22 BEHIND THE IoT CURTAIN
Everyone's ready for the Internet of Things—but what exactly is it?

34 SECTION GRADUATE
GET OFF MY GRID
Smarter, cheaper, greener: Nanogrids

38 SECTION RESEARCH
PREPARE FOR TAKEOFF
The math behind self-flying drones

48 SECTION ALUMNI
THE MORE THINGS CHANGE...
Class Notes from around the world

2015
A YEAR IN REVIEW

ANNUM

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

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

ENGINEERS LEARN BY DOING.

Our new Energy Systems Lab puts the latest technology in the hands of our students, so they can do what's never been done before.

To design systems robust enough to handle unpredictable energy demands, engineers need to get as close to the real thing as possible. In The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, we've completely overhauled our Energy Systems Lab to give both our undergraduate and graduate students hands-on experience with power, communications and smart-grid systems operating at specs they'll find in the real world.

We launched this three-stage renovation in summer 2014. In the first phase, we replaced our power supplies and switchgear with state-of-the-art equipment. In summer 2015 we installed 29 custom-designed workstations that are reconfigurable for a range of student lab experiments, and unique among Canadian universities.

In summer 2016 we will undertake our final phase, adding modular hardware and software for communications and control, offering students front-line experience in designing and programming next-generation smart-grid infrastructure. We are giving our students the power to design the future.

To learn how you can help us realize the final stage of our Energy Systems Lab renovation, contact Joanna Gray at joanna.gray@ecf.utoronto.ca, 416-946-0372 or donate.utoronto.ca/engineering

BOUNDLESS

PHOTO BY RAINA+WILSON

Everyone dreams of their future—we're busy building ours

It's 2015, and we still don't have levitating hoverboards. We do, however, have the nascent Internet of Things—IoT for short. Self-driving cars, intelligent refrigerators, smart hockey helmets: IoT is supposed to make it all possible. But what does an IoT world really look like?

The term has been thrown around for the past several years, building up to become the trendiest of trends. But here in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE), we're busy creating the real technologies—in electronics, communications and photonics—that will constitute the IoT of the future and figuring out what we would do with this large network of connected objects. The cover story of this issue, *Behind the IoT Curtain* on page 22, explores the research that will soon allow us to realize many of the life-altering applications that are now just dreams and hype.

Hype is something we try to avoid—we value evidence, results. But no one can say that the reputation growing around the University of Toronto's Supermileage Team, led by ECE's own Mengqi Wang, is just hype. Always a contender but never previously on top, this year the team brought home a gold medal for racing the most fuel-efficient vehicle at the 2015 Shell Eco-Marathon Americas competition in Detroit. Flip to page 18 for the story of their white-knuckle ride.

Our department maintains its position as the top-ranked ECE program in Canada, a direct result of the excellence of our students, faculty, staff and alumni. Name the problem, be it social, technical or economic, and there is likely an ECE graduate working in a boardroom or operating room to solve it. We strive to keep up with what our alumni have been doing since leaving these halls, and to help former classmates stay in touch with each other. In that spirit, we are launching a new alumni Class Notes section with this issue—turn to page 48 to catch up on news from ECE graduates of all ages, all over the world.

This magazine offers a constellation of stories that reflect how we spent the year. As we turn to face the year ahead, I welcome your feedback on our progress and our future direction. I 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 2015

EDITOR
Marit Mitchell

ART DIRECTOR
Katina Constantinou,
Sugar Design

PHOTOGRAPHY & ILLUSTRATION
Raina+Wilson
Borja Bonaque

COPY EDITORS
Kelly Hayward
RJ Taylor
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

8
If You Build it, They Will Compute
IBM's supercomputing consortium lets ECE researchers tackle challenges from cloud computing to cancer care

12
The One-Step Journey
Four electrical engineering undergrads teamed up to build a powered leg brace for a fellow UofT student

18
Green Gearhead
Mengqi Wang and the UofT Supermileage Team built the most fuel-efficient prototype vehicle to race in 2015. They have plans to do better

16
Climb Every Mountain
Afshin Rezayee is a hands-on hardware guru and serial entrepreneur. Why he's always happy, but never satisfied

22
Behind the IoT Curtain
Everyone's ready for the Internet of Things—but what exactly is it?

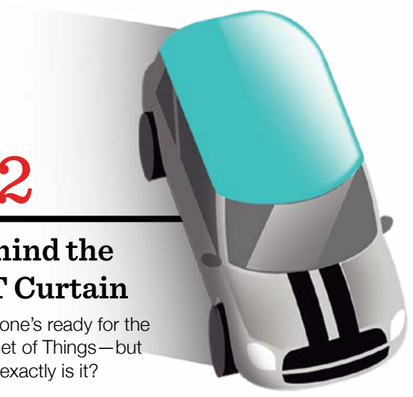

SECTIONS

29
UNDERGRADUATE
Blue Sky Down Under
A team of undergraduates led by ECE's Zhe Gong represented Canada in a gruelling race across the Australian Outback, powered only by the sun

> The goal was clear, if ambitious: to design, build and race the world's most efficient gasoline-powered vehicle at the annual Shell Eco-Marathon Americas. p. 18

34
GRADUATE
Smarter Solar on a Smaller Scale
Researchers want to help you save some solar power for a rainy day, without breaking the bank

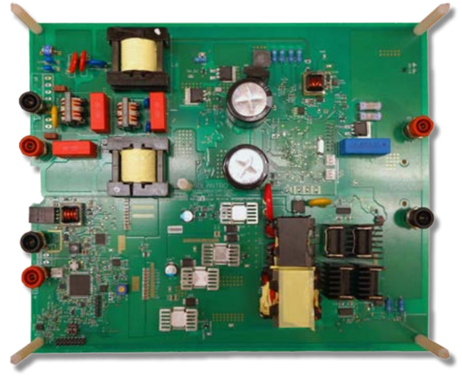

37
RESEARCH
Designing for the Data Deluge
Wei Yu wins E.W.R. Steacie Memorial Fellowship

41
INDUSTRY
Where Are They Now?
What three electrical engineers have learned from their crash course in the real world of entrepreneurship

45
ALUMNI
SkuleTalks, a New ECE Video Series
An accessible way to catch up on the ideas that matter to today's leading engineering thinkers

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. Home to 99 professors, 639 graduate students and 1,512 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.

Among our faculty are 25 Fellows of the IEEE, nine Fellows of the Royal Society of Canada, 11 Fellows of the Canadian Academy of Engineering, and four E.W.R. Steacie Fellows, a prize awarded to the brightest and most promising scientists and engineers across the country. ECE is a hotbed of research commercialization, with more than 160 inventions disclosed and 40 patents filed over the past five years—by far the most productive department for new inventions at the University of Toronto. Read on to find out what makes us the top-ranked ECE department in Canada.

QS World University Rankings 2015

#1 *in* Canada
Electrical & Electronic Engineering

#13 *in* North America
Electrical & Electronic Engineering

#27 *in* the World
Electrical & Electronic Engineering

#1 *in* Canada
Computer Science & Information Systems

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

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

1,512 Undergraduate students
639 Graduate students
99 Professors, including Emeritus
63 Post-docs
51 Admin & tech staff
14 Research associates
12 Visiting professors

25
Chair titles held by ECE faculty in 2014–2015, including Canada Research Chairs, Endowed Research Chairs, Industrial Research Chairs and U *of* T Distinguished Professor Chairs

$59,243
Average salary earned by ECE students on PEY placements in 2015-2016

161 INVENTIONS OUT OF ECE SINCE 2010

40 PATENT APPLICATIONS OUT OF ECE SINCE 2010

80 PERCENTAGE OF ECE INVENTION DISCLOSURES THAT INCLUDE A STUDENT INVENTOR

11 STARTUPS OUT OF ECE SINCE 2010

TIMES HIGHER EDUCATION WORLD UNIVERSITY RANKINGS 2015–2016

Engineering & Technology Universities

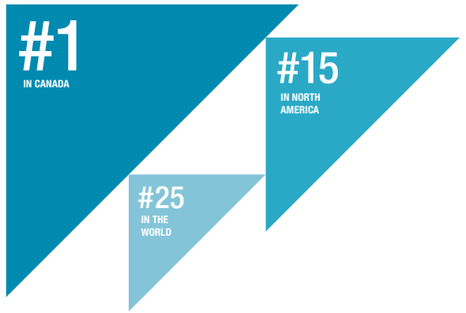

ACADEMIC RANKING OF WORLD UNIVERSITIES 2015

Engineering/Technology & Computer Sciences

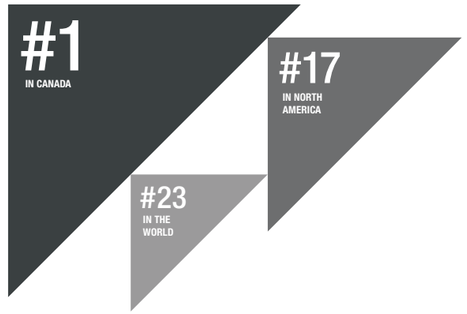

BY MARIT MITCHELL

IF YOU BUILD IT, THEY WILL COMPUTE

ECE RESEARCHERS HAVE LONG FORGED INDIVIDUAL COLLABORATIONS WITH IBM—BUT THE CREATION OF A SUPERCOMPUTING CONSORTIUM HAS ACCELERATED THE PACE OF PROGRESS AND HELPED BRIDGE THE ACADEMIC-INDUSTRY DIVIDE

IT WAS ONE OF THOSE BACK-OF-A-NAPKIN MEETINGS THAT, AT EACH RECOUNTING, SOUNDS A BIT MORE LIKE A MOVIE SCENE: IN WINTER 2010 TOP BRASS FROM THE UNIVERSITY OF TORONTO, WESTERN UNIVERSITY AND IBM SAT DOWN TO DINNER AT A RESTAURANT IN DOWNTOWN TORONTO. ONE BIG QUESTION SAT IN THE MIDDLE OF THE TABLE: WHAT WAS MISSING FROM THE CANADIAN RESEARCH AND INNOVATION LANDSCAPE?

The group was interested in accelerating the research already underway at universities across the country: advancing health care; inventing new technologies for clean energy; improving water quality and monitoring; building livable, stable, sustainable cities; making big data more meaningful. Across the disparate set of goals, a common theme emerged: the single thing all these ambitious projects really needed was computing power, and lots of it.

"We had hundreds, if not thousands, of bilateral research collaboration agreements—a professor and an IBM researcher working together," says Allen Lalonde, senior executive of IBM Canada's Research & Development Centre and a member of the Industry Advisory Board to The Edward S. Rogers Sr. Department of Electrical & Computer Engineering. "But that wasn't necessarily getting leveraged, it wasn't always strategic, and it certainly wasn't very efficient or practical."

In April 2012, the proverbial switch was flipped on the Southern Ontario Smart Computing Innovation Platform—SOSCIP for short, pronounced 'so-sip'. The research and innovation consortium connects professors with the latest IBM high-performance computing tools, analytics software and research expertise. The University of Toronto is now home to SOSCIP's IBM Blue Gene/Q supercomputer, Canada's fastest supercomputer.

"IBM has been working with UofT, and with ECE, since the earth cooled," says Stephen Perelgut, who was IBM Canada university relations manager both before and after SOSCIP's launch. "SOSCIP was born out of trying to figure out how to do what we do more effectively."

Professor Cristiana Amza was one of those early players, having partnered with IBM for more than a decade. Her research has tracked the evolution from housing data in industrial and privately hosted datacentres, to abstracting everything into the cloud.

"Everyone has a cloud infrastructure in operation, but some offerings are pretty basic," says Amza. "We are working toward adding some of the guarantees back in that existed previously, with large amounts of storage in industrial offerings and in the private cloud...in order to do that, there has to be some re-architecting of the whole service."

Her group leverages hundreds of terabytes of data that IBM collects daily from its datacentres and combs through them, looking to improve the architecture of today's cloud systems, and for early warning signs of trouble. They are also creating an 'alarm' program to digest all this low-level information and alert operators of imminent failures, where exactly to look for problems, and how best to respond when a component goes down.

The University of Toronto is now home to SOSCIP's IBM Blue Gene/Q supercomputer, Canada's fastest supercomputer.

"So far there have been fun things going on in the cloud, but if you really want to run a business leveraging the cloud, and I think there's a push toward that, then of course there will be a big outcry if there's any data loss that's very important to businesses," says Amza. "Right now, users are moving just like those cartoons where there's a character walking on air—then they notice they're off the plank. So we're extending the plank while they're still walking on air."

Big shifts in technology often evoke the feeling of walking on air—will this approach prove solid, or will its adopters plummet? Right now global companies, including IBM, Microsoft and Intel, are investing in the idea of boosting their datacentres by introducing reconfigurable hardware in the form of field-programmable gate arrays (FPGAs).

ECE at UofT is home to world leaders in the FPGA field, and IBM is working with Professor Vaughn Betz and PhD candidate Jeff Cassidy to test out a new SOSCIP platform that links its POWER8 CPUs with an Altera Stratix V FPGA via their Coherent Accelerator Processor Interface (CAPI). (Fun fact: before joining ECE as a professor in 2011, Betz was Altera's senior director of software engineering and one of the architects of the Stratix family of products and CAD tools—so it's a safe bet that Cassidy has been thoroughly schooled on the ins and outs of both the hardware and software).

Cassidy is applying the POWER8-CAPI system to an area that is just starting to fully realize the benefits of high-performance computing: cancer treatment. He's developing models to predict how photons of light will behave inside human tissues, to improve the prospects of using photodynamic therapy as a treatment for head and neck cancers. Access to the brand-new CAPI hardware is provided through SOSCIP, and he also works closely with IBM researchers in Austin, Texas and Rochester, Minnesota.

IBM has been working with UofT, and with ECE, since the earth cooled. SOSCIP was born out of trying to figure out how to do what we do more effectively.

Thanks to significant new investments by the Federal Economic Development Agency for Southern Ontario and IBM, SOSCIP has been able to double the number of academic members in the consortium and expand its computing platforms.

