

24 INTO THE DEEP

Will artificial intelligence transform medicine as we know it?

34 SECTION GRADUATE

OFF THE CHARTS

A portal for patient records

38 SECTION RESEARCH

APP-Y AND HEALTHY

The automation of medicine

48 SECTION ALUMNI

ITEMS OF NOTE

Alumni notes from around the world

2016

A YEAR IN REVIEW

ANNUM

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

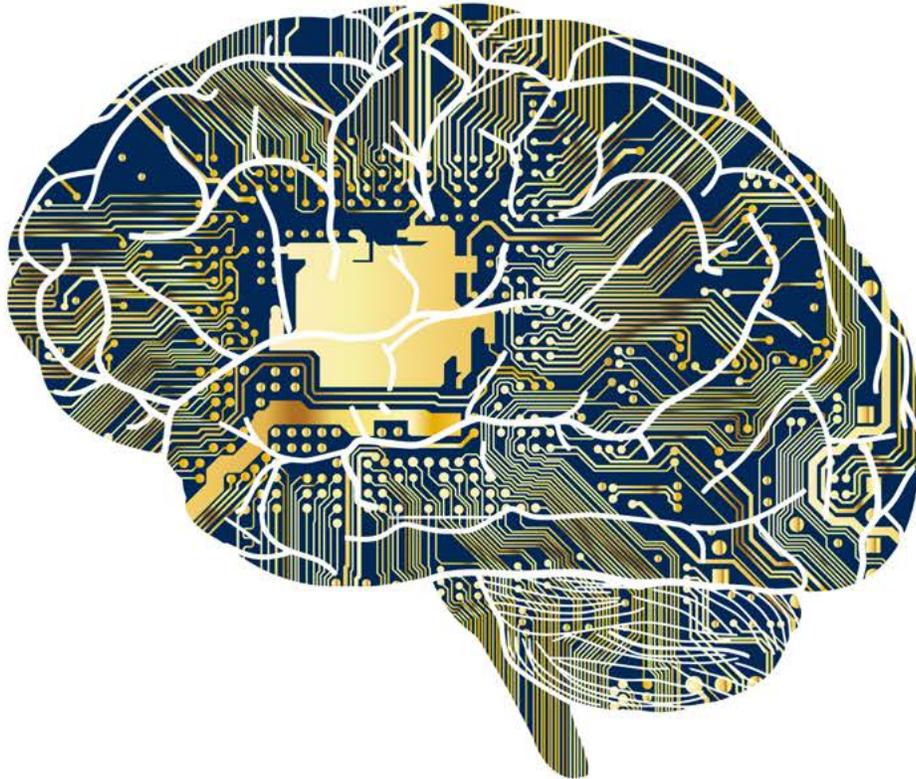
TAGCCCCTGACCATTGCAACGTAACGACTGACCCTTGACGATCGCGTAA
CCGTAGCCGTAGACCATTGCAACGTAACGGCAACGACGTACAGCTACCT
TAAAACCTGACCATTGCAACACGTACAGCTACCTAACCTGACCATTGCAA
CGTACCTGACCATAACGTACAGCTACCTAACCTTGCAACGCCATTGCA
GACCATTGCAAACCATTTGCAACGTAACGACTGACCCTTGACACGTAAACGT
CCGTAGCCGTAAACGTACAGCTACCTGACCATTGCAACGTAACGGCAACG
TAAAACCTGACCATTGCAACACGTACAGCTACCTAACCTGACCATTGCAA
CGTACCTGACCATTGCAACGCCATTGCAACAGCCATTGAACGCCATCAT
GCTACCTAACCTGACCATTGCAAACCATTTGCAAACGTAAGCAAACAGTACA
GACGTACCTGACCATTGTAACAGCTAGGTTAGCTTCGACGTACAGCTGAAA
TAGCCGATCGCGTAAACCTGACCATTGCAAACGTAACCGTTGACCGACTG
CCGTAGCCGTAAACGTACAGCTACCTGACCATTGCAACGTTGGCAAACAACG
TAAAACCTGACCATTGCAACACGTACAGCTACCTAACCTGACCATTGCAA
CGTACCTGACCATTGCAAACGCCATTGCAACACGTACAGCTACCTAACCT
GACCATTGCAAACCATTTGCAACGTAACGACTGACCCTTGACACGTAAACGT
CCGTAGCCGTAAACGGACCATTGCAACGTAACGGCAAACGTACAGCTACCT
TAAGACAGCTACCTAACCTGACCATTGCAACACTCACCTGACAA
CGTACCTGACCATTGCAAACGCCATTGCAAACAGCTTGCCAAAACGCCAT
GCAAACAGTACAGCTACCTAACCTGACCATTGCAAACCATTTGCAAACGTAA
GACGTACCTGACCATTGTAACAGCTAGGTTAGCTTCGACGTACAGCTGAAA
TAGCCGATCGCGTAAACCTGACCATTGCAAACGTAACGACTGACCCTTGAC
CCGTAGCCGTAAACGTACAGCTACCTGACCATTGCAAACGTAACGGCAAACG
TAAAACCTGACCATTGCAACACGTACAGCTACCTAACCTGACCATTGCAA
CGTACCTGACCATTGCAAACGCCATTGCAACACGTACAGCTACCTAACCT
GACCATTGCAAACCATTTGCAACGTAACGACTGACCCTTGACACGTAAACGT
CCGTAGCCGTAAACGTACAGCTACCTGACCATTGCGTTAACGGCAAACG
TAAAACCTGTACAGCTACCTAACCTGACCATTGCAAACCATTTGCAACACG
CGTACCTGACCATTGCAAACGCCATTGCAAACAGCCATTGCAAACGCCAT
GCAAACAGTACAGCTACCTAACCTGACCATTGCAAACCATTTGCAACGTAA
GACGTACCTGACCATTGTAACAGCTAGGTTAGCTTCGACGTACAGCTGAAA
GCTACCTAACCTGGACCATTGCAAACCATTTGCAAACCATTTGCAAACAGTACA
CCGTAGCCGTAAACGTACAGCTACCTGACCATTGCAAACGTAACGGCAAACG
CCTGACCATTGCAAACACGTACAGCTACCTAACCTGACCATTGCAAATAAAA
AACGCCATTGCAAACAGCCATTGCAAACAGCCAT
CCATTGCAAACGTAACACAGTACAGCTACCTAAC
TACAGCTAGGCGTACAGCTGAAATTAGCTTCGA



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



IT'S JUST WHAT THE ENGINEER ORDERED.



Neural networks, closed looped control, the transmission of information through electrical signals: it turns out neuroscience and electrical and computer engineering have more in common than you might think.

An emerging field uniting medicine and engineering is neuromodulation, whereby implanted electrical devices stimulate areas of the nervous system to relieve symptoms and treat complex neurological disorders. Electrical and computer engineers are well-positioned to make significant contributions to the field of medicine and to neuromodulation in particular. Right now, researchers in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering are using big data, integrated circuits and electrode arrays to help find solutions to complex medical problems.

We're working on creating a Centre for Neuromodulation to develop, test and implement systems to be used to control and modulate brain, spinal cord and peripheral nerve function.

To learn how you can support this effort, contact Jennifer Lancaster at jennifer@ecf.utoronto.ca or 416 978 7210

BOUNDLESS

The future is here

Artificial intelligence is very much a reality — whether we are cognizant of its omnipresence or not. Machines are presenting us with tailored advertisements while we browse the internet, analyzing our financial portfolios before suggesting stocks to invest in and adjusting the temperature on our thermostat when we enter our living rooms. But behind every machine that “thinks for itself” is a very real person designing hardware, writing algorithms and developing software.

It’s no secret that artificial intelligence systems are being used to engage customers and understand their behaviour — look no further than your virtual assistant or that customer service chat box on your favourite website. But how will artificial intelligence alter other aspects of our lives? The cover story of this issue, Into the Deep on page 24, explores how one researcher in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) is using artificial intelligence techniques to unlock the secrets of the human body.

Predicting the future is no easy task. Predicting the future of technology is nearly impossible. Who would have guessed even 15 years ago that 140 characters of text could impact an election or that the taxi industry would be transformed by a company that does not own a single cab. We are not here to predict the future, but in ECE the future is here: our undergraduate students are some of the brightest in the country and they are developing our technological future as



PHOTO BY RAINA WILSON

you read these pages. Meet three of them who are helping us to recruit the next cohort of stellar students on page 12.

Our relationships with our industry partners keep us informed of industry trends, connect us to real-world applications for our research and help us to prepare the next generation of electrical and computer engineers to meet the demands of the future. One such example of our department’s industrial collaboration is our ongoing partnership with Huawei. Various research groups within the department are working with Huawei on an array of topics: from 5G cellular networks to silicon photonics. Learn more about ongoing research collaborations with Huawei on page eight.

This magazine offers you a panoramic view of the students, faculty, alumni and facilities that collectively help our department maintain its position as the top-ranked ECE program in Canada and among the best in the world. I hope you will reconnect with the department and share your views on the future of ECE, whatever the future may hold — you may reach me directly at chair@ece.utoronto.ca.

FARID N. NAJM,
PROFESSOR & CHAIR

EDITOR

Jessica MacInnis

ART DIRECTOR

Katina Constantinou,
Sugar Design

PHOTOGRAPHY & ILLUSTRATION

Raina+Wilson
Kerry Shaw
Matt Chase

COPY EDITORS

Kelly Hayward
Avivah Wargon

EDITORIAL INQUIRIES

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

Visit us online:

www.ece.utoronto.ca

Join the conversation:



ANNUM is published annually by the communications office of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering. All material is copyright © 2016, and may not be reprinted without the express written permission of the author.

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 Reproductions.



8



Information Huawei

ECE professors are collaborating with Huawei on advanced communications technology research from FPGAs to 3D optical circuits

12

Special Envoys

Current undergraduate students act as ambassadors to prospective students and alumni



16

From A to B to CD4

James Dou is helping people with HIV manage and monitor their disease in remote areas around the world



20

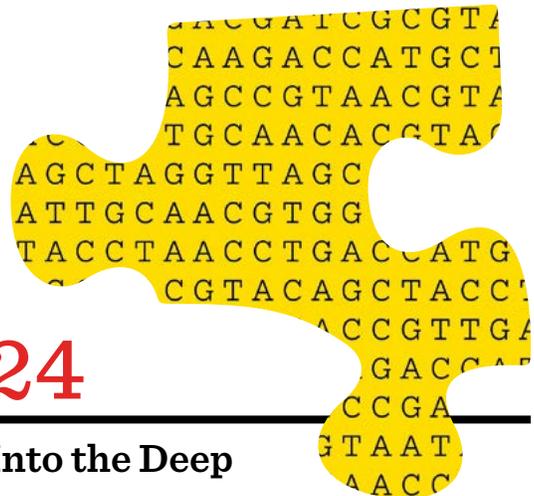
Employee Number 32

Rami Rahim's career has taken him from a small startup to a publicly traded multinational business but kept him at the same company for nearly 20 years

24

Into the Deep

Can artificial intelligence techniques teach us about the human body?



