture-related projects to improve FPGAs’ speed, area, power and ease of use. A first project comprises redesigning the interconnect structures within FPGAs to improve power efficiency, through approaches such as charge recycling, capacitance and glitch reduction. A second project considers coarse-grained reconfigurable fabrics, which are less flexible than traditional FPGAs, yet for specific applications, deliver superior speed, power and area performance. A third direction relates to synthesizeable FPGA fabrics (implemented in standard cells) that can be incorporated into larger systems-on-chip.

Balmain, Keith

WWW.WAVES.UTORONTO.CA/PROF/BALMAIN

RF Phenomena in Magnetized Plasmas

Magnetized plasmas that are encountered both in space and in the laboratory are highly anisotropic media which have a large impact on RF antennas operated in such media. One effect of a transmitting antenna in a magnetized plasma is the occurrence of luminous discharges such as those observed during the 1995 OEDIPUS-C sounding rocket experiment as well as in subsequent experiments carried out in our laboratory. In the OEDIPUS-C experiment, luminous discharges were observed close to the transmitting antenna early in the flight while Argon thrusters driving the separation of two sub-payloads were running. The formation of these RF-discharges as well as the unique shape of the resulting glow patterns are studied in order to better understand the processes that led to their occurrence. In the laboratory experiment, intended as a scaled-down version of the OEDIPUS-C RF antenna, similar RF-discharges were observed that extended as glowing band-shaped beams directed along the ambient magnetic field and extending all along the laboratory plasma column. Research consists of the evaluation of data from the OEDIPUS-C experiment and computational simulations of both the rocket experiment and the laboratory experiment, taking advantage of the vastly improved computer hardware and software available now, as opposed to 1995. Eventually, improved laboratory experiments with emphasis on relating the computed and measured phenomena will be designed and carried out.

▲ Simulation of RF Tag Interrogator Units underneath Urban Rail Trains

RF tag interrogation systems are used for position determination in the automatic control of urban rail trains. Such systems consist of a Tag Interrogator Unit (TIU), which is attached to an antenna mounted underneath the rail car and transmits interrogation pulses to tags installed between the rails. The TIU operates in a complex electromagnetic environment that contains some components that are highly conductive, such as the rail car itself, the wheel bogie sets, the rails and rail spikes going into the ground. Other components have lower but still significant conductivities and permittivities different from that of free space, such as the ground, the ballast and the sleepers (ties). In the design of a positioning system it is important to gain an understanding of the influence of these components on the RF field distribution underneath the rail car, especially in the plane where the tags are located. Finite-element software is used to simulate the RF fields under the train in the presence of all components and particularly to predict the electric field distribution, most importantly in the plane of the tag and the signal reflected by the tag. The results can be used to better understand the response of the TIU, leading to improvements in both the determination of the train position and the design of such TIU systems.

Bardakjian, Berj

HEART.IBME.UTORONTO.CA/~BERJ/BERJ

▲ ⬟ Bioengineering of the Brain

The main themes of the research are within the general field of neural engineering and, in particular, bioengineering of the brain. The purpose is to (1) characterize both normal and pathological brain electrical rhythmic activities and (2) anticipate, then abolish, the pathological electrical rhythmic activities in the brain, such as epileptic seizures. The approach is to characterize the spatiotemporal relations of the electrical rhythmic activities in neuronal populations and use cognitive devices to classify the dynamical features of the biological neural networks in the brain. The developed cognitive devices will be implemented as low-power hardware to be incorporated into the biological neural networks in a closed feedback loop. This will be used to provide implantable devices as therapeutic tools for brain disorders.

Betz, Vaughn

WWW.EECG.UTORONTO.CA/~VAUGHN

■ ⬣ Improved FPGA Architecture and CAD

My team seeks to find both better architectures and better Computer-Aided Design (CAD) tools for a type of integrated circuit — Field-Programmable Gate Arrays (FPGAs). FPGAs are a type of computer chip that can be reprogrammed to perform any function. As the cost of creating chips with billions of transistors has risen to $100 million, most applications cannot justify a custom-fabricated chip and instead are best served by a reprogrammable chip. Our research seeks to find the best “architectures” for FPGAs — what function blocks should they include and, perhaps even more importantly, how can we best programmably interconnect the huge number of function blocks in modern FPGAs? We are investigating how the circuitry of FPGA programmable interconnect should be modified to take best advantage of the latest (22 nm and below) process technologies. We are simultaneously investigating a radical change to FPGA on-chip communication in which we augment conventional FPGA programmable interconnects with packet-switched networks on chip. This fundamentally raises the level of abstraction of communication on the chip, but requires new CAD tools, which we are also developing, to automate this new and different design flow. We also seek to find new algorithms and computer-aided design tools to allow FPGA designs to be completed more quickly and to run at higher speeds and make more efficient use of the chip. In particular, we are looking at how to make highly scalable placement and routing algorithms that can handle the latest chips, which contain billions of transistors, in a reasonable run time. Finally, we are also investigating new areas in which FPGAs can accelerate computation.
A key application in this area we are currently researching is the modelling of photodynamic cancer therapy. By simulating the paths of millions of photons in complex human tissue, we seek to determine the best arrangement of fibre-optic probes to destroy a tumour (and minimize the exposure of healthy tissue) by activating a photosensitive catalyst only in a local area of the patient’s body.

Broucke, Mireille

WWW.CONTROL.UTORONTO.CA/~BROUCKE

▲ Control for Complex Specifications

The field of Systems Control has traditionally been focused on steady-state control specifications in the form of stabilization and tracking. The goal of the project is to develop a theory of control for complex specifications, in particular enabling systematic methods of design and control of the transient phase of a dynamic system. These complex specifications may include safety and liveness specifications, logic-based specifications and temporal specifications. Problems of control with complex specifications arise in all of the disciplines that apply Systems Control, ranging from robotics to process control.

▲ Patterned Linear Systems

Complex dynamic systems that are made of a large number of simple subsystems with simple patterns of interaction arise frequently in natural and engineering systems. There is generally no overarching theory that explains the phenomena exhibited by such systems. We have introduced a class of linear control systems called patterned systems, which mathematically capture the structure of a collection of identical subsystems with a fixed pattern of interaction between subsystems. The project involves developing a control theory for patterned systems and our approach is based on the geometric theory of linear control systems. The aim of our study is to determine if patterned systems may provide a template for the development of a more unified framework for dealing with systems, typically distributed, which consist of subsystems interacting via a fixed pattern.

Brown, Stephen

WWW.EECG.TORONTO.EDU/~BROWN

■ ⬣ CAD and Architecture for FPGAs

My research is focused on many different aspects of field-programmable gate array technology, including the design of the chip architectures and the algorithms that are used to implement circuits in these devices, as well as applications of FPGAs. In addition to my faculty position at the University of Toronto, I maintain an active involvement in the Altera Toronto Technology Centre, where I provide direction for the University Program that is offered by Altera. By combining my involvement in both the University of Toronto and Altera, it has been possible to develop research results that are both interesting from the academic point of view and of practical use when implemented in an industrial-quality CAD tool. My current research effort is in the area of CAD flows for FPGA devices and specifically in the development of new methods of efficiently compiling high-level language code (such as C code) into circuits that can be implemented in FPGA devices.

Chan Carusone, Tony

WWW.ISL.UTORONTO.CA

● ⬣ ■ Digitally Assisted Analog Front Ends

When sensing signals in the physical world, undesired interfering signals often arise with an amplitude much larger than the targeted signal of interest. For example, the signals received by an antenna include many channels with widely varying signal strength. In sensor interfaces, the sources of interference are manifold, including power supply noise, electrical noise from neighbouring circuits, and even mechanical sources of interference. Digital signal processing may be used to ultimately extract our target signal, but this requires digitization of both the desired information-bearing signal, and the dominant sources of interference, thus placing considerable burden on the analog front end and analog-to-digital converter (ADC). We seek to perform mixed analog/digital cancellation of the interference, thus relaxing the requirements on analog amplifiers, filters, and ADCs. We apply our work to problems in wireless communication and sensor interfaces.

● ■ ⬣ Energy-Efficient I/O for Supercomputing

Supercomputing infrastructure has reached an industrial scale. The video, search and cloud computing services on which our modern economy relies are delivered from warehouse-sized facilities housing thousands of individual servers. The energy consumption of these facilities is enormous, with significant economic and environmental consequences. It is therefore very frustrating that these distributed supercomputers are designed to be most energy efficient when operating under peak load, when in fact most of the time these facilities operate at 10–50% of their peak utilization in order to ensure robust service provisioning. Under these normal-use conditions, they operate at 15–65% of their peak energy-efficiency. Our research strives to improve the energy-efficiency of warehouse-scale computers under all normal-use conditions. Specifically, the interconnections within these distributed computing environments are targeted. Without progress on the digital input/output (I/O) circuits at either end of these interconnects, their share of overall server power will increase to over 50% in the next decade. We develop I/O integrated circuits and subsystems that are capable of near-zero energy consumption when idle, sub-nanosecond wakeup times, and the ability to intelligently and autonomously scale power consumption over a decade or more of bandwidth, from a few Gb/s per lane up to 50–100 Gb/s per lane. Leveraging our lab’s past work on burst-mode clock and data recovery circuits and low-power injection-locked oscillators, we utilize nanoscale CMOS technologies to create integrated circuit prototypes of these subsystems. Considering a 220 TWh annual energy budget for supercomputers worldwide, and the importance of Canada’s ICT sector, we are very excited about the impact our work will have.

● ⬣ ■ Highly Integrated Optical Transceivers

Optical fibre is already the dominant communication medium for high data rates over long distances. However, there is increasing interest in the use of optical fibre for communication over shorter distances. For example, in rack-mounted computing and storage environments, where the cost of operating the equipment over its lifetime now exceeds its initial purchase cost, optical communication becomes increasingly attractive at data rates of 25+ Gb/s. At these data rates, the losses inherent in communication over copper cables cause it to consume more power than optical communication, impacting energy costs. Optical fibre’s thin diameter permits better airflow (hence, reduced cooling costs) and easier maintenance than copper cables. Fibre’s immunity to electromagnetic interference is attractive for automotive and other harsh environments. Moreover, optical fibres can be routed in tight bundles with much less crosstalk than copper wires, making it a scalable medium. Even in consumer applications demanding multi-Gb/s throughput, optical cables are attracting increasing interest because of their light weight, flexibility and thin diameter. To exploit the fundamental advantages of optical communication in these areas, we develop highly-integrated, dense and low-power optical transceiver circuits. We prototype our developments in the most advanced integrated circuit technologies available.

● ⬣ ■ Ultra-Short-Reach Chip-to-Chip Communication

Our capacity for digital communication continues to increase as we integrate more and more functionality into fewer electronic components. Integration enables lower-cost, smaller-size systems with lower power consumption. Unfortunately, our ability to integrate complete communication systems onto a single chip today remains limited by small but finite defect rates during chip fabrication, which limit the maximum number of transistors that can be reliably integrated. In addition, the need to combine different fabrication technologies to realize, for example, a high-density memory alongside high-speed transmitter and receiver circuits, makes integration challenging. Emerging dense interconnect technologies may offer a path forward. These technologies place multiple chips in close proximity and connect them with densely packed wires that may be less than 1 cm long. Our long-term vision is to use these ultra-short-reach (USR) links to interconnect multiple chips so seamlessly that system performance can

scale without bottlenecks. We are developing demonstration platforms for USR links, including transceiver circuits that are extremely small and consume very little power.

Cheng, Hai-Ling Margaret

WWW.IBBME.UTORONTO.CA/FACULTY/CORE-FACULTY/H-L-M-CHENG

Cellular Magnetic Resonance Imaging for Non-invasive Cell Tracking and Monitoring

The ability to see cells inside a living body can transform how we detect and diagnose disease and monitor treatment. From differentiating healthy from unhealthy cells, to tracking therapeutic cells that are injected into the body, cellular imaging is an active biomedical research area. Yet, when we think of cellular imaging, we usually think of looking at samples under a microscope. This research program strives to develop a similar capability to look at cells but using magnetic resonance imaging (MRI) for non-invasive,deep-tissue penetration in a living subject. Our goal is to improve detection sensitivity and specificity to make MRI the technology of choice for non-invasive human cellular imaging.

Magnetic Resonance Imaging for Tissue Engineering

This broad research program aims to advance the capabilities of non-invasive magnetic resonance imaging (MRI) to enable critical advances in tissue engineering and regenerative medicine. Specific focus is given to MRI on a physiological, cellular and molecular level to tackle difficult tissue-engineering problems such as angiogenesis and cell-based therapy.

Quantitative Magnetic Resonance Imaging of Microvascular Physiology

The microvasculature, or the nutritive blood vessels in our body, is essential to maintaining tissue health. Loss of proper microvascular function underlies a wide range of conditions, including cancer, stroke and cardiovascular disease. This is why being able to evaluate the health of these small blood vessels is important for diagnosis and assessment of treatment effect. Yet, despite the ability of current advanced technologies to probe microvascular function, such as blood flow, in a non-invasive manner, it remains very difficult to detect the early signs of microvessel dysfunction. In this research program, we are developing new non-invasive imaging methods based on magnetic resonance imaging to find those early changes sooner than we can today, in the hope that early detection will enable early intervention for improved outcome.

Chow, Paul

WWW.EECG.TORONTO.EDU/~PC

Internet-Scale Memory Systems

With the vast amount of data accessed and stored using the Internet, new memory architectures are required to host the data that can provide low latency access, low power dissipation and a compact form factor. Current systems use the collected main memory of a cluster of high-end servers for an application that does not need the computation power of such systems. This project explores the intrinsic requirements of such systems without the constraint of using common compute server platforms. Field-Programmable Gate Arrays will be used to develop more flexible, application-specific and novel architectures for building Internet-scale “Big Data” memory systems.

Programming Models and Architectures for Reconfigurable and Heterogeneous Computing Systems

This research investigates approaches to computing using systems of multiple, heterogeneous computing devices. The heterogeneity addresses the need for special-purpose accelerators that provide performance or other efficiencies, such as more efficient energy usage. A key focus is the use of Field-Programmable Gate Arrays (FPGAs), a form of configurable hardware. Such systems can be found in an embedded device or in high-performance computing systems. Important issues being addressed are better methods for programming, testing and debugging and system architectures. Much of the research is driven by applications. One aspect is to work with users of high-performance computing facilities and help them to improve performance through better algorithms and the use of accelerators implemented with FPGAs and/or GPUs.

Virtualizing FPGAs for Cloud Computing Platforms

Field-Programmable Gate Arrays (FPGAs) are a programmable hardware resource that can be used to build application-specific hardware accelerators for many types of computation and network processing. In a cloud-based platform, the hardware of the computing resources is abstracted from the user in the form of a virtual machine that looks uniform to the user regardless of the physical hardware platform. This project explores how to place FPGAs into this virtualized environment when using Open Stack. The prototyping platform is the SAVI Networks platform (savinetwork.ca) where real applications are being built, including the demonstration of hardware extensions to Open Flow that minimally affect performance compared to using software virtual machines.

Davison, Edward

WWW.CONTROL.UTORONTO.CA/PEOPLE/PROFS/TED

Control of Large-Scale Decentralized Systems

Our research is focused on the control of large-scale systems, where only limited information about the overall system is available to the control agents of the system. Such systems occur often in modern industrial society, for example, in chemical engineering, electrical power systems, aerospace systems, transportation systems, building temperature control systems, large flexible space structures and pulp and paper control systems, as well as in other areas such as management science and biological systems. Problems areas that immediately arise from large-scale systems are current areas of research: decentralized control, intelligent control, fault-tolerant control and the control of unknown systems. A direct application of this research is presently being applied to the control of large flexible space structures, earthquake-resistive building structures, and electric power systems with particular focus on micro-grid systems and spinal cord injury patients.

Dawson, Francis

WWW.ELE.UTORONTO.CA/~DAWSON

Improving Energy Efficiency of Energy Conversion Processes

The general research interests are in the area of modelling systems powered by electrical energy. At the component level, the current focus is on developing improved models that can describe the electric and thermal fields in electrochemical storage devices. The objective is to use reduced order multiphysics models to develop energy management controllers that can extend the life of an energy storage device. At the system level, the objective is to determine the system architecture and control philosophy that lead to an optimal integration of energy storage devices and power converters, subject to a specific generating and electrical load profile. Other areas of interest include the modelling of thermoelectric and piezoelectric devices.

Draper, Stark

WWW.ECE.UTORONTO.CA/PEOPLE/DRAPER-S

Exploiting Feedback to Architect Streaming Digital Communication Systems for Short Delays and High Reliability

In this project we reexamined the architectural thinking that underlies digital communication systems. This architecture was not designed with increasingly important real-time delay-sensitive streaming and collaborative applications in mind. Applications such as high-end video conferencing, vehicular networks, machine-to-machine communications, and the coordination and fast reconfiguration of distributed systems such as factory robots, demand high-reliability real-time data delivery under strict deadlines. In preliminary work we have shown how to realize astonishing improvements in the reliability of communications at short delays by smartly incorporating receiver-to-sender “feedback” into streaming data systems. In this project we will continue to develop the fundamental theory and will also develop the error-correcting codes and decoding algorithms required to implement these ideas in practice. Finally, we will develop a wireless testbed that will consist of a number of wireless devices in which we can prototype our new architecture and algorithms.

Large-Scale Linear Programming Decoding via the Alternative Direction Method of Multipliers

When binary linear error-correcting codes are used over symmetric channels, a relaxed version of the maximum likelihood decoding problem can be stated as a linear program (LP). This LP decoder can be used to decode at bit-error-rates comparable to state-of-the-art belief propagation (BP) decoders, but with significantly stronger theoretical guarantees. However, LP decoding when implemented with standard LP solvers does not easily scale to the block lengths of modern error-correcting codes. In this project we draw on decomposition methods from optimization theory, specifically the Alternating Direction Method of Multipliers (ADMM), to develop efficient distributed algorithms for LP decoding. The key enabling technical result is a nearly linear time algorithm for two-norm projection onto the parity polytope. This allows us to use LP decoding, with all its theoretical guarantees, to decode large-scale error correcting codes efficiently. Our approach has the potential to solve longstanding issues of great industrial importance such as the “error-floor” problem of low low-density parity-check (LDPC) codes; the existence of which has slowed the adoption of these state-of-the-art codes into applications requiring ultra-low error rates such as magnetic storage.

Re-architecting Last-Level Caches for Low-Voltage Operation

Power management is a first-order priority in the design of modern processors. Dynamic voltage/frequency scaling (DVFS), wherein operating voltage is lowered in step with reduced computational demand, is one of the most successful and widely adopted power reduction techniques. However, increased process variability with technology scaling imposes limits on the minimum operating voltage. Below this minimum, large-scale memory structures such as the last-level cache (LLC) cannot be guaranteed to operate reliably. In this project we combine techniques from error-correction coding with architectural insights to redesign LLCs to improve low-voltage performance. Our initial results show that joint optimization of device size, redundancy, and amount of error-correction can yield significant savings in chip area (up to 27% reduction in LLC area for a minimum operating voltage of 600mV in 32 nm technology). We use these insights to design a novel heterogeneous cache architecture that dynamically adjusts the available cache size to match real-time computational demands. By combining larger cells for lower-voltage operation with smaller cells for higher-voltage high-performance operation, the heterogeneous design provides an additional 15%–20% reduction in LLC area at negligible average runtime increase.

The Privacy/Security Trade-off across Jointly Designed Biometric Authentication Systems

In the area of secure biometrics, work has been done to build an information theoretic framework characterizing privacy and security of single biometric systems. People have worked extensively on designing such systems, some cryptographic in nature, and others tied to error-correcting codes. However, there is still little known about security and privacy across multiple jointly designed systems. This work will focus on the privacy/security trade-off across multiple “secure sketch” biometric systems. Secure sketch is a type of biometric system architecture related to error-correcting codes where a system is characterized by a parity-check matrix over a finite field, or equivalently by a subspace of a vector space over that same field. Given a set of systems (a design), we introduce worst- case measures of privacy leakage and security breach

in case a subset of the systems becomes compromised. It turns out that more secure designs are necessarily less private and vice versa. We study the trade-off between privacy and security by relaxing a restricted version of the problem, by studying the algebraic structure of the problem, and by formulating graph theoretic questions. These approaches generate bounds on achievable privacy/security pairs.