"What we wanted to do was deliver enterprise-level, really scalable solutions that the largest, most advanced and complex customers of the world were using," says Lalonde. "Why shouldn't universities and researchers and students have access to that?"

The experts who will carry this joint industry-academic work forward are the new fleet of IBM post-doctoral fellows. Now in its fourth year, this unique program sees IBM hire PDFs and place them with partner professors. The researchers spend their days at one of the SOSCIP member academic institutions, but as IBM employees, they have access to IBM mentors, professional development opportunities, as well as access to all the company's intellectual property. The program has been so successful, IBM is replicating the model in other provinces.

"I think it's absolutely been a slam-dunk success," says Lalonde. "The walls came crashing down and people started going back and forth across them—we're really working on the same stuff, together." A

BY MARIT MITCHELL
PHOTOGRAPHY BY RAINA+WILSON

THE ONE-STEP JOURNEY

WHEN FOUR ELECTRICAL ENGINEERING UNDERGRADS TEAMED UP TO BUILD A POWERED LEG BRACE FOR A FELLOW U OF T STUDENT, THEY WEREN'T SURE HOW FAR THEY COULD GET IN JUST EIGHT MONTHS. THEY SURPRISED BOTH THEMSELVES AND THEIR CLIENT—AND SET FUTURE PROJECTS IN MOTION

"People's minds jump right to Iron Man when they see this," says Thomas Garside, "but I was always more inspired by Robocop."

Every robo-clad hero needs an origin story, and this is Garside's.

It begins with a young man from Sault Ste. Marie who left his hometown to undertake a master's degree at the University of Toronto, and the group of people who rallied around his idea to invent something completely new.

"We knew we were looking for a biomedical project," recalls Elizabeth Sumitro, who joined electrical engineering undergrads Lakmini Perera, Kayatri Rangarajan and Sanjana Seerala to form the core of Garside's team. The four discovered a capstone project description posted by Professor Paul Yoo, who is both a member of the biomedical engineering group in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering and a core faculty member in the Institute of Biomaterials and Biomedical Engineering (IBBME).

The challenge: build a powered leg brace for a fellow student, Garside, who has cerebral palsy. Because his hamstrings are tight, it's difficult for Garside to fully extend his lower legs and stride forward. The powered brace would assist Garside in extending his lower leg into his next step. The team needed to build the brace to spec, test and deliver the product to client—all in under eight months, the duration of the fourth-year capstone course ECE496.

Garside had long been ruminating on his idea for a powered brace, but had met many dead ends in trying to see it built. "I'd approach professors in my Faculty and others, and say, 'Look, you don't know me, and you don't have time to get to know me, but I have this very personal issue with my mobility and can you help me solve my problem, even if it doesn't necessarily help you?'" says Garside. "And that's a very intense thing to do."

FROM L. TO R: KAYATRI RANGARAJAN, SANJANA SEERALA, LAKMINI PERERA, THOMAS GARSIDE AND ELIZABETH SUMITRO

The team of four electrical engineers designed a powered exoskeletal leg brace to help fellow U of T student, Thomas Garside, fully extend his lower leg while walking. They discovered that the human gait is surprisingly hard to mimic.

After arriving at UofT, Garside joined the Engineering Toastmasters public speaking group. In chatting with one of the group members about his idea for a powered brace, the student suggested he pitch it as a capstone project, and that he start with Professor Yoo. Yoo is also the faculty supervisor for the University of Toronto branch of the Tetra Society, a group that aims to help UofT students overcome physical barriers, both on campus and off.

"When Tom came to me with his idea for a project, no one had been brave or crazy enough to try it," says Yoo. "The idea was more than most student volunteers were able to take on. But we threw it out there, and this group bid on it."

Brave and crazy though they may be, Perera, Rangarajan, Seerala and Sumitro are also energetic, eloquent, organized and diligent. Even before school resumed in September, they were meeting with Yoo, Garside and Ken Strauss of Tetra, who acted as their mechanical expert and consultant, hashing out the scope of the work and their respective responsibilities.

"I think we had four alternatives, and we compared objectives of our project against all the alternatives and arrived at this decision that the EMG and foot-pressure sensors were the best combination for us," says Perera.

Here's how the powered leg brace works: EMG (short for electromyography) sensors read the unique signatures emitted by Garside's muscles when he's about to take a step. Pressure sensors on the heels and toes of both feet signal a microcontroller in the brace to drive the motor to fully extend the right leg. An angle sensor located on the brace tracks the leg's angle, extending and retracting in response to the EMG and pressure signals.

Sumitro took the lead on the EMG, Perera on the pressure sensors, Seerala wrote the algorithms that assimilate signals from the sensors in the microcontroller, and Rangarajan worked on both EMG and the microcontroller. By October, they were using the biomedical teaching lab housed in IBBME to get EMG readings from Garside's legs and work them into their design. They also consulted with Professors Tyson Beach and Luc Tremblay, both of the University of Toronto's Faculty of Kinesiology & Physical Education, to learn more about implementing the EMG sensors for controlling locomotion.

"Even with EMG there are a lot of different possibilities in terms of muscles in the leg we could have used. We did some preliminary testing with my signals and discovered that my pattern is not at all like Tom's pattern," says Sumitro. "So we ended up doing the testing with him and found out that the quadriceps gave the strongest readable pattern to let us determine when he was taking a step forward."

It wasn't always smooth striding, especially as pressure built ahead of the final ECE Design Fair.

With just three days until Design Fair, the group was pulling another 12-hour shift in the Design Lab to incorporate their new angle sensor, freshly delivered from the U.S. "We only ordered one component—we should have had a backup, but we didn't realize," remembers Perera. "And while we were testing, it burned. Without it our brace wouldn't be operational, even if we had every other component. We had to order it from the U.S., and it was Friday night."

"I think we called every single shop," says Rangarajan. They came clean to Yoo who helped them hatch a backup plan that relied on video they'd captured during testing, but in a miracle of modern delivery, the angle sensor arrived in less than a day.

"We've popped capacitors, burnt batteries," says Rangarajan. "We had to make sure it worked 100 per cent, so maybe we ran it too many times..."

In the end, it ran when it needed to—to acclaim from course evaluators, Yoo, and most importantly, Garside.

"One thing I did know about engineering design is you can't say 'OK, this is exactly what I want,' because we're developing something that's so new and so cutting edge. We were facing problems that no one had faced before and we didn't know what the exact outcome of what this was going to be," says Garside. "It was really great to feel it all come together and to have a technology that's really tangible."

The team earned the Gordon R. Slemon Design Award in recognition of a project that best combines creativity with good design and execution, and the John W. Senders Award for the imaginative and successful application of the principles of human factors to the design of a medical device.

Although Perera, Rangarajan, Seerala, Sumitro and Garside have now graduated, this is just a first step—another design team is picking up where this one left off, aiming to build a second brace in the mirror-image of the first, benefiting from motion-capture data the team didn't have time to incorporate, and smoothing out the motor control to better follow Garside's gait.

Here, as in any good origin story, creativity and tenacity pay off. A

Another design team aims to build a second brace, in the mirror-image of the first.

When Tom came to me with his idea for a project, no one had been brave or crazy enough to try it.

BY MARIT MITCHELL
PHOTOGRAPHY BY RAINA+WILSON

CLIMBING EVERY MOUNTAIN

SERIAL ENTREPRENEUR AND HANDS-ON HARDWARE GURU AFSHIN REZAYEE ON WHY HE'S ALWAYS HAPPY, BUT NEVER SATISFIED

When his plane began its descent into Toronto in 1999, Afshin Rezayee peered out the window expecting to see mountains. His hometown in Iran is surrounded by them, and he thought this new city would be the same.

"I admit it, I'm not that great with geography," he says. "I knew there were a lot of mountains in Canada but I didn't know that there are not a lot in Toronto." Never one for discouragement, Rezayee turned Toronto into his home anyway—and built his own mountain.

That's what Rezayee does: he builds things. In 2013 he co-founded Kili Technology Corporation—short for Kilimanjaro, for his love of mountains. A fabless semiconductor company, Kili designed chips for payment and secure mobile authentication, and Rezayee was targeting the burgeoning mobile payment market.

"We had contactless reading capabilities on our silicon, we had contact, we had a dual-core processor, we had embedded flash memory," says Rezayee. "At the time, our silicon was probably the only one in the world that had all these features integrated on the same die. We put all these in because we estimated in the future that these needs would come up—and they did."

They even built a slick handheld card reader to house the chip, and their design and execution were so promising that in March 2015 Kili struck a deal with Square, the company started by Twitter co-founder Jack Dorsey.

Square intends to revolutionize how we pay for things. Its hardware—a tiny white dongle—plugs into any smartphone or tablet to turn it into a secure payment terminal. This summer the company announced plans to launch its first contactless card reader, designed not only for credit and debit tap technology, but also to accept Apple Pay directly from mobile devices—Apple unveiled the reader at its 2015 World Wide Developer's Conference in San Francisco.

The deal was the result of Rezayee's hard work at Kili and persistent desire to collaborate. Having developed the chip, designed and built a payment terminal, and developed a sales funnel across Canada, the U.S. and China—all while growing Kili from two to 30 people—Rezayee and Square discussed cooperation.

They opened discussions with Square's business development team, which turned to the possibility of doing a deal. "There's a rule in sales: always say yes," he says now. Conversations got serious, and in March 2015 Rezayee gave his final yes.

Today he's the Canada Hardware Lead for Square, and Kili's office became Square's second Canadian location (its first is in Waterloo). Life is different in some ways—Rezayee spends less time on sales and marketing, and more time on product development. Life is much the same in other ways—he's still motivated to see his silicon in the hands of merchants across the globe, and he's still working.

With two young sons at home, one almost four years old and the other just nine months, Rezayee is busy—but he's always been busy, and probably always will be. He can be classified as that rare breed, the serial entrepreneur, having launched his first company fresh out of his undergraduate degree in 1995. He and his friends designed and built PCBs to control elevators, then PCs before PCs were ubiquitous, then early MP3 sound systems. He earned a master's degree in communications, taught university courses, then shifted gears again and decided to pursue a PhD at the University of Toronto under the supervision of Professor Ken Martin.

"When I came, I told Ken, 'I'm going to do a lot of hard work, I want to finish it fast. Is there any issue with finishing in say, two years, three years?'" remembers Rezayee. "He said no. If you manage to do your publications and your chips and they work, you can finish. That's what he said, and he was right." Rezayee completed his first defence in three years, and produced and tested six chips in that time.

Armed with a PhD, he joined Martin and Professor David Johns' startup company, Snowbush Microelectronics, and saw firsthand its growth and subsequent purchase by Gennum, now Semtech. He then joined the authentication provider SecureKey Technologies Inc., launched SecureKey's chip team, and then spun it off into Kili with co-founder Greg Wolfond.

So he's worked on a lot of different things. The common thread through all of them is his enthusiasm for his work, and the people he works with.

"When you have a great idea, and if you find the money to build it, then you can put together a great team. And great people, they always need to work on great things. So if you give great people a product to work on, let's say five years, 10 years after that they get bored, so they want to work on something new," he says. "If you have great ideas, you should find the right people and you should build it."

Building from the ground up is something he learned from his father, as a boy in his hometown of Isfahan, in the mountains.

"He said something that I really like—he said that when you want to do something, your eyes will tell you that you can't do it. So you get scared by what you see," says Rezayee. "But you don't do it with your eyes, you do it with your hands. So just go and do it! Don't let your eyes scare you, do it with your hands."

BY MARIT MITCHELL
PHOTOGRAPHY BY RAINA+WILSON

GREEN GEARHEAD

THE UNIVERSITY OF TORONTO SUPERMILEAGE TEAM, CO-FOUNDED BY ECE'S MENGQI WANG, TOOK TOP SPOT AT THE 2015 SHELL ECO-MARATHON AMERICAS RACE—BUT SHE'S NOT GOING TO LET THAT SLOW THEM THEM DOWN

On a Sunday in April, clouds cruised calmly past the skyline of Motor City.

On the ground, the scene was anything but calm.

"Drag the car back here!" yelled Mengqi Wang, running full-tilt toward the white fuelling tent next to the start line of the 2015 Shell Eco-Marathon Americas. "There are only two cars in line—if we get there right now, we can still make another run!"

University of Toronto's Supermileage Team surged into action, hauling their fully custom 35-kilogram vehicle from the tent, where technicians had been meticulously measuring the millilitres of gasoline they'd burned on their last run.

Barely trailing the reigning champions from Université Laval, Wang didn't dare to hope that her team could pull out a win—she just wanted one more chance on the track. With the course scheduled to close in mere minutes, driver Kristine Confalone slid into the tiny bullet-shaped car and hit the gas.

THE GOAL WAS CLEAR, IF AMBITIOUS: TO DESIGN, BUILD AND RACE THE WORLD'S MOST EFFICIENT GASOLINE-POWERED VEHICLE AT THE ANNUAL SHELL ECO-MARATHON AMERICAS.

Back in 2013 the entire Supermileage Team consisted of Wang, then a third-year undergraduate in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, and Jonathan Hamway, a mechanical engineering student. Former design-project partners, the two shared an uncompromising love of building things from scratch, moving fast, and never letting inexperience hold them back. Both had considered joining one of the well-established design teams based in UofT Engineering—the Blue Sky Solar Racing Team or Formula SAE Racing Team—but couldn't shake the itch to invent something totally new and run their own show. Hamway discovered that several versions of a supermileage group had existed in years past but had spluttered out. He suggested they resurrect it.

"And I thought, yeah, that's a fantastic idea," says Wang. "It was just the two of us in the beginning, we obviously had no idea how teams worked or what was necessary."

The goal was clear, if ambitious: to design, build and race the world's most efficient gasoline-powered vehicle at the annual Shell Eco-Marathon Americas. How to get there was a little hazier—previous iterations of the club hadn't even managed to ship out to competition.

Wang, Hamway, and their fledgling team, now eight engineers strong, ran straight at that brick wall. "We did a lot of preliminary research and then tried to implement all of it in our first year's vehicle—as many new systems as possible," says Wang. "And I would say multiple systems failed, realistically, the engine being the most crucial of all of them."