SECTIONS

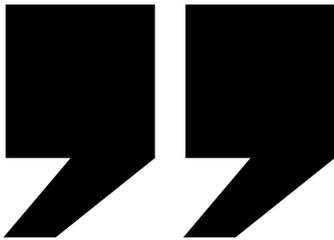


29

UNDERGRADUATE

Empowering Students in ECE's Newly Renovated Energy Systems Lab

New lab equipment provides students and researchers with specs they'd find in the real world



We are effectively making the world a smaller place; we're enabling people to connect and do business and keep informed.

p. 22



34

GRADUATE

Charting New Territory: From Masters To MedChart

James Bateman wants to give patients access to their medical records with the click of a button

38

RESEARCH

App-y and Healthy

Accessing healthcare from the palm of your hand



41

INDUSTRY

From Internship to Industry Partnership

How Ali Sheikholeslami maintains a successful collaboration with Fujitsu Laboratories

45

ALUMNI

Welcome Somen Mondal

Meet ECE's newest member of its Alumni Board of Advisors



OUR STORY

Founded in 1909, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) maintains a proud history of world-leading research and innovative education. We are the top-ranked ECE department in Canada and home to 100 professors, 649 graduate students and 1,513 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 28 Fellows of the IEEE, nine Fellows of the Royal Society of Canada, 13 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 47 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 2016

#1

in
Canada

Electrical & Electronic Engineering

#11

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

#11

in the World

Computer Science & Information Systems

1,513 Undergraduate students

649 Graduate students

100 Professors, including Emeritus

71 Post-docs

51 Admin & tech staff

15 Research associates

8 Visiting professors

161

Inventions
out of
ECE since
2011

47

Patent
applications
out of ECE
since 2011

11

Startups
out of
ECE since
2011

88

Percentage
of ECE
invention
disclosures
that include
a student
inventor

**TIMES HIGHER EDUCATION
WORLD UNIVERSITY RANKINGS 2016–2017**

*Engineering & Technology
Universities*



**ACADEMIC RANKING OF
WORLD UNIVERSITIES 2016**

*Engineering/Technology &
Computer Sciences*



23

Chair titles held by
ECE faculty in
2015–2016, includ-
ing Canada Research
Chairs, Endowed
Research Chairs,
Industrial Research
Chairs and U of T
Distinguished
Professor Chairs

\$55,125

Average salary
earned by ECE
students on PEY
placements
in 2016-2017

INFORMATION

HUAWEI

Huawei is one of the world's largest information and communications technology companies; their products have been deployed in over 170 countries, serving more than a third of the world's population. The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) is one of Canada's largest departments of its kind, with professors working in emerging fields across a number of research areas. So when Huawei's Canada Research Centre in Ottawa sought to focus on the wireline, wireless, optical and IP networking sectors they looked west to the University of Toronto to partner with some of Canada's top researchers in advanced computing and communications technology.



As one of the largest purchasers of Field-Programmable Gate Arrays (FPGAs) in the world, it was unsurprising that Huawei would want to collaborate with professors in ECE — the department is home to some of the top researchers in the field. FPGAs are integrated circuits that can be programmed, and reprogrammed, after manufacturing to implement any digital circuit: they provide flexibility and agility to the end user.

Professors Jason Anderson, Vaughn Betz and Paul Chow are working with Huawei to accelerate computations, improve energy efficiency and reduce cost by leveraging FPGAs. “We envision that in the future, it will be commonplace for FPGAs to stand alongside regular processors in datacentres, where the FPGAs realize hardware accelerators specially tailored to application needs,” says Anderson.

A key aspect of the research relates to the nascent concept of FPGA Virtualization. “While the hardware design to connect an FPGA to a processor is not very difficult, the real challenge is building out all of the technology that would allow us to run a variety of applications on the

FPGA, and switch context between the different applications,” explains Anderson. “This is akin to how a regular processor is able to load and execute different software programs.” The FPGA Virtualization challenge is to develop the hardware and software infrastructure that is needed to make this happen.

Whether you are on Twitter in your living room, online shopping on your lunch break or using your cell phone on a road trip, you’re triggering a data request, causing data to travel around the world, making stops at datacentres in different places that run different jobs at the same time. Professor Baochun Li is working on making this data move faster — and smarter. Li is working with Huawei to address performance issues experienced by big data applications deployed in Wide Area Networks, where large amounts of data are processed across multiple physical machines, in different geographical areas.

A typical job involves shipping about a terabyte of data — that’s about a hard drive’s worth of data. It costs about ten cents per gigabyte to ship data across the datacentres at a rate of somewhere between 50-100 megabytes per second: it’s not only costly but time consuming to ship data from datacentre to datacentre. “So that is something we are trying to solve: how do we run this huge job that is cutting across multiple datacentres in the most efficient way,” says Li. “We are trying to assign these tasks to different datacentres so

We are trying to assign these tasks to different datacentres so we know which datacentre is running which task. If you assign it in the best way possible, it's going to reduce the time for the job to be completed."

we know which datacentre is running which task. If you assign it in the best way possible, it's going to reduce the time for the job to be completed."

It might be an understatement to say that silicon has played an important role in the history of technology innovation, but if you asked Professor Peter Herman what the next chapter in silicon's contributions might be he'd be inclined to say silicon photonics.

An emerging trend, silicon photonics is a technology in which data moves between and within computer chips by light rather than electrical conductors. "For decades there has been this idea to build optical circuits in silicon," says Herman. "This is now commercially viable as a technology base, but the waveguides here are incredibly small... much smaller than optical fibre waveguides and that is the key technology challenge: how do you take waveguides in glass and connect the big glass cables with the silicon on-chip devices."

In this project with Huawei, Herman is working on the age-old problem with a high-tech twist: fitting a square peg into a round hole, or more accurately, fitting glass fibres with silicon chips. "Simply put, the fibre doesn't fit with the silicon chip," says Herman. "So we are working on using lasers to modify the glass to leave these

waveguides behind. What we intend to do is take new forms of glass fibres, each having multiple waveguide cores, and precisely weld or bind the fibres to blocks of glass. A focussed laser would then write a 3D pattern of optical waveguide circuits in a reverse fan-out that will bring fibre light into a tight pattern to the opposite end of the glass where we can match and efficiently connect with the array of waveguide channels on the flat silicon chip." This 3D fibre-to-chip configuration will enable a much higher density of silicon chips to connect and communicate efficiently, essential to handle the explosion of information in our growing datacentres.

For Huawei, these three projects are just the tip of the iceberg: in April of 2016 it was announced that Huawei will invest \$3 million in research projects with the University of Toronto over the next three years. "Partnering with ECE allows us to leverage the talent and network of researchers in the department, and helps us to drive innovation in new and exciting directions," says Scott Bradley, VP of Corporate Affairs at Huawei Canada. "Whether it's in photonics, hardware, cloud computing or many of the other research areas we partner with ECE in, we're excited to see technologies developed here in Toronto used around the world." While 170 countries and a third of the world's population seems like a lot, it turns out that it's a small world after all. ■



FROM L: ROBIN TODD,
ANDREW MAKSYMOWSKY
AND SARINA SIT

SPECIAL ENVOYS

An ECE ambassadorship is a volunteer position held by undergraduate students in the department. They interact with prospective students and their parents, and alumni to share their undergraduate experience, both inside and outside of the classroom. While their backgrounds, areas of study and outside interests may vary, they all share a common sentiment: a passion for electrical & computer engineering

Both his father and grandfather are alumni of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE), but for Andrew Maksymowsky (Year 3 CompE) it was a circuits kit given to him as a seven-year-old that eventually inspired him to apply to ECE. And though the familial ties to ECE are strong and his passion for electrical and computer engineering started when most kids were just mastering addition and subtraction, Maksymowsky's talents are not just confined to the lab. An accomplished violinist, in his first year at U of T, he jumped at the opportunity to put a quartet together to play a gig at the Royal Ontario Museum in Toronto. It was an exciting opportunity to play inside of one of Canada's iconic cultural venues, and he didn't let on to his band that they would be playing in front of royalty: Prince Edward was in attendance that night. When His Royal Highness asked if the quartet was from the Royal Conservatory of

Music the four shook their heads and said they were from the University of Toronto — and not from the Faculty of Music, but from the Faculty of Applied Science & Engineering.

Maksymowsky is one of about 30 ECE Ambassadors — current undergraduate students in ECE who represent the department at recruitment, outreach and alumni events throughout the year. A volunteer position, ECE Ambassadors interact with current students, prospective students and their parents, and alumni. The Ambassadors all share a passion for electrical and computer engineering, and for the department, but they also reflect the variety of students enrolled in ECE.

Maksymowsky's involvement with the U of T solar car team started with a message on his phone. Maksymowsky received a text message from Frank Gu (Year 3 CompE), the managing director of the Blue Sky Solar Racing team, asking him if he wanted to work on a project with the car.

I'm working with talented people who not only know so much but are happy to share their knowledge.

The team's aerodynamic simulations were taking a lot of time, slowing down the team's progress on the design: "Imagine waiting for an hour to open an email: that's what it was like," says Maksymowsky. "I suggested throwing more computing power at the problem, and found a way to distribute the task onto a cluster of 16 machines." The additional computing power has allowed the team to run more simulations and experiment with the design in hopes of achieving the most aerodynamic vehicle shape possible.

Between band practice, course work, solar car commitments and archery, Maksymowsky is also passionate about sharing his ECE experience with prospective students. "I applied to be an ECE Ambassador because I had a really positive experience with the department when I was deciding what university and program to attend," says Maksymowsky. "At the time, it was great to hear from professors but what really helped me to decide was hearing the experience of current students so that is something I wanted to do for others."

The ECE Ambassadors field questions from both prospective students and their parents at campus open houses, visits from high schools and recruitment events. While high school students often have questions about workload, courses and exams, parents are often interested in job prospects for alumni and internship programs.

For Sarina Sit (Year 2 ECE), applying to UofT Engineering was not a straightforward decision. But after a week at the Girls Leadership in Engineering Experience (GLEE) at UofT at the end of grade 12 she knew that engineering was not only on the table, but at the top of the list. "Throughout all of high school I had enrolled in courses that would allow me to apply for life sciences," says Sit. "But I had such an amazing experience at GLEE that I knew I had to give engineering another look." Sit was accepted to UofT Engineering's General First

Year program and hasn't looked back. "As I went through the courses, I found that the computer and electrical courses were my favourite so I pursued those," says Sit. "Right now, I'm leaning towards electrical, but my courses are still a mix."

Sit is involved in the IEEE student branch in their hardware chapter and will next move into their power chapter. As director of events for the student branch, she has helped to organize networking events and hardware hackathons. She has also volunteered with the Faculty of Applied Science & Engineering's Outreach Office where she helps to promote science, technology, engineering and math education to pre-university students.

In the short term, Sit is going through the process of applying to Professional Experience Year (PEY) internship placements and in her final year in ECE, she hopes to find a great team for her year-long capstone project. Long term, she's ready to roll up her sleeves: "I'm keeping an open mind to all of the opportunities out there," says Sit. "I hope to find a position that allows me to follow my interests and passions, and helps me to continue to develop both my technical and project management skills."

The PEY paid internship program is an opportunity for upper-year undergraduate students to work full-time in an engineering position for 12-16 months, and it allows students to make meaningful contributions and gain career momentum before they graduate. Robin Todd's (Year 3 CompE) PEY placement combines two of her big passions: software development and healthcare. She is working on the Drive team at Synaptive Medical, writing code to help develop a neurosurgical robot arm. Though she was always interested in medicine, her first experience

with software or coding occurred just a few weeks before her first year at U of T began. “I had never had any exposure to software before but I took a week-long course on programming the summer before first year,” says Todd. “In high school I had taken music courses instead of the software and coding courses, but now at my PEY placement, I’m working fulltime as a software developer.”

At ECE events for prospective students, Todd welcomes questions about the PEY program: “At my PEY placement, I learn something new every day,” says Todd. “There’s a lot of problem-solving, and I’m working with talented people who not only know so much but are happy to share their knowledge.” Todd

is busy outside of the classroom and the workplace: she has played intramural soccer on the ECE team, was on the inner tube water polo team and is also an avid runner, completing her first marathon in the fall of 2016. Her final year as an undergraduate student is just around the corner, and she’s already thinking about what’s next: “I’ve always considered applying to medical school, but I’m really loving my PEY placement, so I’m considering staying in software engineering after I graduate,” says Todd. “The important thing for me right now is to keep my options open.”