Eleftheriades, George

WWW.WAVES.UTORONTO.CA/PROF/GELEFTH

Artificial Materials (Metamaterials) from Microwave to Optical Frequencies

We are developing paradigm-shift metamaterial devices and subsystems and related technologies at RF/microwave and optical frequencies. Metamaterials are artificially structured media with unusual electromagnetic properties. Such properties include negative refraction, enhanced evanescent waves through resonant amplification and sometimes a negative group velocity. Our vision is to develop metamaterials that can manipulate and control electromagnetic waves, much as conducting wires manipulate the flow of electrons. Both three-dimensional volumetric and surfaces (metasurfaces) metamaterials are being developed. A recent effort concerns the development of ultrathin metasurfaces for wavefront manipulation, such as refraction (bending of incident plane waves or Gaussian beams), lensing and controlled beam formation. Application areas include super-resolution microwave and optical microscopy, detection and sensing, advanced hardware for wireless communications, wireless power transfer, reduction of interference, space technology, radars, defence, solar-cell concentrators, thermophotovoltaics, infrared focal-plane arrays and many more. Examples of devices include small antennas, multi-functional RF/microwave components (including active devices), sub-diffraction imaging lenses and probes (even operating in the far field), ultrathin lenses, invisibility cloaks and related “transformation optics” lenses, plasmonic optical circuits, plasmonic waveguides and nano antennas. Research includes both experimental work and fundamental theory. Our research is supported by several industrial partners, government agencies and laboratories. Graduates from our group are quite successful in securing faculty positions in academia (e.g., U. Michigan, U. Alberta, McGill, U of T and UBC) and industry (e.g., Apple, AMD, Blackberry, Freescale and Motorola).

Enright Jerger, Natalie

WWW.EECG.TORONTO.EDU/~ENRIGHT

Approximation Computing for Communication

Approximate computing explores opportunities that emerge when applications can tolerate error or inexactness. These applications, which range from multimedia processing to machine learning, operate on inherently noisy and imprecise data. As a result, we can trade off some loss in output value integrity for improved processor performance and energy efficiency. Memory accesses are costly both in terms of latency and energy. We are exploring microarchitectural technques that leverage approximation to reduce the cost of data storage and data communication. For example, we explore load value approximation, a novel microarchitectural technique to learn value patterns and generate approximations of the data. The processor can use these approximate data values to continue executing without incurring the high cost of accessing memory, removing load instructions from the critical path. Load value approximation can also be used to inhibit approximated loads from accessing memory, resulting in energy savings.

Interconnection Networks for Heterogenous Multicore Systems

In recent years, embedded and mobile devices have proliferated in society; with each generation, these devices enable new applications including navigation, digital photography and multimedia. With the scaling of transistor features sizes, the number of components that can be integrated onto a single chip continues to grow, allowing increased functionality. Modern chips may contain several processing cores, graphics processors, memory controllers, I/O interfaces, multimedia accelerators and numerous other specialized functional units. Each of these components within a device requires that data be communicated between it and other parts of the system. To facilitate this communication, various on-chip communication structures have been proposed and utilized. Driven by higher requirements for concurrent bandwidth and lower latency, interconnect fabrics in practical system-on-chip (SoC) devices have evolved from a bus architecture to an on-chip network. This project explores the needs of SoC communication, which differ dramatically from those of general-purpose devices; we are developing novel architectures to meet the power and real-time latency and bandwidth constraints of these systems. These novel on-chip network architectures will work symbiotically with the entire memory hierarchy to allow efficient delivery of data. We are developing novel memory controller scheduling algorithms and exploring the use of 3D stacking to ensure sufficient bandwidth in these SoC systems.

Semantically Rich Networks for Many-Core Architectures

Parallel architectures are rapidly becoming ubiquitous. To leverage the computational power of these multiple cores, communication between cores or devices is essential. This project looks at streamlining the communication between cores via on-chip network innovations to increase its efficiency. The most efficient network design — the one that provides the greatest performance at the lowest cost (area and power) — will be one that most closely matches the functionality required by an application. However, customized, application-specific networks are not appropriate for general-purpose many-core architectures since they run many diverse applications. Therefore, we focus on communication behaviours that are evident across a range of workloads, specifically behaviour exhibited by cache coherence protocols. Cache coherence protocols introduce communication overhead and can substantially impact performance, as many of these operations lie on the critical path. Examining several coherence protocols, we note the presence of coherence primitives that use multicast and reduction operations. Our reduction routing combines redundant messages (such as acknowledgements) during their network traversal to reduce network load. This insight and router architecture can be more broadly applied to any many-to-one communication pattern. In addition to many-to-one and one-to-many messages that are common to many coherence protocols, coherence protocols exhibit additional exploitable behaviour. For example, these protocols are characterized by having a mix of long and short messages; cache lines (data messages) represent long messages while coherence requests (control messages) are typically only a few bytes long. Existing theories for deadlock freedom in fully adaptive routing require a conservative virtual channel allocation scheme. We show that this restriction severely limits performance for short coherence control messages. We proposed a novel flow control technique, whole packet forwarding that accelerates the handling of short coherence messages in the network to improve performance and efficiency. An important feature of these solutions is the low hardware overhead they incur. Small hardware modifications and modest additional logic are required to support each of the above-mentioned designs. These optimizations yield significant throughput and latency improvements for a variety of workloads.

Simulation Methodologies for On-Chip Networks

On-die communication fabrics represent a critically important aspect in the design of future many-core computer systems. As systems scale to increasingly large numbers of on-die agents, the on-die communication fabric will factor dramatically into both the performance and the power consumption of future architectures. This research focuses on two challenges in the design of on-die communication fabrics: physically aware performance and area optimization for communication fabrics and uncore, interconnect and system power management. Within these topics, we are specifically exploring solutions to integrate cache coherence protocol traffic analysis within the early-stage on-chip network design space exploration and the integration of protocol-level information into the quality of service and DVFS mechanisms of the on-chip network. These two thrusts will span issues of correctness, energy/performance efficiency and scalability. Current techniques to simulate on-chip networks are either time-consuming or lack accuracy in the resulting performance and power estimates. Our new traffic models will accelerate on-chip network simulation and allow researchers to reach stronger conclusions about system performance at an early design stage. These models accurately capture sharing behaviour and the interaction of dependent messages in the coherence protocol. These models are parameterized to allow a wide diversity of systems to be simulated with rapid turn-around times.

Francis, Bruce

WWW.SITES.GOOGLE.COM/SITE/BRUCEFRANCISCONTACT

Control Theory and Applications

1. Distributed robotics theory: We study the mathematical theory of robot formations. The robots are typically modelled as unicycles with only onboard sensors and no leaders. The objective is to design local motion strategies so that a team of robots performs a coordinated task, such as forming a circle (with Professors Mireille Broucke, Manfredi Maggiore and Luca Scardovi). 2. Infinite lattices of dynamical systems: In studying the formation of a very large number of vehicles, one approach is instead to model an infinite number of vehicles. The question then arises as to what mathematical framework to take so that the latter model correctly describes the behaviour of the former. This leads to the subject of infinite chains or lattices of dynamical systems (with Professor Avraham Feintuch).

Frey, Brendan

WWW.PSI.TORONTO.EDU

Algorithms for Inference and Machine Learning

Dr. Frey’s group develops new inference theories that can be used for probabilistic and statistical inference in large-scale systems, such as those that arise in telecommunications, robotics, genetics, genomics, vision and signal processing. Dr. Frey is co-author of an article that introduced the factor graph and associated sum-product and max-product algorithms (IEEE Trans Info Theory 2001). A factor graph is a method for decomposing high-order probability models into simpler terms, so that the sum-product or max-product algorithm can be used to efficiently perform inference. A search for “factor graph” on Google returns over 40,000 hits. Other methods developed by Dr. Frey and his colleagues include variational methods for inference in large-scale nonlinear Gaussian models (Neural Comp 1999), the “wake-sleep” algorithm for unsupervised learning (Science 1995), cumulative distribution networks (NIPS, UAI 2008) and loopy belief propagation algorithms for low-level vision (CVPR 2000), phase-unwrapping of medical and satellite images (NIPS 2001), exemplar-based clustering (Science 2007) and facility location (AISTATS 2010).

Data Analysis and the Affinity Propagation Algorithm

Summarizing data by identifying a subset of representative examples is important for scientific data analysis and in engineered systems. Such exemplars can be found by randomly choosing an initial subset of data points and then iteratively refining it, but this only works well if that initial choice is close to a good solution. Dr. Frey’s group developed a new method called affinity propagation, which takes as input measures of similarity between pairs of data points. Real-valued messages are exchanged between data points until a high-quality set of exemplars and corresponding clusters gradually emerges (Frey and Dueck, Science 2007). Because of its simplicity, general applicability and performance, the affinity propagation algorithm is widely used in science and engineering. In the past year, an on-line web tool developed by Dr. Frey’s group was accessed over 100,000 times by over 3000 users (unique IP addresses), 600 of which were from Canada. Google returns over 10,000 hits for the search term affinity propagation. Dr. Frey’s method has been applied to solve problems in biology, genetics, genomics, medicine, physics, chemistry, telecommunications, electronics, archeology, economics and social networks.

WWW.GENES.TORONTO.EDU

Deciphering the Human Genetic Code

Despite widespread claims that the human genome has provided a “book of life,” it turns out that it is very difficult to understand how genes stored in the genome encode the actual genetic messages that control the life, death and ongoing activities of the cells comprising all human tissues. In the words of the famous genomics researcher Eric Lander, “Genome: Bought the book, hard to read.” Recently, Professor Brendan Frey and his research team discovered a fundamentally new view of how living cells “read the genome” and use a limited number of genes to generate enormously complex tissues such as the brain. In a paper that was published in the May 6, 2010 issue of *Nature* and featured on its cover, Dr. Frey describes research conducted by his team. They developed a computational technique based on probability, statistics and machine learning and used it to reveal a second level of information hidden in the genome that can account for the exponentially greater complexity required to create a human being. The work of Dr. Frey’s team was reported in the *Globe and Mail*, the *Toronto Star*, on CBC Radio, BBC Radio and in a variety of other national and international news. Dr. Frey leads an ongoing, multi-year project whose goal is to infer the coding mechanisms underlying the regulation of genes. The project involves experimental collaborators from the Centre for Cellular and Biomolecular Research at the University of Toronto, along with international collaborators. Several of the students and postdoctoral fellows who have graduated from Dr. Frey’s lab have subsequently taken faculty positions at leading universities, including UPenn, UNC and Harvard.

Genov, Roman

WWW.EECG.UTORONTO.CA/~ROMAN

Portable, Wearable and Implantable Sensory Biomedical Electronics

We are heading the Intelligent Sensory Microsystems Laboratory at the University of Toronto. Members of our laboratory conduct research on analog and digital VLSI circuits, systems and algorithms for energy-efficient signal processing with applications to electrical, chemical and photonic sensory information acquisition, biosensor arrays, brain-chip neural interfaces, CMOS imagers, parallel signal processing, adaptive computing and implantable and wearable biomedical electronics.

Goel, Ashvin

WWW.EECG.TORONTO.EDU/~ASHVIN

Binary Instrumentation of Operating Systems

A binary instrumentation system enables the monitoring and manipulating of every instruction in an executing binary. Binary instrumentation systems have been used for developing bug-finding and security tools. For example, Memcheck uses binary instrumentation to detect various types of memory errors dynamically, such as accessing memory after it has been freed. We have developed a binary instrumentation system for the Linux operating system. We aim to use this system to develop tools to find memory bugs in the Linux kernel and to harden the kernel against buggy device drivers. This is joint work with Professor Angela Demke Brown of the Department of Computer Science.

End-to-End Data Reliability

The goal of this project is to ensure data integrity in the face of software bugs. Currently, the project is focused on improving the reliability of file-system software. When file systems are buggy, they can cause data corruption and persistent application crashes. We are developing a system that ensures that a file-system disk image will remain consistent in the face of arbitrary file-system bugs. The key idea is to verify all file-system operations that update the disk at run time using a well-defined set of consistency properties. This is joint work with Professor Angela Demke Brown of the Department of Computer Science.

Gulak, Glenn

WWW.EECG.TORONTO.EDU/~GULAK

VLSI for Digital Communications

In the area of digital communications, we have continued to develop several practical ways to improve the performance and implementation of wireless systems that use multiple antennas (MIMO) for improved diversity and capacity. One of the key elements that we have investigated is the subsystem in the baseband responsible for data detection. A key contribution is the creation of an innovation that we call an on-demand K-best algorithm (a breadth-first search technique) whose complexity scales linearly with constellation size. This innovation is key to supporting higher-order modulation schemes such as 64-QAM and 256-QAM systems that will appear in next-generation communication standards, necessary for Gbps performance. We have implemented and tested our algorithm in 0.13 um CMOS and have generated the best-known results published in the literature to date, with respect to data rate, power efficiency and area. Our results have been extended to soft detection and tested with CMOS prototypes for use with iterative FEC decoding schemes. We have also made recent contributions to an important channel preprocessing block found in all MIMO systems, namely that of QR decomposition, a function needed for decomposing the channel matrix. Our key contribution in this area is the development of both algorithms and a 0.13 um CMOS implementation that demonstrates the world’s lowest (best) processing latency. Another area of recent accomplishment is in a channel preprocessing element known as Lattice Reduction, which can be used to mitigate scattering and antenna correlations that exist in practical MIMO systems. Lattice Reduction is a baseband signal processing algorithm to re-orthogonalize the signal space with the objective of improving BER performance. We have developed several algorithmic innovations and the world’s first CMOS prototypes for Lattice Reduction; the concepts developed will be particularly attractive for low-power implementations. Future work focuses on next-generation wireless OFDM baseband signal processing algorithms and their high-performance, low-power CMOS realization.

Hatzinakos, Dimitrios

Automatic Detection System for Tracking of Bees

This is a project in collaboration with Sigma-Eight Incorporation (SEI), in Witby, Ontario. Recently, SEI was approached by the New Zealand Institute for Plant and Food Research Limited (PFR), which is a Crown Research Institute located in Auckland, NZ. PFR is attempting to find ways to enhance yield from bee populations by creating smarter growth-focused pollination systems. To do this, they need to have more information about bees and their foraging methods, which can be obtained by more accurate tracking techniques. PFR have found a transmitter small enough to work for their purposes; however, they lack a receiving system capable of tracking the signal from these transmitters in real time. Various tracking systems have been developed over the years using different techniques such as electronic and mechanical scanning, time of arrival, electronic rotation, etc. Most of these methods presume the signal will be on for a long period of time and have significant strength. This situation is quite different: the signal will be quite weak and will be on for a very short period of time (i.e., pulsed for about 2 to 5 ms every few seconds). The system also has to be cost effective, since the budget available for multiple stations is somewhat limited, as is often the case in research these days. To keep costs down. SEI is hoping to use DSP-based hardware that it has already developed utilizing dual TI TMS320C5509A processors, a separate VHF receiver daughter card and an RS232 serial output card. Phase-related methods may be investigated to determine direction; however, it is likely that the phase distortion in such weak signals will introduce an intolerable level of uncertainty. Also, it may be difficult to control the timing to the precision required. The most likely method will be a signal-strength-related one utilizing several Yagi antennas pointed in different directions. Instantaneous simultaneous detection is possible on several receivers, which could be compared and interpolated to estimate a direction. Several samples could be used in sequence to refine the direction. Other systems have been built in the past utilizing similar techniques; however, they did not have the processing power or hardware necessary to provide simultaneous detection in all directions and comparison of the signal from the same pulse. Comparing successive pulses on different antennas would add increased uncertainty due to likely variations in signal strength from pulse to pulse. This project requires knowledge in the areas of RF antennas and propagation, digital signal processing, digital hardware, real-time systems and C programming. Most development tools are low or no cost, requiring a typical PC for development. A spectrum analyzer, oscilloscope, and test range will also be required for various aspects of the project.

WWW.COMM.UTORONTO.CA/~BIOMETRICS/MEDICAL

Medical Biometrics

The cardiovascular system offers a variety of physiological signals that can be used as biometrics. While modality such as the electrocardiogram (ECG) is still relatively novel, it is increasingly garnering acceptance as a useful biometric tool, because of some unique characteristics. Existing solutions for biometric recognition from electrocardiogram (ECG) signals are based on temporal and amplitude distances between detected fiducial points. Such methods rely heavily on the accuracy of fiducial detection, which is still an open problem due to the difficulty of exact localization of wave boundaries. To avoid fiducial points detection, the signal is processed holistically, using second-order statistics. Our autocorrelation-based method is a very simple and effective approach that does not require any waveform detection. It depends on estimating and classifying the significant coefficients of the Discrete Cosine Transform (AC/DCT) or the Linear Discriminant Analysis (AC/LDA) of the autocorrelation of heartbeat signals. The AC/LDA algorithm has been incorporated into a prototype system developed at the Biometric Security Laboratory (BioSec.Lab), the HeartID. HeartID is a Matlab-based software with various functionalities, such as user enrolment, database handles, security level adjustment and identification/verification modes of operation.

WWW.COMM.TORONTO.EDU/~SPSN

Self-Powered Sensor Networks

The University of Toronto, AD Telecom, SIRADEL and OMESH Networks are partners in developing compelling materials, communication architectures, software and other critical technologies necessary to create self-powered, ubiquitous and wireless ad hoc sensor networks. Substantial benefits will be realized by the citizens of Ontario and by Canadian society in general with the commercialization of a family of products that take advantage of these sensor networks, along with the novel energy harvesting and power generation technologies used to support them. The panoply of envisioned applications include effective, responsible and sustainable monitoring and governance in structural health, disaster relief, and transportation and law enforcement, as well as public safety and security. During our collaborative effort, we will undertake three main tasks: (1) Creation of sensor hardware that employs redundant architectures, fault-tolerant methods and nano-enabled materials to ensure system integrity, minimize sensed false-positives, increase sensor sensitivity and ease interaction with short-range wireless radios. The proposed research will integrate these aspects in a flexible and low-cost hardware framework. Several types of optical, electrochemical and biological sensing techniques will be investigated, including a quantum dots composite-based authentication-at-a-distance architecture with unambiguous authentication and visual association under all weather conditions, such as fog, rain and snow. (2) Creation of system software and middleware for the extraction, processing and characterization of real-time sensed data. One of the unique contributions of this task involves the advancement of innovative mobile social networking technology, which has the secondary benefit of enhancing next-generation voice, video and data transfer in addition to security/privacy methodologies. The University of Toronto will leverage AD Telecom’s extensive, state-of-the-art infrastructure to collect massive amounts of sensor data in order to provide critical functionality for (i) management of inconsistent and uncertain data; (ii) light-weight data integration; (iii) data cleaning and social network analysis; and (iv) various enhanced security functions for device authentication and data protection under a wide range of attack scenarios. (3) Creation of innovative energy-conserving, capture and storage technologies that use novel nanoscale materials, energy harvesting methods

and renewable energy resources to supply consistent power to sustain autonomous sensor networks. The research on self-powered sensor energy systems will focus on five major areas: (i) power conditioning and conservation; (ii) electromagnetic energy harvesting; (iii) solar energy harvesting; (iv) vibrational-thermal energy harvesting; and (v) energy storage. The ultimate target is low-cost, miniaturized, readily integrable, 24/7 energy generation systems that can sustain on the average a few hundred mW pulses, as well as support continuous current draw at mA scale. Several, if not all, of the energy generation techniques will undoubtedly be extended to the next-generation, large-scale wireless technologies to further reduce dependence on fossil fuels and other environment-taxing resources.