At the 2013 Shell Eco-Marathon in Houston, Texas, the team finished eight of 10 laps around the track, just shy of a complete run.

"We got through technical, safety, all of that—we were qualified to race, we just didn't finish a run," says Wang. "Our projected mileage was 626 miles per gallon back then, which was a lot better than we expected, but by no means a winning number." Hamway, an engine buff, was particularly bothered by the engine's unreliable performance and determined to fix it.

That summer the team returned to Toronto and regrouped. Hamway decided to junk their off-the-shelf engine—a 35cc leaf-blower with a ripcord that the driver had to pull across her chest to get started—in favour of designing and building something completely custom. Ever ambitious, the team also decided to replace their clunky aluminum chassis with a monocoque aerobody—supporting the structure by distributing loads across its shell—which, of course, they knew nothing about.

"We realized that getting the right surface finish is a much trickier process than we had anticipated, which we didn't know until we finished that first car and looked at the body, and we thought to ourselves, 'This looks like crap!'" Wang remembers.

It may not have been the shiniest car on the track, but it ran. Back in Houston in 2014 the team finished second, behind Université Laval, with official mileage of 2,712 miles per gallon (86.7 millilitres per 100 km). They also took home the Technical Innovation award for their custom engine and the Penzoil Tribology award for its efficiency, each worth $2,000. People were starting to notice the University of Toronto team.

The little black bullet pulled cross the finish line, marking the close of the Gasoline-Prototype category at the Shell Eco-Marathon Americas in Detroit. The team from the University of Toronto poured into the white tent to have technicians measure their car's fuel consumption and temperature, but the sampling and recording systems had already been switched over in preparation for the next category, Urban Concept.

It would take some time to get the final results, and besides, Wang had heard that Université Laval, the heavy favourites, had improved in their last run to more than 3,300 miles per gallon. It didn't look good.

By Spring 2014 both Wang and Hamway had graduated, and Wang launched into a master's under the supervision of ECE Professor Wai Tung Ng. She was now the only member of the original eight guiding the team, and they had work to do, bringing down the car's weight and improving the aerobody. They also had a new goal in sight: take first place at the next Shell Eco-Marathon Americas.

But Wang wasn't happy with just the one goal, which she viewed as attainable—she wanted to look two years ahead, to think bigger. She wanted to see the team, now 25 members strong, break into the Battery Electric category, the second-largest field of competition behind Gasoline Prototype.

"I want to make sure that before I completely hand off the team to the new generation, they have the experience of designing a car from scratch, facing the same challenges we faced, really understanding what it means to design a system from nothing."

"It took us two years to learn, and we weren't impeded by pre-existing knowledge. Which does, ironically, seem to be a problem," says Wang. "Whereas if you're really doing it from scratch then you have to fully understand literally every nut and bolt in the car, you have to know why you picked that one. And maybe the reason isn't a very good one—like it just happened to be the cheapest—but at least you know that that's the reason and not because that nut or bolt somehow increases the vehicle's efficiency by twenty per cent."

On the streets of downtown Detroit, Wang and Prashanth Murali Krishna, the team's new mechanical lead, gave a handful of interviews about how it felt to come in second. "We're just happy to improve on last year's performance," Wang told a news crew. They joined the rest of the crew in the UofT paddock, where half of the team was squealing and the other half begging for calm.

"I'm reallllllly nervous," said Murali Krishna.

Then a technician in a white lab coat walked over with the news: 3,421 miles per gallon—the equivalent of 68.7 millilitres per 100 kilometres. That's less than five tablespoons of gas to drive from the foot of Yonge Street in Toronto to Barrie, Ontario, and it was enough to lock up first place. The University of Toronto Supermileage Team paddock exploded.

Wang simultaneously finished her master's and started a PhD with Ng in Fall 2015, and the Supermileage Team keeps moving the finish line a little farther ahead.

"I know for a fact that our car could, and I would even say easily, beat the North American record if we had gotten all the systems working sooner," says Wang. "Just in terms of potential, the car even exactly as it is now, if we get all the systems working, I know, I know because we're so close, we can definitely beat the North American record."

Mengqi Wang never lets a little success slow her down. A

BY MARIT MITCHELL
ILLUSTRATION BY BORJA BONAQUE

BEHIND THE IoT CURTAIN

Everyone's ready for the Internet of Things—but what exactly is it? Researchers in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering are building the backbone that may determine its form

Tell an auditorium of 300 tech-savvy listeners that the Internet of Things is upon us, and you'll see 300 heads nod in agreement. Ask those 300 to define the Internet of Things, and you'll get 300 different answers.

The term 'Internet of Things'—IoT for short—has reached peak buzziness, permeating predictions about future tech in fields from health care to transportation. The information-technology research firm Gartner placed it at the tippy-top of its most recent Hype Cycle, at the 'Peak of Inflated Expectations.'

But what exactly is it? The promises are grand and far-reaching: self-driving cars that make smarter decisions and ease traffic congestion, more efficient distribution and monitoring of electricity across the grid, wearables that stream blood glucose levels directly to the doctor's office. Over the next five to 10 years, myriad technical problems need solutions if IoT is to crawl out of Gartner's 'Trough of Disillusionment' and up the 'Slope of Enlightenment' to realization. Which problems are solved first, and how, may determine the shape of the thing we will come to know as IoT.

"Internet of Things is not one thing—it's a family of technologies," says Professor Alberto Leon-Garcia. "There's the sensing devices, the computing, the communications, and until recently all these have evolved independently. The whole trick to seeing economies of scale is to have fairly uniform technology across mass markets. When that happens, costs drop dramatically, capabilities improve steadily, and that's when you get out of the trough."

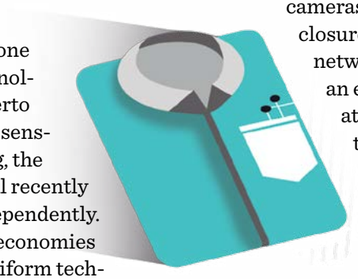

What we know for sure: we will need to handle more data, faster, and with low power.

Four years ago, Leon-Garcia launched a national NSERC Strategic Network called SAVI—Smart Applications on Virtual Infrastructure. The idea was to design and build a next-generation Internet, including a fully-functional testbed and all the applications to run on it.

"Part of the challenge is how to create standards that inter-operate, or create platforms that abstract or virtualize the different types of sensors so they can inter-operate on a common platform," says Leon-Garcia. "In SAVI we look at everything top-down—all the infrastructures are there to support applications."

The genius of SAVI is in its design—it was built from Day 1 to be a software-defined infrastructure, coordinating cloud computing, FPGAs, storage, networking, resource allocation and analytics under a single management system.

Built on top of the SAVI testbed is the Connected Vehicles and Smart Transportation (CVST) project. CVST aggregates data from highway cameras, road incident alerts, road closures, transit systems, bike-sharing networks, border crossings and even an experimental drone stationed at Downsview in north Toronto, to give a comprehensive view of where people are moving in the Greater Toronto and Hamilton Area—the vibrant mass of humanity dubbed the GTHA.

"The question is how to integrate the information that sensors can provide into applications that do useful things. It could be controlling emissions to deal with greenhouse gas emissions, it could be controlling traffic, it could be controlling infrastructures, factories, use of electricity, use of water. And now I'm talking about smart cities, what I call smart infrastructures," says Leon-Garcia. "To me a key benefit of IoT is in enabling global knowledge of these very large infrastructures like cities—that combined knowledge allows you to make intelligent decisions toward universally desirable goals: things like reducing carbon emissions, or reducing travel times."

Reducing travel times for cars is one thing, but we'll also need to reduce travel times for data—and lots of it.

Finding new ways to send more bits down existing pipes—along our current fibre-optic infrastructure—is easier said than done. For starters, we don't even know the fundamental carrying capacity of fibre-optics (known as the Shannon limit after American mathematician and engineer Claude Shannon, father of the field of information theory).

Professor Frank Kschischang is tackling the complex non-linear mathematics of fibre-optic communications systems, trying to find out how close we are to this unknown upper bound. Part of the problem is inherent to signalling with light: when you send intense pulses of light down a fibre, the optical power changes the fibre's index of refraction—in other words, the signal itself alters the channel through which it moves. A higher index of refraction means the light slows, and the signal becomes distorted. This problem can only be described using non-linear math, and it's not straightforward.

"The way that fibre optic communications evolved was to import ideas that had already been pioneered and debugged from wireless communications, throw them into fibre and see how it worked," says Kschischang. "But now we are bumping into this non-linearity and it's causing an opportunity for a rethink."

So far, handling non-linearity has been treated as a tough signal-processing problem—restoring badly distorted data, rather than finding ways to prevent its distortion in the first place. Kschischang is collaborating with researchers in academia and industry on both theoretical and experimental work to overcome non-linearities when sending

Internet of Things is not one thing—it's a family of technologies. There's the sensing devices, the computing, the communications, and until recently all these have evolved independently.

high volumes of data over thousands of kilometres, and at extremely high speeds. His work has implications for long-distance links between Toronto and Tokyo, on a local scale such as within a metropolitan area, and for the hyper-local, such as connecting computers inside of a single datacentre.

"I think there are very few people who know the information theory and the non-linear physics. So we're kind of starting a new breed of people who try to know both, and that's very interesting," he says.

To translate that physics and theory into reality, we need new hardware and devices—and Professor Joyce Poon is designing them.

Today, an email sent from Tokyo to Toronto crosses the ocean as light. But in order to route it to the correct recipient, it needs to be converted from light to electronic signals—then to the ones and zeros that all digital devices interpret. That conversion, from optics to electronics, is incredibly energy-expensive. As data streams constantly from a future network of temperature sensors, light sensors, motion sensors, chemical sensors, pressure sensors, and sensors not yet invented, it's an untenable extra step that Poon aims to make obsolete.

> Building the backbone of the Internet—the core that will allow for sensing, streaming and storage—is an engineering problem.

"I'm trying to answer the question of what you do when there's so much data, and you still have to process it, and route it," she says. Her group is working to design an integrated photonic circuit to route information optically, moving data faster while saving enormous amounts of energy.

"The idea is that as we move to higher and higher data rates coming into each port, it's going to be very difficult to do this switching electronically," she says. "What we found is that we couldn't use existing generic technology to do it, so we developed a brand new platform. If everything works out, these chips will replace the guts that are inside switching equipment right now."

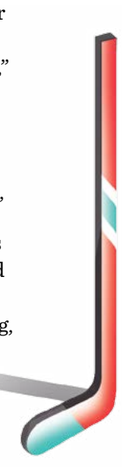

Improving the physical layer in terms of speed is one thing, but Professor Ashish Khisti is worried about security. "When I think of the Internet of Things, I think of RFIDs—tiny tags on everything, pinging simple signals to readers," he says. "But if there are millions or billions of devices communicating constantly, the levels of protection currently being offered are really scary."

Khisti is designing new hardware-based technologies for better security and privacy, such as multi-antenna arrays that concentrate signals to favour the user of interest, and protocols that make eavesdropping much harder.

We've arrived on the brink of the Internet of Things thanks in large part to the semiconductor industry—our lives are now powered by fast, cheap and tiny digital devices in a cornucopia of forms. But 50 years after Gordon Moore predicted that the number of transistors we could pack onto an integrated circuit would double every 1.5 to two years, the big chip-makers are bravely facing a world in which that may no longer be true.

"These are very interesting times because nobody knows what the future of semiconductor technology is going to look like," says Professor Sorin Voinigescu. "The last 50 years have been unique in the history of mankind—no other industry has ever had 40 per cent growth year after year. We're at a fork in the road."

In one direction: finding new materials that allow for continued transistor shrinkage, albeit possibly at a slower pace. In another: tailoring existing CMOS transistor technology, which the entire electronics industry is set up to produce, to IoT and healthcare applications.

"Everything depends on this CMOS technology continuing along this 'forever shrinking' path," says Voinigescu. "The implications for other industries are huge—search, artificial intelligence, robotics—it's all going to saturate."

After his PhD Voinigescu worked in industry for several years, and when he joined the ECE faculty in 2002, focused his research mainly on communications integrated circuits. "But I've gone back to devices in the last three years, precisely because I saw that something had to be done."

This year, in collaboration with networking and communications companies such as Ciena in Ottawa and Finisar in California, his group became the first to design and demonstrate the world's fastest trans-impedance amplifier in silicon, transmitting at 120 gigabytes per second. For comparison, that's four times faster than the current standard, laying groundwork for future generations of fibre-optic networks. They're now working to add more superfast components to build a complete 120Gbaud communications system.

"All the technologies needed for IoT are already in place," says Voinigescu. "But now you need that higher data rate for the Internet backbone to support it, and you need higher computer power with lower power consumption—otherwise the earth will melt."

Building the backbone of the Internet—the core that will allow for sensing, streaming and storage—is an engineering problem. But IoT, in its final form, will be defined by the intelligence built around that backbone.

"IoT is really two sides of a coin," says Leon-Garcia. "One side is the gathering of the information, with unprecedented volume and granularity—but the other is about how do you process that information, and how do you make intelligent decisions based on it?"

In 2025, the auditorium of 300 techies may have a ready definition for IoT. What that definition is depends on the intelligence of our decisions today. A

Undergraduate Study

TRAINING THE NEXT PROBLEM SOLVERS

Electrical and computer engineering power the technological advances that define modern life. An undergraduate degree from The Edward S. Rogers Sr. Department of Electrical & Computer Engineering opens doors to any career imaginable, from app development to artificial intelligence, mechatronics to medicine, and every field 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 biomedical engineering, mechatronics and systems control, software and hardware engineering, energy systems, digital and analog electronics, electromagnetics, photonics, and communications.

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**.

16:1
Undergraduate student-to-faculty ratio in ECE—one of the lowest in U*of*T Engineering

Undergraduate degrees awarded, 2014-2015

ElecE *128*
CompE *129*
Total *257*

BLUE SKY SOLAR REPRESENTS CANADA AT 2015 WORLD SOLAR CHALLENGE

In Fall 2015 a team of U*of*T Engineering undergraduates travelled more than 15,000 kilometres, just for the chance to race another 3,022 kilometres.