Todd might be the only runner in the group, but they are all chasing big dreams. **A**



I APPLIED TO BE AN ECE AMBASSADOR BECAUSE I HAD A REALLY POSITIVE EXPERIENCE WITH THE DEPARTMENT WHEN I WAS DECIDING WHAT UNIVERSITY AND PROGRAM TO ATTEND.

**ANDREW
MAKSYMOWSKY**



FROM A TO B TO CD4

PhD student James Dou's passion for photonics has taken him from inside ECE's labs to helping people living with HIV manage their disease in remote areas around the world — with a stop in the boardroom along the way

You'd be forgiven for not knowing how large a flow cytometer is, but take PhD candidate James Dou's word for it when he says it's too big to carry in a backpack.

Ubiquitous in hospitals and research institutions, flow cytometers are used to analyze particles — cells are suspended in fluid and passed through a laser to detect the presence of certain biomarkers. Used in various forms since the early 1970s, flow cytometry is routinely utilized in the diagnosis and monitoring of various health disorders; blood cancers are often diagnosed with the help of flow cytometers and they are an important tool in tracking the health of patients diagnosed with HIV. As for their size: flow cytometers are about as big as that deep freezer in your basement.

When Dou arrived at the University of Toronto as a MASc student under the co-supervision of professors Stewart Aitchison and Peter Herman, he worked on developing novel tools for cancer research — using lasers to fabricate lab-on-a-chip biophotonic devices for single cell analysis. As he embarked on his PhD under the supervision of Professor Aitchison he had already begun to determine the best application and commercialization potential for a microfluidic lab-on-a-chip device. Dou worked with U of T's Innovations & Partnerships Office to look into a number of potential applications for the lab-on-a-chip device from a market perspective, and the possibilities were seemingly endless. “We looked at blood testing, cancer diagnosis, agricultural analysis, food safety

and many others,” says Dou. “It was a good exercise to go through, but it was also kind of painful — you almost felt like it was a platform that could be used to do too many things and it was difficult to pinpoint exactly which one to focus on.”

Dou entered the proverbial spin cycle: the technological push of his original research led him to pursue

commercialization but the market pull led him away from cancer and single cell analysis and back to the lab to find an application that was both technologically sound and had market viability. After connecting with people involved in HIV testing and diagnosis, it was clear that there was not only a market need for portable CD4 testing, but a global health need as well.

CD4 cells are a type of white blood cell that are integral in protecting the body from infection: the higher your body’s CD4 count is, the better you can fight infections. In people with HIV, CD4 cells are attacked by the virus and often destroyed. This impedes the body’s response when it detects viruses and bacteria. CD4 counts are used to help measure a patient’s disease progression and can help to determine the efficacy of the patient’s medication regimen.

“A patient’s CD4 count indicates how well their immune system is working. CD4 levels are a direct measure of the immune system; it’s an indirect measure of how severe the HIV virus is,” says Dou. “The World Health Organization recommends patients with HIV have their CD4 levels counted every three to six months in initial stages of the treatment and every six to 12 months once patient status has stabilized — these tests are very useful for HIV monitoring.”

Flow cytometry has been the standard method of CD4 testing and patients with HIV in most areas of North America and Europe have relatively easy access to flow cytometers and the technicians who operate them. But for those in rural and remote areas of the world, CD4 testing is no easy undertaking. First, there must be access to a health care practitioner who is trained to draw blood. The blood must then be transported to the closest flow cytometer — sometimes hundreds of kilometres away. An expert must then analyze the results while patients wait days or even weeks for the results. For those who have come long distances to have their blood drawn, it may be difficult to return, days or weeks later, to learn the results.

When we first started ChipCare, less than twenty-five percent of people in Africa who needed access to this testing actually had access to the testing — there is a really big gap in meeting this need.



These technical challenges and commercial opportunities for portable CD4 testing led Dou, Professor Aitchison, and Rakesh Nayyar, an expert in flow cytometry, to found a company called ChipCare to develop a portable cytometer called a polyvalent analyzer: a handheld device about the size of a tissue box. The device was designed not only to be portable, but to withstand environmental challenges found outside of a laboratory setting, including temperature fluctuation, dust and limited or intermittent access to electricity.

As for the process the patient undergoes, it is similar to a home blood glucose test: the patient's finger is pricked and a drop of blood is placed on a transfer pipette and then onto the custom cartridge designed by Dou and the team at ChipCare. A mechanical pump controls the sample flow within the cartridge, and a reagent pellet, integrated into the cartridge during manufacturing, mixes with cells within the blood: these labelled biomarkers cling to the CD4 cells. The cartridge is then inserted into the polyvalent analyzer and light is shone at the sample, the fluorescence of the cells is measured and CD4 cells are counted. The results are available on the digital display of the polyvalent analyzer within minutes.

While Dou expected the technical challenges of developing a point-of-care device for testing CD4 levels he also faced challenges related to working on a product designed for remote health care from a big city: Toronto. "It is definitely a challenge," says Dou. "Whether it's about business matters, market development,

financing — even the technical development and engineering: many people you talk to don't appreciate the challenges of health care in a remote setting." Despite these challenges, ChipCare has had early successes: in 2012 Dou and Aitchison were named two of U of T's Inventors of the Year, and in 2013 ChipCare received a \$2.05 million investment from a collaborative group of angel investors, the Government of Canada and Grand Challenges Canada. By 2015 the company had closed \$5.045 million in Series A financing. Dou expects the device will receive regulatory approval in early 2017, and will be available in the field soon after. "The initial launch of the device and technology will be primarily in Africa because there is a large HIV-positive population and right now there is insufficient access to CD4 testing there to meet the demand," says Dou. "When we first started ChipCare, less than twenty-five percent of people in Africa who needed access to this testing actually had access to the testing — there is a really big gap in meeting this need."

For Dou, his professional goals are lofty for both the short and long term: "There are about 36 million people living with HIV around the world—in the short term I hope we can help them manage and monitor their disease. As for the future, I'd like to explore other applications of this device, whether it is for rural and remote health care diagnostics and monitoring, or otherwise," says Dou. "That, and I look forward to successfully defending my thesis." ■

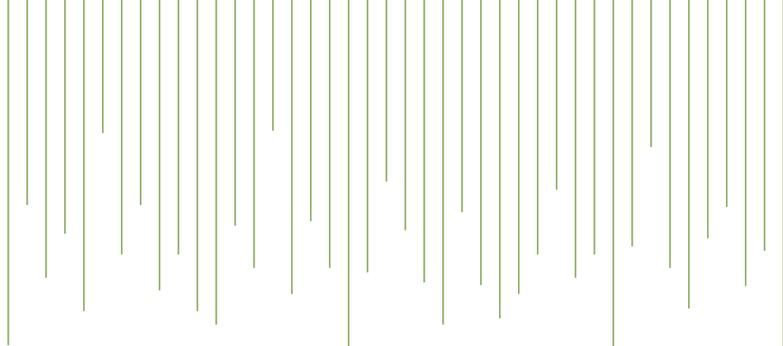
In 2013 ChipCare received a \$2.05 million investment from a collaborative group of angel investors, the Government of Canada and Grand Challenges Canada. By 2015 the company had closed \$5.045 million in Series A financing.

EMPLOYEE NUMBER 32

It's what start-up fairy tales are made of: an unproven company makes its 32nd hire, a junior engineer, to work on the very first product the company will bring to market. That product then revolutionizes the core router market and propels the company from small start-up to a publicly-traded, multinational company designing, manufacturing, configuring and selling routers, networks, switches and software. And as for employee number 32? He's now the CEO.

BY JESSICA MACINNIS
PHOTOGRAPHY BY KERRY SHAW





WE ARE EFFECTIVELY MAKING THE WORLD A SMALLER PLACE; WE'RE ENABLING PEOPLE TO CONNECT AND DO BUSINESS AND KEEP INFORMED.

Rami Rahim has built a career on chasing what he calls those 'ah-ha' moments — the euphoric feeling of figuring out a particularly challenging problem. And while pursuing that feeling has kept him busy, it's also kept him at the same company for nearly twenty years.

The alumnus of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering is the CEO of Juniper Networks, but he was first hired to work on the company's initial product, the M40 Core Router. The timing for both Rahim and for Juniper couldn't have been better: "We introduced the M40 in 1998, when the internet was on the cusp of explosive growth," says Rahim. "So telecommunications companies had no choice but to bet on a small unproven start-up because there was nothing else on the market that could keep up with the scale that they were seeing in their networks."

The success of the M40 put Juniper Networks on the map, but the focus then turned to proving that the company was not a one-product wonder. "We needed to prove that we were a company that could continue to innovate and continue to deliver products that satisfy the scale and performance of the large, global networks around the world," says Rahim. "Essentially, we were able to do that and that's why we are where we are as a company today." The drive for innovation was important to Juniper Networks' growth, but it is also a key reason why Rahim has stayed at the company for so long. Rahim sees the company's mission as more than just producing routers and switches. "We are effectively making the world a smaller place; we're enabling people to connect

and do business and keep informed," says Rahim. "We're helping to deliver online education; we're bringing clean tech energy to third-world countries and assisting researchers in solving very important problems."

Solving those challenges would certainly be an 'ah-ha' moment on a grand scale, but Rahim cut his teeth on solving problems in ECE. Rahim graduated from electrical engineering in 1994. "I remember learning about concepts in dynamics and quantum mechanics that were truly very difficult to grasp," says Rahim. "Those 'ah-ha' moments when you would find yourself sitting in the library for hours, and then have a breakthrough in understanding, were some of the best moments for me at U of T." His fourth-year thesis was on optoelectronics — a topic that interested him, but at the time he didn't really have an appreciation for its practical uses in industry. "It's almost fateful that today I am in communications. I'm the CEO of a large networking company and in this industry optoelectronics happens to be one of the most important areas of technological innovation that will help us to keep up with the demands on networks around the world."

Though his days of developing products, writing code, and solving on-the-ground technical issues are mostly behind him, Rahim keeps his technical skills sharp, partly out of a lifelong interest, but mostly because he feels that there is a significant advantage to executives having a technical background. "In a high-tech company in particular, having the technical foundation gives you the ability to ask the right questions to make the decisions that involve both a business and technical element," says Rahim. "It helps

you detect when there are issues in execution, or in decision making, that would otherwise be difficult without the foundation.”

Rahim sees some trends in engineering tied very closely to emerging trends at Juniper Networks. Machine learning, cyber security and big data are some of the areas Rahim predicts growth in, both in the field of electrical and computer engineering generally, and at Juniper Networks.

I think self-driving cars will be a reality in the not-so-distant future, says Rahim. I think a similar concept can be applied to networking – making networks run themselves as opposed to requiring a lot of manual, error-prone work by humans. At Juniper, we affectionately call it ‘self-driving networks’ – this is where I see machine learning being applied to networking in a lot of interesting ways.

While there is no shortage of challenges in the networking industry today, Rahim is keeping his eye on a challenge further down the pipeline: kids who can code. “We need more technical talent in both Canada and the United States to fuel inno-

vation, and to keep these countries competitive in a global environment,” says Rahim. “It needs to start at a young age by demonstrating to our children how important technology is and by building a solid technical foundation at an early age.”

Rahim was the guest speaker at Convocation in the spring of 2016 and shared some of the advice that has helped him to build his career. To ECE’s graduating class, Rahim offered the following: go beyond your comfort zone, work with people who inspire you, finish what you start, be kind and work hard. While uncertainty might be disconcerting to new graduates, he discouraged the graduating class from getting too caught up in setting out their five- and ten-year plans and reminded them: “You are engineers – nothing is impossible.”