Helmy, Amr S.

PHOTONICS.LIGHT.UTORONTO.CA/HELMY/RESEARCH

Infrared and THz Semiconductor Laser Chips

The coherent radiation afforded by lasers fuels numerous applications ranging from medicine to material processing and telecommunications. In particular, semiconductor lasers offer a form factor, efficiency and portability that have fuelled innovations in all industrial sectors. A new class of semiconductor lasers has been recently developed with even more novel features. Semiconductor lasers enabled by Bragg reflection waveguides (BRW) are essentially one-dimensional photonic bandgap structures that are doped in a p-i-n profile, where light is guided by Bragg reflectors with light propagating parallel to the epitaxial layers. Bragg reflection lasers enable the realization of high-power single-mode lasers and amplifiers with larger mode volumes, higher gain coefficients and stronger mode discrimination than their counterparts. Moreover, this class of novel lasers empowers applications related to nonlinear frequency conversion in monolithically integrated optoelectronic integrated circuits. This research focuses on using this class of lasers to develop high-performance single-mode lasers for realization of electrically injected monolithic optical parametric oscillators. These chip-based sources can provide continuous coverage of spectral regions, which are not accessible by other technologies such as quantum cascade lasers. Examples of niche applications served by this unique platform include sources for environmental and biomedical sensing elements in the 1–9 μm window and chip-based THz spectroscopy sources. These sources play a pivotal role in enabling high-resolution, high-sensitivity chemical sensing and environmental monitoring applications due to their superior tuneability and spectral brightness. For example, certain molecules containing carbon-hydrogen bonds have infrared footprints within the 2–3 μm spectral window. In addition, H2O exhibits significant absorption features around 2.5 μm, which can be used for in-situ combustion measurements of moisture and temperature. This spectral window is not covered with quantum cascade lasers and can benefit from a broadly tunable coherent source. We were able to recently demonstrate such a source using the devices described above. In addition we are also able to develop sources with no moving parts to cover extended regions of the spectrum using a single device. As an example, efforts are underway in the group to cover the 7–11 μm window of radiation to be able to test for most known explosives.

Monolithic Microwave Photonics and THz Pulse Sources with Ultralow Phase Noise

Widely tunable, stable photonic-based microwave and pulse sources are indispensable for applications in numerous fields such as telecommunication systems, radar systems and modern metrology. These sources are often bulky and require a stable RF signal source, which increases the system's cost and complexity. In addition the attributes of these sources are usually limited by the characteristics and in particular the bandwidth of the RF sources utilized. Recently we introduced a novel, simple method to generate an optical clock with wavelength tunability. The beating signal generated by two single-mode lasers causes the modulation of the gain saturation of an SOA that is placed inside a ring laser cavity. This technique is particularly versatile in comparison to its counterparts; the repetition rate is controlled by the frequency difference between the two CW light sources, overcoming the bandwidth limitation of other techniques, which require a RF source. In addition, the operating wavelength is tuned by sweeping the central wavelength of the band-pass filter. Moreover, this new technique is cost-effective and provides the possibility for hybrid integration as it is comprised of semiconductor chips that can be heterogeneously integrated on an Si platform. This research enables optical pulse-trains to be generated in an all-optical setting based on gain-induced four-wave mixing in semiconductor optical amplifiers. A unique advantage of this versatile approach is the optical control it affords of the repetition rate, which can be tuned by controlling the frequency difference between the various light sources employed in the setup. Using novel designs, we are able to drastically improve upon the stability of all optical techniques through injection locking. Robust and low-phase noise pulse generation in the 100s of GHz has been measured with a line-width ~ 1 Hz with no need for RF sources, optical stabilization or optical feedback. Recently, we successfully achieved injection locking using 10 MHz optical frequency comb source, while utilizing an external cavity to eliminate the residual modes. Injection locking using such a low (10 MHz) optical frequency comb source enables and provides more flexibility for numerous applications.

PHOTONICS.LIGHT.UTORONTO.CA/HELMY

Monolithic Quantum Photonic Devices and Circuits

Photons, the particles of light, play a pivotal role in the emerging area of quantum information science, such as optical quantum computing and quantum cryptography. However, these futuristic-sounding technologies only exist in specialized labs; practical commercial systems are not available to date. One of the reasons these technologies have not moved into practical settings is that they need to be implemented using bulky components that are not portable and are sensitive to vibrations. Current technologies required to produce the building blocks of quantum systems do not allow a high level of integration of these components. Those include devices for the generation, manipulation and detection of paired photons that are entangled. These entangled photons are an essential building block for quantum systems as required by quantum mechanics. For example, the mainstream technique for producing entangled photon pairs is to use a strong laser beam to hit a nonlinear crystal.

With a probability of 10^–10, a photon in the laser can be converted into a pair of entangled photons. Such a system is extremely inefficient and very energy consuming. Besides, the whole system, including the laser, nonlinear crystal, mirrors and lens, etc., takes up a great deal of space on an optics table. The optical setup requires delicate construction and is sensitive to the external environment. Thanks to recent breakthroughs by our group, fully integrated, portable and robust entangled photon sources were made possible using mainstream semiconductor technologies. We successfully demonstrated the generation of entangled photons from a semiconductor chip. This electrically powered, alignment-free chip is specially engineered, which not only increases the photon generation efficiency compared to the bulk crystal counterparts, but also makes the integration with other optical components possible. Meanwhile, our group has been working on engineering the generated photon properties on the same platform to enable quantum computing, material spectroscopy, quantum and, imaging among other applications. Ultimately, our techniques will allow the entire photon generation and manipulation processes to take place on the same chip, which could be a big step towards a practical, commercial quantum computer and other quantum information processing systems.

PHOTONICS.LIGHT.UTORONTO.CA/HELMY/NANOPHOTONICS

Nanophotonic Devices and Networks

Nanophotonics research in the group focuses on novel hybrid plasmonic-based optical devices for integrated optical circuit and signal processing. With enhanced light–matter interaction, sub-wavelength footprint, and minimal parasitics, plasmonic devices are promising candidates for next-generation optoelectronic components that can help alleviate the latency and power dissipation bottlenecks in current VLSI technologies. Utilizing plasmonic devices to enable dense, on-chip optical integration, our research also explores reconfigurable hybrid plasmonic network structures consisting of 2D arrays of intersecting waveguides. By manipulating the network topology and therefore the interference between plasmonic waves propagating among the interconnected waveguide junctions, the network can be engineered to support different spectral responses at various output ports. Moreover, by introducing functional materials such as polymer or 2D alternatives with gate-variable properties, the dispersion characteristics of the structure can be reconfigured via biasing. We are currently investigating programmable plasmonic components for optical switching and signal processing. To reduce the computational cost associated with numerical simulations, our research also involves constructing analytical models for these plasmonic network structures. Specifically, we analyze the power distribution within intersecting plasmonic waveguides through impedance analysis and a scattering matrix model. This allows the output response of networks with any arbitrary topology to be encapsulated into closed-form expressions that do not require numerically extracted parameters. Finally, although plasmonic components can provide the field confinement necessary to support dense integration, the significant light attenuation due to Ohmic damping fundamentally limits the device performance. Our research looks to overcome this challenge through long-range, coupled plasmonic waveguide structures. By engineering the field symmetry across a common metal layer within coupled waveguide systems, we have demonstrated that the Ohmic dissipation can be minimized regardless of waveguide configuration or material platform. As a result, a radically improved attenuation–confinement trade-off can be achieved in comparison to common types of plasmonic waveguides proposed to date. This design approach provides a powerful tool for developing a broad range of plasmonic devices such as modulators and photodetectors with small footprint and low insertion loss.

PHOTONICS.LIGHT.UTORONTO.CA/HELMY/RESEARCH

Nondestructive Analysis of Liquid-, Gas- and Aerosol-Phase Nano- and Biomaterials in Optofluidics Using Optical Spectroscopy

Conducting Raman spectroscopy in hollow-core optofluidics such as photonic crystal fibres (HCPCFs) results in significant Raman intensity enhancements compared to direct sampling in cuvette. This platform can be used as a useful method for ultrasensitive detection of vibrational modes of chemical and biological molecules. The enhancement technique in all liquid-core waveguide platforms is mostly based on their use as a waveguide to confine both the liquid and the optical field over a long distance, and the degree of enhancement attained for a specific solution depends on the physical parameters of the waveguide. The great potential of hollow-core photonic bandgap optofluidics for optical sensing originates from the increased light–matter interaction volume and efficient accumulation of the Raman scattering along the extended length of the waveguide. The well-confined excitation interacts directly with the sample molecules while propagating along the length of the waveguide and Raman scattering can be efficiently excited along the fibre's entire length. In our research we utilize different optofluidic techniques for enhancing the retrieved Raman/FTIR signal of nanomaterials in liquids, gases and aerosols. Unprecedented details in analyzing various nanostructures and biological molecules utilizing optofluidic fibres such as photonic crystal fibres (PCFs) in Raman spectroscopy have been achieved. Techniques and applications that combine surface-enhanced Raman spectroscopy (SERS) with optofluidic-assisted Raman spectroscopy to enable nanomolar sensitivity of nanolitre volumes are also being examined. Recently a detailed, non destructive characterization of CdTe nanoparticles was carried out using Raman spectroscopy for solutions with QD concentration of 2 mg/mL, which is similar to their concentration during the synthesis process. Our platform allows clear vibrational modes corresponding to the structure and interactions of the QDs to be observed. These vibrational modes include those of the CdTe core, Te defects, CdSTe interface, thiol agent and carboxylate-metal complexes. These modes are correlated with the crystallinity of the QD core, interfacial structure formed upon stabilization, QD–thiol interaction mechanisms, water solubility of the QDs and their potential bioconjugation abilities.

Herman, Peter

PHOTONICS.LIGHT.UTORONTO.CA/LASERPHOTONICS

▲ 3D Laser Fabrication: Enabling Nano-Optics for the Nanosciences

Nanoscience and nanotechnology define significant trends today that seek to exploit the new physical laws encountered as the structures and devices we make become smaller and smaller. Most attention is on the quantum effects that dramatically alter the electrical, magnetic, optical, and mechanical – virtually all – properties of materials in surprising but exceedingly useful ways. Optical physics is playing a significant role in both these trends. Near-field, multi-photon and confocal microscopy and laser tweezers are opening the frontiers of nanoscience by probing and manipulating individual nanostructures, while laser lithography is a nanotechnology used in high-volume manufacturing of nanosize transistor gates in microelectronic chips. As this race to shrink the world goes forward, optical physics is also evolving with its own set of challenges to understand and to harness the new optical phenomena in nanostructures much smaller than the wavelength of light. Nano-optics has emerged as the new discipline that promises new optical materials (photonic bandgap crystals, metamaterials, plasmatronics) to guide light at dimensions below conventional diffraction limits or probe the electron wavefunction of protein molecules with powerfully enhanced optical resolution. To this end, the proposed NSERC program seeks to invent a new means of laser optical beam delivery that will facilitate the fabrication of 3D nano-optical systems. Near-field and phase-shifting techniques will be exploited in multi-level diffractive optical elements to design "intensity defects" within 3D periodic interfering laser patterns. Photosensitive optical materials exposed to these modified laser patterns will see nano-optic devices precisely embedded at the critical points of a 3D periodic lattice to enable the nanofabrication of compact 3D photonic crystal circuits, 3D optical-domain metamaterials and nanofluidic chromatography sensors for cell proteomics. This significant extension of laser holography promises a powerful advance in nano-optics and defines a new paradigm for high-volume manufacturing — contactless 3D nanomolding — of significance to Canada's optics, biophotonics and nanotechnology industries.

▲ Intelligent Beam Control for Ultrashort Laser Manufacturing of Photonic and Biomedical Microsystems

The symphony of colours and light flashes generated during laser machining attest to the dramatic undulating physics evolving rapidly as material is transformed through various states of matter, heated to sunlike temperatures, shocked to explosive pressure and finally ejected at supersonic velocity. Understanding and controlling such complex phenomena is a major science challenge and, regrettably, too poorly understood to benefit today's industry as it attempts to steer wanton laser processes into reliable nanoscale manufacturing methods for a new generation of smart medical devices, biosensors and telecommunication products. The proposed program aims to improve the fundamental understanding of laser interactions at the forefront of "burst" ultrafast laser processing and self-focusing "filamentation" machining — effects first discovered by our group. Our aim is to turn this understanding into "intelligent" laser control methods that can manage the highly nonlinear light interactions in transparent materials and possibly open a new direction for three-dimensional manufacturing. A novel burst generator provides tailored laser profiles to build up heat accumulation and annealing effects to counter shock and other collateral damaging effects. State-of-the-art delivery systems with self-learning algorithms for spatiotemporal beam shaping are to be developed for dynamic focusing into transient plasma and defect centres that promise to control the size, position, stress and morphology in various laser machining directions. Powerful "5-D" spectroscopic and phase-contrast microscopy tools will uniquely harvest the rich optical signature of the laser physics to offer real-time monitoring as optical and microfluidic devices take shape. Femtosecond laser filamentation is a new opportunity for deep penetration machining and stress-induced scribing of transparent media like flat-panel display, silicon wafers and lab-on-a-chip devices. The program aims to deliver new laser diagnostic and control systems for 3D manufacturing of telecom optical circuits, optical fibre assemblies, smart medical catheters, optical sensors and other high-value photonic systems for our Canadian partners.

Hum, Sean

WWW.WAVES.UTORONTO.CA/PROF/SVHUM

● Advanced Handset Antennas for LTE-A and 5G

The project is exploring the development of advanced handset antenna for next-generation standards such as LTE-A and 5G. A novel design technique, based on identifying the characteristic modes of the handset chassis, is used to design compact antennas providing a multitude of capabilities, including high decorrelation between antenna ports, which is important for improving the capacity of the handset, using technologies such as MIMO (multi-input/multi-output) and carrier aggregation; adaptive capabilities to optimize the capacity of the handset in a changing environment (induced by movement or user handling); and high isolation between ports (for alleviating filtering requirements in the RF front end of the handset). The result is an advanced antenna design supporting a multitude of functionalities in a single aperture.

● Multibeam Reflectors for Satellite Applications

Realizing multibeam apertures on satellite platforms while satisfying volume, mass, and cost requirements is extremely challenging, resulting in a wide array of competing architectures for multibeam systems employing one-feed-per-beam. Reflectarrays offer numerous advantages in this space, including low mass, low profile and, importantly, the ability to provide fine control over the scattered signals at different frequencies and polarizations. This creates a compelling case for investigating the use of this architecture in multibeam satellites. This project is exploring this question for practical telecommunication systems, with support from ESA.

⬢ Reconfigurable Leaky-Wave Antennas

This project is exploring the creation of very-low-profile 1D and 2D apertures which can be efficiently fed using an internal feed. As such, they can potentially offer the gain and flexibility of reconfigurable reflectors and lenses without the bulk associated with those architectures. Instead, radiation is produced through purposefully tailored leakage from surface waves travelling within the aperture. Active devices can be embedded in the aperture to provide full phase control of scatters embedded within the aperture. The result is a flat and potentially low-cost beamforming platform that can be used in a variety of applications, such as RADAR, satcom-on-the-move, rail signalling, and others.

● Reconfigurable Wideband Spatially Fed Arrays

This project is developing wideband spatially fed arrays, such as reflectarrays and array lenses, for aerospace applications. The goal is to create very flat and thin antenna apertures that can be electronically scanned that exhibit much wider bandwidths than are possible with conventional implementations of these architectures. At the same time, spatially fed architectures provide a high-performance, cost-effective alternative to traditional phased arrays. Applications include point-to-point communication systems, satellite systems, radar, and remote-sensing systems.

● Transparent Reflectarrays for Satellites

The goal of this project is to develop optically transparent reflectors based on reflectarray technology, which can be placed over solar panels on a satellite to save space and leverage the area afforded by solar panels for additional use as a high-gain antenna aperture. This project is being pursued jointly with the University of Toronto Institute for Aerospace Studies Microsatellite Technology Centre (MSTC) to test new microwave technologies on emerging microsatellite systems.

Iizuka, Keigo

WWW.KEIGO-IIZUKA.COM

▲⬟ Omni-Focus Video Camera

Our major achievements during the recent past were the invention of two novel types of distance-mapping video cameras. The first invention, called the Axi-Vision Camera, is a distance-mapping camera that is based on the combined principles of time of flight and modulated light illumination. Television programs produced by using the Axi-Vision Camera have been broadcast from NHK, Japan. In a contest sponsored annually by *Optics & Photonics News* of the Optical Society of America, the paper on the Axi-Vision Camera was selected as one of the most significant scientific accomplishments described in a refereed journal in 2002. The Axi-Vision Camera was commercialized by NHK Enterprises, Japan and the first unit was sold for $400,000. We received the 2003 Fujio Frontier Award in recognition of our leading-edge research and development of the Axi-Vision Camera. The second invention, called the Divcam (short for Divergence Ratio Axi-Vision Camera), is a distance-mapping camera that utilizes the universal decay rate of the illuminating light with distance. The Divcam is lightweight, compact, portable and reliable, has a fast response and is low cost: a U.S. Patent was filed on the Divcam and later extended to an international patent through the Patent Corporation Treaty. The omni-focus video camera, which needs the information of distance, was invented as a natural extension of the Divcam. Its invention was reported by various news organizations and magazines, including Fox News Network in the U.S. Some media even stated that the omni-focus video camera would revolutionize the global camera industry. Recently, the omni-focus video camera was used to obtain a super deep 3D image. The article "Super Deep 3D Images from a 3D Omnifocus Video Camera" highlighted this achievement: the image appeared on the cover of the February, 2012 issue of the journal *Applied Optics*.

Iravani, Reza

WWW.ELE.UTORONTO.CA/PROF/IRAVANI/IRAVANIMAIN

▰ Control, Operation and Energy Management of AC and DC Microgrids

This RD&D project addresses the challenges, strategies, solutions and technologies for monitoring, protection, control, and operation of (1) utlity-grade urban, rural, and remote AC microgrids subject to the high depth of penetration of renewable and alternative energy resources, and (2) DC microgrids.

▰ Modelling, Analysis, Control, Protection, and Operation of High-Voltage Direct Current (HVDC) Systems and Technologies

This research addresses challenges in the development of analytical and time-domain simulation models and control/protection strategies/algorithms for optimal operation of the interconnected AC power system that imbeds HVDC links and HVDC grids, mainly for large-scale integration of wind and solar power, and energy storage.

▰ Real-Time Hardware-in-the-Loop (RT-HIL) Simulation of Integrated HVDC-AC Power Systems and Microgrids

This work includes research and development of analytical and real-time simulation strategies and the corresponding hardware/software tools for the analysis, control, protection and operation of (1) large interconnected AC power systems that imbed overlay High-Voltage Direct-Current (HVDC) grids, and large-scale wind and solar power plants; and (2) microgrids with a high depth of penetration of distributed generation and storage units.

Jacobsen, Hans-Arno

WWW.MSRG.ORG

⬢ ACC — AspeCt-oriented C

AspeCt-oriented C (www.AspeCtC.net) implements an aspectoriented extension to C and offers one possible language design for an aspect-oriented C language. AspeCtC is open source, released under GPL. AspeCt-oriented C is a research project conducted by the Middleware Systems Research Group at the University of Toronto. ACC enables aspect-oriented software development with the C programming language. AspeCt-oriented C consists of a compiler that translates code

written in AspeCt-oriented C into ANSI-C code. This code can be compiled by any ANSI-C compliant compiler, for example GCC. AspeCt-oriented C is a proposed language design and compiler. ACC serves as one viable AspectC language design. AspeCt-oriented C ships with a set of Compiler Tools that help to use ACC as part of larger development projects, either to integrate aspects and ACC compiler into larger builds or to organize new software development builds with aspects in mind. The objective of the AspeCt-oriented C project is to build a robust compiler to support aspect-oriented programming with C. ACC achieves this by building on proven aspect-oriented language designs for other languages, most notably the AspectJ language for aspect-oriented programming with Java.