The Blue Sky Solar Racing Team, led by Managing Director Zhe Gong (ElecE 1T4+PEY), had spent the last two years designing, building and testing their entirely custom solar-powered vehicle, Horizon. Blue Sky was the only Canadian team to enter the 2015 World Solar Challenge, a gruelling race across the Australian continent, from Darwin in the north to Adelaide in the south, powered only by the sun.

"We're incredibly happy, the challenges we faced in the last two years made it seem improbable that we would be able to accomplish what we did," said Gong. "It also meant a lot to represent Canada on the international stage. This year's competition was the toughest we've ever seen and we'll definitely be working hard to climb the ranks in two years. Look out for us!"

The team met its goal of completing the full course without any breakdowns, help or 'trailering', in a time of 47 hours, 39 minutes and 39 seconds.

PEY Profile

Chandini Chandrabalan (ElecE 1T7+PEY)
Technical Services Intern, Protection and Control at Hydro One in Toronto

PHOTO BY ROBERTA BAKER

Q: WHY DID YOU CHOOSE PEY?

A: I decided to pursue an internship through the PEY program because beyond the perspective I was getting at school, I thought it would be an ideal platform for me to add to my experience. By extending what I've learned in the classroom to applications in the real world, I hoped to explore various fields within ECE in a new light. I was also eager to experience working in a corporate environment for an extended period of time, and wanted to work on meaningful projects, build a work ethic and really get to know my fellow co-workers.

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

A: In the Protection and Control department at Hydro One, I work with both the protection relays that support the main transmission operation of Hydro One as well as the control equipment which allows for remote control of the many transformer stations across Ontario. Initially, I assisted in a cyber-security project which was mandatory for Hydro One to be compliant with industry standards. Recently, I led a team of three fellow co-op students, and with the supervision of senior engineers, our primary task is to review upgrades to transformer stations. Upgrades include changing from electromechanical relays to microprocessor relays, or designing protections and control devices for a new station, and updating a central database in keeping with company standards.

Q: ARE YOU WORKING TOWARD A MAJOR PROJECT OR MILESTONE?

A: In addition to the regular work of the department I mentioned, I'm hoping to get involved in projects that focus on special protections for high-priority stations, as well as projects that aim to ease the process of transferring information from the design engineers at headquarters to the controllers at the main control centre. Outside of the technical work internal to my department, I'm also a part of Hydro One's Co-Op Committee, which acts as a liaison between the full-time staff and the students to enhance the co-op experience at the company. Through the committee, I'm helping to coordinate programs such as a mentorship program between full-time employees and students, a series of notable speakers, career panels and conferences, and various other professional-development and social events.

Q: WHAT HAVE YOU LEARNED ABOUT INDUSTRY OR WORKING LIFE THAT YOU'LL BRING BACK TO SCHOOL?

A: The most important aspect of industry life that I've realized is the ability to be accountable for my work. I've found that in school, students mostly operate as individuals, and we alone face the consequences of our action or inaction. In industry, I'm seeing for the first time that my work is a piece of a bigger picture, and why it's important that I present accurate work in a timely manner. Taking initiative and responsibility for the tasks I do allows me to actively collaborate with people with different areas of expertise. When returning to school in the fall, I hope to adapt the way I learn at school to how I now learn new material at work. With a new sense of the type of work that is performed in different industries I can paint a bigger picture, making the learning process at university a much more active and involved experience. I'm also eager to return to school and work with students from different fields of engineering, as well as arts and science, where I can contribute my knowledge and skills toward a common goal such as a capstone project.

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 2015–2016, 252 ECE students are working in PEY placements at 84 companies located across Canada and around the world, earning an average salary of $59,243. To learn more, visit **uoft.me/ece-pey**.

UNDERGRADUATE ENROLMENT, 2011–2012 TO 2015–2016

	PART TIME	FULL TIME	TOTAL
2011–2012	256	1096	**1352**
2012–2013	224	1142	**1366**
2013–2014	259	1133	**1392**
2014–2015	272	1208	**1480**
2015–2016	313	1199	**1512**

ECE PEY PLACEMENTS, 2011–2012 TO 2015–2016

	ElecE	CompE	TOTAL
2011–2012	134	71	**205**
2012–2013	102	79	**181**
2013–2014	133	77	**210**
2014–2015	156	70	**226**
2015–2016	160	92	**252**

$104,116
Top salary earned by an ECE student on PEY placement in 2015-2016

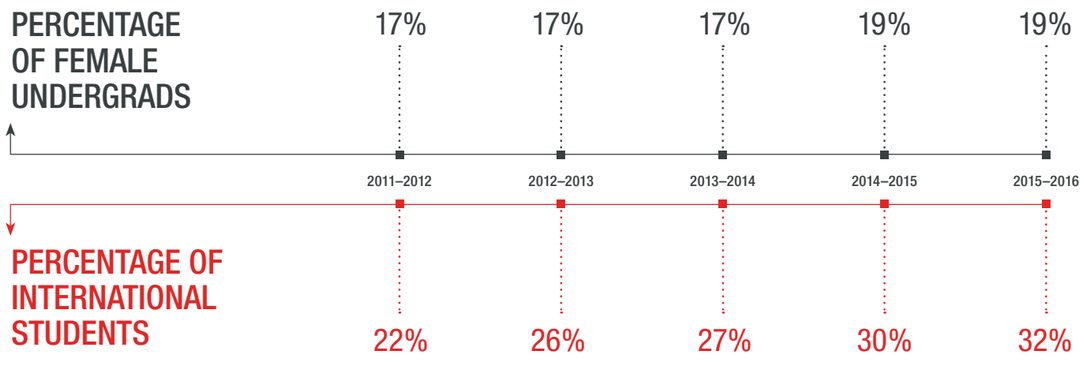

Graduate Study

AN INTERNATIONAL DESTINATION

Throughout its 106-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/ece-gradadmission**.

$2.4M
in graduate student scholarships awarded by ECE for 2015-2016

OUR GRADUATE STUDENTS BY FIELD OF STUDY

	MASC	PHD	TOTAL
BIOMEDICAL ENGINEERING	4	7	**11**
COMMUNICATIONS	32	58	**90**
COMPUTER ENGINEERING	39	62	**101**
ELECTROMAGNETICS	12	15	**27**
ELECTRONICS	15	34	**49**
ENERGY SYSTEMS	20	22	**42**
PHOTONICS	16	37	**53**
SYSTEMS CONTROL	7	10	**17**

Mario Badr named one of U*of*T's top teaching assistants

Mario Badr, a PhD candidate in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, took home a 2015 TA Teaching Excellence Award from the University of Toronto's Teaching Assistant Training Program (TATP). One of just four recipients out of more than 6,000 teaching assistants across the entire University, Badr has taught Engineering Strategies and Practice I & II, Computer Fundamentals, Computer Organization and Computer Architecture.

"I love TAing, I love teaching," says Badr. "This semester I tried to incorporate more games into my tutorial for Engineering Strategies and Practice—I asked my students to take marshmallows, string and spaghetti and work together to build the tallest freestanding structure. I think by doing this type of work early in the course, it sets them up well for the teamwork that comes later."

The TATP Teaching Excellence Award was created in 2003 to recognize the outstanding contributions of teaching assistants at the University of Toronto, and seeks to value the work of TAs who regularly inspire and challenge undergraduate students. The awards committee considers the TA's knowledge of his or her subject area, communication skills, organizational skills, demonstrated enthusiasm, and ability to provide students with effective feedback, as well as testimonials from both students and faculty supervisors.

Badr counts his supervisor, Professor Natalie Enright Jerger, among his teaching mentors, and draws inspiration from the teaching styles of UofT Engineering Professors Jonathan Rose, Peter Weiss and Jason Bazylak.

"I guess motivation for teaching is something I've always had," says Badr. "I've been taking the PPIT program, Prospective Professors in Training. I want to keep that door open, and if I could become a professor then that's something I would like to do."

ONGOING PARTNERSHIP WITH SOLANTRO AIMS TO BRING A LITTLE SUNSHINE INDOORS

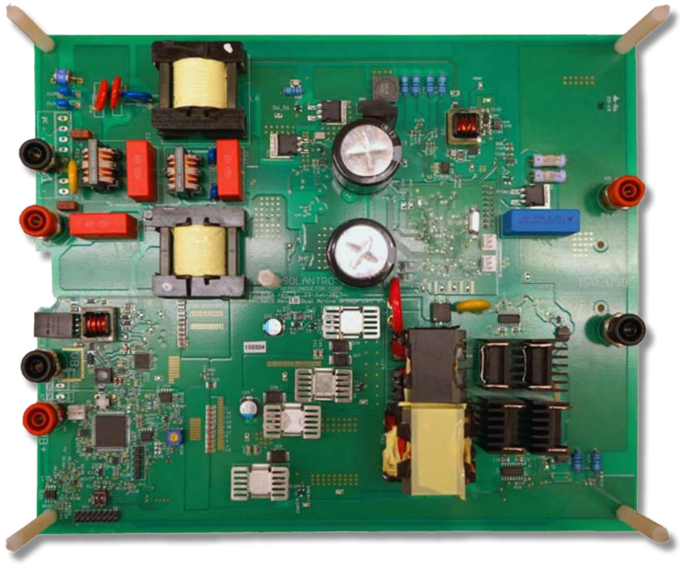

The inverter Poshtkouhi designed in collaboration with Canadian semiconductor company Solantro.

Shahab Poshtkouhi wants to help you save some solar power for a rainy day. For his PhD, carried out under the supervision of Professor Olivier Trescases, he worked closely with Ottawa-based Solantro, a company that designs custom semiconductors specifically for solar applications. With Solantro, Poshtkouhi designed and built a nanogrid—an independent system that captures solar energy, stores and distributes it, all on a tiny scale just right for your home.

For the eco-conscious consumer, purchasing a custom solar array and storage system, or an electric vehicle, can cost many thousands of dollars. “For the average home owner, or a small business owner like a farmer, that upfront cost is prohibitive,” says Poshtkouhi. “Our main goal is to bring that cost down, and to build a system that’s modular and expandable.”

Solar panels on the farmer’s roof, or out in a field, generate energy in the form of direct current, or DC power. But when she plugs in her toaster, that farmer is pulling energy in the form of alternating current, or AC. To convert DC to AC you need a special circuit called an inverter—not something most of us have lying around. Poshtkouhi spent summer 2015 working at Solantro to design an inverter to interface DC solar panels and batteries with the AC nanogrid. He and Solantro plan to demonstrate these modular nanogrids in early 2016.

“I hope that someday you’ll be able to walk into a hardware store and buy this entire system—solar panels, inverters, batteries, and smart distribution all in one—and install it in your home, easily and cheaply,” says Poshtkouhi.

Enrolment in Research-stream Programs 2015–2016

MASc 145

PhD 245

TOTAL 390

TOTAL GRADUATE ENROLMENT, 2011–2012 TO 2015–2016 (HEADCOUNT)

	FALL 2011	FALL 2012	FALL 2013	FALL 2014	FALL 2015
MASC	153	174	167	173	145
PHD	212	236	249	253	245
MENG	114	155	140	169	249
TOTAL	479	565	556	595	639

GRADUATE DEGREES AWARDED, 2010–2011 TO 2014–2015

	2010–2011	2011–2012	2012–2013	2013–2014	2014–2015
MASC	61	57	56	70	71
PHD	39	34	37	27	35
MENG	43	25	56	85	76
TOTAL	143	116	149	182	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 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 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 globe. Learn more how research in ECE is shaping the world at **uoft.me/eceresearch**.

$311K – 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, 2014-2015

INSTITUTE OF ELECTRICAL & ELECTRONICS ENGINEERS, CONTROL SYSTEMS AWARD // **PROFESSOR BRUCE FRANCIS**
INSTITUTE OF ELECTRICAL & ELECTRONICS ENGINEERS, FELLOW // **PROFESSORS DEEPA KUNDUR, BAOCHUN LI**
ROYAL SOCIETY OF CANADA, FELLOW // **PROFESSOR BRENDAN FREY**
CANADIAN ACADEMY OF ENGINEERING, FELLOW // **PROFESSOR TED SARGENT**
ENGINEERING INSTITUTE OF CANADA, JULIAN C. SMITH MEDAL // **PROFESSOR ALBERTO LEON-GARCIA**
ENGINEERS CANADA AWARDS, MEDAL FOR DISTINCTION IN ENGINEERING EDUCATION // **PROFESSOR JONATHAN ROSE**
AMERICAN SOCIETY FOR ENGINEERING EDUCATION, TOP 20 UNDER 40 // **PROFESSOR MICAH STICKEL**
ALFRED P. SLOAN FOUNDATION, SLOAN RESEARCH FELLOWSHIP // **PROFESSOR NATALIE ENRIGHT JERGER**
UNIVERSITY OF TORONTO, UNIVERSITY PROFESSOR // **PROFESSOR TED SARGENT**
UNIVERSITY OF TORONTO, DISTINGUISHED PROFESSOR IN DIGITAL COMMUNICATIONS // **PROFESSOR FRANK KSCHISCHANG**
UNIVERSITY OF TORONTO, INVENTOR OF THE YEAR AWARDS // **PROFESSORS PARHAM AARABI, RICHARD COBBOLD**

WEI YU WINS E.W.R. STEACIE MEMORIAL FELLOWSHIP

Your smartphone may keep getting smarter, but its network is struggling to keep up.

PHOTO BY NSERC

Demand for fast, cheap and plentiful data continues to surge, but wireless communications infrastructure is reaching the limits of what it can provide to users—unless we can find more efficient ways to engineer our networks.

Professor Wei Yu is doing exactly that—his influential work tackles the design and optimization of wireless communication systems. In 2015 he was awarded an E.W.R. Steacie Memorial Fellowship from the Natural Sciences and Engineering Research Council.