In Rahim’s view, the next generation of engineering graduates will face limitless challenges. For example, there is no end in sight to the amount of network capacity that humans will consume, and it’s no longer just about connecting people, but things: “By 2020, more than 50 billion devices around the world will be connected,” says Rahim. As with any challenge Rahim has encountered over the course of his career, the more complex, the better: “Engineers are very motivated by solving difficult problems,” says Rahim. “If problems are too easy, they become less interesting.” Rahim’s professional journey has been tied to solving technical and business challenges at Juniper Networks for almost two decades. As for why he’s stayed: “Boredom is a problem I’ve never had here.” ■

TAGCCCCTGACCATTTGCAACGTAACGACTGACCGTTGACGATCGCGTAACCGTA
CCATTGCAACACGTACAGCTACCTAACCTGACCATGCAAGACCATGCTAGCCAT
CGTACCTGACCATAACAGTACAGCTACCTGACCCGTAGCCGTAACGTACAGCTA
AACAGCCATTGAACGCCATCATTAGTACCTGACCATTGCAACACGTACAGCTAC
GCAACACGTACAGACGTACCTGACCATTGTACAGCTAGGTTAGCTTCGACGTACA
TGCCGTAGCCGTAACGTACAGCTACCTGACCATTGCAACGTGGCAACAACGTAG
ACCATTGCAACGCCATTGCAACACGTACAGCTACCTAACCTGACCATTGCAGACC
CCGTAGCCGTAACGGACCATTGCAACGTAAACGGCAACGTACAGCTACCTACACG
CTAGCCGATCGCGTAACCTGACCATTGCAACGTAAACGACTGACCAGTTGACTGAT
TAACGTACCTGACCATTCAACACGTACAGCTACCGACCATTGCAGACCATTGCAA
GCAACGCGTACCTGACCATTGCAACGCCATTGCAACAGCTTGCCGATCACGCCA
CACGTACAGCTACCTAACCTGACCATTGCAGACCATTGCAACGTAAATAAGACAGC
TTAGCTTCGACGTACAGCTGAGAGCAACACGTACAGCTACCTAACCTGACCATTG
TGCGATCCGCCATCCGTAGCCGTAACGTACAGCTACCTGACCATTGCAACGTAA
CATTAAAGTCCTGACCATTGCAACCATTTGCAACACGGCTACCTAACCTGGACCAT
TGACCATTGCAACGTAAACGGCAACGCCTGACCATTGCAACACGTACAGCTACCT
TGAGATTAGCTTCGAGCACTGACCATTGCAGACCATTGCAACGTAGACACGTACA
CATTAAAGTCCTGACCATTGCAACATGCAAGCTAACGTACCTGACCATTCAACACG
ACGCAGCTACCTAACCTGACCATTGCAAGACCATTGCAGACCATTGCAACGTAAACG
CGTACCTGACCATAACAGTACAGCTACCTGCTACCTGACCATTGCAACGTAAACG
AACAGCCATTGAACGCCATCTACCTAACCTGACCATTGCAGACCATTGCAACGTAA
GCAACACGTACAGACGTACCTGACCATTGTACAGCTAGCCGATCGCGTAACCTGA
TGCCGTAGCCGTAACGTACAGCTACCTGACCATTGCAACGTGGCAACAACGTGC
ACCATTGCAACGCCATTGCAACACGTACAGCTACCTAACCTGACCATTGCAGACC
CCGTAGCCGTAACGGACCATTGCAACGTAAACGGCAACGTACAGCTACCTACACG
CTAGCCGATCGCGTAACCTGACCATTGCAACGTAAACGACTGAOCGTTGACTATG
ACCGACCATTGCAGACCATTGCAACGTAAACGACTACGTAAACGTCCGTAGCCGTA
CCATTGCAACCGTAGACACGTACAGCTACCTAACATTGCAACGTAAACGACTGACC
GCAACAGCTTGCCGATGACGCCATGACCGTACCTGACCATTGTACAGCTAGGTTAG
CACGTACAGCTACCTAACCTGACCATTGCAGACCATTGCAACGTAAATAAGACAGC
TTAGCTTCGACGTACAGCTGAGAGCAACACGTACAGCTACCTAACCTGACCATTG
TGCGATGCGCCATCCGTAGCCGTAACGTACAGCTACCTGACCATTGCAACGTAA
CCATTGCAACACGTACAGCTACCTAACCTGACCATTGCAAGACCATTGCAGACCAT
CGTACCTGACCATAACAGTACAGCTACCTGACCCGTAGCCGTAACGTACAGCTA
AACAGCCATTGAACGCCATCATTAAAGTCCTGACCATTGCAACACGTACAGCTAC
GCAACACGTACAGACGTACCTGACCATTGTACAGCTAGGTTAGCTTCGACGTACA
TGCCGTAGCCGTAACGTACAGCTACCTGACCATTGCAACGTGGCAACAACGTAA
ACCATTGCAACGCCATTGCAACACGTACAGCTACCTAACCTGACCATTGCAGACC
CCGTAGCCGTAACGGACCATTGCAACGTAAACGGCAACGTACAGCTACCTACACG
CTAGCCGATCGCGTAACCTGACCATTGCAACGTAAACGACTGACCAGTTGACTAAG
TAACGTACCTGACCATTCAACACGTACAGCTACCGACCATTGCAGACCATTGCAA
GCAACGCGTACCTGACCATTGCAACGCCATTGCAACAGCTTGCCCTTGACGCCA
CACGTACAGCTACCTAACCTGACCATTGCAGACCATTGCAACGTAAATAAGACAGC
TTAGCTTCGACGTACAGCTGAGAGCAACACGTACAGCTACCTAACCTGACCATTG
TGCTAGACGCCATCCGTAGCCGTAACGTACAGCTACCTGACCATTGCAACGTAA
CATTAGTACCTGACCATTGCAACCATTTGCAACACGGCTACCTAACCTGGACCAT
TGACCATTGCAACGTAAACGGCAACGCCTGACCATTGCAACACGTACAGCTACCT
TGAGATTAGCTTCGAGCACTGACCATTGCAGACCATTGCAACGTAGACACGTACA
CATTAGTACCTGACCATTGCAACATGCAAGCTAACGTACCTGACCATTCAACACG
ACGCAGCTACCTAACCTGACCATTGCAAGACCATTGCAGACCATTGCAACGTAAACG
CGTACCTGACCATAACAGTACAGCTACCTGCTACCTGACCATTGCAACGTAAACG
AACAGCCATTGAACGCCATCTACCTAACCTGACCATTGCAGACCATTGCAACGTAA
GCAACACGTACAGACGTACCTGACCATTGTACAGCTAGCCGATCGCGTAACCTGA

GCCGTTAGACCATTGCAACGTAACGGCAACGACGTTACAGCTACCTTAGTACCTGA
TGCAACGTAACGACTGACCGTTGACACGTAACGTTAACCTTGCAACGCCATATC
CCTGACCATTGCAACGTAACGGCAACGCGTACCTGACCATTGCAACGCCATTGC
CTAACCTGACCATGCAAGCTACCTAACCTGACCATGACAGACCATTGCAACGTA
GCTGAGATAGCCGATCGCGTAACC TGCAACGTTAGACCGTTGACCGAC
TACCTGACCATTGCAACACGTTAC ACCTGACCATGCAACGTTACCTA
ATTGCAACGTAACGACTGACCG CGTGCAACTCACCTGGCATGC
TACAGCTACCTAACCTGACCAT CGTACAGCTACCGACCGTTGA
CCCTGTACAGCTACCTAACCTG AACGCCATTAACTTGTGCG
CGTAACGACTACGTAACGTTCCGT GTACAGCTACCTGACCACAACG
TGACGTACCTGACCATGTACAGCT CTTGACGTTACAGCTGAGAGCAA
TACCTAACCTGACCATGCAACATT GTACCTGACCATGTACAGCTAGG
CAGACCATTGCAACGTTAACGTTAC GCAACGCCATTGCAACAGCCAT
CGGCAACG AAGTCGC
GCAGACCA AGCTACC
AACCTGAC GTACAGC
GCTACCTA ACGTAGA
TACAGCTA AATGCAA
ACTGACCG ATTGCAC
GCAACGCG TGCGACT
GACACGTA ACAGTGC
CCATTGCA GATACGT
AACCTGAC GTACCTA

BY JESSICA MACINNIS
ILLUSTRATION BY MATT CHASE

INTO THE DEEP

Thirty years ago, neural networks were dismissed as an impractical approach of artificial intelligence. Today, neural networks are seeing a revival in a field called deep learning. While the applications of deep learning are seemingly endless, Professor Brendan Frey is using it to transform medicine as we know it

CGACTG CAACTC TGGCATGC
TACAGCTACCTAACCTGACC CCGTTGAACTCGTACATGC
ACCTGTACAGCTACCTAACCGCCATTAACTAACGGTCA
CGTACAGCTACCTGACCACA CGAGCACTGACCATGCAGA
GTTGACACGTTATTGCAACGTTGACCATTGCAACGCCATT
CTTCGACGTTACAGCTGAGAG GTAACGTTACGTTGACACAGC
TACCTAACCCCATGC GACCAT CAGCTAGG
CAGACCATT CAGCCAT
CGGCAATA TCGCGTA
TGCAACGTTCCATATC
CCTGACCA CCATTGC
CTAACCTG AACGTA
GCTGAGAT GACCGAC
GTCCTGAC GTACCTA
ATTGCAAC GGCATGC
TACAGCTA CCGTTGA
TCCTGTAC TGTTGCG
CGTAACGACTACGTAACGTTCCGT GTACAGCTACCTGACCACAACG
TGACGTACCTGACCATGTACAGCT CTTGACGTTACAGCTGAGAGCAA
TACCTAACCTGACCATGCAACATT GTACCTGACCATGTACAGCTAGG
CAGACCATTGCAACGTTAACGTTAC GCAACGCCATTGCAACAGCCAT
CGGCAACGCGTACCTGACCATT GCAACAGCCATTGCTAGACGC
GCAGACCATTGCAACATTGCAA TAGCCGTTAACGTTACAGCTACC
AACCTGACCATGCAATAGTAGA ATGTACAGCTAGGCGTTACAGC
GCTACCTAACATTGCAACGTTAAC TTGACACGTTATTGCAACGTTAGA
TACAGCTACCGACCATGACCGTAG GTACAGCCTTGTGCGAATGCAA
ACTGACCGTTGACACGTTAACCTTGGCAACGCCATATCTGACCATTGCAC
GCAACGCGTTACCTGACCATTGCAACGCCATTGCAACGTTCTAACCTGCGACT
GACACGTTACAGACACAGCTACCTAACCTGACCAGACGACTGCAACGTTACAGTGC
CCATTGCAACGTTAGACCGTTGACCGGTTAGCTTCGTTACAGCTGAGATACGT

When 86-year-old May Ashworth book-ended her Google search with “please” and “thank you,” her grandson was not the only millennial to get a laugh out of it — his screenshot of her search went viral when he tweeted it. Google’s official Twitter account even responded. While Ashworth is likely not the only person to assume that there was a human sitting at Google’s headquarters assigned to answer her query, she might be the politest. It is within Ashworth’s lifetime that artificial intelligence has progressed from science fiction to a fact of life.

Artificial intelligence has been a part of our lexicon for decades, often conjuring visions of malevolent robots, cyborg law enforcers, or a pair of squabbling androids. While there is no universally accepted definition, artificial intelligence is most generally accepted as the umbrella term under which we get computers to perform tasks as well as — or better than — a human can. From self-driving cars to AlphaGo to May Ashworth’s Google search, artificial intelligence is no longer in the realm of science fiction. Words like machine learning and deep learning have become ubiquitous: sorting the spam in your inbox, recognizing your face in photos on Facebook and suggesting the next show to binge-watch on Netflix. While we might take these conveniences for granted, or like Ashworth, not even realize they exist, it is becoming difficult to imagine a life without them. But what exactly are they, and what is the next frontier?