WWW.SITES.GOOGLE.COM/A/GAPP.MSRG.UTORONTO.CA/RESEARCH-EQOSYSTEM

eQoSystem: Towards Declarative Distributed Applications

The eQoSystem project (eqosystem.msrg.org) seeks to simplify the development and management of business processes deployed on a distributed Service Oriented Architecture (SOA). The target architecture is an enterprise system with distributed services coordinated by application workflows or business processes. Declarative goals, specified in Service Level Agreements (SLAs), are used to assist in the development of such applications and to automate the monitoring, deployment and resource provisioning tasks. The eQoSystem project is conducted by the Middleware Systems Research Group (MSRG) at the University of Toronto and is a collaboration involving IBM Toronto and NSERC.

WWW.MSRG.UTORONTO.CA/PROJECTS/PADRES

The PADRES ESB — Events and Services Bus

PADRES (padres.msrg.org) is an open-source, enterprise-grade event management infrastructure that is designed for large-scale event management applications. Ongoing research seeks to add and improve enterprise-grade qualities of the middleware. The PADRES system is a distributed content-based publish/subscribe middleware with features built with enterprise applications in mind. These features include (i) intelligent and scalable rule-based routing protocol and matching algorithm; (ii) powerful correlation of future and historic events; (iii) failure detection, recovery and dynamic load balancing, and (iv) system administration and monitoring. As well, the PADRES project studies application concerns above the infrastructure layer, such as (i) distributed transformation, deployment and execution; (ii) distributed monitoring and control; (iii) goal-oriented resource discovery and scheduling, and (iv) secure, decentralized choreography and orchestration. A publish/subscribe middleware provides many benefits to enterprise applications. Content-based interaction simplifies the IT development and maintenance by decoupling enterprise components. As well, the expressive PADRES subscription language supports sophisticated interactions among components and allows fine-grained queries and event-management functions. Furthermore, scalability is achieved with in-network filtering and processing capabilities. The PADRES research project is conducted by the Middleware Systems Research Group (MSRG) at the University of Toronto and is a collaboration involving various industry partners and Canadian funding agencies.

Johns, David

WWW.EECG.TORONTO.EDU/~JOHNS

Advanced Interface Circuits for MEMS Technology

Micro-ElectroMechanical Systems (MEMS) refer to tiny devices that combine micrometre-scale mechanical devices with micro-or nanoscale electronic circuits to sense physical quantities. Some recent examples of commercial applications for MEMS are pressure sensors used as microphones in devices such as cellphones and hearing aids, inertial sensors used in airbag deployment as well as positional control in hand-held games and cellphones and gyroscopes used for image stabilization for cameras as well as angular velocity measurement in hand-held games. This research program investigates new circuits and architectures that will significantly improve MEMS power dissipation as well as improve accuracy performance. With improved accuracy, new applications can be developed that are not otherwise possible. For example, a highly accurate inertial sensor can be used to track position by integrating acceleration to obtain velocity and then integrating velocity to determine distance travelled.

Joy, Mike

Current Density and Conductivity Imaging with MRI

In 1989 I initiated a research program whose goal was to create images of the electrical current density (CD) inside the body based on Magnetic Resonance Imaging (MRI). In the next five years this research was widened to include the imaging of tissue electrical conductivity. This work has resulted in two novel techniques, Current Density Imaging (CDI) and Current Density Impedance Imaging (CDII). The imaging of tissue conductivity has been a recurring objective since the 1930. Today, the best-known method is Electric Impedance Tomography (EIT). EIT measures currents and voltages on the skin and relates them to possible conductivity distributions in the body. Unfortunately EIT is a very ill-posed problem and consequentially has poor resolution at depth. This severely limits its effectiveness. The novelty of CDI and CDII is that MRI can be used to measure the magnetic fields arising from internal electric current density and that the conductivity can be accurately computed from these fields. This has been verified in my lab and internationally. This distinguishes CDI and CDII from EIT. The unanswered question is; “Can CDI and CDII give rise to a new technique that is more medically useful?” To be medically useful these measurements must be safe, accurate, of high temporal and spatial resolution and clinically feasible. Presently there are no methods that meet these requirements. We have used CDI in live animals (5 kg pigs) and spatial resolution of 2 cm and temporal (gated) resolution of 10–20 ms. CDII is accurate when conductivity is isotropic. Since tissues are typically anisotropic we are presently testing a technique combining MRI diffusion tensor imaging (DTI) and CDII (DT-CD-II). These MRI sequences are distinct from those developed by others in that CDI is accurate and does not depend on the nature of the tissues in which the current flows. The consequence is that we must physically rotate the tissues being imaged. In the long term I wish to remove or mitigate this requirement. In summary, my most recent success has been to be the first to image anisotropic conductivity.

Kherani, Nazir

WWW.ECF.UTORONTO.CA/~KHERANI

High-Efficiency Silicon Photovoltaics

The objective of this project is to research, develop and integrate a set of thin-film technologies that will lead to prototype demonstration of high-efficiency silicon photovoltaic solar cells. The novelty of the research lies in the development and integration of unique, production-worthy technology elements which will ultimately make it possible to attain the lowest cost per watt peak (Wp) of silicon-based photovoltaic solar cells. Silicon offers one of the highest photovoltaic energy conversion efficiencies. This property, along with its stability, abundance, environmental compatibility and technological maturity, make silicon a prime material for photovoltaics. However, the challenge today is cost. Cost can be reduced by decreasing production and material costs and by increasing energy conversion efficiency. The unique technology elements comprising this project are (1) development of high-efficiency solar cell concepts with the objective of producing the greatest quantity of solar electricity per gram of silicon (i.e., thin silicon); (2) use of low-temperature, high-quality thin-film synthesis techniques with the objective of implementing low-thermal-budget and high-production-rate processing; and (3) integrated development of PV cell concepts, photon harvesting techniques and production processes compatible with the drive to continually reduce the silicon absorber thickness.

Micro-Power Sources and Sensors

The objective of this project is the development of micro-power sources and micro-sensors for various sensing applications with the ultimate realization of self-powered sensors. Research on micro-power sources focuses on the development of continuous vibrational piezoelectric energy harvesters using (1) ambient vibrational sources and (2) tritium-occluded-in-silicon beta-emitting sources. In the latter case, integration of the beta source enables an autonomous vibrational energy generator. These vibrational energy harvesters are based on aluminum nitride piezoelectric material. Research on sensors focuses on development of an aluminum nitride–ultrananocrystalline diamond platform to synthesize surface acoustic wave nano-transducers operating at GHz frequencies. Recent research has demonstrated resonant frequencies and velocities exceeding 10 GHz and 10,000 m/s, respectively.

NICE Composite Materials

Nano-Integrated Carbon-Enveloped (NICE) composites are being developed (1) as smart coatings for energy conservation in building applications and (2) as novel materials for photonic applications. NICE composites, based on diamond-like carbon film, which is a silicon-compatible material, have thermal, mechanical, optical and electrical properties that can be tailored over extremely wide ranges, yielding a versatile material for photonic, optoelectronic and micro-electro-mechanical systems (MEMS) applications. The objective of this project is to demonstrate NICE composites as a viable platform material for the development of smart coatings for building energy applications and to explore their viability as passive and active rare-earth base photonic materials.

Photonic Crystal–Photovoltaics

The aim of this project is to investigate photonic crystal–photovoltaic integrations with the aim of creating high-efficiency, economic, third-generation solar cells. The novelty of the research lies in innovative integrations of nanomaterials and thin-film semiconductors. As thin-film crystalline or nanocrystalline silicon solar cells are made thinner, light trapping at wavelengths near the absorption edge becomes increasingly important (e.g., absorption lengths are 10 mm and ~1 mm for wavelengths of 800 nm and 1100 nm, respectively). Upon applying perfect random scattering on an incident silicon surface with a lossless back reflector, a maximum path length enhancement of ~50 is expected, though in reality the actual value is closer to 10. However, much larger path length enhancement factors, on the order of 103 to 104, are required to effectively absorb the longer wavelength light. An alternative approach is light localization through the application of photonic crystals. Photonic crystals (PC) are periodic dielectric structures that affect the behaviour of electromagnetic waves similar to periodic potentials in semiconductor lattices that affect the behaviour of electron waves. Through the exploration of a range of nano-integrations, we have recently proposed a novel class of transparent conducting porous nanocomposite films amenable to a variety of device applications. One application involves the use of a selectively transparent and conducting photonic crystal as an intermediate reflector for efficiency enhancement of tandem thin-film silicon micromorph solar cells. Other application areas include bifacial PV, LEDs and catalytic process applications.

Khisti, Ashish

WWW.COMM.UTORONTO.CA/~AKHISTI

Low-Delay Communication Systems for Streaming Media

We investigate theoretical foundations and practical architectures of communication and compression techniques optimized for low-latency applications such as conferencing and cloud computing. It turns out that traditional methods that separate compression and error correction into different modules are far from optimal when end-to-end latency is considered. Furthermore, the instantaneous dynamics of the communication channel play a fundamental role in the ultimate performance limits of low-latency systems. Therefore both the theoretical approaches and resulting architectures for low-latency communication systems are radically different from traditional approaches to reliable communication systems. The proposed project tackles this challenge in collaboration with Hewlett Packard Laboratories.

Kschischang, Frank

WWW.COMM.UTORONTO.CA/FRANK

Energy of Decoding

The capacity of an additive white Gaussian noise channel depends on its signal-to-noise ratio (SNR); the greater the SNR (i.e., transmitter energy), the greater the capacity. It is known that by using error-correcting codes of very long block length, code performance at transmission rates up to the channel capacity can be achieved. However, operation near the “Shannon limit” requires complicated encoding and decoding algorithms, which can themselves consume considerable energy (particularly at the decoder). This can amount to a large fraction of the total energy used by the communication system. Using Thompson’s VLSI model, this project reconsiders the coding problem when encoding and decoding energy is taken into account (in addition to the traditionally accounted-for transmitter energy).

Fibre-Optic Communication Using the Nonlinear Fourier Transform

Fibre-optic transmission systems are evolving at a rapid pace towards achieving greater spectral efficiencies. Coherent detection is supplanting noncoherent detection and polarization multiplexing and advanced modulation schemes are being implemented. Today’s high-speed electronics enable very sophisticated signal processing and coding to be applied, even at extremely high data rates, yet there is a significant gap between what has so far been practically achieved and what is known to be achievable in theory. In this work we study information transmission techniques based on the nonlinear Fourier transform. The nonlinear Fourier transform (NFT), a powerful tool in soliton theory and exactly solvable models, is a method for solving integrable partial differential equations governing wave propagation in certain nonlinear media. The NFT decorrelates signal degrees-of-freedom in such models, in much the same way that the Fourier transform does for linear time-invariant systems. In the proposed communication scheme, which can be viewed as a nonlinear analogue of orthogonal frequency-division multiplexing commonly used in linear channels, information is encoded in the nonlinear frequencies and their spectral amplitudes. Unlike most other fibre-optic transmission schemes, this technique deals with both dispersion and nonlinearity directly and unconditionally without the need for dispersion or nonlinearity compensation methods. Much work remains to be done, however, in translating this theoretical idea into practice.

Spatially Coupled Algebraically Decodable Codes for High-Speed Data Transmission

Optical fibres support very high-speed communication channels (hundreds of Gbits/s per wavelength) and designing error-control coding schemes that can correct channel errors at such high speeds is a daunting task. This research investigates one promising family of codes, so-called spatially coupled algebraically decodable codes, for such applications. This family includes “staircase codes,” a hardware-friendly class of codes with excellent code performance. Our ongoing research is investigating methods to incorporate soft-decision information and to combine coding with higher-order modulation.

Kundur, Deepa

WWW.COMM.UTORONTO.CA/~DKUNDUR/RESEARCH

A Cyber Security Impact Analysis Framework for the Electric Smart Grid

The scale and complexity of the smart grid, along with its increased connectivity and automation, make the task of cyber protection particularly challenging. Recently, smart grid researchers and standards bodies have developed technological requirements and potential solutions for protecting cyber infrastructure. However, grid protection remains daunting to asset owners because of resources limitations. Important questions arise when identifying priorities for design and protection: Which cyber components, if compromised, can lead to significant power delivery disruption? What grid topologies are inherently robust to classes of cyber attack? Is the additional information available through advanced cyber infrastructure worth the increased security risk? The goal of this project is to develop a framework to assess the impact of cyber attacks on the electric smart grid. Our approaches borrow from mathematical principles from control and communication theory to identify new vulnerabilities stemming from the use of cyber infrastructure and the relative physical impact of cyber attacks. One outcome is a vulnerability analysis tool that can be employed by smart grid stakeholders to identify critical cyber infrastructure that must be prioritized for system hardening.

Cyber-Physical Protection of the Smart Grid

The emerging smart grid represents an engineering system with tightly coupled and coordinated cyber and physical components. The close interaction of such diverse components may lead to emergent system behaviours and new forms of vulnerabilities. However, opportunities may also exist through the coupling to improve system survivability in the face of faults and attack. This research program pioneers the development of a modelling and analysis methodology for cyber-physical smart grid systems by harnessing the power of dynamical systems frameworks. Through integration of mathematical tools from the fields of nonlinear dynamical systems, graph theory and game theory, we aim to address timely and important system operation, control and security problems influenced by the needs of electric power utilities. The work will provide timely design insights and instruments essential for developing more reliable, secure and survivable smart grids. Solutions for resilient smart grid development and operation are just emerging and the proposed research provides a necessary framework to better assess, redevelop and prioritize them. Moreover, this research helps to reinforce the synergy among communication, computation, economic and electricity networks, fostering an important interdisciplinary view of the emerging smart grid. The ability to build resilient smart grid systems will provide commercial and environmental benefits by facilitating widespread adoption of smart grid infrastructure, revolutionizing the electricity marketplace and reducing our society’s ecological footprint.

Kwong, Raymond

WWW.CONTROL.UTORONTO.CA/~KWONG

3D Conformal Thermal Therapy of Soft Tissues for the Treatment of Localized Cancer Using MRI-Controlled Ultrasound Therapy

MRI-guided ultrasound therapy is a powerful method of cancer treatment in which ultrasound energy, guided by magnetic resonance imaging, is used to coagulate a target region of tumour. This kind of treatment has been developed as a non-invasive alternative to conventional therapies such as surgery and radiation, which often lead to long recovery times with high complication rates. Successful application of this technology for treatment of localized cancer depends critically on the ability to deliver ultrasound energy to a targeted region of the affected organ, while avoiding thermal damage to surrounding structures. This requires accurate control over spatial and temporal deposition of energy to regulate the temperature. The goal of this research is to develop advanced control strategies to enhance the treatment effectiveness of the ultrasound therapy system.

Dependability and Security in Control and Multimedia Systems

Control and multimedia systems have become increasingly sophisticated and complex. Failures in these systems can lead to large financial losses or even catastrophes. For control systems, our research combines advanced tools from control and artificial intelligence to detect failures or discover previously unknown faults. We integrate diagnostic information to reconfigure control systems so that they are dependable even when failures occur. We seek to make multimedia systems more secure by designing new strategies to embed forensic information that protects copyrights, is resilient under content manipulation attacks and deters piracy.

Lehn, Peter

WWW.ELE.UTORONTO.CA/~LEHN

Power Electronics to Enable More Sustainable Electrical Energy Networks

Professor Lehn’s research lies in the area of medium- and high-power applications of power electronics to form more reliable, cost-effective and sustainable electrical energy systems. Of specific interest is the development of converter systems and network architectures for low-cost, low-loss integration of wind, solar and energy storage resources, including plug-in hybrid/electric vehicles. Improving robustness and power quality of the electrical grid via intelligent control of power-electronically interfaced sources and loads is a major focus. Research into high-power applications revolves around exploitation of established and emerging HVDC and FACTS technologies to improve utilization and stability of power transmission systems.

Leon-Garcia, Alberto

WWW.MET.UTORONTO.CA/ALG

Application Platforms and Smart Infrastructure

We are developing systems for the control and management of resources in power utility grids, transportation systems and cities. These systems leverage the collection of state information using a vast array of sensors. We use a service-oriented approach to resource management that extends methodologies from cloud computing and we extend and apply graph theoretic models for the flow of demand in these resource networks. We are assessing the effectiveness of current supercomputing systems in the control and management of future smart infrastructures.

Autonomic Service Architecture

We are developing an architecture for a new network and service management and control system that largely manages and controls itself and is able to accommodate a multitude of existing and future applications, thus promising to be highly cost efficient and flexible. This ideal self-regulating management and control system would be responsive to ever-changing demands and even equipment failure and would autonomously regulate and optimize configurations of data flow, be able to protect itself from harmful impact — and even have the capabilities to self-heal.

Connected Vehicles and Smart Transportation

The ORF Research Excellence Project on Connected Vehicles and Smart Transportation is a collaborative project between industry, government and academia to develop an information gathering and sharing platform to enable smart applications for transportation and transit in the public and private domains. The CVST system leverages the sensing capabilities of mobile devices and public-sector sensors to provide real-time state information. Users can then make decisions that reduce travel time, increase productivity and reduce energy consumption and vehicle emissions.

Design of Converged Communications and Computing Infrastructure

We consider the design of optical networks that can provide connectivity for future datacentres that can interconnect several million servers. On this scale, power consumption and space for Ethernet cabling become severe challenges. We are designing optical switching fabrics that combine WDM and burst switching to provide extremely high capacities with high spatial compactness and low power consumption.

Green Networking

In this project we are developing resource management techniques for the operation of computer networks that are energy efficient and that have low carbon emissions. We use graph theoretic methods to synthesize flow and capacity assignment as well as topology designs that are conducive to green networking.

Green Telco Cloud

In this project we are investigating the migration of the telecom service provider infrastructure into a green cloud computing infrastructure. We model and experimentally assess the performance of existing and future services using cloud computing. We focus in particular on services that depend on wireless access networks.

WWW.SAVINETWORK.CA

NSERC Strategic Network on Smart Applications on Virtual Infrastructures

The NSERC Strategic Network on Smart Applications on Virtual Infrastructures is a partnership between Canadian industry, academia, government, education research networks and high-performance computing centres. SAVI is developing a virtualized converged computing and communications infrastructure that can support the rapid deployment of large-scale distributed applications. A key innovation in SAVI is the notion of a smart edge that complements remote datacentres to build an extended cloud. A major goal of SAVI is to develop a national testbed to support experimentation in future internet protocols and architectures as well as future large-scale applications.

Optical Networks for Ultrascale Datacentres

We are designing optical networks that can provide connectivity for future datacentres that can interconnect several million servers. At this scale, power consumption, footprint and space for Ethernet cabling become severe challenges that can be addressed through the deployment of optical multiwavelength transmission and switching technologies.

Levi, Ofer

WWW.BIOPHOTONICS.UTORONTO.CA

Optical Biosensors and Biomedical Imaging Systems

Our research interests include developing biomedical imaging systems and optical biosensors based on semiconductor devices and nanostructures and their application to biomedical diagnostics, in vivo imaging and study of biomolecular interactions. The goal of our work is to integrate sensor components into miniature functional biosensors and apply them to novel biology and biomedical applications. As such, our research is interdisciplinary and includes semiconductor device physics, optics, micro- and nanofabrication, chemistry and applications in biomedical diagnostics, cancer studies and neurobiology.

Li, Baochun

IQUA.ECE.TORONTO.EDU

Airlift: Video Conferencing as a Cloud Service Using Inter-Datacentre Networks

It is typical for enterprises to rely on services from cloud providers in order to build a scalable platform with abundant available resources to satisfy user demand and for cloud providers to deploy a number of datacentres interconnected with high-capacity links across different geographical regions. We argue that multiparty video conferencing, even with its stringent delay constraints, should also be provided as a cloud service, taking full advantage of the inter-datacentre network in the cloud. We present Airlift, a new protocol designed for the inter-datacentre network, tailored to the needs of a cloud-based video conferencing service. Airlift delivers packets in live video conferences to their respective destination datacentres with the objective of maximizing the total throughput across all conferences, yet without violating end-to-end delay constraints. To make the optimization problem easier to solve, Airlift uses intra-session network coding and the notion of conceptual flows. A real-world implementation of the Airlift protocol has been developed, which shows that our new protocol design performs substantially better than state-of-the-art peer-to-peer solutions.