"I'm passionate about information theory, which is the mathematical foundation of modern digital communications," he said. Yu is currently investigating novel ways that base-stations and smartphones in a radio-access network may cooperatively transmit to and receive information from each other, in order to enhance signal quality and to reduce interference for wireless data access. His discoveries have impact on the network architecture, transceiver design, and network deployment for future generation wireless cellular services.

The award allows researchers who show extraordinary promise to focus on their research, relieved of teaching and administrative duties for a two-year term. Each Fellow receives a research grant of $250,000.

"Being able to focus on research 100 per cent of the time is hugely important," said Yu. "The Steacie Fellowship makes it possible for me to recruit the best students and post-docs to my research program to help develop the next generation of wireless technologies."

ECE INVENTIONS DISCLOSED OVER THE PAST FIVE YEARS

In cases of cross-disciplinary inventions, disclosures are expressed as a fraction

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

PROPORTION OF TOTAL NSERC FUNDING FOR ELECTRICAL AND ELECTRONICS ENGINEERING AT LEADING CANADIAN UNIVERSITIES, 2010-2011 TO 2014-2015

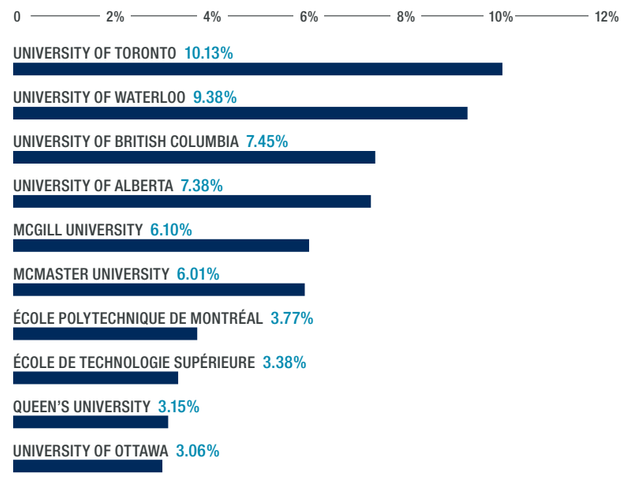

It's a bird—it's a plane—*it's an algorithm!*

Drones don't fly themselves—and neither do fighter jets, commercial airliners, or space shuttles. But as the era of the self-driving car dawns, why not design a self-flying plane?

That's exactly what control theorists at the Institute for Robotics and Mechatronics (IRM) are working toward. Their research, a collaboration between Professors Mireille Broucke of ECE, Angela Schoellig of the University of Toronto Institute for Aerospace Studies, and PhD candidate Mario Vukosavljev, aims to design control systems for flying quadrocopters—lightweight drones with four rotors.

Control systems are the 'brains' of machines, algorithms that allow them to 'behave' the way we want them to and often accounting for the unpredictable, such as responding to variations in environment or even mid-air collisions. The team's current control methodology considers safety, speed, and live response of the quadrocopter—next, they'll work on teaching their drone some fancy flying, such as slalom patterns and flips.

This is just one project underway in the Institute for Robotics and Mechatronics. Formed in 2010, IRM unites thinkers from across engineering disciplines and industry to develop more intelligent robots and devices to advance medicine and healthcare, rescue and exploration, military and security, and fields as yet unexplored. Learn more at **irm.utoronto.ca**

PHOTO BY ADRIAN ESSER, DYNAMIC SYSTEMS LAB, UTIAS

ECE RESEARCH FUNDING 2009–2010 TO 2013–2014

	FEDERAL	PROVINCIAL	INDUSTRY	OTHER	TOTAL $
2009-2010	8,637,522	3,948,866	1,481,958	2,878,488	**16,946,834**
2010-2011	9,239,607	3,201,271	2,238,062	3,187,226	**17,866,166**
2011-2012	10,403,537	1,569,027	2,960,981	3,192,836	**18,126,381**
2012-2013	9,727,505	2,481,833	2,355,784	3,347,918	**17,913,040**
2013-2014	11,498,219	3,662,706	3,145,172	4,110,224	**22,416,321**

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. The latest of these, on the topic of 'Information Security in a Connected World', showcased our research leadership in the areas of cloud security, security of cyber-physical systems, and protecting medical biometrics, and attracted participants from the fields of aerospace, business, medicine and more. Learn more about upcoming events at **uoft.me/eceevents**.

ECE accounts for half of all U*of*T Engineering startup companies and invention disclosures

WHERE THE MONEY GOES:

OUR 2015-2016 FINANCIAL PICTURE AT A GLANCE ($ MILLION)

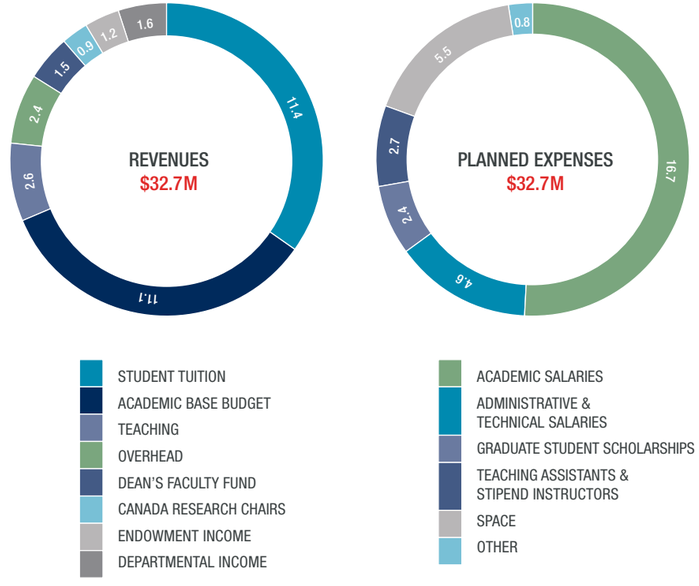

Where are they now?

Illuster Technologies Inc.

ECE alumni Miad Fard, Mehrad Mashayekhi and Richard Medal took the technology they developed for their fourth-year design project and founded Illuster Technologies Inc. in 2013. The trio and their unique hardware and software platform for teaching analog electronics were profiled in the last issue of *ANNUM*.

ANNUM caught up with co-founder Miad Fard to hear what three electrical engineers have learned from their crash course in the real world of entrepreneurship.

Q: WHAT'S DIFFERENT FROM 12 MONTHS AGO?

A: In the past year we have re-branded, re-launched our website, and struck a distribution partnership with one of the world's leading educational technology providers. Our original plan was to market and provide customer services on our own, but that would limit us to a very local portion of the market as we continued to grow enough to handle sales on a global scale. This would of course delay our ambitions of changing the way electronics is taught around the world. By bootstrapping and forming partnerships for marketing and service, we not only fulfill our ambitions, but we can shift our focus to what we excel at: designing and engineering labs to teach electronics.

Q: WHAT IS YOUR FOCUS NOW, AND IN THE NEAR-TERM?

A: Illuster is moving toward a more defined identity. When we first started the company, we were coming at it as engineers and our scope was a bit out of focus. As we gain experience on the business side, we start to see what we actually bring to the table. We are still focused on bringing a modern take on teaching electronics to the classroom, but we also want to provide development tools for enthusiasts and the electronics industry in general. We feel that the same way coding and computational literacy have made their way into the mainstream, eventually electronics will as well and we want to be providing tools to teach and develop electronics to the public once that time comes.

Q: HOW ARE YOU BALANCING THE DEMANDS OF A STARTUP COMPANY WITH OTHER PROFESSIONAL AND PERSONAL OPPORTUNITIES?

A: All three of us are trying to grow both professionally and intellectually as the company grows. We feel it's important to make sure you have a healthy relationship with your company, not only financially but personally as well. You always hear about passionate entrepreneurs sinking their life savings into their company, but a more overlooked scenario is when a young company eclipses the professional growth of its founders. It is important to be as realistic as you are passionate. You do have to sacrifice more time this way, but that has never been a problem for us; our mentality has always been: "Make sacrifices now. Get rewarded later."

175 Employers Hiring ECE Students on PEY Over the Past Five Years:

2MPOWER
360INCENTIVES.COM
407 ETR CONCESSION COMPANY LTD.
A THINKING APE
ABB SWITZERLAND LIMITED
ACUITY
ADEXA INC.
AERCOUSTICS
AESO (ALBERTA ELECTRIC SYSTEM OPERATOR)
AFFINITY SYSTEMS LIMITED
AGFA
AGFA GRAPHICS NV-BELGIUM
ALTERA
AMAZON
AMD
AMERICAN EXPRESS CANADA INC.
ANALOG DEVICES INC.
APPLE
ARISTA NETWORKS
ARUP
ASSETMINE INC.
ATS AUTOMATION TOOLING SYSTEMS INC.
AUTOLIV ELECTRONICS CANADA INC.
AVIYATECHNOLOGIES INC.
BANK OF MONTREAL
BELL CANADA
BELL MOBILITY
BELL TV / VISIONMAX
BIG VIKING GAMES
BLACKBERRY LIMITED (FORMERLY RIM)
BLUECAT NETWORKS
BMW GROUP CANADA
BOMBARDIER AEROSPACE
BROADCOM
BRUCE POWER
BURO HAPPOLD
CANADA HEALTH INFOWAY INC.
CANADA POST CORPORATION
CANADIAN NATURAL
CANCER CARE ONTARIO
CASEWARE INTERNATIONAL INC.
CAST SOFTWARE INC.
CELESTICA INC.
CGI GROUP INC.
CIENA
CIRBA INC.
CISCO SYSTEMS INC.
CITY OF BRAMPTON
COMMUNICATIONS AND POWER INDUSTRIES CANADA INC.
DELOITTE & TOUCHE LLP
DEMONWARE
DEPARTMENT OF NATIONAL DEFENCE
DESIRE2LEARN
DOUBLETHINK INC. (MYBLUEPRINT)
DRDC TORONTO
EASTERN POWER LTD.
ECAMION
ENERSOURCE
ENVIRONMENT CANADA
ENVISION MOBILE (FORMERLY VIZIO MOBILE INC.)
EPSON
ERICSSON CANADA INC.
ESNA TECHNOLOGIES
EVENTMOBI
FINANCEIT CANADA INC.
FIO CORPORATION
FIXMO
FU TIN BUSINESS CO. LTD.
GAZOOMOBILE
GE CANADA
GENERAL ELECTRIC
GENERAL MOTORS OF CANADA LTD.
GENESYS VG INC.
GLOBALIVE COMMUNICATIONS
GOOGLE
GUOMAO GROUP CO. LTD.
HH ANGUS
HITUPP INC.
HONDA
HONEYWELL
HP CANADA
HUSKY ENERGY
HUSKY INJECTION MOLDING SYSTEMS LTD.
HYDRO ONE
IBM CANADA LTD.
IESO
IMPERIAL OIL
INTEGRATED DEVICE TECHNOLOGY (IDT)
INTEL CANADA
INTELLIWARE DEVELOPMENT INC.
INTUIT CANADA
ITG
KAPSCH TRAFFICCOM CANADA INC.
KIJIJI
KNOWROAMING
LABATT BREWERIES OF CANADA
LOBLAW COMPANIES LTD.
M&E ENGINEERING
MAGNA CLOSURES
MARIN SOFTWARE
MARVELL TECHNOLOGY GROUP
MAXXIAN
MCW CONSULTANTS LTD.
MEDIAADX
METROLAND DIGITAL
MICROSEMI CORPORATION
MINISTRY OF CHILDREN AND YOUTH SERVICES
MINISTRY OF TRANSPORTATION
MODIFACE
MOLD-MASTERS LTD.
MOTOROLA SOLUTIONS/PSION
NATIONAL INSTRUMENTS ELECTRONICS WORKBENCH GROUP
NATIONAL RESEARCH COUNCIL CANADA
NEXJ SYSTEMS INC.
NTHGEN SOFTWARE
NVIDIA
OMNIVEX CORPORATION
ONTARIO FINANCING AUTHORITY
ONTARIO MINISTRY OF EDUCATION
ONTARIO MINISTRY OF HEALTH & LONG-TERM CARE
ONTARIO POWER AUTHORITY
ONTARIO TEACHERS PENSION PLAN
ONTARIO POWER GENERATION
PEARSON EDUCATION
PERSPECSYS
PHILIPS (CHINA) INVESTMENT COMPANY
POINTCLICKCARE
PSION INC.
QNX SOFTWARE SYSTEMS
QUALCOMM CANADA
QUANTA TECHNOLOGY
QUICKPLAY MEDIA
RBC CAPITAL MARKETS
RED HAT CANADA LTD.
RESEARCH IN MOTION
RL SOLUTIONS
ROGERS COMMUNICATIONS
ROYAL BANK OF CANADA
RUGGEDCOM INC.
SANDVINE
SAP CANADA INC.
SCICAN
SCOTIABANK
SEMTECH
SIEMENS
SMITH AND ANDERSEN CONSULTING ENGINEERING
SOTI
SOUTHPAW TECHNOLOGY INC.
SUNNYBROOK HEALTH SCIENCE CENTRE
SUNWELL TECHNOLOGIES INC.
SURFEASY
SYMANTEC CANADA LTD.
TECK RESOURCES LTD.
TEKNION FURNITURE SYSTEMS
TELUS
TEMENOS
THALES RAIL SIGNALLING SOLUTIONS
THE HOSPITAL FOR SICK CHILDREN
TIMEPLAY
TORONTO HYDRO
TORONTO TRANSIT COMMISSION
TRANSCANADA
TRAPEZE GROUP
TXIO
TYCO ELECTRONICS CANADA LTD.
UBER TECHNOLOGIES INC.
UNIVERSITY HEALTH NETWORK
UPSTREAM WORKS SOFTWARE
V SEMICONDUCTOR INC.
VENNSA TECHNOLOGIES INC.
WATTPAD
WEALTHFRONT
WINMAGIC INC.
XAGENIC INC.
XTREME LABS INC.