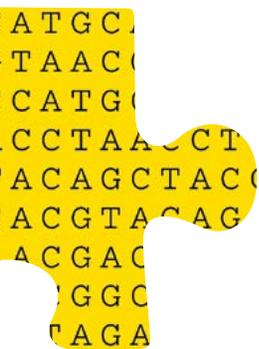
Machine learning, a branch of artificial intelligence, is the term used to describe how computers can analyze data and make subsequent decisions or determinations based on them. Conceptually, artificial neural networks are an interconnection of primitive functions that each mimic a biological neuron. Deep learning uses multiple layers of these artificial neural networks, rearranging the connections between them after exposure to new data. Effectively, these deep learning systems “learn” much like a human does: through interactions and experience. It is these layers of neural connections that make a deep learning system, well, deep: the system delivers better performance with more layers.

Professor Brendan Frey, in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, studied under University of Toronto

Computer Science Professor Geoff Hinton, whom many consider to be ‘the father of deep learning,’ focusing on computational vision, speech recognition and text analysis while he was completing his PhD. “Geoff was very inspirational because he believed that neural networks could be trained in such a way that they could automatically ‘understand’ the data, instead of just memorizing examples and linearly interpolating, which common machine learning systems, such as support vector machines, do,” says Frey. “You can trace a lot of innovations in the field back to Geoff.” Together, they authored one of the first papers on deep learning, published in *Science* in 1995.

When Frey joined ECE as professor in the early 2000s, his research interests soon pivoted from vision recognition, speech recognition and text analysis when he was told that his unborn child had a genetic problem — suddenly identifying cats in YouTube videos seemed trivial. The diagnostician could only offer Frey ambiguous information: it could be a big problem or it could be nothing. For Frey, this was the catalyst in changing the focus of his research. “I learned two important lessons: first, that I needed to work on something that would more profoundly change society, such as by helping people that were in my situation, and second, that there was a lot of value to patients in producing computational systems that could manage probabilities and reduce uncertainties,” says Frey.

It was around this time that the human genome had been sequenced, and Frey realized that the deep learning techniques he had been using in vision and speech applications could do what would be impossible for a human to do — understand and interpret the three billion letters of the human genome. Frey’s research group continued to work on vision and speech applications, but began to focus more and more on developing machine learning methods that could understand the human genome. He foresaw that deep learning techniques could create more robust medical systems, because they can balance different sources of information and avoid over-specialization, or memorization, of individual cases. “It’s like an undergraduate student doing well on the exam but not really understanding the material,” says Frey. “What you really want is a student who can understand the material



and you can pop a surprise quiz on them and they do well — it means they know what is going on, and they aren't just studying to ace the exam questions.”

While the computational capacities for deep learning were continuing to expand, advances in medicine were opening doors to combine the two. “The genome had been sequenced in 2002 and at the time, and even now, many researchers focused on gathering genome data and correlating it with disease labels.” Instead, Frey focused on using deep learning to train computer systems that can understand what happens inside a cell when DNA is altered. In turn, these systems can then be used to detect and fix disease — much more challenging than applications in vision, text analysis and speech processing. It's one thing to train a computer to identify a stop sign in a photo — this is a task easily completed by most humans. But it's another thing entirely to train a computer to do something that no human can do: not only interpret the three billion letters of the genome — but make sense of it. Frey is bridging what he calls the “genotype-phenotype gap” by using deep learning to understand the text of the genome and how errors in it can lead to disease phenotypes.

Genome biology is ripe for deep learning: it produces reams of data, and pretty soon you will be able to sequence your DNA for less than a trip to the grocery store. It's fitting that the computational system Frey uses to understand how genetic variation leads to disease is itself based on a biological

system. But artificial neural networks are not new: they have been around since the early days of AI. They were largely dismissed because their use was limited by slow computers and limited data sets. “When I was a grad student in the early 90s, training a neural net took too long,” says Frey. “It was very hard to be creative; it could take months to get results, so the focus was on theory.” Today, computers are faster, data sets are larger and more easily accessible, and graphics processing units (GPUs) are utilized to accelerate computing. These three technological advances have enabled deep learning experiments to be run in minutes, and applications to grow in recent years at an unprecedented rate.

Frey's research team published a series of papers in *Nature* and *Science*. Then, in 2015, he founded Deep Genomics, building a team with expertise in both machine learning and computational biology. In November of 2014, Deep Genomics received \$5 million in seed funding and launched their first product, SPIDEX, which provides a set of genetic variants and predictions for their effects on human splicing across the entire genome.

In 2015, there was a panel at the premier conference for machine learning. Some of the giants in the field including Demis Hassabis from DeepMind, Yoshua Bengio of the Université de Montréal and Yann LeCun of Facebook were asked what the next big thing would be in machine learning. The answer was nearly unan-

It's another thing entirely to train a computer to do something that no human can do: not only interpret the three billion letters of the genome — but make sense of it.

imous: medicine. The Human Genome Project determined the sequence of all the genes in the human genome — we now have the blueprints to the human body. But how do we read them, and what do they mean? “It's one thing to have a framework to describe how things work,” says Frey. “It's another thing to have a reliable, scalable and trustworthy system that can make predictions on what's going to happen.” While neural networks were dismissed just a few decades ago, they may hold the key to unlocking our understanding of ourselves. Deep learning reduces the need for humans to ‘teach’ computers to complete a task, but will deep learning techniques enable computers to teach us to understand the human genome? If you ask Professor Frey, the answer is yes. ■

Undergraduate Study

*TRAINING THE
NEXT PROBLEM
SOLVERS*

Behind advances in technology, business, medicine and beyond, you'll find electrical and computer engineers using math and science to solve problems and create innovative new products. 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. **A**

16:1
Undergraduate
student-to-
faculty ratio
in ECE

Undergraduate
degrees awarded,
2015-2016

ElecE 159
CompE 123
Total 282

EMPOWERING STUDENTS IN ECE'S NEWLY RENOVATED ENERGY SYSTEMS LAB

In The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, students have long studied the generation, transmission and distribution of electrical power in the Energy Systems Lab. In this lab, undergraduate students studying electrical engineering receive hands-on experience with power and communications systems at specifications found in the real world.

With the support of Dean Cristina Amon, the Energy Systems Lab underwent major renovations between 2014 and 2016 to improve safety, modernize power distribution infrastructure and ensure students are at the forefront of energy systems education. The electrical switch gear has been replaced, new DC and AC power supplies installed, and advanced machine experimental benches are now available to students. The old switch gear was dismantled and hundreds of metres of electrical cabling were replaced. New custom-designed workstations were installed, along with fibre-optic network cabling — the lab has the facilities to study the control and management of micro grids, smart grids and also integration of renewables into modern utility networks.

As our energy needs and technological advances continue to change the way we power our world, the Energy Systems Lab offers our students and researchers access to state-of-the-art equipment with specs they will find in the real world.





CRACKING CODE

U of T's Engineering Outreach office designs and delivers programs to promote science, technology, engineering and math (STEM) education to pre-university students. In the summer of 2016, ECE's Karen Cerullo (CompE Year 3) joined a team of U of T Engineering students who were tasked with developing and delivering curriculum to teach students from grades three to eight how to code. After developing and beta testing their workshops over the course of the summer, they launched a two-day pop-up camp at U of T. Karen's beta testing and pop-up coding camp meant that close to 300 kids had a coding experience in the summer that was directly linked to school curriculum. Earlier in the year, Karen helped to deliver a series of coding workshops at InnovateU — Canada's largest one-day STEM event for kids — enabling 1400 students to have a coding experience on campus. ANNUM sat down with Karen to hear how she spent 2016 with some future electrical and computer engineers.

Q: WHAT WAS YOUR SUMMER POSITION?

A: I was hired by the Engineering Outreach office and I was on a team of seven U of T Engineering students. Our job was to find a way to integrate coding education into curriculum that already exists at the elementary school level and develop workshops that would teach coding. For example, how can we teach coding with science, with math or with language arts? So we developed and delivered about five workshops: it's actually really easy to teach a kid how to code or how to read binary numbers, they are open-minded and very committed to learning something new.

Q: WERE THERE ANY EXCITING BREAKTHROUGHS THAT YOU SAW WITH THE STUDENTS?

A: I think one of the most exciting things was not just teaching the students how to code, but teaching them to use the coding mindset to solve any problem, whether in the classroom or otherwise. One of our workshops didn't use a computer or tablet at all. It was fun to see them apply what they learned on the computer screen to a totally different model.

Q: HOW DOES WHAT YOU WERE TEACHING ELEMENTARY STUDENTS INTERSECT WITH WHAT YOU ARE LEARNING IN ECE?

A: One of our workshops was entirely based on dynamic programming, which was really cool because

it's actually something I am learning right now in third year. Even though it's a third year engineering concept, we were able to break it down and teach it to grade three students. The idea is that you take a complicated problem and break it into a bunch of smaller problems and work backwards to solve the bigger problem — I think that's a really important skill to learn.

Q: WHAT SURPRISED YOU THE MOST?

A: I think just the rapid progression of the students' understanding over the course of just a few days. For many of them, the first day was their first exposure to anything to do with coding and just a few days later they were creating apps and writing their own programs. One of the programming languages we used with them is called Scratch and by the end of their time with us some of them were teaching me things on it! It really felt like we were all learning at the same time.

Q: WHAT'S NEXT FOR THIS PROGRAM?

A: Right now, it's running as a Saturday program. The team that developed the workshops over the summer alternate Saturdays and we teach all of the content we had developed. We're finessing the workshops a bit and I'd love to see it grow into a larger-scale program next summer.

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 2016–2017, 226 ECE students are working in PEY placements at 86 companies located across Canada and around the world, earning an average salary of \$55,125. To learn more, visit uoft.me/ece-pey.