IQUA.ECE.TORONTO.EDU/SPOTLIGHTS/GESTUREFLOW

GestureFlow: Streaming Multi-Touch Gestures

To support collaboration among multiple users in real time, we propose that gestures are streamed in a broadcast fashion from one user to all participating users. Streaming gestures themselves, rather than application-specific data, makes it possible to optimize the design and implementation of a gesture broadcast protocol that can be reused by any gesture-intensive application that needs to support multiparty collaboration. We believe this is a more elegant and reusable solution that serves the needs of an entire category of gesture-intensive applications. Once received, a gesture stream can be rendered in real time by a live instance of the same application on a receiver. To take such broadcast of gestures a step further, we believe that multiple gesture broadcast sessions need to be supported concurrently, so that any participating user can be the source of a gesture stream. In this work, we have designed GestureFlow, a new gesture broadcast protocol specifically designed for multiple concurrent broadcast sessions of user gestures. We point out that gesture streams typically incur low yet bursty bit rates, unlike traditional media streams. They do pose unique challenges, as gesture streams need to be received with the lowest possible delay and packet losses are not tolerable. In our design of GestureFlow, we use network coding and present a detailed design that takes advantage of inter-session network coding to support low latencies across multiple broadcast sessions.

IQUA.ECE.TORONTO.EDU/SPOTLIGHTS/MULTI-RESOURCE-PACKET-PROCESSING/

Multi-Resource Packet Processing in Software-Defined Networks

Queueing algorithms determine the order in which packets in various independent flows are processed, and serve as a fundamental mechanism for allocating resources in a network appliance. Traditional queueing algorithms make scheduling decisions in network switches that simply forward packets to their next hops, and link bandwidth is the only resource being allocated. In modern network appliances, e.g., middleboxes, link bandwidth is no longer the only resource shared by flows. In addition to packet forwarding, middleboxes perform a variety of critical network functions that require deep packet inspection based on the payload of packets, such as IP security encryption, WAN optimization and intrusion detection. Performing these complex network functions requires the support of multiple types of resources, and may bottleneck on either CPU or link bandwidth. For example, flows that require basic forwarding may congest the link bandwidth, while those that require IP security encryption need more CPU processing time. A queueing algorithm specifically designed for multiple resources is therefore needed for sharing these resources fairly and efficiently. One of the challenges we have investigated is the fairness-efficiency trade-off. The fairness-efficiency trade-off shown in the example above generally exists for multi-resource packet scheduling, but it has received little attention. Existing multi-resource queueing algorithms focus solely on fairness. However, for applications with a loose fairness requirement, trading off some fairness for higher efficiency and higher throughput is well justified. In general, depending on the underlying applications, a network operator may weigh fairness and efficiency differently. Ideally, a multi-resource queueing algorithm should allow network operators to flexibly specify their trade-off preference and implement the specified trade-off by determining the “right” packet scheduling order.

IQUA.ECE.TORONTO.EDU/SPOTLIGHTS/MATCHING

Optimizing Datacentre Operations with Practical Complexity

The unprecedented growth of mega-datacentres, in which hundreds of thousands of machines are assembled to process a massive amount of data for Internet-scale services, has been driving the evolution of computing. Designing algorithms to optimize datacentre operations is thus imperative. At the same time, the scale of the infrastructure calls for novel approaches to reduce the complexity of the solutions and make them practical. In this project, our research objective is to resolve the tussle between optimality and practicality in designing algorithms for datacentres. First, for a single datacentre we have designed Anchor, a resource management system that effectively allocates server resources to virtual machines. Instead of being optimal, Anchor is designed to be flexible and practical and uses a unified mechanism to support diverse allocation policies expressed by operators and tenants. It abstracts performance goals as preferences and uses a novel stable matching algorithm to solve the matching problem efficiently. We have also studied the problem of workload management for multiple centres distributed over a wide geographical area, where it is possible to go for both optimality and practicality. I propose to exploit the geographical diversity to reflect the electricity and bandwidth price difference at different locations and ISPs and develop a novel distributed algorithm to solve the large-scale optimization problem with faster convergence than that of traditional methods.

Liang, Ben

WWW.COMM.UTORONTO.CA/~LIANG

Broadband Multimedia Communication in the Mobile Environment

Multimedia content is the single most influential factor driving the need for increased mobile network capacity and device capability. The proliferation of cloud-based content distribution services and video social networking applications will severely stress the existing mobile systems. The inherent heterogeneity of both wireless access technologies and mobile devices enables rich and ubiquitous multimedia services, but it also significantly complicates system design. We are interested in investigations into system optimization, resource management and algorithm design to create innovative technologies for multimedia communication in the mobile environment. Promoting a synergistic approach, we work at the interface between mobile access, broadband communication and distributed-system technologies.

WWW.COMM.UTORONTO.CA/~LIANG/RESEARCH

Fair Resource Scheduling in Large-Scale Networked Systems

The principles of network science permeate wide-ranging applications such as communications, cloud computing, power grid management, transportation and biology. A central issue is how to effectively share network resources among competing agents. We are interested in developing new theories and practices for fair resource scheduling in large-scale networked systems. Examples of our investigation include cloud computing economics, distributed smart grid control, and multi-resource fair scheduling.

Heterogeneous Data Communication for Mobile Cloud Computing

Two revolutionary technologies, cloud-based computing and smart mobile devices, have fuelled the emergence of a new mobile cloud computing paradigm. On the one hand, cloud computing centres, accessible through the Internet, produce shared pools of always-on computing resources such as software, CPU and storage. On the other hand, smart mobile devices have evolved, driven by the drastic rise of hardware, application and wireless communication capabilities, to become ubiquitous tools for both content consumption and content creation. These two new technologies perfectly complement each other, with cloud servers providing the engine for computing and smart mobile devices naturally serving as human interface and untethered sensory inputs. This research targets the large-scale heterogeneous communication and networking architecture expected to serve as the backbone of the emerging mobile cloud computing paradigm. We envisage a seamless global system of computing, communication and applications, supported by a synergistically operated mobile cloud computing system, incorporating hybrid macro cloud centres, micro cloudlets, and smart mobile devices. Topics of investigation include mobile computation offloading, virtual machine placement and allocation, and cooperative joint communication-computation.

Resource Management and Optimization in Wireless Networks

In next-generation heterogeneous wireless networks, the increased number of networked devices and the broadband nature of application demands will increase the need for efficient resource sharing. The goal of this research is to

develop fundamental theories, communication algorithms and networking protocols for efficient allocation of spectrum, hardware, and power in high-throughput wireless networking environments. Topics of our investigation include cooperative communication, small-cell networks, interference management, stochastic optimization and dynamic resource allocation.

Lie, David

WWW.EECG.TORONTO.EDU/~LIE

■ Computer Systems Security

My research goal is to make computer systems safer and more reliable. With the degree to which computing has permeated our lives, from mobile smartphones to ubiquitous cloud computing, it is crucial that this infrastructure that we rely so heavily on be secure and reliable. I take a variety of approaches to achieving this goal, including techniques using operating systems, computer architecture, formal verification and networking. I like building prototypes with my students to demonstrate our ideas. My current areas of focus are building more secure smartphones and cloud computing systems.

Liebeherr, Jorg

WWW.COMM.UTORONTO.CA/~JORG

● Enabling Heterogeneous Self-Organizing Machine-to-Machine Networks

Machine-to-machine (M2M) communications have created a new paradigm for creating distributed applications that connect previously unseen numbers of intelligent devices. M2M applications must satisfy demands for low cost, scalability, and low overhead, and must be able to operate over a mix of different communication systems. Thus, M2M applications can greatly benefit from self-organizing approaches to networking, because of their ability to adapt to changes of network topology, traffic mix and service requirements. However, the use of self-organizing design principles for heterogeneous M2M systems remains largely unexplored. The objective of this research project is to harvest the potential of self-organizing networks for M2M communications. The ability to have large-scale networks that can be deployed instantaneously and inexpensively creates opportunities for simpler and more resource- and cost-efficient networking, and may lay the foundation for innovative technologies. We develop new theoretical and practical approaches to realize self-organizing M2M communications in three areas: (1) scalable routing and name resolution over collections of wired and wireless networks; (2) real-time performance monitoring for network optimization; and (3) dynamically loadable and flexible traffic control algorithms for M2M networks. The project carries the proposed solutions from theory to implementation and provides proof-of-concept prototypes.

● Hybrid Networks for Safety-Critical Mobile Communication Systems

Application-layer service overlay network solutions developed over the last decade have enabled the deployment of network services not natively available in the Internet, such as content delivery systems, broadcast video delivery and distributed directory services. Nodes in an ASON generally communicate in a peer-to-peer fashion without a requirement for servers or data centres. In a collaboration with Thales Canada Transportation Solutions (TCTS), we develop service overlay network solutions in support of safety-critical communication in train control systems. Specifically, we will use service overlay networks to establish hybrid networks, where an existing (wireless and wired) infrastructure network is enhanced by a mobile SON provide alternate communication paths. The main challenge of the project is to satisfy latency requirements of the train control system. Solutions from this project will reduce the need for infrastructure of the rail signalling system in train tunnels, as well as increase the availability of communication between train cars and the back end infrastructure.

Liscidini, Antonio

WWW.EECG.TORONTO.EDU/~LISCIDIN

⬢ Smart Power Optimization for Wireless Transceivers

Mobile Internet access has become very popular with the introduction of 3G and 4G networks that offer high-speed wireless connections. Until now, this mobile revolution has been driven by the possibility of having low-cost mobile terminals with Internet access, enabling ICT applications in education, health, government, banking, environment monitoring and business. Although several “smarter” phones with multi-standard capabilities have been introduced, the path towards a universal mobile radio is far from smooth. Nowadays smartphones are still extremely expensive compared to simple phones and have a battery life limited to a couple of days. The main reasons for these limitations are the use of dedicated transceivers for each standard supported and the ever-increasing demand for better performance and thus faster communication. These two factors nullify all attempts to reduce power dissipation and the overall bill of material. The main idea of this project is to apply the concept of reconfigurability, enabling mobile terminals to dynamically and autonomously adapt to changing environmental conditions and reducing their energy consumption. There are many examples in nature of dynamic fitting of performance to changing boundary conditions, since it represents the best way to achieve maximum efficiency in highly complex systems. Almost all ecosystems are based on this principle, which allows them to evolve while minimizing energy dissipation.

⬢ Ultra-Low-Power Transceivers for Wireless Sensor Networks

Wireless communication represents one of the most important revolutions of the last century. Although initially based only on star-mesh networks (e.g. cellular), at the end of the 1990s some wireless systems started to also adopt peer-to-peer (P2P) architectures, Wireless Sensor Networks (WSNs) being a prime example. These systems do not require base-stations since they are formed by autonomous short-range wireless nodes. All these nodes monitor and control the environment defining the working area by their spatial distribution. Since the high density of units makes the system more flexible and relaxes the sensitivity of the single receiver, in ZigBee network performance is exchanged with the possibility of enabling long-lasting and cheap devices. Unfortunately the target of a large-scale diffusion of WSNs was partially missed due to difficulty in realizing both long battery life and a high level of system integration in order to minimize the costs of the single device. Recently, with a consolidation of technologies like MEMS, the possibility of energy harvesting and the evolution of compact energy storage cells, industry interest in WSNs is rising again. The goal of this project is to realize a transceiver with average power consumption below 100 uW to operate from harvested energy, sustaining an autonomous short-range communication to enable an ultra-low-power wireless sensor network. In this case low power consumption and low costs will be achieved by combining the functionality of several building blocks. Indeed, recycling bias current and devices is the prime strategy to minimize area, power consumption and complexity of the transceiver. Furthermore, minimizing the overhead associated with each start-up-wake-up cycle will maximize the efficiency of the node.

Lo, Hoi-Kwong

WWW.COMM.UTORONTO.CA/~HKLO/QRNG

● High-Speed Quantum Random Number Generator

Quantoss is a high-speed quantum random number generator (QRNG) prototype, which is a joint effort of Mars Innovation, the Department of Electrical & Computer Engineering and the Department of Physics at the University of Toronto. It generates truly random numbers from the quantum phase noise of a laser. For more information about the technology, please refer to the following articles: B. Qi, et al., *Opt. Letters*, 35, 312–314, (2010); F. Xu, et al., *Opt. Express*, 20, 12366–12377, (2012). We plan to develop the prototype further to make it compact and low cost. We also plan to develop the software for the classical post-processing.

WWW.COMM.UTORONTO.CA/~HKLO

● Measurement-Device-Independent Quantum Key Distribution

Quantum cryptographic systems are, in principle, unconditionally secure. In practice, quantum hacking has emerged as a key challenge to their security. To foil quantum hacking, we have recently proposed an entirely new approach — measurement-device-independent quantum key distribution (MDI-QKD) — that could “short-circuit” all detector security loopholes. In other words, the system will be automatically immune to all detector side channel attacks. This is remarkable because it means that commercial QKD detection systems would no longer require any special security certifications and, in fact, they can even be manufactured by a malicious eavesdropper. We will build a prototype MDI-QKD system and demonstrate its robustness against detector side channel attacks. We will also develop the theory of MDI-QKD and take into full account various imperfections in real-life devices. Our work will allow us to use our enemy in quantum cryptography.

● Quantum Cryptography: From Theory to Practice

We seek to build high-speed (> 1 Gbit/s), unbreakable, secure communication systems based on quantum mechanics. “The human desire to keep secrets is almost as old as writing itself.” With the advent of electronic businesses and electronic commerce, the importance of encryption for secure communications is growing. Standard encryption schemes are based on unproven computational assumptions. In contrast, quantum code-making offers perfect security in communication, based on the laws of physics. Our goals are to dramatically improve both the performance and the security of practical quantum key distribution systems. We do so through system building and studying hacking strategies and countermeasures.

Maggiore, Manfredi

WWW.SCG.UTORONTO.CA/~MAGGIORE

▲ Formation Control in Multi-Vehicle Systems

This research, performed in collaboration with ECE Professor Luca Scardovi, aims at developing strategies to control rigid formations of a large class of vehicles. The vehicles in question are propelled by a thrust vector and possess an actuation mechanism that induces torques about the three body axes. Examples include quadrotor helicopters, vertical take-off and landing (VTOL) aircrafts, underwater vehicles and satellites. The challenge in this research problem is that each vehicle can sense its relative displacement, orientation and velocity with respect to nearby vehicles, but doesn’t know its absolute position or orientation. Yet, using this limited information, the group of vehicles should co-operate to achieve a rigid formation.

▲● Formation Control of Nanosatellites

NASA and the European Space Agency have proposed the deployment of nanosatellite clusters to create a platform for scientific observation of the universe. The idea is to launch nanosatellites to orbit with each satellite carrying a mirror and to assemble the cluster in a rigid formation. The result would be a large orbiting telescope with unprecedented resolution and range. One of the key challenges in deploying such a telescope is the development of formation control algorithms. The electric actuators used to propel nanosatellites (electric thrusters) produce very low thrust with low resolution. These two factors, combined with tight specifications on the accuracy of the control task, make formation control particularly difficult. This research, in collaboration with Professor Chris Damaren at UTIAS, aims at developing a formation control methodology that takes into account the characteristics of electric thrusters and solves the formation control problem with the required accuracy.

▲ Virtual Constraints: A New Paradigm for the Control of Motion

The traditional approach to making robots perform complex motions relies on a hierarchical decomposition of the control task — motion planning at the high level and reference tracking at the low level. This approach has proven to be inadequate

in complex motion control problems such as locomotion in multilegged robots or flight in birdlike robots. This research aims at developing a new paradigm for motion control. This paradigm is based on the concept of virtual constraint — a constraint on the states of a control system that does not physically exist, but can be enforced via feedback control. The literature demonstrates that this idea has been used to induce stable walking in biped robots and we believe it can be used to emulate the flight of birds and insects and the swimming of fish and, more generally, to induce complex behaviours in robots.

Mann, Steve

WWW.EYETAP.ORG/RESEARCH/MEDR.HTML

Augmented Reality Will Never Work and That's Why We Need Mediated Reality

Augmented reality, whether through hand-held iPhone applications developed in our lab and elsewhere, or by eyePhone (electric eyeglasses), has been shown to be problematic because it causes information overload. What we've learned is that an older concept called "mediated reality" overcomes these problems. We've developed various mediated-reality iPhone apps as well as eyeglass apps, etc., that help people see better and find their way better. It emphasizes the fundamentals of physics, computer science and engineering. It is also closely coupled with the undergraduate and graduate course ECE516: wearcam.org/ece516/

WWW.INTERAXON.CA

Brain-Computer-Interaction (BCI) and EEG-Based Cyborg Technologies

Brain-computer-interaction (BCI) systems developed as part of wearable computing and cyborg technologies have been widely deployed in industry. Our work was showcased at the Vancouver Olympics as part of Ontario House and continues to be adapted into various products around the world. BCI based on the Chirplet Transform (http://wearcam.org/chirplet.htm) has been the subject of a recent PhD thesis and a number of research papers. See www.eyetap.org/publications

WWW.WEARCAM.ORG/COMPARAM

Comparametric Equations and High Dynamic Range (HDR) Imaging

High Dynamic Range imaging has many applications, such as in electric eyeglasses. On the pure-math side, there's the theory of comparametric equations. On the practical side, there are applications in extending the dynamic range of imaging devices such as electric eyeglasses, portable cameras and cellphones. See "Comparametric Equations with Practical Applications in Quantigraphic Image Processing," *IEEE Transactions on Image Processing*, vol. 9, no. 8, pp. 1389–1406, Aug. 2000, which you can download from www.eyetap.org/publications

WWW.EYETAP.ORG

EyeTap Electric Eyeglasses, Personal Safety Devices and Systems

The EyeTap electric eyeglasses cause the eye itself to become both a camera and a display for computer-mediated reality that achieves augmented reality but also goes beyond it, not only augmenting but also modifying. The wearable face-recognizer puts virtual name tags on people, etc. The mediated vision helps people see better and find their way better, and generally improves their personal safety. See www.eyetap.org

Lifeglogging: Lifelong Videocapture

Since early childhood I've been wearing a computer system that captures my life. In the 1990s I miniaturized this into a necklace with fish-eye lens and various sensors (wearcam.org/neckcam.htm) and presented this work to Microsoft as the Keynote Address of CARPE in 2004. Microsoft has subsequently manufactured a similar product called SenseCam. Other companies such as DARPA, HP Labs and Nokia have also been building on this lifeglogging work. The work is known by many other names such as lifelogging, lifeblogging, CARPE, or lifestreaming. We now have a community of more than 80,000 "cyborgs" online and research continues into the mobile multimedia iPhone apps, as well as versions built inside the eye sockets of the blind.