ECE CORPORATE PARTNERS, 2011-2015

A.U.G. Signals Ltd.
ABB AB
Advanced Micro Devices Inc.
AEG Power Solutions Inc.
Alcatel Canada Inc.
Altera Corporation
Asahi Glass Co., Ltd.
Bell Canada
BLiNQ Networks Inc.
Brammo Inc.
Broadcom Corporation
Carinthian Tech Research Institute
Christie Digital Systems Canada Inc.
Ciena Canada Inc.
Digital Predictive Systems Inc.
Diros Technology Inc.
E. I. du Pont Canada Company
eCamion Inc.
Ericsson Canada Inc.
Exar Corp.
Finisar Corporation
Fuji Electric Co., Ltd.
Fuji Electric Systems Co., Ltd.
Fujitsu Laboratories Ltd.
Fujitsu Labs of America Inc.
Gener8 Inc.
Genia Photonics Inc.
Gennum Corporation
GO Lighting Technologies Inc.
Hatch Ltd.
Hewlett-Packard Company
Huawei Technologies Co., Ltd.
Hydro One Networks Inc.
IBM Canada Ltd.
Intel Corporation
International Business Machines
IPPLEX Holdings Corporation
ITS Electronics Inc.
IVG Fiber Ltd.
Kapik Inc.
Kinectrics Inc.
Lattice Semiconductor Ltd.
Lockheed Martin Canada
Mark IV Industries Corp.
Mattson Technology Canada
Maxim Integrated Products Inc.
Meta North Inc.
Mircom Technologies
Mitsubishi Electric Research Lab.
Morgan Solar
National Semiconductor
Neurochip Corp.
NXP Semiconductors Netherlands B.V.
OneChip Photonics Inc.
Ontario Power Authority
Opus One Solutions Energy Corp.
Polaris Industries
QD Solar Inc.
Qualcomm Canada Inc.
Qualcomm Technologies Inc.
Quanser Inc.
Rambus Inc.
Raytheon Canada Limited
Redline Communications Inc.
Research in Motion Ltd.
Robert Bosch Corporation
Semiconductor Research Corporation
Sendyne Corp.
Silicon Mitus Inc.
SINTEF Energi AS
Solana Networks
Solantro Semiconductor Corp.
Solar Ship Inc.
Taiwan Semiconductor Manufacturing
TELUS
Telus Mobility
Texas Instruments
Thales Canada Inc.
Toronto Electric Ltd.
Toshiba Corporation
Total American Services Inc.
Ultra Electronics
Unisearch Associates
Varilume Lighting Inc.
ViXS Systems Inc.
Wurth Elektronik eiSos GmbH&Co. KG
Xilinx Inc.
Xogen Technologies Inc.
Zentrum Mikroelektronik Dresden AG

$13.1 million
Corporate funding to ECE over the past five years

HOW TO MULTIPLY YOUR MONEY

In ECE, we leverage industry financial support with federal, provincial and institutional matching to enable millions in industry-partnered research per year. A $52,500 industry contribution to a research project grows to $202,500—a 4x leverage on the initial investment.

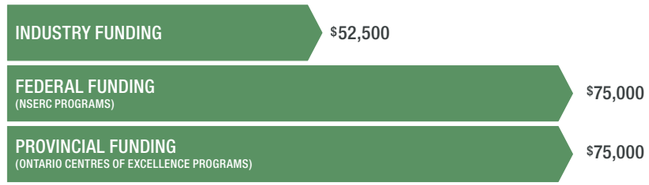

The outcome: a joint project with the support of world-renowned researchers, expert administration, project management and motivated graduate students.

Contact us to explore the vast potential of partnering on research and technology development.

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

ILLAN KRAMER, PHD
DIRECTOR OF CORPORATE, GOVERNMENT & INTERNATIONAL PARTNERSHIPS
ILLAN.KRAMER@UTORONTO.CA
416 978 6990

ECE INDUSTRY ADVISORY BOARD

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering formed its Industry Advisory Board in 2014. The board provides ECE with objective advice about the department's role in contributing to technological and economic development, and engineering practice. Its members are:

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, FORMERLY DIRECTOR OF SOFTWARE ENGINEERING // ALTERA
SONG ZHANG, DIRECTOR, TECHNOLOGY PLANNING & PARTNERSHIP // HUAWEI CANADA

Alumni

STRENGTH OF THE ECE ALUMNI NETWORK

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—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 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.

Total number of active alumni worldwide:
11,381

ECE ALUMNI BOARD OF ADVISORS:

ALAN BOYCE
ElecE 7T8
Senior management consultant, entrepreneur

JOHN EAST
ElecE, MBA, UC Berkeley
Former CEO of Actel Corporation

ALEX GRBIC
CompE 9T4, MASc 9T6, PhD 0T
Director, Product Marketing at Altera

CATHERINE LACAVERA
CompE 9T7, JD/MBA UofT
Director, Litigation at Google

NATASHA LALA
CompE 9T8
Chief of Staff at OANDA Corporation

ALEX SHUBAT
ElecE 8T3, MASc 8T5,
PhD Santa Clara, MBA Stanford
Former CEO of Goji Food Solutions Inc., executive advisor, entrepreneur

PHOTO BY ROBERTA BAKER

Catch the premiere of SkuleTalks, a new ECE video series

Stay up-to-date on the latest research underway in ECE with SkuleTalks—short videos that keep ECE alumni current on the big questions in fields from wearable tech to smart-grid security and more. Whether you've just graduated or haven't designed a circuit in 30 years, SkuleTalks are an accessible way to catch up on the ideas that matter to today's leading engineering thinkers.

View fresh SkuleTalks:
uoft.me/skuletalks

KING EMBODIES SPIRIT OF SPRING REUNION

Alex King (ElecE 3T5) was the only University of Toronto alumnus to receive a medal celebrating 80 years since his convocation at this year's Spring Reunion Chancellor's Circle Medal Ceremony.

"It makes me very proud," King said. "Receiving an 80-year medal is pretty special."

The 102-year-old King attended Spring Reunion with his wife, Claire, son, Bill, and daughter, Helen. He lives in Collingwood, Ont. and had not visited the St. George campus in many years—he said that King's College Circle looks similar, but there are many new buildings he has never seen before. When he was enrolled in electrical engineering, all of his classes took place in The Little Red Skulehouse—the original home to U of T Engineering, which was demolished in 1966.

His affiliation with the renowned Class of 3T5 is one he cherishes. He said that alumni from his class met regularly for many years, and were still having reunions when he was well into his 80s. "I had a lot of buddies in those days," he said. "It seemed like we got closer as we got older."

King said that throughout his life he was a "victim of circumstance"—always embracing new opportunities. He feels that it's important that students and young alumni alike embrace the same mentality. "Never hesitate and don't miss a chance," he said.

FALL ALUMNI LECTURE & NETWORKING RECEPTION

ECE alumni from all graduating classes returned to Skule again last Fall to catch up with classmates and former professors. In November 2014, attendees heard Professor Jonathan Rose share stories of his newest interdisciplinary research on mobile applications.








PHOTOS BY JIMMY LU

FOURTH-YEAR & ALUMNI CELEBRATION

Each spring, on the day all fourth-year students finish presenting their final capstone design projects, ECE hosts a reception to recognize this important milestone and toast its imminent graduates. This year, alumna Kia Puhm (CompE 9T4+PEY) shared her memories of ECE and urged the graduating class to nurture the valuable relationships they formed with their classmates here.







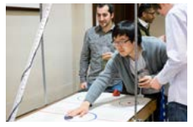
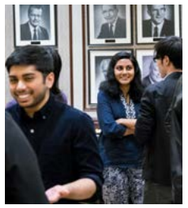


PHOTOS BY VINCENT TSE

SPRING REUNION

This year's Spring Reunion drew graduates from the class of 5T0 all the way up to 1T5 back to ECE's departmental lunch and lab tours on Saturday, May 30, 2015. Attendees compared notes on how the department has changed since they were students, enjoyed a meal together, and visited the state-of-the-art Toronto Nanofabrication Centre and UofT Engineering's Blue Sky Solar Team workshop. Mark your calendars for the next Spring Reunion, planned for May 28, 2016!






CONVOCATION

Following our big spring Convocation ceremony in June, ECE hosted a lunch to congratulate the classes of 1T5 and 1T4+PEY. Despite heavy skies, spirits were light as newly minted alumni and their families joined faculty and staff at the celebratory event on campus.








PHOTOS BY JIMMY LU

CLASS NOTES

6T4

JOAN HIRDLE (NÉE ALEXANDER)
ELECE 6T4

Over 50 years since graduation and my memories of ECE are pretty dim. We had to do a small program on one of the first university computers. It had paper tape input. Every time you corrected a mistake the paper tape dupe machine would dupe a correct input wrong. Of course I went into computers. I married an Englishman and we moved to Australia. We had two sons. While they were young I did contract work and then started back into full time work at 50. I retired at 66. I am now so busy I don't know how I found time to work. I play a lot of competition bridge and enjoy doing things on my iPad.

LEN KLOCHEK
ELECE 6T9, MASC 7T2

Retired in 2011, after 25 years of teaching at Seneca College, and 15 years spent in industry.

Keep in regular contact with classmates Brian, Hillar, Marty, Neil, and Van.

JAMES (JIM) WELCH
ELECE 6T9, MASC 7T4, JD 1982

Briefly—work at Nebraska Med. Center 2+ years, work at local Power Co. 2+ years, obtain Juris Doctorate, test into Nebraska Legal Bar, Professional Engineers, Notary and into US Patent Bar. Obtain US DOE Grant, successfully fabricate Single Device CMOS at UNL with help from U of T, obtain many patents on results. Obtain 300+ patents for clients, (eg. J.A. Woollam Co., Nebr. Univ. Profs. and Independents, as well as for own inventions). Have one of the best Scientology collections in the world, and continue to study it. One hobby is working on old cars. Practice as Prof. Engineer some, but mostly as Patent Attorney for 30+ years. Still single (as one secretary gal said—who would put up with you—Indeed!). Am easing toward retirement, but clients won't let me! (See USPTO, LinkedIn and Avvo Websites for more.)

6T9

JAY BEATTIE
MASC 6T9, PHD 7T5

For now, I will share with you this photograph taken circa 1968. It is of eight graduate students in Electrical Engineering at U of T. I am one of them and the rest were visiting my cabin on Gloucester Pool in the Severn River in Ontario. I don't remember all their names so see if alumni readers can identify them. We were mostly in the Power Systems area working with Stuart Robertson and Wasyl Janischewskyj.

7T6

PETER ARATO
ELECE 7T6, MASC 7T9

Undergrad UofT EE in the early 70's was a wonderful balance of learning and socializing. I was fortunate to experience every aspect to the fullest, including being a first year class rep, secretary of the engineering society, a sometime Lady Godiva member, and still somehow learn my profession well enough to enter the MASc program. My time at the Institute of Biomedical Engineering was far more focused on academics and research with only a few necessary diversions such as the annual Slave Auction fundraiser. It was the most enjoyable yet stressful time in my life, and I hope eventually when I retire, I look forward to returning to University and continue my education.

7T8

MICHAEL MIRSKY
ELECE 7T8, MENG 8T5

Memories of weekends at EUT pounding away at the keypunch machines to get that FORTRAN program working right! Misery was dropping your deck of punchcards! Remember the old original Sandford Fleming 126 before it burned down? It was a huge lecture hall decorated by punch card airplanes stuck in the ceiling that had been launched by rubber bands! After graduating, I walked across the street along with a good number of my classmates to join Ontario Hydro. After a wonderful 35 year career with Ontario Hydro/Ontario Power Generation, am now retired living in Israel. Regards to all my classmates!

7T9

DAVID WAECHTER
ELECE 7T9, MASC 8T2, PHD 1987
AT CARLETON UNIVERSITY

David Waechter (7T9) is releasing a book this year about his late father's career. His father, Ralph William Waechter, was a University of Toronto graduate in aeronautical engineering (4T8). The first chapter of the book discusses Ralph Waechter's graduating class and the early work on aeronautics at the university. A later chapter discusses how Ralph once brought his son David along to the University of Toronto on a Saturday afternoon to watch as he ran his deck of computer punch cards through the card reader (at the engineering annex building) for a special project. Little did he know that some years later, the same son would be running his own punch cards through the same system as an EE student. The book is titled "*Flight Test: The Avro Arrow and a Career in Aeronautical Engineering*" (available through www.innerscale.com).

8T0

SIMON DODGE
ELECE 8T0

Principal Engineer, zSeries Security Engineering/Strategy, WellsFargo Bank. Moved to Atlanta, Georgia, USA in 1985, married 1991. Specialist in zSeries security controls (pka mainframes). Anticipating traveling in retirement (hopefully soon).

8T5

AKIKO OGAWA
ELECE 8T5, MASC 9T0

Been living and working in Japan since my MASc graduation. Not much new to report there, except for the recent relocation to an office building along the Tokyo coast.

My fondest memories of ECE include being the backstage crew for Skule Nite, volunteering at engineering pub nights, summer softball games during grad school, playing cards in the Engineering cafeteria and studying for exams at the Robarts Library with my classmates.

8T6

NICOLAS DANIELS
ELECE 8T6, MASC 8T9

Hi everyone from Elec Eng 8T6 and all the former Dukes of Devonshire, particularly "East House Beast House". I'm retired now with cash for life, had a great career in the public service with CSEC...
Regards: Nick *(Jack, Scratch)* Daniels

8T9

JOHN LAZAROU
ELECE 8T9

Just moved back to Toronto last year after 16+ years of working in Finance in New York.

9T0

SAM AHUJA
ELECE 9T0, MASC 9T3

I shifted into software engineering... who would have thought the principles of engineering and problem solving I learned in power systems could be applied to software engineering and running a business in general. I have fond memories making friends and surviving through engineering together!

LAWRENCE LAI
ELECE 9T0

Now in Hong Kong.

9T1

TIN-SUN (TIM) LEE
ELECE 9T1+PEY

Proud dad of two boys with loving wife. Lived in the city & downtown before moving to Markham.

Worked at Nortel, Celestica and currently at Creation Technologies.

Built a career as an in-circuit test (ICT) engineer for EMS companies and I also support modern applications of communications and instrumentation functional/system testing in a manufacturing environment.

I have fond memories of my years at UofT and of the fantastic PEY (16 month) experience at IBM. I wish all the best to our alma mater, to all who have attended and to future graduates.