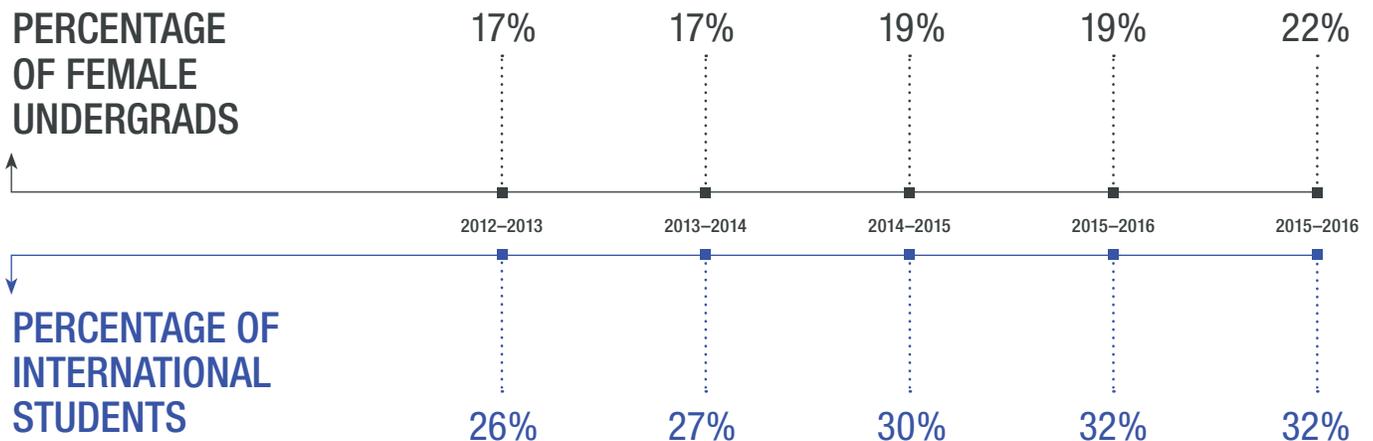
\$101,318
 Top salary earned by an ECE student on PEY placement in 2016-2017

UNDERGRADUATE ENROLMENT, 2012–2013 TO 2016–2017

	PART TIME	FULL TIME	TOTAL
2012-2013	224	1142	1366
2013-2014	259	1133	1392
2014-2015	272	1208	1480
2015-2016	313	1199	1512
2016-2017	283	1230	1513

ECE PEY PLACEMENTS, 2012–2013 TO 2016–2017

	ElecE	CompE	TOTAL
2012-2013	102	79	181
2013-2014	133	77	210
2014-2015	156	70	226
2015-2016	160	92	252
2016-2017	135	91	226



Graduate Study

*AN INTERNATIONAL
DESTINATION*

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, com-

puter engineering, electromagnetics, electronics, energy systems, photonics, and systems control. They are supervised by our faculty of 79 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. **A**

\$2.2M
in graduate student
scholarships
awarded by ECE
for 2016-2017

OUR GRADUATE STUDENTS BY FIELD OF STUDY

	MASC	PHD	TOTAL
BIOMEDICAL ENGINEERING	6	6	12
COMMUNICATIONS	30	55	85
COMPUTER ENGINEERING	41	60	101
ELECTROMAGNETICS	14	16	30
ELECTRONICS	23	33	56
ENERGY SYSTEMS	15	23	38
PHOTONICS	20	40	60
SYSTEMS CONTROL	9	9	18



Mohamed Abdelfattah receives TA Teaching Excellence Award

Mohamed Abdelfattah, a PhD candidate in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, was awarded a 2016 TA Teaching Excellence Award from the University of Toronto's Teaching Assistant Training Program (TATP). This award, conferred to just five recipients across the University in 2016, recognizes the outstanding contributions of teaching assistants at the University of Toronto. The TATP Teaching Excellence Award also seeks to value the work of TAs who challenge and inspire undergraduate students.

"I know firsthand the benefit of having great TAs and instructors," says Abdelfattah. "TAing is not just about grading exams and supervising classrooms — it's about enabling undergraduate students to learn how to connect the dots and develop skills they will take with them as they embark on their careers." Abdelfattah has taught Computer Organization, Digital Systems and most recently was head teaching assistant to Professor Vaughan Betz in a course entitled Communication and Design.



CHARTING NEW TERRITORY: *FROM MASTERS TO MEDCHART*

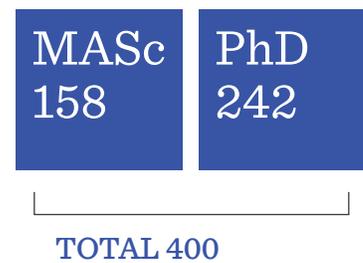
James Bateman has spent the last few years making things smaller. As a MASc student under the supervision of Professor Sorin Voinigescu, Bateman worked on nanoelectronic devices. “Transistors are getting faster and smaller, but what is the limit on existing silicon technology?” says Bateman. “The research we conducted tried to find this limit using a first principles approach while also investigating novel alternative materials.”

As the devices he was working on got smaller, the binder of medical information related to a family member’s care was getting larger. “Travelling between appointments, clinics and hospital stays when my father-in-law was ill was a complicated process,” says Bateman. “My mother-in-law carried around a binder full of test results, prescriptions and appointment details — she found she could provide better care when she had access to all of his medical information at any given time.” While Bateman was working on his thesis, he was also thinking about how he could make that overflowing binder smaller and more accessible.

Bateman’s personal experience led to the creation of MedChart, which provides patients with electronic access to their medical records. By the time he defended his thesis MedChart had already received seed funding and was providing patients and caregivers access to medical records through their cloud-based patient portal. “I think providing patients with access to their own records is empowering,” says Bateman. “We are now working with major healthcare providers, like hospitals and clinics; there’s a demand for easy access to personal healthcare records.”

As for the connection between his research and his work with MedChart: “Both deal with complex systems — and finding ways to make these systems more efficient,” says Bateman. But not everything is getting smaller: “We’re excited to see the company grow: by 2017 patients across Canada will have access to their personal health records through MedChart.”

Enrolment in Research-stream Programs 2016–2017



TOTAL GRADUATE ENROLMENT, 2012–2013 TO 2016–2017 (HEADCOUNT)

	FALL 2012	FALL 2013	FALL 2014	FALL 2015	FALL 2016
MASC	174	167	173	145	158
PHD	236	249	253	245	242
MENG	155	140	169	249	249
TOTAL	565	556	595	639	649

GRADUATE DEGREES AWARDED, 2011–2012 TO 2015–2016

	2011–2012	2012–2013	2013–2014	2014–2015	2015–2016
MASC	57	56	70	71	70
PHD	34	37	27	35	47
MENG	25	56	85	76	97
TOTAL	116	149	182	182	214

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 4,500 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 about how research in ECE is shaping the world at uoft.me/eceresearch. 

**\$348K – 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 **8**
PHOTONICS **9**
SYSTEMS CONTROL **5**
EMERITUS **21**

TOTAL 100

PIERO TRIVERIO WINS ONTARIO EARLY RESEARCHER AWARD

Simulation techniques can be applied to electronic design, electrical grids, and medicine



PHOTO BY KIWOS

There are complex networks all around us — the network of tiny wires that exchange information in our smartphones, the electrical grid that delivers power to consumers and the human vascular system that carries blood to and from the heart. Professor Piero Triverio and his team develop innovative mathematical models to predict how these networks will function. For example, engineers are striving to make sure that the millions of components in a smartphone exchange information quickly and reliably. “As we push the performance of our electronics, the number of phenomena that need to be taken into account as we design them increases,” said Professor Triverio. “Predicting how components exchange information is a difficult but

crucial step. We build models and algorithms to do just that.” Triverio and his team have also applied their models to predict electricity flow in power grids, and to study blood flow in human arteries. In 2016, Professor Triverio was awarded an Early Researcher Award from the Government of Ontario. These awards are conferred to scientists in the first five years of their academic career and provide them with funding to help build their research teams. “Our work is just one example of what mathematics can do to improve the day-to-day lives of people,” said Triverio. “Whether it is helping to make cellphones faster and smaller or proposing ways to make the diagnosis of medical issues safer and less invasive.”

SELECTED AWARDS WON BY ECE FACULTY, 2015-2016

CANADIAN ACADEMY OF ENGINEERING, FELLOWS
Professors [Nazir Kherani](#) and [Deepa Kundur](#)

INSTITUTE OF ELECTRICAL & ELECTRONICS ENGINEERS, FELLOW
[Professor Dimitrios Hatzinakos](#)

NSERC, BROCKHOUSE PRIZE FOR INTERDISCIPLINARY RESEARCH
[Professor Ted Sargent](#)

COMMITTEE FOR THE STATUS OF WOMEN IN COMPUTING RESEARCH, BORG EARLY CAREER AWARD
[Professor Natalie Enright Jerger](#)

INSTITUTE OF ELECTRICAL & ELECTRONICS ENGINEERS, INFORMATION THEORY SOCIETY, AARON D. WYNER DISTINGUISHED SERVICE AWARD
[Professor Frank Kschischang](#)

UNIVERSITY OF TORONTO, CONNAUGHT INNOVATION AWARDS
Professors [Stewart Aitchison](#), [Roman Genov](#), [Glenn Gulak](#)

UNIVERSITY OF TORONTO, PRESIDENT'S TEACHING AWARD
[Professor Jonathan Rose](#)

UNIVERSITY OF TORONTO, INVENTION OF THE YEAR AWARDS
[Professor Brendan Frey](#)

ECE INVENTIONS DISCLOSED OVER THE PAST FIVE YEARS

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

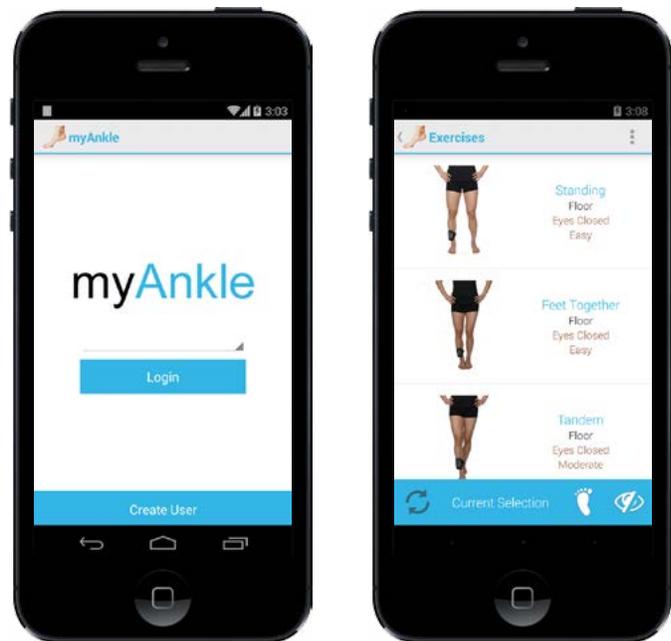


App-y and Healthy

Fifteen years ago, if you sprained your ankle you might reach for your phone to call your doctor to book an appointment. Today, you might reach for your phone to measure your balance, improve your physiotherapy and track your recovery.

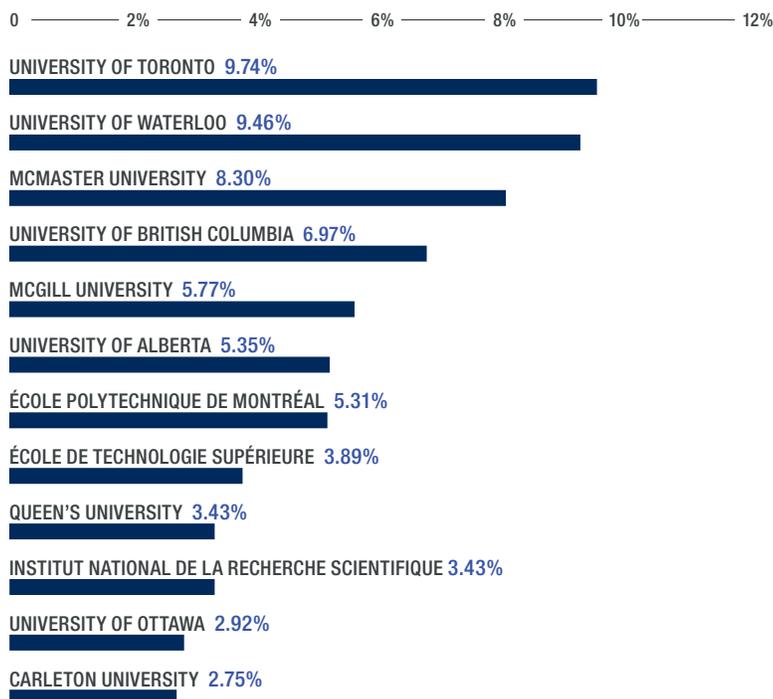
myAnkle, a mobile application, does just that; using a smartphone's accelerometer, the app is able to detect the stability of a user's leg and ankle and track their progress over time as they complete exercises to regain strength and balance. The first iteration of myAnkle was developed in a course that matches graduate students from across U of T to students with graduate level programming skills from The Edward S. Rogers Sr. Department of Electrical & Computer Engineering and the Department of Computer Science and is taught by Professor Jonathan Rose.