Musical Instruments and Other Human–Machine Interface Inventions

This research looks at innovative human–machine interaction based on arrays of air jets, or the like, to create volumetric tactile input devices such as air typing, in which the fingers move through space and interact with air currents. We've also developed a computer keyboard that has no moving parts. In place of each key is a finger hole, supplied by a system that detects restrictometric parameters of air flow of waste air from the CPU fan. The research is based on Karman Vortex shedding across bluff bodies, like the shedder bars in the hydraulophone (underwater pipe organ) and turbulences as a form of input and output medium. This research goes beyond what's possible with the hydraulophone; take a look at this video and then imagine the possibilities when we miniaturize it and use air instead of water: vimeo.com/14018088

WEARCAM.ORG/ABSEMENT/EXAMPLES

Physics-Based Modelling Using Presement and Absement

Velocity is the time-derivative of position or displacement; differentiating once more gives acceleration. But what happens when you take the time-integral of displacement? The result is something called absement. Integrating again gives absity. Integrating once more gives abseleration. Absement, absity and abseleration arise in fluid flows. For example, the amount of water flowing through a valve is the absement of how open the valve is, i.e., the time-integral of the openness. Other examples of absement arise in hydraulophonic sound production (sound from vibrations in water: see www.wearcam.org/absement/

Mojahedi, Mo

WWW.MOGROUP.UTORONTO.CA

Engineering the Electric and Magnetic Dispersive Responses of Artificial Media

Many of our modern conveniences are the consequence of our ability to control and modify the behaviour of naturally occurring materials and to design and manufacture artificial materials and systems with novel properties. In electromagnetic theory, the behaviours of materials and systems are characterized according to the so-called "dispersive effects." Depending on the researcher's area of interest and expertise, he or she may use different terminologies such as delays, indices or velocities to characterize the same dispersive effects. Despite these different nomenclatures, fundamental and important relations exist among the various delays, indices and velocities. The dispersion engineering paradigm formulates our attempts to control and manipulate these various delays, indices or velocities — the dispersive effects — by synthesizing artificial materials and designing novel systems. These systems in turn allow us to control and manipulate the amplitude and phase of voltage or current waveforms and/or electromagnetic pulses in order to achieve a desired outcome. For example, the paradigm of dispersion engineering has been used to demonstrate unusual behaviours such as negative or superluminal group delays and negative refractions. In addition to scientific interest in such unusual behaviours, dispersion engineering has been used to design more functional microwave devices such as broad band phase shifters, efficient antenna arrays and interconnects with reduced latency, to name a few.

Nanoplasmonic and Nanophotonic Devices

Performance of computers is expected to eventually reach its fundamental limits in terms of speed, bandwidth, power consumption and electromagnetic interference. The problem lies partly in the degrading performance of electrical interconnects. Unlike transistors, in which functionality increases with miniaturization, the functionality of electrical interconnects degrades substantially with miniaturization. One suggestion is to replace the electrical interconnects with optical interconnects, which do not suffer from signal latency, limited bandwidth or high power consumption compared to their electrical counterparts. However, there is a major problem with optical interconnects and waveguides. The optical mode size, and hence the device size, are approximately proportional to the operational wavelength. In other words, while transistors with dimensions of approximately 50 nm are common today, the micron size of optical devices makes their integration with electronics difficult. Surface plasmon polariton (SPP) — surface waves at the interface between a metal and dielectric — may provide a solution. These plasmonic waveguides, like optical interconnects, have small latency and large bandwidth but, unlike the optical interconnects, they can easily be miniaturized. However, plasmonic waveguides have their own challenges. Chief among these are (1) large propagation losses, and (2) lack of various efficient and integrated plasmonic devices such as polarizers, directional couplers and bends, to name a few. In order to overcome the losses associated with SPP while maintaining a small device size, our group was among the first to propose a hybrid plasmonic waveguide (HPWG). The HPWG can be viewed as an optimized structure exhibiting a compromise between loss and mode size. Moreover, fabrication of our HPWG is compatible with the existing silicon technology. Our HPWG can be used as a building block for the next generation plasmonic devices such as TM- and TE-pass polarizers, polarization independent couplers and other novel components.

Moshovos, Andreas

WWW.EECG.TORONTO.EDU/~MOSHOVOS

Bandwith-Efficient DRAM Controllers in Noncoherent Systems

Embedded and mobile hand-held devices have been proliferating, enabling applications that were impossible or cumbersome with the big iron machines of the past. Each new generation of these devices offers more capabilities, enabling new applications: While early mobile devices were capable of simple tasks and low bandwidth communication, today's devices offer many more capabilities such as multimedia, navigation, digital photography, etc. As their capabilities increase, novel applications such as health monitoring will be possible. For these possibilities to materialize, mobile and embedded systems need to become more powerful while maintaining reasonable uptime. A mobile system today contains several compute engines that are all supported by an external memory device. As the computation needs increase, more data needs to be fed to these engines. The link between these engines and the memory is the memory-controller. The memory controller can greatly affect how much data the memory can provide, and at what energy cost. This work will develop memory controller technologies that will boost data feeding capabilities while taking power into account. The goal is to develop the memory-controller technology that will be used in future-generation mobile devices in support of more demanding applications, while allowing the device to stay on for longer periods of time.

Exploiting Multimegabyte On-Chip Memory Hierarchies

Several technology and application trends favour chip multiprocessor (CMP) architectures, which integrate multiple processor cores, a memory hierarchy and interconnect onto the same chip. CMPs could be used for commercial servers and for end-user systems, as they can support both multiprogram and parallel/multithreaded workloads. They can also be used as the building blocks for shared multiprocessors (SMPSs). Designing high-performance and power-aware memory hierarchies and interconnects is imperative for CMPs in order to meet the memory demands of multiple processors and applications while not exceeding power constraints. Continuing application trends towards larger memory footprints, multiprogram workloads and the ever increasing speed gap between on-chip and off-chip memory compound to put further pressure on the on-chip memory hierarchy and interconnect. Furthermore, on-chip integration presents us with new trade-offs and opportunities for optimizations that need to be exploited to deliver the expected performance/watt. Additional opportunities are provided by stack-die and on-die DRAM technology that may be used to incorporate multimegabyte caches. The key questions addressed by this research are (1) How do we manage these multigigabyte

caches; are the techniques currently used still adequate, or is there room or need to rethink these decisions? and (2) Can we exploit this tremendous wealth of on-chip storage to further optimize performance beyond what is possible by simply caching instructions and data? Accordingly, the proposed research comprises two thrusts: The first considers the use of coarsegrain tracking to achieve performance that is otherwise not possible with conventional cache management techniques. The second exploits the on-chip caches to store program metadata in addition to instructions and data. Program metadata is information collected at run time about program behaviour that can be used to anticipate and optimize for future program demands.

FPGA-Friendly Processor Architectures for Irregular Applications

Our assertion is that, as embedded applications evolve, some of them will exhibit irregular behaviour. We have demonstrated that conventional soft processors are inefficient for this purpose and have proposed FPGA-friendly designs for various processor structures. We are looking at conventional processor implementations developed for ASIC implementation, identifying inefficiencies when these are implemented on an FPGA. We are proposing alternative organizations that are FPGA-friendly instead.

Power-Aware Cache-Based Structure Design

Computing devices comprise processing elements that process digital information and memory elements for storing digital information. Because of technological constraints, memory tends to be significantly slower than the processing elements it supports. Accordingly, virtually all modern computing devices employ caches, which are additional small and fast temporary memories that serve to accelerate most references to the otherwise slow memory elements. In recent years, power dissipation has emerged as an additional critical design constraint in computing device design; power limits performance for all devices and uptime for portable ones. Low power dissipation and performance are at odds: high performance typically comes at the price of high power dissipation. Caches account for a significant portion of total power dissipation (e.g., 25% to 45% in a modern processor). Because of trends in semiconductor technology, their power dissipation is expected to increase in relative terms. Accordingly, there is a need for techniques to reduce their power dissipation while maintaining performance and usability. In addition, virtualization is emerging as a key technology for future server systems. Caches will play an important role in virtualization, as they can accelerate access to memory from devices without intervention from the processor (this is necessary for adequate performance). Accordingly, there is a need to understand and develop caching mechanisms to support this aspect of virtualization. In addition, this project investigates the caching and communication architecture for fused CPU and graphics processor systems.

Smartphone and Mobile Platform Architecture

Cellphones are changing the way we interact with each other, access information and do business. Just a few years ago, cellphones offered just voice calling and short text messaging. Today cellphone capabilities parallel those of recent, high-end desktop systems, while they include several additional capabilities to communicate and interact with the physical world, such as embedded cameras and touch, position and acceleration sensors. New applications are continuously emerging including image-based searching, speech recognition and translation. All of this is possible because smartphones incorporate significant power. However, providing this power is a very challenging task, because smartphones must operate using limited energy sources while maintaining a reasonable manufacturing cost and a relatively small physical size. A key mechanism for improving smartphone capabilities is computer architecture that studies how to build smartphones given the available manufacturing technologies, while taking into consideration the applications that these phones will use. Computer architecture faces continuous challenges for two reasons: (1) the properties of the underlying technology change significantly over time and (2) so do the applications. This work seeks to understand smartphone applications and to identify opportunities to improve smartphone architectures, leading to next-generation smartphones. The primary target of the work is applications that acquire, manipulate and use images and video in smartphones. Expected benefits include increased compute capability and functionality, novel imaging applications, improved energy efficiency and reduced cost for smartphone platforms.

Nachman, Adrian

WWW.ECE.UTORONTO.CA/PEOPLE/NACHMAN-A

MRI-Based Impedance Imaging

This ongoing project seeks to image electric properties of tissue with novel use of Magnetic Resonance Imaging apparatus. It is joint research with Professor Mike Joy's laboratory, where Current Density Imaging was first invented. Recent progress includes the first electric conductivity images of the heart in live animals and the first method to image anisotropic conductivities (such as those in muscles and brain tissue).

Spatio-Temporal Analysis of Multicontact Nerve Cuff Recordings

This project, in collaboration with Professor Jose Zarrifa's laboratory at the Toronto Rehabilitation Institute, seeks to find new inverse problem approaches for extracting information from peripheral nerves. Applications include neuroprostheses for individuals with neurological injuries and amputations, as well as a valuable tool for basic neuroscience research.

Najm, Farid

WWW.EECG.UTORONTO.CA/~NAJM

Power Grid Verification

With increased power dissipation and reduced supply voltage, modern large microprocessor chips draw over 150 amperes from the external supply! These levels of current are unprecedented in microelectronics and are a key challenge for design. Apart from the design issues of delivering a well-regulated low-voltage supply at such high current, a key problem for chip designers is to make sure that the increased

voltage drop and/or rise (due to IR drop and/or Ldi/dt drop) in the on-chip power/ground grid does not lead to functional failures. Another major problem is designing the grid so that the grid metal branches do not suffer from electromigration failures. We are aware of at least two industrial instances (a DSP core and a large microprocessor) where the chip had to be redesigned because functional failures on silicon were caused by current-induced noise on the power grid. However, checking the grid node voltages and branch currents is very time-consuming and expensive, so that it is often done incompletely or not at all. We are developing efficient techniques for verifying that the voltages and currents of the power/ground grid are safe and within user specifications and, if the grid is found to be unsafe, for redesigning and optimizing the grid to achieve safety.

Ng, Wai Tung

WWW.VRG.UTORONTO.CA/~NGWT

Smart Power Integration and Semiconductor Devices

Our research group is focusing on the integration of power devices, smart power integrated circuits and power management systems. Our group has worked extensively in the development of CMOS-compatible HV fabrication processes for automotive and consumer applications in the 40–100V range. We also have ongoing collaborative projects with our industrial research partners to develop discrete and integrated power MOSFET and silicon- and SiGe-based BiCMOS fabrication processes for smart power ICs and wireless applications respectively. In recent years, we have focused on the design and implementation of VLSI power management circuits. This includes the demonstration of integrated soft-switching topology with predictive dead-time control and a practical DVS (Dynamic Voltage Scaling) system in 2004. We are also working on the integrated DC-DC converters with digital control. One of our aims is to prove that a digital controller is a viable approach to portable power management. We were able to demonstrate an integrated DC-DC converter with dynamically adjustable power transistor size for power conversion efficiency optimization at ISPSD'06 for the first time. In 2010, we introduced a superjunction power FinFET at IEDM 2010. This is exciting work toward enabling the next-generation FINFET CMOS fabrication technology to be compatible with the implementation of smart power ICs. Power management is currently an area of intense activity. Our group is focused on the integration of the controller with power output stages. We currently have activities in all-digital on-chip temperature sensors for thermal management applications, gate driver circuits with dynamically adjustable driving strength to suppress EMI and improve efficiency, and dead-time control circuits, as well as integrated class-D power amplifiers.

Pavel, Lacra

WWW.CONTROL.UTORONTO.CA/~PAVEL/LP_RESEARCH2

Decentralized Optimization and Game Theory

We are working on decentralized dynamic optimization from mathematical problem formulation to algorithm design. The optimizing agents could be nodes in a network, channels in a link or network or even autonomous robots in a group formation. We consider either a game theoretical framework or an optimization framework. In a game theoretical framework, agents or players are endowed with an individual cost function to be optimized and the aim is to achieve a Nash equilibrium, whereby no player has an incentive to deviate from its action. In an optimization framework, our work considers a number of agents that co-operate to estimate the minimum of the sum of their locally known cost functions. These agents are to dynamically adjust their actions, in response to their individual cost and the analogous decisions made by neighbouring agents (nodes), a consensus-based idea.

Dynamic Optical Network Control and Self-Optimization

We seek to create new algorithms for automatic, dynamic network self-optimization by using system theoretical and control methods and incorporating both energy efficiency and transmission performance criteria. These algorithms will be implemented in protocols for self-management and will allow on-demand wavelength capacity to be set up, reconfigured and readjusted with minimal human intervention. The system theoretical approach of our research will lead to scalable tools and techniques that take into account the full interaction between the various layers in a dynamic adaptive network and ensure a robust network performance.

Phang, Khoman

WWW.EECG.UTORONTO.CA/~KPHANG

Friends of Design

Friends of Design is a network to promote communication within the ECE department at the University of Toronto. Our goal is to have contacts in each area of ECE willing to direct inquiries to appropriate experts within the department and the network. Inquiries are welcome from students, faculty and staff, as well as undergraduate students, outside faculty, academe and industry and alumni. Find us on Facebook at https://www.facebook.com/groups/266464955455/

Plataniotis, Konstantinos N. (Kostas)

WWW.DSP.UTORONTO.CA

Affective Signal Processing: Unravelling the Mystery of Emotions

Emotion plays an important role in our daily activities and greatly influences many areas, such as learning, decision making and interaction with others. Our decisions and courses of action are adapted according to the emotional cues we receive while interacting with others. This allows the exchange of information to be much smoother and more effective. Integrating the emotional states of a user into a human–mobile interface will provide a user-centric experience that enables the interaction to be more intuitive, flexible and efficient. We are proposing an affective signal processing system that enables real-time analysis, tagging and inference of cognitive-affective mental states from facial video and

EEG recordings. This framework combines vision-based processing of the face (e.g., a frown or smile) with EEG predictions of mental states (e.g., interest or confusion) to interpret the meaning underlying EEG and facial signals over time.

Privacy-Enhancing Face Recognition

This research encompasses novel ideas in security, biometrics, privacy and smart data management principles. It creates a radically new digital asset distribution paradigm where privacy-enhancing solutions are used to minimize risks to privacy, strengthen regulatory oversight and promote public confidence.

Privacy-Protected Video Surveillance

The proposed development is a privacy protection system for video surveillance. It protects the personally identifiable visual information of subjects appearing in video surveillance footage by performing reversible encryption on the corresponding pixel regions (e.g., the face). This is a unique and effective privacy-enhancing solution that can be applied immediately after video capture, but is reversible using a secret key, thus negating the need to store the unprotected original video footage in case an incident investigation occurs.

Signal and Image Processing for Stereoscopic Cameras, Biometric Sensors and Laser Radar Applications

We support development in the areas of (1) signal-image processing for 3D imaging applications in the frequency regimes of visible and infrared stereoscopic camera systems and biometrics sensors; (2) image and data fusion for these multisensor systems; and (3) implementing the signal-image processing developments for biometrics sensors (e.g., face -tracking features using stereoscopic cameras, vital signs from EEG and ECG sensors) and imaging aid systems for landing helicopters in poor visibility.

Poon, Joyce

WWW.PHOTON.UTORONTO.CA

Integrated Photonics for Communications and Computing

We invent, design, fabricate and measure integrated photonic devices and circuits, such as electro-optic transceivers and optical switches, for communications and computing. Our unique strength is the breadth of technologies we access. We partner with collaborators in industry, academia, and research institutes from around the world to use the most sophisticated electronic-photonic integration platforms. Our photonic devices and circuits are implemented in the following material systems and platforms: silicon-on-insulator (SOI); indium phosphide on SOI; silicon nitride on SOI; indium phosphide; and correlated electron materials (vanadium dioxide). Our goal is to demonstrate integrated photonic-electronic devices and circuits that are ultra-low-power, high-speed, and compact, to meet the needs of communication and computing systems of the future.

Integrated Quantum Photonic Circuits

We are using silicon-on-insulator (SOI) platforms available through foundry services to design and implement photonic integrated circuits for quantum physics experiments and quantum information applications. The goal of the research is to integrate the typically large, tabletop-sized quantum experiments into single chips. We are presently designing and testing entangled photon pair sources based on spontaneous four-wave mixing, quantum interference experiments, high-performance optical filters and polarization controllers and modulators for quantum cryptography.

Prodic, Aleksandar

WWW.ELE.UTORONTO.CA/~PRODIC

Power Management and Integrated Switch-Mode Power Supplies

Low-to-medium switch-mode power supplies (SMPSs), used in cellphones, computer systems, communications, vehicles, medical devices and other applications that consume power from a fraction of a watt to several kilowatts, have traditionally been controlled by analog means. This is mostly due to operation at high switching frequencies and requirements for low-power cost-effective implementation. As such, they suffer from limited flexibility and are not best suited for integration with modern digital systems. Our research has developed enabling technologies for implementing digital controllers in high-frequency low-power SMPSs and is currently focusing on fully utilizing the digital control advantages as well as on the development of novel converter topologies.

Qian, Li

WWW.ECF.UTORONTO.CA/~QIANLI

Fibre-Optic Sensing

We utilize photonics technology to create instrumentation for fibre-optic sensing and metrology. Our frequency-shifted interferometry technique has been demonstrated to have a variety of applications, such as dispersion measurement; fibre length measurements; multipoint optic sensing for stress and/or temperature sensing (used in civil structures); multipoint chemical gas sensing for environmental monitoring as well as industrial monitoring in hazardous environments; and liquid level sensing in cryogenic environments (required in space applications). Virtual-reference interferometry (VRI), a technology developed by one of our graduate students, has been commercialized in a successful start-up company.

Nonlinear Optical Devices: Ultrafast Switching and Frequency Conversion

Many photonic devices rely on nonlinear optical properties of materials. For example, ultrafast switching devices that operate in the 100 GHz range and beyond utilize the ultrafast nonlinearity of optical materials. They can be widely used in high-speed data communication and signal processing. We are developing a sophisticated model for nonlinear optical materials that possess ultrafast and resonant optical nonlinearity and will use it to design compact, ultrafast, optical logic gates for signal processing. Nonlinear optical devices are also used for frequency conversion, which has wide application in lasers and optical communications. We are developing fibre-based frequency converters that would create new types of lasers in the wavelength region (mid IR) that is difficult to access by conventional means. Such lasers may be used for environmental sensing and biomedical applications.

Quantum Communication

Quantum physics introduces revolutionary ideas that enable the creation of new tools and methods unimaginable previously. For example, in communication, quantum technology offers unbreakable communication security. Transforming fundamental quantum concepts into practical tools is, however, not without considerable challenges. We are developing advanced engineering tools, devices and systems that utilize quantum concepts and implement them using existing fibre-optic technologies. These include a specialized optical homodyne detector for quantum key distribution, a fibre-based entangled photo pair source and fibre-based quantum key distribution systems.

Rose, Jonathan

WWW.EECG.UTORONTO.CA/~JAYAR/

Automatic Interconnect Synthesis and Optimization for FPGAs

The circuitry that connects computing modules on Field-Programmable Gate Arrays has become more complex, and the creation of it is one of the slowest and most difficult parts of digital circuit design on FPGAs. The goal of this project is to make all forms of interconnect design easier to do, by automating the creation of interconnect, and making it easier to change the connectivity. Ultimately, we wish to optimize the interconnect demands of an application circuit simply by responding to higher-level performance requirements on each logical link, such as bandwidth and latency constraints. We hope to use this methodology on both small units of computation (found within a compute unit) and the more common inter-unit communciation.