9T5

TONY ORSI
ELECE 9T5, MASC 9T8

Great memories of my time at Skule in the ECE and IBBME programs where I learned lots of useful information that I am now using in my practice of Intellectual Property Law. Would love to hear from others. Reach me at torsi@bereskinparr.com

9T7

KEVIN GIOR
MASC 9T7

This year, I established an international solution trading company, Aewon Solutions Corporation, in Tokyo, Japan. Please feel free to contact me if you have any innovative solutions to offer in Japanese and Korean market. Looking forward to hearing from you soon.
kevin.gior@aewon-solutions.com

0T2

KENAN HUSKOVIC
ELECE 0T2

Enjoying life and work at MMM Group Limited—a Canadian multi-disciplinary Consulting Engineering firm (mmmgrouplimited.com). Reach out to me if you are interested and want to discuss opportunities: huskovick@mmm.ca

0T3

MATHEW SZETO
COMPE 0T3+1

Currently working in London, UK at RBC Capital Markets on the Institutional Equities Trading desk.

JUNQIAN ZHANG
ELECE 0T3+PEY

I really miss the numerous days and nights we spent in BA common room, working on projects, assignments and exams.

0T6

BRIAN CHAU
ELECE 0T6+PEY

IP Lawyer at Norton Rose Fulbright.

BUSHRA JAVAID KHAN
ELECE 0T6, MENG 1T2

One marriage, a master's degree, one toddler, two pregnancies, and nine years of working in the chip manufacturing industry as a design verification engineer has been incredible to say the least. When I graduated from ECE with a bachelor's I was hopeful, proud and full of energy. I didn't know much about what exactly I wanted to do but my suggestion to all the new grads out there is do something you love. You will never regret it. I love my work and it makes it easier to balance work and family. Follow your heart. ECE thanks for giving me so many opportunities!

Jot a class note! Email your news to eceinquiry@utoronto.ca

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

Distinguished Lectures Series
2015 - 2016

Vincent Poor
Princeton University
Sept. 17, 2015

Fundamental Limits on Information Security and Privacy

Kristin Pettersen
Norwegian University of Science and Technology
Oct. 22, 2015

Snake Robots: A Robotic Solution for Firefighting, Search and Rescue, and Subsea IMR Operations

Steven Low
California Institute of Technology
Nov. 12, 2015

Control and Optimization of Smart Grid

Michal Lipson
Columbia University
Nov. 19, 2015

Silicon Photonics: The Optical Spice Rack

Susan Hagness
University of Wisconsin - Madison
Jan. 28, 2016

Imaging Tissue and Treating Cancer with Microwaves

Boris Murmann
Stanford University
Feb. 25, 2016

Mixed-Signal Circuits for the Data-Driven World

Ranu Jung
Florida International University
March 17, 2016

Closing the Loop: Nerves, Machines and Interfaces

Kathryn McKinley
Microsoft Research
April 8, 2016

Programming the World of Uncertain Things

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.

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

This Directory contains a listing of all research projects currently underway in ECE. Projects are listed alphabetically by principal investigator, and categorized by topic or field using six distinctive icons.

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		•		•		
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		•				
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. 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 them. 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 image 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, of copyright violations, video compression, and intelligent image resizing, as well as a broad range of other important applications. Although different approaches to understanding visual information have been explored in the past, one of the most promising 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 (but 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

■ Automatic Performance Tuning for GPUs

Graphics Processing Units (GPUs) have been effectively used to accelerate many applications. These many-core processors deliver a performance an order of magnitude higher than multicore cores at a fraction of the power. However, GPUs require application developers to restructure, or optimize, their application codes to exploit the underling GPU architecture. These optimizations are tedious to apply, may or may not benefit performance and interact with one another in non-intuitive ways. The goal of this project is to ease this burden on programmers by developing compiler-based automatic performance tuning of GPU applications. We explore two approaches. The first attempts to correlate application code features with the effectiveness of the optimizations for a large set of training applications in order to build a machine-learning model. The compiler uses this model to decide what optimizations to apply. In a second approach, the space of possible optimizations is incrementally explored using heuristics to determine a good set of optimizations to apply. The project also builds the necessary compiler and run-time support for auto-tuning.

■ Directive-Based Programming and Optimizations for GPUs

This project develops high-level programming models for GPUs. More specifically, it involves the design and implementation of *hi*CUDA, a directive-based language for GPUs. The language facilitates programming GPUs through simple directives added to the sequential code while maintaining the well-adopted CUDA/OpenCL programming models, and does so with no penalty to performance. We developed a prototype hiCUDA compiler, which we have released to the public domain at www.hicuda.org. We are currentlyextending this infrastructure to support directive -based optimizations for OpenCL and CUDA kernels.

■ High-Performance FPGA Overlays

Field Programmable Gate Arrays (FPGAs) offer massively parallel resources that, if exploited by application developers, can deliver high levels of performance. However, the widespread use of FPGAs to accelerate applications is hindered by (1) their low-level programming abstraction that requires expertise in hardware design—expertise that application developers often lack; and (2) the long development cycles associated with FPGA design tools, to which software developers are not accustomed. In this project, we design, implement and evaluate overlays (FPGA circuits that are in themselves programmable) that can make the use of FPGAs by software developers easier. We designed and prototyped an overlay architecture that projects the software model of pipelined dataflow graphs (DFGs). Instances of this overlay

architecture can deliver performance in the gigaflops range that scales with FPGA resources and are fast and easy to program. We are current exploring compiler-based solutions for automatic extraction of DFGs of applications; determining the best overlay instance for a given application; extending the design of the overlay to multiple FPGA devices; and exploring a just-in-time compilation framework for dynamically and transparently translating binary code into overlay circuits.

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; and 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 that 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

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 FPGA speed, area, power and ease of use.

LEGUP.EECG.TORONTO.EDU

High-Level Synthesis of Hardware Circuits from Software Programs

High-level synthesis (HLS) raises the abstraction level for hardware design by allowing a software program to be automatically synthesized into a hardware circuit. HLS aims to offer the flexibility and ease of use associated with software, along with the speed and energy advantages of customized hardware. Both of the major FPGA vendors have been investing heavily in HLS in recent years, and there has been much research on the topic in academia as well. State-of-the-art HLS is nearing the point where software engineers are able to design hardware, with the quality of the HLS hardware produced being comparable to human-crafted implementations. LegUp is a high-level synthesis tool under active development at the University of Toronto. LegUp accepts a C program as input and automatically compiles the program to a hybrid architecture comprising a processor (a soft-core MIPS or a hardened ARM) and custom hardware accelerators. Results show that LegUp produces hardware solutions of quality comparable to that of commercial high-level synthesis tools. LegUp is open source and freely downloadable (legup.eecg.toronto.edu), providing a powerful platform that can be leveraged for new research on a wide range of HLS and hardware/software co-design topics.

Bardakjian, Berj

WWW.IBBME.UTORONTO.CA/FACULTY/CORE-FACULTY/BERJ-L-BARDAKJIAN/

Bioengineering of the Brain

The main themes of the research fall 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, 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 cognitive devices developed will be implemented as low-power hardware to be incorporated into the biological neural networks in a closed feedback loop. These will provide implantable devices to be used 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 that 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

ISL.UTORONTO.CA

● ■ ⬣ Digitally Assisted Analog Front Ends

When sensing signals in the physical world, undesired interfering signals with an amplitude much larger than the targeted signal of interest often arise. 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: 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

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 using magnetic resonance imaging (MRI), for non-invasive, deeptissue penetration in a living subject. Our goal is to improve detection sensitivity and specificity so as 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) for enabling critical advances in tissue engineering and regenerative medicine. Work focuses specifically on 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/

■ ⬣ FPGAs for Cloud Computing Platforms—Virtualization and Applications

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 regardless of the physical hardware platform. This project explores (1) how to place FPGAs into this virtualized environment when using Open Stack; and (2) how to scale up heterogeneous programming environments to work in a large-scale heterogeneous environment. The prototyping platform is the SAVI Networks platform (savinetwork.ca) where real applications are being built to drive the development of the programming environment as well as to characterize the behaviour and performance of large-scale systems.

■ ⬣ 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.

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 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 used in the control of large flexible space structures, earthquake-resistive building structures, and electric power systems with particular focus on microgrid systems and spinal cord injury patients.

Dawson, Francis

WWW.ELE.UTORONTO.CA/~DAWSON

Improving Energy Efficiency of Energy Conversion Processes

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 for 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 600 mV 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 run-time 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 crypto-graphic 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 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/MAIN.HTML

Engineered Materials (Metamaterials) from Microwave to Optical Frequencies

We are developing paradigm-shift metamaterial devices and subsystems, and related technologies from RF/microwaves to optical frequencies. Metamaterials are engineered materials 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, satellite communications, radar, 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 have been quite successful in securing faculty positions in academia (e.g., UMich, UAlberta, McGill, U of T and UBC) and industry (e.g., Apple, AMD, Google, 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.

Interconnect Solutions for Interposer-Based Systems

Silicon interposer technology (“2.5D” stacking) enables the integration of multiple memory stacks with a processor chip, thereby greatly increasing in-package memory capacity while largely avoiding the thermal challenges of 3D stacking DRAM on the processor. Systems employing interposers for memory integration use the interposer to provide point-to-point interconnects between chips. However, these interconnects only utilize a fraction of the interposer’s overall routing capacity, and in this work we explore how to take advantage of this otherwise unused resource. We are exploring general approaches for extending the architecture of a network-on-chip (NoC) to better exploit the additional routing resources of the silicon interposer and to take advantage of new opportunities afforded by the interposer. We propose an asymmetric organization that distributes the NoC across both a multicore chip and the interposer, where each sub-network is different from the other in terms of the

traffic types, topologies, the use or non-use of concentration, direct vs. indirect network organizations, and other network attributes. Through experimental evaluation, we show that exploiting the otherwise unutilized routing resources of the interposer can lead to significantly better performance.

■ 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 into the early-stage on-chip network design space exploration and 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 turnaround times.

Francis, Bruce

WWW.SITES.GOOGLE.COM/SITE/BRUCEFRANCISCONTACT

▲ Retired from teaching and research, now writing textbooks

I have retired from teaching and research and am now writing textbooks:
(1) *Flocking and Rendezvous in Distributed Robotics*
(co-authors are Bruce Francis and Manfredi Maggiore).
(2) *A Brief Course in Digital Signal Processing*
(co-authors are Bruce Francis and Deepa Kundur).

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 online web tool developed by Dr. Frey's group was accessed over 100,000 times by over 3,000 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 daughtercard 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 with 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 includes 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) lightweight 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 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 Lasers

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 has fueled innovations in all industrial sectors. A new class of semiconductor lasers with even more distinctive features has been developed recently. 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 in comparison to 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 tunability and spectral brightness. For example, certain molecules containing carbon-hydrogen bonds have infrared footprints within the 2–3 μm spectral window. In addition, H_2O 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 numerous fields of applications such as in 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 an RF source. In addition, the operating wavelength is tuned by sweeping the central wavelength of the bandpass filter. This new technique is also cost effective and provides the possibility for hybrid integration as it consists 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 and no need for RF sources, optical stabilization or optical feedback. Recently, we successfully achieved injection locking using a 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 to produce 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 a big space on an optics table. The optical setup requires delicate construction and is sensitive to 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, and quantum 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 has been achieved. Techniques and applications to 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 attests 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 sciencetific 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 in order 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. It will then be possible 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. This technique uses 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 and 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 embed 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 aspect-oriented 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 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 is an open-source, enterprise-grade event management infrastructure that is designed for large-scale event management applications. Ongoing research seeks to add to 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) refers 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

IBBME.UTORONTO.CA/FACULTY/EMERITUS-FACULTY/MIKE-JOY

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 1930s. Today, the best-known method is Electric Impedance Tomography (EIT). EIT measures currents and voltageson 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 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 of 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 Co-simulator Platform for Cyber Security Analysis of Microgrid Systems

Effective modelling and simulation of complex power system disturbances, especially those stemming from intentional cyber attack, represents an open engineering research and development problem with recent national focus. The Northeast Blackout of 2003 and the December 2013 large-scale power outage in the Greater Toronto Area clearly demonstrate the vulnerability of the Canadian grid to incidental and natural disruptions; given the increasing dependence of power systems on communications and computation, intentional cyber attack would thus have potential for great devastation. Simulators are a cost-effective and safer alternative to conducting experiments with prototype or real systems. Experiments can also be executed faster than in real time for efficient what-if analysis. Thus, tools for modelling and simulation of smart-grid systems are of paramount importance to power system stakeholders for judicious planning and preparedness for contingencies. Challenges stem from the need to develop intelligent models of cyber-physical interdependencies within emerging smart-grid systems, the facility to portray realistic and meaningful cyber attacks, and the ability to balance precision, scale and complexity. This three-way collaborative project amongst Professor Deepa

Kundur, University of Toronto, L'Institut de recherche d'Hydro-Québec (IREQ) and Opal-RT Technologies, Inc. aims to develop a cyber-physical co-simulator platform for the purpose of studying the impacts of cyber attacks on emerging microgrid systems. The results of the project will benefit a major microgrid project within IREQ by providing a framework to test communication and control strategies. Furthermore, the software integration insights arising from the research will be transferred to Opal-RT to equip them with knowledge to better support their existing and future clients.

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 resource 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.

Online Seizure Detection Using EEG-Based Features of Neural Rhythms Measured from a Wireless Headset

In this project we study signal processing algorithms for the early real-time detection of seizures. Our algorithms make use of EEG-based signals measured from a portable headset developed at Avertus Epilepsy Technologies. The algorithms must address interference of the EEG signals to identify the most critical features for detection. This project is in collaboration with Prof. Berj Bardakjian of ECE and IBBME at the University of Toronto.