Professor Rose, a renowned expert in Field Programmable Gate Arrays (FPGAs), pivoted his research focus from integrated circuits to the automation of medicine. He is exploring and developing the nascent field of computerized support for medicine; and the creation of apps like myAnkle is just one way electrical and computer engineers can contribute to the field of medicine. Rose is also focused on the connection to the underlying hardware of mobile devices: "As mobile devices continue to evolve and the Internet of Things continues to expand, there's a huge opportunity to make medicine more accessible to people in all corners of the world."



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

PROPORTION OF TOTAL NSERC FUNDING FOR ELECTRICAL AND ELECTRONICS ENGINEERING AT LEADING CANADIAN UNIVERSITIES, 2011-2012 TO 2015-2016



ECE RESEARCH FUNDING 2010-2011 TO 2014-2015		FEDERAL	PROVINCIAL	INDUSTRY	OTHER	TOTAL \$
		2010-2011	8,972,287	3,184,950	2,213,062	3,491,863
2011-2012		10,180,802	1,569,027	2,910,981	3,484,571	18,145,381
2012-2013		9,811,957	2,481,833	2,305,784	3,465,717	18,065,291
2013-2014		11,618,423	3,652,656	3,117,251	4,122,724	22,511,054
2014-2015		14,700,926	4,975,436	3,260,799	2,115,017	25,052,178

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.

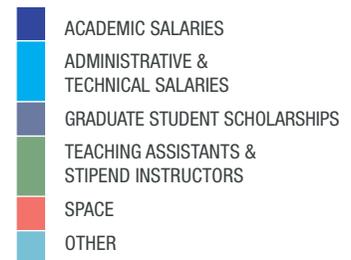
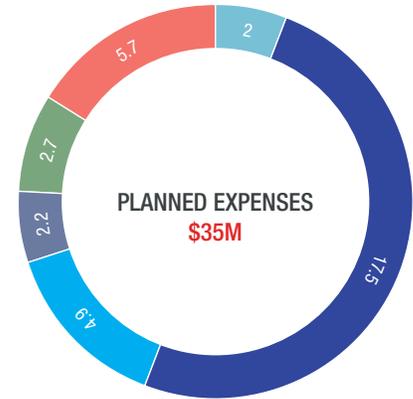
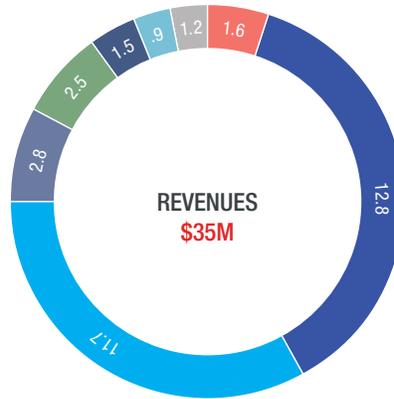
We enjoy active partnerships with more than 85 industry funders and collaborators. Faculty members from ECE gathered

in summer 2016 to discuss the departmental research vision for the next five to 10 years and identify ways in which professors and researchers can work with industry to tackle challenges collaboratively. Attendees considered the greater societal and economic benefits that could stem from research in areas such as data, machine learning and neuromodulation. Learn more about upcoming events at uoft.me/eceevents. 

ECE accounts
for 37% of all
UofT Engineering
invention
disclosures

WHERE THE MONEY GOES:

OUR 2016-2017
FINANCIAL PICTURE AT
A GLANCE (\$ MILLION)



RESEARCH PARTNERSHIP SPOTLIGHT

When Professor Ali Sheikholeslami was a graduate student, he embarked on an internship with Fujitsu Laboratories. Nearly twenty years later, and now a professor in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, his research partnership with Fujitsu Laboratories remains active. *ANNUM* sat down with Professor Sheikholeslami to hear about his collaboration with Fujitsu Laboratories and how both the partnership and the research have developed over the course of two decades.

Q. YOU ARE A MEMBER OF ECE'S ELECTRONICS GROUP; TELL US ABOUT YOUR RESEARCH AREAS.

A. My research group works in the areas of analog and digital integrated circuits, high-speed signaling and VLSI memory design (including FeRAM, SRAM, CAM, and MRAM).

Q. DESCRIBE SOME OF THE MAJOR AREAS OF FOCUS THAT YOU HAVE WORKED ON WITH FUJITSU LABORATORIES.

A. I've worked with Fujitsu Laboratories

on three major areas: the first was on non-volatile memory. This is where information is stored to a device and can later be accessed even when power has been disconnected. Today, we see this technology used a lot in contactless smart cards. Since there is no battery on these cards, the information is held in non-volatile memories. When you hold your card near a terminal at a train station to pay for your ticket, a small amount of electromagnetic energy is transferred to the card to power up the memory and to read/write its information. The second area of focus has been on spin electronics, where we use the spin of electrons — instead of their charge — to store and retrieve information in what is called spin-transfer torque magnetoresistive random-access memory or STT-MRAM. The other area that I have worked with Fujitsu Laboratories is on high-speed chip-to-chip communications.

Q. HOW LONG HAVE YOU BEEN WORKING WITH FUJITSU LABORATORIES?

A. I began working with Fujitsu Laboratories as a graduate student in the late 1990s and since then, the partnership has been ongoing. Over the years, more

than 20 graduate students from my research group have been involved in research projects with Fujitsu Laboratories and several of them have gone to Fujitsu Laboratories for internships, similar to the internship I had when I was a graduate student. Together, we have published more than 30 papers — it's been a really productive partnership.

Q. WHY HAS THIS PARTNERSHIP BEEN SO SUCCESSFUL?

A. It's been successful because it's really been a true collaboration; we are constantly communicating to understand the problems and challenges, and how we can work together to find solutions. Over the years, we've worked closely on some really interesting problems and we've been able to make progress on them because of the two-way nature of the partnership. Whether I approach them with open problems that I think they can help with or vice versa, we've had very open lines of communication and I think the length and success of the research partnership reflects that.

204 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 ADS
ADEXA INC.
AERCOUSTICS
AEROFIS
AFFINITY SYSTEMS LIMITED
AGFA GRAPHICS NV-BELGIUM
ALAZZAM CONSULTING
ALTERA
AMAZON
AMD
ANALOG DEVICES
APPLE
ARCHITECH
ARISTA NETWORKS
ARUP
ASSETMINE INC.
AUTOLIV ELECTRONICS
CANADA INC.
AVIYA TECHNOLOGIES INC.
BANK OF MONTREAL
BELL MOBILITY
BIBLIOCOMMONS INC.
BIG VIKING GAMES
BLACKBERRY LIMITED
(FORMERLY RIM)
BLUECAT NETWORKS
BMW GROUP CANADA
BOMBARDIER AEROSPACE
BURO HAPPOLD
CANADA HEALTH INFOWAY INC.
CANCER CARE ONTARIO
CASEWARE INTERNATIONAL INC.
CAST SOFTWARE INC.
CELESTICA INC.
CGI GROUP INC.
CIBC
CIENA
CIRBA INC.
CISCO SYSTEMS INC.
CITY OF BRAMPTON
COMMUNICATIONS AND POWER
INDUSTRIES CANADA INC.
CONAVI MEDICAL INC.
CONTAX INC.
DAC GROUP
DELL
DELOITTE
DEMONWARE
DESIRE2LEARN
DOUBLETHINK INC. (MYBLUEPRINT)
EASTERN POWER LTD.
ECAMION
ENERSOURCE
ENVIRONMENT CANADA
ENVISION MOBILE
EPSON
ERICSSON CANADA INC.
ESNA TECHNOLOGIES
EVENTMOBI
EVERTZ MICROSYSTEMS LTD
FINANCEIT CANADA INC.
FIO CORPORATION
FITBIT
FIXMO
FLIPP CORPORATION
FU TIN BUSINESS CO. LTD.
GENERAL ELECTRIC CANADA

GENERAL MOTORS OF
CANADA LTD.
GENESYS CANADA
LABORATORIES INC.
GOOGLE
GUOMAO GROUP CO. LTD
HH ANGUS
HONDA
HONEYWELL
HP CANADA
HUSKY ENERGY
HUSKY INJECTION MOLDING
SYSTEMS LTD
HYDRO ONE
IBI GROUP
IBM
IBM CANADA LTD.
IESO (FORMERLY ONTARIO
POWER AUTHORITY)
IMPERIAL
INDIGO BOOKS AND MUSIC
INFINERA
INTACT FINANCIAL CORPORATION
INTEL
INTELLIWARE DEVELOPMENT INC.
INTERFACEWARE INC.
INTUIT CANADA
ITG
KAPSCH TRAFFICOM
CANADA INC.
KIJJI
KNOWROAMING
LABATT BREWERIES OF CANADA
LOBLAW COMPANIES LTD.
M&E ENGINEERING
MAGNA
MARIN SOFTWARE
MARVELL TECHNOLOGY GROUP
MAXXIAN
MCW CONSULTANTS LTD.
MEDIAADX
MERCATUS TECHNOLOGIES INC.
MICROSOFT
MINISTRY OF CHILDREN AND
YOUTH SERVICES
MINISTRY OF TRANSPORTATION
MODIFACE
MONERIS
MONGODB
MORGAN SOLAR
MORGAN STANLEY
NATIONAL BANK FINANCIAL
MARKETS
NATIONAL INSTRUMENTS
NATIONAL RESEARCH
COUNCIL CANADA
NEXJ SYSTEMS INC.
NVIDIA CORPORATION
OMNIVEX CORPORATION
ONTARIO FINANCING
AUTHORITY (OFA)
ONTARIO MINISTRY OF EDUCATION
ONTARIO MINISTRY OF HEALTH &
LONG-TERM CARE
ONTARIO POWER GENERATION
ONTARIO TEACHERS PENSION
PLAN (OTPP)
ORACLE
PALANTIR TECHNOLOGIES INC.
PEARSON EDUCATION
PERASO TECHNOLOGIES, INC
PERISCOPE CAPITAL INC
PERSPECSYS

PHILIPS (CHINA) INVESTMENT
CO., LTD.
PIVOTAL
PRECISIONHAWK
PRINCESS MARGARET
CANCER CENTRE
PROOFPOINT INC.
PSION INC.
QNX
QUALCOMM
QUANTA TECHNOLOGY
QUICKPLAY MEDIA
QUICKTAPSURVEY
RAKUTEN INC.
RAPYUTA ROBOTICS CO., LTD.
RBC
RED HAT CANADA LTD.
RESEARCH IN MOTION
RL SOLUTIONS
RUBIKLOUD TECHNOLOGIES INC.
SANDVINE
SAP CANADA INC.
SAP LABS, LLC
SCICAN
SCOTIABANK
SECURITY COMPASS
SEMTECH
SIEMENS
SMITH & ANDERSON
SOTI INC.
SOUTHPAW TECHNOLOGY INC.
STRATEGIC MAPPING INC.
SUNCOR ENERGY
SUNNYBROOK HEALTH
SCIENCES CENTRE
SUNWELL TECHNOLOGIES INC.
SURFEASY
SYMANTEC
SYNAPTIVE MEDICAL
SYNOPSYS
TECK RESOURCES LTD.
TEKNION FURNITURE SYSTEMS
TELUS
TEMENOS
THALES GROUP
THE HOSPITAL FOR
SICK CHILDREN
THEREDPIN.COM REALTY INC.
TIMEPLAY
TOFUN MARINE SUPPLY INC.
TORONTO HYDRO
TORONTO TRANSIT
COMMISSION (TTC)
TRADEREV. (FORMERLY
NTHGEN SOFTWARE)
TRANSCANADA
TRAPEZE GROUP
TUNNELBEAR INC.
TXIO
TYCO ELECTRONICS CANADA LTD.
UBER TECHNOLOGIES, INC.
ULTIMATE SOFTWARE
UNIVERSITY OF TORONTO
UPSTREAM WORKS SOFTWARE
V SEMICONDUCTOR INC.
VENNSA TECHNOLOGIES INC.
WATTPAD
WEALTHFRONT
WINMAGIC DATA SECURITY
XAGENIC INC.
XILINX INC.
XLV DIAGNOSTICS INC
YELP
ZEBRA

ECE CORPORATE PARTNERS, 2012-2016

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.
RBC
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.
Havelaar
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 BV.
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

\$13.6 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.