WWW.EECG.UTORONTO.CA/~JAYAR/ECE1778

Creative Applications for Mobile Devices

Mobile smartphones have given rise to an explosion of creativity over the past few years. There have been exciting, inspiring and incredibly useful software apps in the areas of medicine, music, psychology, senior support, banking, cooking, global health, exploring, travel, shopping, games and many more fields. These applications have only just scratched the surface of the potential of mobile devices. As our understanding of how mobile technology can be used, many new possibilities will occur to each of us. As new hardware sensors and other capabilities are added to the phones, ever more applications will become possible. The purpose of this research/graduate course is to build a collaborative environment of creativity for new applications of mobile devices. Graduate students from all disciplines at the University of Toronto are invited to take the course for credit. This course has been taught four times, and will be available for the fifth time beginning in January 2015. It is primarily a project-based course in which the goal is to produce a working app by the end of the course. Projects will be done in groups of two or three. Students with computer programming skills will be matched with those from non-programming backgrounds to do projects in the latter students' disciplines.

WWW.EECG.UTORONTO.CA/~JAYAR/CIMSAH/RESEARCH-PROJECTS.HTML

Eye Diagnosis: Perimity

As our eyes age, we can suffer from macular degeneration, often due to glaucoma. If an eye doctor suspects you have this problem he or she will send you to have your 'visual field' measured by doing a perimetry test. The standard device for doing this costs roughly $30,000 and is very heavy. The goal of this project is to build a perimetry device using an standard Android tablet. If it works, then people around the world will be able to do this test at home, at their leisure. This work is done in collaboration with Professor Moshe Eizenman and Dr. Graham Trope at the Toronto Western Hospital.

Eye Tracking on Mobile Devices

The goal of this project is to bring low-cost and low-energy eye gaze estimation to mobile devices. The technology, based on Professor Moshe Eizenman's many years of experience in eye-gaze estimation, will allow tablets and phones to know where on the screen a person is looking. This project has an immense number of applications in psychiatry, education and commerce.

Mobile App for Smoking Cessation

The Ontario government spends a significant amount of money on funding clinics to help people stop smoking. They do this because every dollar spent here saves two dollars by preventing the illnesses related to smoking. The Nicotine Dependence Clinic on College Street, just south of U of T's Downtown Toronto campus, has been helping people quit smoking for many years, and does research on how to do this best. In collaboration with that clinic (Dr. Peter Selby and his research team), we seek to build a mobile application that will help people quit smoking — by recording their habits, helping remind them of the reasons to quit, and perhaps alerting them to imminent trigger situations.

Sargent, Edward

WWW.LIGHT.UTORONTO.CA

A Biochip for Gene-Based Disease Detection

We are building integrated circuits for the detection of a panel of biomarkers that indicate the early onset of specific types of disease. We configure nanostructured electrodes on a conventional integrated circuit; functionalize these electrodes with a nucleic acid probe having a sequence complementary to the target molecules of interest; and sensitively detect

hybridization when it occurs. We are applying the chip to the early detection of cancer and to the sensitive and rapid detection of 'superbugs' such as MRSA at the point of need.

Low-Cost High-Efficiency Photovoltaics

We seek to create low-cost high-efficiency solar cells. Our approach employs colloidal quantum dots — semiconductors that are synthesized and processed in the solution phase and that, through quantum size-effect tuning, allow the sun's full spectrum to be absorbed.

Sarris, Costas

WWW.WAVES.UTORONTO.CA/PROF/SARRIS

Advanced Radio Propagation Modelling for Next-Generation Rail Signalling Systems

The public need for rail transportation safety can be effectively served by precise train control systems, enabled by advances in wireless technologies. Communications-based train control (CBTC) systems aim to provide reliable, wireless rail signalling and train navigation via a number of access points (transponders), which cover the entire area of the railway network. A critical safety mission of CBTC system planning and installation is to ensure that the number and position of access points will maintain wireless connectivity for the trains. While lack of coverage in a cellular communication network may result in dropped calls or slow data speeds, a similar effect in a CBTC system may have much more serious consequences for passenger safety. As a result, a detailed radio survey, whereby wireless propagation measurements are carried out over the entire railway network, precedes the installation of CBTC systems. Conducting a radio survey requires significant resources (time put in by qualified personnel and funds), while the line remains out of service. Often, the pressure to complete the survey results in overestimating the necessary number of access points. This redundancy results in higher installation and maintenance costs that reduce the competitiveness of CBTC solutions. This project is focused on the development of a powerful software package that can significantly accelerate the radio survey, using advanced propagation modelling techniques to optimize the distribution of access points for CBTC systems. The project will develop a comprehensive modelling framework for radio-wave propagation in complex railway environments, validated through measurements.

Multi-User Wireless Power Transfer

The principle of inductive coupling offers a route to energy transfer for very small distances. On the other hand, the engineering of coupled magnetic resonances between capacitively loaded loops has led to the recent demonstration of energy transfer over a distance of about two metres. Subsequently, several research groups have focused their efforts on the possibility of wireless charging of hand-held devices from a medium distance. This research faces two particularly important challenges: first, the rapid degradation of energy transfer efficiency with distance and, second, the safety limitations on exposing the general public to magnetic fields. The twin challenges of efficiency and safety are intimately connected, as higher efficiency implies that lower magnetic field amplitude is needed. In this project, a new direction in wireless energy transfer will be investigated, based on near-field antenna beamforming and the emerging technology of electromagnetic meta-materials. Also, new algorithms will be developed for the creation of multiple beams that can enable the concurrent charging of multiple devices.

Stochastic Computational Electromagnetics

Research on computational electromagnetics has been dedicated to the simulation of arbitrarily complex yet well-defined structures. However, several cutting-edge research areas, notably plasmonics and nanotechnology, employ devices that are increasingly subject to fabrication process variability. Moreover, while electromagnetic simulators are now able to model large-scale wireless propagation problems, they are still limited by the inherent statistical variability of indoor and urban environments. In general, the development of powerful electromagnetic simulation tools that effectively incorporate statistical uncertainty is bound to have a far-reaching impact on the pace of technological advancement with respect to grand research challenges such as the design of low-cost yet efficient solar cells, the development of biomedical instrumentation for cancer detection and treatment and wireless service planning. The current state of the art in scientific computing under stochastic uncertainty is based on post-processing data from repetitive simulations. Not surprisingly, this approach has existed for years and it is too time consuming to incorporate in a typical engineering design cycle. In other words, while the level of complexity and the significance of modelling uncertainty are constantly rising, the relevant modelling tools have remained fundamentally the same. This project is aimed at closing this gap, in order to meet the challenge of modelling statistically variable electromagnetic structures and fields with applications in plasmonics, biomedical hyperthermia and wireless communications. Our approach is focused on the fundamental reformulation of field solvers to embed statistical uncertainty in a computationally efficient manner.

Scardovi, Luca

WWW.SCG.UTORONTO.CA/~SCARDOVI

Analysis and Control of Complex Interconnected Systems

It is well recognized that control has proven to be an essential ingredient in almost every engineering system, ranging from power and automotive systems to space missions, and that feedback is a key element in many natural phenomena, ranging from molecular pathways in living organisms to ecological systems. Recent years have witnessed an increasing interest in systems that are composed of (possibly many) interconnected units. As a whole, those systems often exhibit one or more features that cannot be predicted from the properties of the individual parts. These properties (called emergent behaviour) are not an attribute of any single entity: they are irreducible and are generated by their interconnection. Emergent behaviour can lead to surprising and useful phenomena such as memory, intelligence and self-organization in cells, but can also have disastrous consequences. Examples include the spread of infectious diseases, neuronal synchronization disorders in the brain, collective motion in bacteria, and locust swarms. It is therefore of great interest to understand the principles behind the emergence of such properties and investigate methods of controlling them. The control and systems-theory paradigm is natural in this context, but unfortunately "off-the-shelf" techniques are not always appropriate for such complex systems. In the present research effort, we propose to overcome these limitations by developing new principles and methodologies that go beyond classical stability and regulation theory. Future applications range from the domain of biological networks to the domain of complex man-made systems and include closed-loop control of neuronal synchronization, analysis and control of synthetic biological circuits, and coordination in autonomous sensing networks, among others.

Sheikholeslami, Ali

WWW.EECG.UTORONTO.CA/~ALI

Circuits for Spin Electronics

Spin Electronics (or spintronics) is a new and emerging field of science and engineering that exploits the spin of electrons, in addition to their charge, for the purpose of information storage, transport and manipulation. The ultimate aim of research in spintronics is the discovery and invention of new devices, such as spin transistors and their integration into semiconductor technology to create better functionality and performance at lower cost and complexity. The purpose of this research is to explore circuit techniques for spin-based devices that are suitable for nonvolatile memory applications, replacing conventional memory technologies such as DRAM, SRAM, Flash and EEPROM. The basic structure of a spin-based memory cell is a magnetic tunnel junction (MTJ) that consists of two ferromagnetic layers separated by a thin layer of insulating material. One of the two ferromagnetic layers is a thick layer whose magnetization is fixed. The other one is a thin layer, also called a free layer, whose magnetization can be switched so that its direction is either parallel or anti-parallel to that of the fixed layer. This corresponds to storing a digital 1 or 0 in the cell. Reading the stored bit is achieved by examining the resistance of the cell. This resistance is low for the parallel state and high for the anti-parallel state. There are several challenges in the operation of the MTJ device that must be addressed before the spin-based memory can compete favourably against Flash. Our goal in this research is to devise circuit techniques to circumvent the devices shortcomings and ease the requirements for the underlying technology.

High-Speed Wireline Signalling

This research targets circuit design for high-speed chip-to-chip signalling, backplane signalling and optical communication. This includes circuit designs for the transmitter (such as the design of MUX, equalization and driver) and the receiver (such as adaptive equalization and clock and data recovery). At speeds beyond multi-Gb/s, even a few inches of a PCB trace acts like a transmission line and as such exhibits f requency-dependent attenuation, signal reflection, crosstalk and timing jitter. The goal of circuit design in this area is to compensate for the channel attenuation, reduce signal reflections and reduce crosstalk and timing jitter so as to reduce the bit error rate (BER) of the communication link while using less than a few mW per Gb/s operation. In the past few years, we have been able to contribute to this research through the design of ADC-based receivers that allow for extensive signal equalization in the digital domain. Our latest work in this area was presented at ISSCC 2013 where we presented a 10Gb/s blind baud-rate receiver using an ADC front end. In the past few years, we have also contributed to the design of non-data-aided equalization techniques and to burst-mode CDRs. Moving forward, there are still many challenges in the area of high-speed signalling as demand for signalling speeds of 28Gb/s and beyond grows. These data rates impose stringent requirements on both the channel equalization and the power budget for these links. We strive to address these challenges in the near future.

Smith, Peter W. E.

WWW.ECF.UTORONTO.CA/~UPL

Ultrafast Photonics

My teaching and research interests have centred on the study of ultrafast photonic and nonlinear optical effects in materials, and the development and characterization of ultrafast optical devices for all-optical switching and signal processing. A number of materials systems have been investigated, including bulk and composite semiconductor materials, semiconductor optical amplifiers, organic polymers, inorganic crystals and colloidal semiconductor nanocrystals. We have found that with suitable preparation and treatment, many of these materials can be made to exhibit large optical nonlinearities with very rapid (picosecond or sub-picosecond) response times. Such materials will form the basis for a new generation of ultra-rapid all-optical signal-processing devices. These devices, because they operate at ultrafast rates in the optical domain, would eliminate the 'electronic bottleneck' that limits the capacity of current-day data communications systems.

Sousa, Elvino

Autonomous Infrastructure Wireless Networks

The research focuses on our vision for 4G wireless networks. This vision stresses the deployment aspects of the physical layer, including features that allow the physical layer to autonomously configure itself after the deployment of base stations or access points by users. The base stations and access points are deployed by users in a random manner and the network infrastructure is built and grows organically. This approach to wireless network deployment will greatly reduce the cost of base stations and access points and result in networks with much greater capacity, which is required for the emerging broadband wireless services. This vision encompasses the current industry developments referred to as femtocells. The work is also related to what is referred to as cognitive radio.

Stumm, Michael

WWW.EECG.TORONTO.EDU/~STUMM

System Software Performance Optimizations

Our primary objective is to make improvements to operating systems that significantly enhance kernel and application performance. Currently, we are primarily targeting multicore-based systems. Our general approach is to exploit Hardware Performance Counters (HPCs) that today are an integral part of all processors, and use them to measure overheads and identify bottlenecks online and in real time. We then feed the information gathered from the HPCs to the system resource managers so that they can make informed decisions on how best to use system resources from a performance point of view. Using this approach, we recently introduced exception-less system calls to Linux that improved the throughput of Apache by over 100% without any modification to Apache and we improved the throughput of MySQL by 40%.

Tate, Joseph (Zeb)

WWW.ELE.UTORONTO.CA/~ZEB

Online Parameter Estimation for Wind-Driven, Doubly Fed Induction Generators

The rapid increase in variable generation technologies such as wind and solar power throughout many nations' power grids has the potential to significantly reduce reliability. To ensure this does not happen, power companies must run a multitude of simulations that identify potential problems before they occur in the real system. The ability of these simulations to accurately inform decisions is only as good as the models being used and the lack of confidence in dynamic models is one of the main problems associated with these new-generation technologies. For several reasons — such as model order reduction to make simulations tractable, the reluctance of generator manufacturers to release detailed models or parameter sets and the relatively high installation of wind generation — the accuracy of wind-generator models is becoming increasingly important to planning engineers. This project is looking at ways to use ambient wind power generator measurements (i.e., without introducing artificial stimuli to the system) to determine wind parameter models for use in simulation-based studies. Thus far, we have been exploring the performance of various nonlinear parameter estimation schemes, in particular the Extended and unscented Kalman Filters (EKF and UKF respectively), to select the most appropriate algorithm for this application. We have seen via simulation with high-bandwidth sampling that, for a relatively simple DFIG model, the UKF is generally superior to the EKF in both robustness and speed of convergence, confirming the UKF's superior performance when applied in other disciplines. Two extensions are currently being investigated: first, the impact of reduced bandwidth sampling (e.g., using 30–60 samples per second) and second, testing of the UKF estimator with AC measurements.

Phasor Measurement Unit Data Characterization and Compression

Phasor measurement units (PMUs) are the primary smart grid component being added to the North American transmission network (i.e., the high-voltage network used for large, inter-area power transfers). One of the main reasons these units are being introduced is to enable wide-area situational awareness and control of the power grid. These applications will require substantial investments in cyber-infrastructure and this research project is looking at ways to both characterize PMU data and use this characterization to achieve high levels of data compression. Results obtained thus far indicate that accounting for the unique characteristics of PMU data can lead to significantly higher lossless compression ratios in comparison to those of generic lossless compressors.

Wind Impact Metrics for Short-Term Power Grid Operations

One of the main challenges associated with the increasingly widespread introduction of wind generators is figuring out ways to control their inherent variability. While operators have always had to deal with uncertainty in electricity utilization, the availability of generation resources has traditionally been either controllable and/or known in advance. As the supply mix moves more towards variable generation resources such as wind and solar power, operators will have to learn ways to anticipate problems and take corrective actions in order to maintain system reliability. This research focuses on ways to quantify and visualize the potential impact of wind generator variability over short time horizons (e.g., four hours in the future), so that operators can have a better understanding of potential problems on the network. Because the potential impacts on the grid depend heavily on both the levels of wind generation and their distribution throughout the system, most of the work thus far has focused on development of accurate ARMAX models that account for the non-independence of wind generators' outputs. Once these models have been developed, the next stage of this project will focus on formulation and calculation of metrics that use the forecast statistics to highlight bout potential grid problems and suggest appropriate preventive controls.

Taylor, Joshua

WWW.ELE.UTORONTO.CA/~JTAYLOR/INDEX.HTML

Gaming in Modern Electricity Markets

Ten years ago, California's power system was rocked when energy traders manipulated vulnerabilities in electricity markets. The resulting California Electricity Crisis culminated in blackouts and economic losses of millions of dollars. This was but one of many widespread examples of participants in power markets exploiting design flaws, leading to increased risk of physical failures as well as unnecessarily high electricity rates for end users. Considerable analysis and experience have yielded power markets that, while still vulnerable, have not experienced egregious abuses in recent years. As we shift our dependence onto renewable energy sources, energy storage and demand response resources, power markets are changing to reflect the changing physical landscape. This will inevitably introduce new vulnerabilities, which could potentially lead to new disasters like the California Electricity Crisis. It is therefore imperative that power markets be systematically designed to induce fair and honest participation among market participants. In this project, we apply game theoretic tools to assess the vulnerabilities of power markets, and use mechanism design to develop countermeasures that ensure market participants do not have incentives to game the system. For example, by examining the equilibrium of a dynamic game model of energy storage markets, we can see if tactical behaviours can lead to poor social outcomes. Tools likes the Vickrey-Clarke-Groves mechanism enable us to make such tactical behaviours unattractive by imposing auxiliary payments like an upfront tax.

Learning to Manage Electrical Loads

It has never been possible to provide electricity with 100% reliability; this will become even more of an issue as we increase our reliance on volatile renewable sources of energy like wind and solar. Demand response programs incentivize loads to modify their electricity consumption to accommodate uncertainty in the power supply. For example, an office building may receive a reduced electricity rate for allowing their air conditioning to be shut off a few times per year, relieving a stressed power system on the hottest, most demanding days of the summer. Demand response has many advantages like low infrastructure cost and fast response times, but presents a number of new challenges because the number of electric loads dwarfs the number of traditional generation resources, and the characteristics of each load are fundamentally uncertain. For example, the state of a load may change as a function of weather, evolving hardware components or the people who use it. In a demand response program, each time a load is utilized, new information about it becomes available. In this project, we investigate how load aggregators can improve their capabilities by factoring learning into their demand response algorithms. The problem is both very large in scale and high-dimensional in its uncertainty, necessitating approximations with the highest tractability.

Trescases, Olivier

WWW.ELE.UTORONTO.CA/~OT

Battery Management for Electric Vehicles

Despite numerous technological innovations, the proliferation of EVs in Canada is primarily limited by the range and cost of today's vehicles. Reducing the cost and extending the range of EVs is a major multidisciplinary challenge faced by the global automotive industry. Advances in lightweight materials, battery chemistry, battery management and power electronics are needed to meet future customer expectations and convert entire fleets from gasoline to EV technology. Another major hurdle in the widespread acceptance of EVs is the uncertainty about the lifetime and reliability of the battery pack, especially in the harsh Canadian climate. This has delayed the adoption of Lithium-Ion (Li-ion) based battery technology until very recently, despite vastly superior energy density compared to the Ni–MH batteries used in the first generation of hybrid vehicles. Making better use of the energy capacity by increasing the system efficiency is the key to reducing the overall size and cost of the EV battery.

Regenerative braking (Regen) is often used in electric vehicles to capture kinetic energy that is otherwise wasted in the brake pads when the vehicle comes to a stop. Instead of simply applying the mechanical brakes during deceleration, an EV equipped with a Regen system uses the motor as a generator in order to transform mechanical energy into stored charge in the battery. Even the latest lithium-based batteries have a relatively poor ability to quickly absorb energy without affecting long-term performance. The maximum output power of modern Li-ion batteries is typically at least three times higher than the maximum input power. Repeatedly using Li-ion batteries to both absorb this large negative burst of power during Regen and provide a large positive power burst during acceleration can significantly raise the pack temperature and accelerate aging. Automotive-grade Ultracapacitors (UCaps) have recently been developed as an energy storage technology to complement batteries. Commercial Ucaps have input and output power densities on the order of 12 kW/kg, which is at least one order of magnitude higher than that of Li-ion batteries. On the other hand, the 6 Wh/kg specific energy of these Ucaps is at least 10 times worse than that of Li-ion batteries, leading to the concept of using a hybrid storage system consisting of a smaller Li-ion battery and a Ucap. Using this approach, the battery serves purely as an energy tank, while the Ucap is sized to meet the surge input and output power requirements. Effectively managing the energy flow between the Ucap, the battery and the motor requires new power electronic topologies and advanced control schemes. The main goal of this project is to develop new models, control schemes and power electronic converters to extract the maximum performance from modern EV energy storage systems.