Smart Grid: Cyber-Physical Operation, Security and Quantum Technology

Future electricity generation and distribution networks will be enormously sophisticated. They will incorporate the ever-increasing renewable sources of energy (e.g., solar, wind, hydro) in a most cost-effective and energy-efficient way; they will need to be robust and resilient in the event of a natural disaster or a malicious attack; and they will enable competing firms to share the same physical and cyber infrastructure while protecting the private usage information of millions of individuals. In short, the networks will have to employ advanced sensing, computation, and communications mechanisms to enable efficient and flexible generation, transmission, and distribution of power through a grid—a Smart Grid. We propose to build a research infrastructure platform to study the control and communication of a large-scale smart grid. A salient feature of the proposed infrastructure is the incorporation of a physical quantum network, reserved for the most critical part of the smart-grid communication, where security, authenticity, and latency of information exchange are critical to system monitoring and control. The project consists of four carefully crafted components: (i) the development of a hardware-in-the-loop real-time simulator platform for large-scale smart grids employing in-house FPGA-enabled Real-Time Digital Simulator (RTDS) systems that provide fast and reliable study of complex power systems; (ii) the design of a real-time heterogeneous smart-grid communication network simulator that enables the dynamic study of practical communication constraints and a variety of cyber attacks at different network layers along with a defense-in-depth approach for protection; (iii) the construction of an information-theoretically secure physical quantum key distribution (QKD) network, consisting of commercial QKD systems as well as a newly proposed measurement-device-independent QKD system, to be custom built in-house; and (iv) the overall integration of the first three components to produce a dynamic real-time large-scale bed test, the first of its kind, facilitating unprecedented insights into smart-grid integration and operation.

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.NAL.UTORONTO.CA/RESEARCH

Connected Vehicles and Smart Transportation

The ORF Research Excellence Project on Connected Vehicles and Smart Transportation is a collaboration 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 that enables users to make decisions that reduce travel time, increase productivity and reduce energy consumption and vehicle emissions. A live portal showing the state of traffic in the Greater Toronto Area is available at cvst.ca.

Design of Next-Generation Smart Infrastructures and Service Platforms

The convergence of three technologies—cloud computing, software-defined networking, and machine-to-machine (M2M) communication—provide an opportunity to create application platforms that offer unprecedented technical capabilities, scalability, energy efficiency, security, flexibility, and economics. Cloud computing provides on-demand computing power for applications at unprecedented price points. Software-defined networking allows flexible network equipment to be tailored to the needs of applications. M2M communication allows sensor and control devices to be attached to the Internet in support of new applications. In combination, these three technologies enable applications that are smart in the sense of being aware of context as well as the state of the environment. These applications will be used to manage resources in smart infrastructures (transportation, power, water, air quality, buildings, etc). They will also provide novel services, such as a personal assistant that is aware of context, activity schedule, and the goals of an individual. We view computing and communications resources as organized in multiple tiers, with remote massive data centres at the heart of the cloud, a smart edge providing telecom services, and then a fog consisting of gateways and sensor and actuator devices. Creating and deploying applications across this multitier cloud is challenging because of the variety of diverse technologies and incompatible protocols. Our research focuses on the design of management systems that span the cloud, smart edge, and fog to coordinate the allocation of resources to applications.

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.

NSERC Strategic Network on Smart Applications on Virtual Infrastructures

The NSERC Strategic Network on Smart Applications on Virtual Infrastructures (SAVI) 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. SAVI has designed and deployed a seven-node national testbed to support experimentation in future internet protocols and architectures as well as future large-scale applications. The SAVI testbed is federated with the U.S. GENI testbed for network innovation.

Levi, Ofer

WWW.BIOPHOTONICS.UTORONTO.CA

Optical Biosensors and Biomedical Imaging Systems

Our main fields of interest include biophotonics and semiconductor optical devices, and in particular, the development of miniature optical biosensors and biomedical imaging systems enabled by semiconductor optical devices and nanostructures. We seek to design, fabricate and use miniature optical imaging systems for portable microfluidics diagnosis systems and for in vivo applications such as optical brain imaging and continuous monitoring of tissue kinetics. Our research is divided into two main categories: (i) miniature biosensors for optical sensing inside Lab-on-a-Chip microfluidics chips; and (ii) biosensors and optical imaging systems for portable imaging inside the body (in vivo imaging) for neural imaging, tissue imaging, and cancer studies applications. Our lab collaborates with research labs and with physicians at Toronto's hospitals, to translate our studies and apply the imaging systems we develop to disease monitoring and patient care.

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.

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 co-operative 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 co-operative 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 produced 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 on 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 datacentres. 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 to 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, which enabes mobile terminals to dynamically and autonomously adapt to changing environmental conditions and reduces their energy consumption. There are many examles 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-lasting 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.

●▲ Wireless Sensor Network for Steelmaking Optimization

The optimization of steelmaking processes is based on multiple sensors placed around the furnace to monitor several parameters such as gas concentrations, temperature and vibrations. These sensors have to survive in the very harsh environment of the furnace in order to provide reliable measurements, which are used for safety, quality, and environment and process control. Currently the sensors are bulky and connected by cables, leading to frequent mechanical damage due to the hot particles and equipment moving near the furnace. Furthermore, the vast numbers of sensors, in many cases redundant due to the problems explained above, need to be periodically maintained to provide reliable safety and process information. The proposed project has two goals. The first one to investigate the possibility of replacing the cables with a wireless bridge based on an ultra-low-power transceiver. The transceiver will be implemented using CMOS technology, making it compact, low cost and with a long battery life. The second goal is to demonstrate the possibility of combining the transceiver with the sensors required to monitor the steelmaking process, creating a fully integrated wireless sensor node. With a large number of these low-cost wireless sensors, it will be possible to build a wireless sensor network around the furnace, magnifying the sensing capabilities and improving reliability due to the absence of cables.

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 theory, 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 could 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 into 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. The work 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 broadband 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 of imaging 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.HTM

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.

Energy Optimization Algorithms in Railway Networks

Concerns over carbon emissions, climate change and sustainability are motivating global efforts to reduce energy consumption in transportation systems. Reducing energy consumption in railway networks is an important component and goal of such efforts. This requires advances in dynamic real-time optimization. Yet, real-world applications of mathematical optimization techniques are not widespread in the railway industry. This research project addresses this gap and is focused on real-time optimization of railway networks towards minimizing energy consumption. Specifically, the objective of this project is to formulate a framework and the associated mathematical algorithms for dynamically re-adjusting timetables in a multiple-train railway network such that the total electrical energy consumed is minimized and the utilization of produced regenerative energy is maximized.

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, 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.

Photonic Devices and Systems for New Applications

We are designing photonic devices and systems for new areas of applications, such as quantum information, neuroscience, surgical guidance, and NewSpace. We are extending the capabilities of photonic technology by developing new tools and techniques. This research is at an exploratory stage. Please contact the PI for more details.

Prodic, Aleksandar

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

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). We are currently extending this technique to vibration sensing. Another technology we developed is called the virtual-reference interferometry (VRI). It was developed by one of our graduate students, and has been commercialized by a successful start-up company. It is used to characterize the dispersion of optical fibres, waveguides, and devices.

Quantum Optics and 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 our group's fibre-optic technologies. These include fibre-based entangled photon pair sources and fibre-based quantum key distribution systems. Our entangled photon source based on a periodically-poled optical fibre is the world's simplest polarization-entangled photon pair source, and we are the first to use such a technology. It provides pure (>1000:1 signal-to-noise ratio), broadband (>100 nm), highly polarization entangled (>99% interference visibility) photon pairs. Such sources are being used for reconfigurable multi-user quantum key distribution systems. We will further explore the technology's potential in quantum sensing.

Rose, Jonathan

WWW.EECG.TORONTO.EDU/~JAYAR/SOFTWARE/GENIE/

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—we will do this 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 have demonstrated the use of a new tool, called GENIE, that can help designers make efficient interconnect at the fine-grain level (within an IP core) and at the coarse-grain level (between IP cores).

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

Creative Applications for Mobile Devices

Mobile smartphones have given rise to an explosion in 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-grows, 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 for mobile devices. Graduate students from all disciplines at the University of Toronto are invited to take the course for credit. 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: Perimetry

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. So far, tests on people with normal vision have yielded good results. 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 Smoking Cessation 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) has been helping people quit smoking for many years, and does research on the subject of how to do this best. Our collaboration with that clinic (Dr. Peter Selby and his research team) is seeking to build a mobile application that will help people quit smoking—by recording their habits, reminding them of the reasons to quit, and perhaps alerting them to imminent trigger situations.

Sargent, Edward

WWW.LIGHT.UTORONTO.CA

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. We are also exploring new materials including perovskites, promising optoelectronic materials that exhibit impressive photovoltaic performance. Finally, we are studying systems of quantum dots embedded into perovskites, and these exhibit remarkable optoelectronic properties traceable to their atom-scale crystalline alignment.

On-Chip Gene and Protein Analysis of Cells and Bacteria

We create integrated circuits for the detection of panels of biomarkers (nucleic acids, proteins, and small molecules) that indicate the early onset of specific types of disease. We have also recently developed the means to capture rare cells, such as circulating tumour cells, on an integrated circuit, enabling subsequent biomolecular analysis. 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.

Quantum Dots for Light Emission Applications

We develop new ultra-efficient light emitters that reduce energy consumption and achieve improved colour purity relative to existing display and projection light sources. We use colloidal quantum dots engineered to provide ultra-narrow light emission in the green and red under blue excitation.

Solar fuels

Our investigations focus on converting solar energy into stored chemical energy, mimicking the natural photosynthesis process. We develop highly active catalysts for the splitting of water to hydrogen and for the reduction of CO_2 to carbon-based products (CO, methane, methanol) through electrochemical processes using electricity from solar photovoltaic cells. We also work to develop efficient photocatalysts for the direct production of solar fuels from water and CO_2 via a sunlight-driven photocatalysis process. The ultimate aims of our projects are to generate clean and renewable solar fuels and to close the carbon cycle.

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 the 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 purposes 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 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 with Flash. Our goal in this research is to devise circuit techniques to circumvent the device's shortcomings and ease the requirements of 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 both 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 act like a transmission line and as such exhibit frequency-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 cross talk 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 10-Gb/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 28 Gb/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 research 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 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 better understand 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 developing 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 potential grid problems and suggest appropriate preventive controls.

Taylor, Joshua

WWW.ELE.UTORONTO.CA/~JTAYLOR

Gaming in Modern Electricity Markets

Ten years ago, California's power system was rocked when energy traders manipulated vulnerabilities in its 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, 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 issue is becoming more pronounced as we increase our reliance on intermittent renewable sources of energy like wind and solar. Demand-response programs incentivize users 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 its 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 because 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 the development of tractable approximations with rigorous performance guarantees.

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 are a major multidisciplinary challenge for 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 power burst of 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 SMPS 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. Maximum power point tracking (MPPT) is performed on a PV array 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.

Non-invasive Assessment of Aortic Coarctation via Computational Fluid Dynamics

Patients with aortic coarctation suffer from a severe narrowing of the aorta. Currently, diagnosis is done with catheterization, an invasive procedure with a 1-in-1,000 risk of patient death or serious injury. In collaboration with physicians at the Sunnybrook Health Sciences Centre and the Toronto General Hospital, we are developing a non-invasive alternative. The geometry of the patient's aorta is first reconstructed from medical images. Then, a computational fluid dynamic simulation is used to calculate the pressure gradient across the coarctation, and assess disease severity in a non-invasive way. Computational methods are expected to bring substantial innovation in medicine, improving patient care and reducing costs.

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 upon 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 circuits and antennas.

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 since the earliest stages of design. Through our models, designers can maximize product reliability and performance without resorting to costly prototyping. This research activity is of immediate interest to 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 molecularfactors 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 [Ca^2+]. 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 at the Canadian National Institute for the Blind (CNIB). The system has a tracking and navigation feature 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 (VANETs) 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 parts 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 that 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 premier conferences for its impact on the community.

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 100 Gb/s. 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 Gbaud. Optical modulators require 3–6V electronic signal swing for proper operation. The large voltage swing and the very broadband operation from DC to over 80 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 multi-chip solution. In addition, because of the relatively large resolution (7–8 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 novel large swing (>3V differential), multi-bit (6–8 bit) 55-nm BiCMOS and 28-nm SOI CMOS DAC and ADC topologies for 16-QAM, 64-QAM and OFDM optical transmitters operating at 1 Tb/s. Several record-breaking front-end building blocks of the ADC and DAC recently developed in our group and operating at 120 Gb/s with over 90 GHz bandwidth were presented at IEEE BCTM and IEEE CSICS in October 2015.

▲ 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 convertors 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. Initial experimental results have demonstrated one of the critical building blocks of the ADC, the fastest track-and-hold amplifier reported to date, with a clock frequency exceeding 90 GHz and 40-GHz input bandwidth.

▲ Atomic-Level and 2D Crystal Electronic Devices

This exploratory research focuses on the simulation, design and fabication of novel nanoscale metal nanowire transistors. The ultimate goal is the physical implementation of sub-7 nm gate length transistors, beyond the 2030 ITRS horizon.

● ▲ 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 GHz. SiGe BiCMOS and SOI CMOS technologies will be used and compared for lowest power operation. The main applications are in autonomous navigation of vehicles and drones, distance measurements, and touchless gesture control of small wearable devices and IoT connected devices.

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

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

▲ Advanced Design of Peripheral Nerve Interfaces

We are developing novel interfaces that are aimed at improving the clinical use of neurostimulation technologies, which include functional electrical stimulation and electrical neuromodulation. Projects include the design of (1) a minimally invasive method of electrically activating the peripheral nervous system (enhanced transcutaneous electrical nerve stimulation, eTENS), and (2) peripheral nerve electrodes that can achieve long-term recording of neural activity.

⬟ Electrical Neuromodulation for Bladder Dysfunction

The objective of this project is to develop therapeutic platforms for treating urinary dysfunction 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.

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 co-operative transmission techniques to enhance the capacity, coverage and reliability of wireless cellular networks. The goal is to develop novel coordinated signal processing, resource allocation and network optimization techniques for 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 co-operate in transmitting and receiving signals, and the heterogeneous architecture where remote radio units are deployed within the cellular structure to enhance coverage. This research 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 will now 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; and (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 premier 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 allowing 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.