High-Frequency Digitally Controlled DC-DC Converter ICs

As the world faces unprecedented environmental challenges, energy efficiency and power management have taken centre stage. Switched-mode power supplies (SMPSs) are the key enabling technology for efficiently delivering the tightly regulated supply voltages required by today's modern mixed-signal (digital+analog) integrated circuits (ICs) and systems. The SMPS acts as the interface between the energy source, such as a battery, and the load ICs. A typical SMPSs uses a combination of high-speed, low-resistance semiconductor switches, energy storage components, sensors and control circuits to regulate one or more output voltages in the presence of disturbances. State-of-the-art SMPSs have a power conversion efficiency above 90%. The resulting low heat dissipation allows multiple SMPSs to be integrated with their load circuits into a single IC. The clear trend in SMPS research is toward adaptive digital control loops, increased integration within system-on-chip (SoC) applications, higher efficiency over the full operating range and higher switching frequency, resulting in smaller energy storage components. The long-term goals of the proposed research are to make tomorrow's power management systems smaller, more efficient, more robust and more reliable, while reducing electromagnetic interference (EMI) and environmental impacts. The research focuses on new high-frequency control schemes, system-level optimization, thermal management, low-power mixed-signal circuits and power MOSFET optimization.

Power Converters for High-Efficiency LED Lighting

Electric lighting accounts for approximately 11% of the world's total power consumption. The development of cost-effective power management circuits for compact fluorescent bulbs (CFLs) has led to drastic improvements in lighting efficiency. While CFLs are clearly an improvement over archaic incandescent bulbs, they are difficult to dim, they contain poisonous mercury and the chromatic properties of the light are less than ideal. High-brightness (HB) LED modules are rapidly emerging as a promising candidate to replace CFLs in numerous lighting applications, since HB-LEDs are mercury free, scalable and can be easily dimmed. The main goal of the project is to develop smart digital switched-mode power supplies (SMPSs) to optimize the lighting efficiency and chromatic properties of HB-LED systems using closed-loop thermal and optical feedback. The targeted controller can efficiently regulate the electrical-to-optical energy conversion process under a wide range of operating conditions.

Power Electronics for Photovoltaic Applications

Solar energy has long been recognized as one of the most abundant forms of clean renewable energy. Countless research efforts around the globe are contributing to the steady decline in the cost of photovoltaic (PV) power, with the promise of reaching grid parity in the near future. This is a complex target, as the prices of conventional energy sources are constantly in flux and heavily dependent on government subsidies. The penetration level of solar power is rapidly increasing in most developed countries because of government incentives and multidisciplinary technological advances. The exponential growth of PV technology presents tremendous opportunities for all companies in the semiconductor supply chain, ranging from discrete power devices to mixed-signal control ICs. Performing maximum power point tracking (MPPT) on a PV array is used to continuously optimize the total harvested power under time-varying temperature and illumination fluctuations. It has been demonstrated that performing distributed MPPT (DMPPT) on a per-panel or even per-cell basis, instead of using a single MPPT controller across the entire PV array, can substantially improve the total system efficiency under partial shading conditions. The main goals of this project are to quantify the benefits of DMPPT for different levels of granularity and to develop new high-efficiency power electronic converter topologies and control schemes for both monocrystalline silicon and multijunction III-V PV systems.

Triverio, Piero

WWW.WAVES.UTORONTO.CA/TRIVERIO

Electromagnetic Transients in Power Distribution Networks

We develop numerical models and algorithms to predict electromagnetic transients in power distribution networks. Lightning, faults and switching activity can induce fast transients on power networks, potentially compromising grid stability. As network complexity increases, because of the penetration of renewable sources and distributed generation, predicting these phenomena becomes more and more challenging. Our techniques provide a fast way to investigate broadband transients in large power networks made by overhead, underground and submarine cables.

Modelling and Simulation of Complex Systems

Numerical techniques for the simulation of complex systems are a strategic asset in many scientific and industrial projects. However, computational complexity is often a big issue. Our group develops techniques to generate compact models for highly complex components based on system identification and model order reduction. Models can be extracted from high-fidelity simulations or experimental results and enable a fast simulation of large-scale systems. This approach has been applied to the design of high-speed circuitry and to the thermal simulation of 3D integrated circuits with liquid cooling.

Signal Integrity and Electromagnetic Compatibility Engineering

Signal integrity and electromagnetic compatibility issues, like crosstalk and interference, are a major concern in the design of electronic systems. We develop efficient mathematical models, seamlessly compatible with mainstream design tools, to predict and minimize these issues from the earliest stages of design. Through our models, designers can maximize product reliability and performance without resorting to cos tly prototyping. This research activity is of immediate interest for the microelectronic, automotive and aerospace industries.

Truong, Kevin

APEL.IBBME.UTORONTO.CA

Computational Tools for Protein Sequences, Structures and Networks

Cells are composed of protein signalling networks that perform biological functions such as regulating cell growth or catalyzing biochemical reactions. As a result, the malfunction of proteins often causes human illnesses such as Alzheimer's disease, heart disease and cancer. My long-term research goal is to create synthetic protein signalling networks that will allow us to one day manipulate cell biology with the same precision as we do electrical circuits and computer networks. To accomplish this goal, my proposal will focus on developing computational tools for studying protein sequences, structures and signalling networks. First, to infer the function of a protein sequence, the Smith Waterman (SW) algorithm is used to find its similarity to proteins of known function. As sequence databases grow larger, faster sequence comparison approaches are required, such as using accelerated field-programmable gate array (FPGA) hardware. To make the FPGA solution more affordable, I will develop FPGA hardware for accelerating the SW algorithm using fewer resources while maintaining a comparable speed. The next task is to study the protein signalling kinetics within cells. Fluorescent protein biosensors are powerful tools, but the design of these biosensors often consists of trial and error. Using a computational tool to model the conformational space of protein biosensors, I improved the design, but the tool was not quantitative. To address that problem, I will include molecular factors that select preferred biosensor conformations. Last, to design synthetic protein networks or model larger existing networks, I will develop a computational tool for simulating the spatial and temporal kinetics of protein signalling networks. The sum of this work will yield insights into protein sequences and their networks that will ultimately aid in developing therapies for human illnesses.

Live Cell Imaging and Control of Caspase Kinetics Using Engineered Proteins

Over the past decade, members of the caspase family of proteases have been extensively studied for their critical role in apoptosis. The caspase family displays rich spatial and temporal kinetics in living cells, such as cascading activation and differential subcellular expression. While such characteristics confound many biosensor designs, they accentuate the strengths of fluorescent protein biosensors. By employing the principle of fluorescence resonance energy transfer (FRET), protein biosensors can be created to image the kinetics of caspase activation in living cells. Furthermore, we can control the exact moment that caspase activation occurs within the cell using an inhibitory protein of caspase that is engineered to be switchable on [Ca2+]. This goal will be accomplished by achieving three things: (1) targeting caspase biosensors to subcellular organelles ; (2) imaging caspase cascades in living cells; and (3) finally, engineering proteins to control caspase activation based on XIAP (an X-chromosome-linked inhibitor of apoptosis protein) and a Ca^{2}+ binding protein called calmodulin (CaM). This work will pioneer designs for engineered proteins that will provide new tools for fundamental studies in cell biology.

Valaee, Shahrokh

WWW.COMM.UTORONTO.CA/~VALAEE

Localization of Wireless Terminals in Indoor Environments

Location-based services (LBS) are emerging as new applications on mobile phones. In LBS, the main challenge is to locate the user, especially in indoor and covered areas where GPS service is not available or has unacceptable errors. In this research we estimate the location of a mobile phone using the strength of signals arriving from Wi-Fi access points. We have designed and developed the system on three Wi-Fi-enabled phones and PDAs and have tested it in an office building at the University of Toronto, in a shopping mall in north Toronto and in the Canadian National Institute for the Blind (CNIB). The system has a tracking and navigation system that uses voice instruction to help visually impaired individuals find their way in indoor environments.

Wireless Communications in Vehicular Environments

In 2000, an estimated 6,394,000 motor vehicle crashes were reported to the police in the U.S. Based on a series of in-depth investigations of police reports and on-scene investigations, human factors were found to be the definite cause of 70.7% of the crashes. Most of those accidents could have been prevented if devices that allowed vehicle-to-vehicle (V2V) and vehicle-to-roadside (V2R) communication had been installed in cars. The goal of our research is to create vehicular ad hoc networks (VANET) that can be used to enhance safety on roads and to provide telematic services such as information on road conditions and traffic congestion and mapping.

Veneris, Andreas

WWW.EECG.UTORONTO.CA/~VENERIS/ANDREAS VENERIS.HTM

CAD for VLSI Verification, Debugging, Testing and Synthesis

The semiconductor industry has products reaching all aspects of commercial and consumer markets domestically and internationally. The rapid growth of this industry in the past 30 years is attributed partly to advances in the Electronic Design Automation industry community and Computer-Aided Design (CAD) tools to assist engineers designing such complex high-performance devices. The research of our group centres on the development of CAD software that expedites the verification, debugging, synthesis and testing of computer chips. Our research has been funded by major government and industrial partners and has won awards at premiere conferences for its impact on the community.

Venetsanopoulos, Anastasios

WWW.DSP.TORONTO.EDU

Digital Signal/Image Processing, Digital Communications, Biometrics and Biomedicine

In digital signal/image processing our work is contributing to both fundamental and applied areas. New digital filters, such as linear, nonlinear, multidimensional and multispectral, both fixed and adaptive, have been introduced.These filters have been applied to geophysics, imaging and financial data, and to biomedicine, including radiology, tomography, mammography and MRI. In the area of telecommunications, we have dealt with the problems of sonar transmission through fading dispersive random channels and image compression, focussing on techniques called perceptually lossless techniques. We have also explored biometrics and multimedia.

EN.WIKIPEDIA.ORG/WIKI/ANASTASIOS_VENETSANOPOULOS

Research in Telecommunications, Signal and Image Processing, Multimedia and Biometrics

We have four research topics in this area: (1) multimedia (image compression, image and video retrieval); (2) digital signal/image processing (multichannel image processing, nonlinear, adaptive and M-D filtering, knowledge-based processing and recognition, 3D imaging, and biomedical applications); (3) telecommunications; and (4) biometric research.

Voinigescu, Sorin

WWW.EECG.TORONTO.EDU/~SORINV

▲ ⬢ 200 GS/s DACs and ADCs for Optical Transceivers with QAM and OFDM Modulation

As internet traffic continues to increase exponentially because of the explosion of mobile multimedia devices, there is renewed demand for electronic circuits and optoelectronic systems that can operate at serial data rates in excess of 50 Gb/s. Current OOK- and QPSK-modulated optical transmitters are limited to 50-Gb/s bit streams because of the limitations of optical filters and optical fibres. This is one of the reasons for the widespread adoption of QPSK modulation in the 110-Gb/s Ethernet standard. Rather than transmitting and receiving a single 110-Gb/s NRZ bit stream, four NRZ bit streams, each at 28 Gb/s, are employed to modulate two orthogonally polarized optical carriers in-phase and in-quadrature. The next generation of optical fibre standards will inevitably imply the deployment of both higher-level modulation schemes such as 16-QAM and OFDM and of higher serial bit streams at and beyond 112 Gb/s. Optical modulators typically require 3ï¿½5V electronic signal swing for proper operation. The large voltage swing and the very broadband operation from DC to over 55 GHz are beyond the best performance reported for digital-to-analog converters in nanoscale CMOS and SiGe BiCMOS technologies. One option would be to use a low-voltage swing DAC followed by a very large voltage-swing, large-gain, broadband linear amplifier fabricated in III-V technology. This is an expensive multichip solution. In addition, because of the large resolution (7ï¿½10 bit) required, it is almost imperative that the DAC directly drive the optical modulator to avoid signal distortion. This proposal seeks to research and develop large swing (>3V differential), multibit (6ï¿½10 bit) BiCMOS-55-nm and 28-nm SOI CMOS DAC topologies for 16-QAM, 64-QAM and OFDM optical transmitters. One such SiGe BICMOS DAC, developed in our group and operating at 50 GS/s with 6 Vpp differential swing, was presented at IEEE IMS in 2012 and won the best student paper award. An NSERC CRD grant was recently approved to extend this research to the investigation of 200 GS/s ADC topologies for 1Tb/s fiber-optic receivers.

▲ ■ 56 GS/s 7-bit DAC and ADC

The research focuses on architectures and physical implementation of low-power 56-GS/s, 7-bit digital-to-analog and analog-to-digital for next-generation optoelectronic transceivers with multilevel modulation formats. The circuits will be fabricated in the world's most advanced 55-nm SiGe BiCMOS and 28-nm FDSOI technologies.

▲ Atomic-Level and 2D Crystal Electronic Devices

This exploratory research focuses on the simulation, design and fabrication of novel nanoscale metal nanowire transistors. The ultimate goal is the physical implementation of subï¿½10 nm transistors.

▲ ⬢ High-Efficiency mm-Wave Transmitter Array

This research focuses on novel transmitter array architectures based on high-efficiency direct mm-wave digital modulators. The goal is to deliver 64 QAM modulated transmitters with over 65% power-added efficiency and over 36, 30 and 26 dBm of output power at 45 GHz, 94 GHz and 138 GHz, respectively. New circuit topologies, IQ correction techniques and free-space power combining and modulation schemes are explored. Both SiGe BiCMOS and 45 nm SOI technologies are being investigated for the practical implementation.

▲ Low-Power mm-Wave Distance Sensor

Ultra-low-power single-chip mm-wave distance sensors and active tags will be investigated and demonstrated in silicon at 60 GHz and 80 G.

Wong, Willy

WWW.INDIVIDUAL.UTORONTO.CA/WILLY

⬟ Sensory Neuroengineering

My interests are in the areas of neuroscience, signal processing and sensory engineering, with particular application to speech and hearing. My work lies at the intersection of biomedical, computer and communications engineering. Students working in my group typically have an interdisciplinary approach with interests in signal processing, algorithm design, modelling and psychology. Currently I have three active project areas: (1) Theoretical studies of sensory information processing. We are developing a physics-based model of the process by which the sensory system transmits information from the periphery to the brain. The goal of this project is to elucidate the underlying physical laws that govern the process of sensation. (2) Biomedical signal processing. We work with a team of doctors and surgeons, using signal processing and data mining techniques to better understand and diagnose neurological disorders like schizophrenia or major depressive disorder. We are also engaged in invasive brain studies for rehabilitation applications. (3) Acoustic and speech processing. We focus here on problems related to speech and hearing loss and the tools that we can develop to aid those with communication disorders.

Wonham, Murray

WWW.CONTROL.UTORONTO.CA/~WONHAM

▲ Supervisory Control of Discrete-Event Systems

Our research is on supervisory control of discrete-event systems, that is, logic control of systems described in a framework of automata and formal languages. We focus on system architecture and 'intelligent' computing techniques as a means of confronting state-space explosion and exponential complexity. Specifically, architectures include monolithic (as a 'worst case'), refined to decentralized, distributed, and hierarchical and their combination as heterarchical system organizations; while computing includes modelling and algorithmic development using state-tree structures, an adaptation of state charts. Existing

applications by ourselves and others include manufacturing workcells, chemical engineering processes, automobile push-button systems and industrial diagnostic systems.

Yoo, Paul

WWW.IBBME.UTORONTO.CA/FACULTY/CORE-FACULTY/PAUL-YOO

⬟ Electrical Neuromodulation for Bladder Dysfunction

The objective of this project is to develop a neural prosthesis for restoring bladder function in individuals with neurological or idiopathic disorders. This involves (1) the investigation of neural mechanisms underlying various stimulation-evoked bladder reflexes; (2) development of minimally invasive methods of electrically engaging these neural pathways; and (3) the clinical translation of these technologies in patients with chronic urinary dysfunction.

⬟ Neural Interface for Vagus Nerve Stimulation Therapies

Vagus nerve stimulation is an emerging neurostimulation modality that is currently used clinically to treat epilepsy and (more recently) chronic heart failure. However, the widespread and reliable use of this technology has been hampered by the limited ability to activate and record neural activity from this compound nerve trunk. This project aims to develop innovative peripheral nerve interface technology designed for implementing closed-loop control of neuromodulation therapies.

⬟ Novel Electrode Platform for Obstructive Sleep Apnea

Although CPAP remains the primary means of treating the symptoms of obstructive sleep apnea, neurostimulation technology (e.g., hypoglossal nerve stimulation) has emerged as an effective alternative for patients noncompliant with long-term CPAP therapy. The primary objectives of this research are to (1) develop minimally-invasive neurostimulation technology for treating symptoms of obstructive sleep apnea; and (2) uncover the pathological mechanisms that contribute to the clinical presentation of repeated airway collapse during sleep.

Yu, Wei

WWW.COMM.UTORONTO.CA/~WEIYU

● Design and Optimization of Next-Generation Wireless Cellular Networks

Prof. Wei Yu's research program aims to use novel cooperative transmission techniques to enhance capacity, coverage and reliability of wireless cellular networks. The goal is to develop novel coordinated signal processing, resource allocation and network optimization techniques in the design and analysis of interference management and mitigation methods for future wireless networks. Interference mitigation is expected to become a crucial task in wireless system design as future networks become more densely deployed, frequencies more aggressively reused and the network topologies increasingly heterogeneous. Prof. Wei Yu's research program focuses on two network architectures in particular: the cloud radio-access network (C-RAN) architecture where base-stations cooperate in transmitting and receiving signals and the heterogeneous architecture where remote radio units are deployed within the cellular structure to enhance coverage. Prof. Wei Yu's research program aims to advance the state-of-the-art in the theoretical capacity analysis of wireless networks and to impact the design philosophy, standards development and evolution of future-generation wireless networks.

Yuan, Ding

WWW.EECG.TORONTO.EDU/~YUAN

■ Understanding and Improving the Availability and Performance of the Cloud

With the massive growth of internet connectivity, software is increasingly being delivered by internet services — a model commonly known as 'cloud computing'. Such online services are almost entirely implemented as distributed systems, and their availability (i.e., uptime) and performance (i.e., latency and throughput) are critical as system outages now will simultaneously affect all users. Our research goal is to improve the availability and performance of these cloud-based distributed systems. In particular, our group focuses on (1) understanding the real-world failures that occur in these systems; (2) developing practical techniques to detect the defects and thus prevent the outages from occurring; (3) developing practical tools to debug the performance of such systems. Our initial investigations have been successful: we have detected hundreds of new bugs in some of the most widely used distributed systems, including Hadoop, HBase, Cassandra, Hive, etc. Many of these problems have already been fixed by the developers since we reported them. Our techniques have also resulted in multiple papers presented at OSDI, the premiere conference on operating systems.

Zhu, Jianwen

WWW.EECG.TORONTO.EDU/~JZHU

■ Compiling Software to Silicon

My research focuses on highly automated methods of designing high-volume systems-on-chip, as well as field-programmable gate arrays (FPGAs). The automation is achieved by solving a series of optimization problems that ultimately allow the design of integrated circuits to be as easy and as fast as programming software in C/C++, while not compromising circuit performance. More specifically, these methods can help chip design companies to quickly design hardware processing engines for today's smartphones. These engines, such as video codecs, graphic processing units and security and baseband processors, are the key differentiators that allow smartphones to provide a rich user experience within the tight energy budget of batteries. Similarly, these methods, by utilizing FPGAs, can help telecom and IT infrastructure equipment vendors improve the performance and lower the cost of equipment that will power the next-generation internet and wireless infrastructure.