\$52,500 *industry investment* → **\$202,500**
leveraged 4x

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.

ALLISON BROWN, PHD
DIRECTOR OF FOUNDATION & CORPORATE PARTNERSHIPS
ALLISON.BROWN@ECF.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:

I-CHENG CHEN, DIRECTOR OF MOBILE ENGINEERING // AMD

SAI-KIT ENG, DIRECTOR OF PROGRAM MANAGEMENT // QUALCOMM CANADA

DR. RON HO, DIRECTOR OF ENGINEERING // INTEL

WALTER KINIO, VP, RESEARCH AND INNOVATION // THALES CANADA, TRANSPORTATION SOLUTIONS

ALLEN LALONDE, SENIOR EXECUTIVE // IBM CANADA RESEARCH & DEVELOPMENT CENTRE

SONG ZHANG, DIRECTOR, TECHNOLOGY PLANNING & PARTNERSHIP // HUAWEI CANADA

Alumni

*STRENGTH OF
THE ECE
ALUMNI NETWORK*

Graduates of ECE can be found around the world: in boardrooms and in classrooms, in mining and in media, at startups and in surgical theatres. Our alumni community is a powerful global network of leaders in a wide array of fields. 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 Jessica MacInnis at eceinquiry@utoronto.ca or 416 978 7997. **A**

Total number
of active alumni
worldwide:
11,536



PHOTO BY ADAM COISH

ECE ALUMNI BOARD OF ADVISORS: WELCOME SOMEN MONDAL

Members of the ECE Alumni Board of Advisors play an important role in informing the department's strategic direction by providing valuable insight into industry trends while advising on alumni engagement and advancement.

This year, ECE alumnus and long-time U of T Engineering Alumni volunteer Somen Mondal joined the board. Mondal graduated from ECE with a BSc in Computer Engineering in 0T2. He served as Co-Founder and CEO of Field ID until it was successfully acquired by Master Lock LLC in 2012. That same year, Mondal received the Faculty of Engineering's Early Career Award. This award is presented to an individual who has distinguished themselves in the profession, the community, the University and other related fields. His leadership at Field ID helped earn the company a spot on the Profit Hot 50 and Deloitte Fast 50 lists, and he was named winner of the Ontario Ernst & Young Entrepreneur of the Year award in 2012.

Somen Mondal is the Co-founder and CEO of AI recruiting company, Ideal. Mondal and his team use artificial intelligence to help employers make precise and efficient high-volume hiring decisions. Companies use Ideal's Intelligent Screening technology to sift through resume noise and instantly identify who to interview.

ECE alumni celebrated at U of T's Arbor Awards ceremony

Created in 1989, the Arbor Awards recognize University of Toronto volunteers for their outstanding personal service to the university.

This year, three ECE alumni received Arbor Awards for their dedication to the university and for their contributions to the experience of U of T students, faculty, staff and alumni. Congratulations Emmy Choi (ElecE 9T5), Catherine Lacavera (CompE 9T7) and Paul Malozewski (ElecE 8T3).

2016 ECE ALUMNI BOARD OF ADVISORS:

JOHN EAST

ElecE, MBA, UC Berkeley
Former CEO of Actel Corporation

ALEX GRBIC

CompE 9T4, MSc 9T6, PhD 0T3
Vice President, Product Marketing
and Planning, Programmable
Solutions Group, Intel Corporation

CATHERINE LACAVERA

CompE 9T7, JD/MBA U of T
Director, Litigation at Google

NATASHA LALA

CompE 9T8
Managing Director, Solutions for
Business at OANDA Corporation

SOMEN MONDAL

CompE 0T2,
MBA Queen's University
Co-founder and CEO of Ideal

ALEX SHUBAT

ElecE 8T3, MSc 8T5,
PhD Santa Clara,
MBA Stanford
CEO Espresa



Fourth-Year and Alumni Celebration

PHOTOS BY ALAN WU

Alumni, professors and graduating fourth-year students got together at the conclusion of their busy Design Fair week to toast the graduating class, celebrate their accomplishments and welcome them to the ECE alumni community. Alumna Rita Gao (ElecE 0T8) delivered the keynote address.



Fall Alumni Networking Lecture & Reception

PHOTOS BY ALAN WU

ECE alumni from all years came back to Skule to catch up with classmates and hear Professor Deepa Kundur share the latest on her research on security of cyber-physical systems.





Spring Reunion

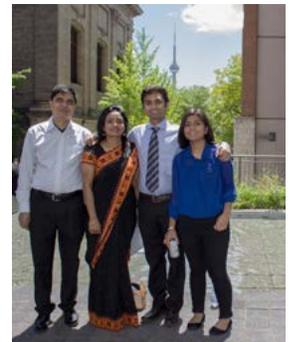
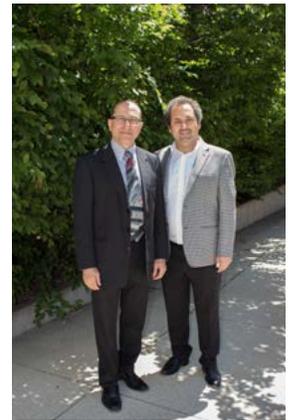
Alumni from 6T0 to 1T6 were welcomed back to campus for our annual ECE lunch and lab tours at this year's Spring Reunion. Attendees visited the newly renovated Energy Systems Lab and joined Frank Gu (Year 3 ECE), the managing director of U of T Engineering's Blue Sky Solar Team for a tour of their workshop. Join us for the next Spring Reunion on June 3, 2017!



Convocation

PHOTOS BY
SHARMINIE VIGNESWARAN

Graduating ECE students heard an inspiring convocation speech from alumnus Rami Rahim (ElecE 9T4) at our Convocation ceremony and then joined their fellow 1T6 and 1T5+PEY classmates at a lunch hosted by ECE. The sun was shining as our newest alumni were joined by their friends, families, faculty and staff to celebrate their accomplishments.



CLASS NOTES

6T6

RUSS JONES
ELECE 6T6, MENG 7T6

A wonderful place to get in on the ground floor, as the study of computing was in its infancy.

7T0

JAY BEATTIE
MASC 7T0

Having spent 13 years in university, it was my privilege to work with some of the best in the field of Power Systems. At U of T, it was Professors Wasyl Janischewskyj and Gordon Slemmon. They shaped my career and I am forever grateful. WJ was extremely knowledgeable in the field of high voltage engineering, and a meticulous teacher of the subject. His assignments were usually due in three weeks, and they needed every minute of that time to complete. In class, the jovial Professor Slemmon (I don't recall that he was called anything but that) would challenge you with his magnetic puzzles. Out of class, he offered wisdom and encouragement.

I graduated from the University of Western Ontario in Electrical Engineering in 1966. I chose Electrical because of something a professor said in a 2nd year electrical class. Basically he paraphrased Michael Faraday. The beauty of electricity is not that it is mystical or mysterious, but that it is under law.

That professor was Stuart Lauchland, founder of the Electrical department at Western and one time president of the APEO. When I graduated, I wanted to do post graduate work under him. But he discouraged the idea. He told me to go somewhere else, so I would get a wider perspective of the subject. Towards that end, he gave me a letter of introduction to U of T. So I wandered down the road to Toronto.

But first, I stopped at Guelph where I worked for Canadian General Electric for a summer. Being a bachelor with nothing better to do and more money than I had ever had before, I took up flying. Over the years since, I got a commercial licence, took up gliding, and set

a Canadian soaring record for altitude. Being drawn to both power flying and soaring, I bought a motor glider and flew it for 20 years until it was destroyed in a windstorm.

I learned about the empirical side of engineering at the University of Toronto where our graduate student quarters were next to the Civil Engineering materials lab. Every half hour or so, the whole building shook as they broke another piece of concrete to find out how strong it was. I suspect they are still doing those experiments.

Not knowing what to do with life, I continued on as a student at Waterloo for another five years. For diversion, I dabbled in student politics, becoming president of the Graduate Student Union and a member of both the University Board and the Senate. I took so long to get a PhD, my entrance exam elapsed and I had to redo in year five.

Back in '66, if you wanted 13 job offers, you took 13 interviews. Things were different in '75. But I did get 3 opportunities: post-doctoral in Scotland, Ontario Hydro Research, or Saskatchewan Power Research. I chose the latter, reasoning that it is better to not be a small fish in a big pond. Besides, I had just got married and didn't have the courage to face the unknown as a post-doc. That union produced two children. Brien has his masters in Computer Science and plays the trumpet in Saskatoon; Caroline is working on her masters in Sociology while raising two children in Winnipeg.

After 41 years, I am now the second longest survivor working for SaskPower collecting both a salary and a pension. (The number one guy is a technologist in charge of switching station maintenance. He has 48 years!) Over this time, I have gone through a succession of demotions. I started two steps removed from the president – I reported to a man who reported to the president. At my nadir, I was 6 steps removed from the president. But things are looking up, I am now only five steps removed, and doing the same job – basically firefight-

ing. I did do some research and got a few patents. For each one, the company paid me \$1, and the patent lawyer got \$30k.

Amongst the many fires, perhaps my most memorable achievement at SaskPower was predicting how they should re-arrange the cables feeding a hotel in Regina following an explosion in the hotel's transformer vault that injured the hotel owner. It was caused by an unbalanced current overload. We shut the hotel down at midnight and proceeded to measure the self and mutual impedances of the 13 wires that fed the hotel (four 3-phase circuits in parallel plus a neutral). You need to know that the impedance of, and hence current in, a wire depends on its position with respect to other nearby wires. At dawn, we had the 85 measurements to complete the impedance matrix which I then fed into a computer program to find which permutation of wires minimized the imbalance. It reminded me very much of the hours I and other students spent churning away on the IBM 1620 in the Galbraith Building from dusk till dawn on those three-week assignments. The 1620 was so slow, we used to say, that you could hear it think. Back in Regina, the crew rearranged the cables according to my calculations and the resulting load was almost perfectly balanced. Like Professor Lauchland said, electricity is under law – first Murphy's then Kirchoff's in this case.

8T9

JOHN LAZAROU
ELECE 8T9



Back to living in Toronto now after 16 years of working for Goldman Sachs in New York City.

9T4

YING HO HOWARD LAU
COMPE 9T4



I am currently working in the banking industry in Hong Kong.

9T6

ANTHONY MAK
ELECE 9T6

Time flies and it has been two decades since I graduated from ECE and first of all, I would like to say hi to my thesis supervisor, Prof. Hatzinakos in the communications group. Since graduation, I have been working in China and Hong Kong and now I am responsible for international business development in a group of testing laboratories which is a not-for-profit organization. Meanwhile I am doing my MBA at The University of Manchester and expecting graduation in 2017. I can be reached via <https://hk.linkedin.com/in/makanthony>.

9T9

MOHAMMED KHALID
PHD 9T9



I am currently an Associate Professor in the Department of Electrical and Computer Engineering at the University of Windsor in Windsor, Ontario. I am also serving as the Graduate Coordinator in my department.

I have fond memories of time spent as a PhD student in Pratt building. The U of T campus and surrounding city neighbourhoods provided lots of opportunities for recreational activities.

I am happy to report that my youngest daughter is following in my footsteps and is enrolled in our undergraduate program as a second year student.

0T8

IKENNA OSUJI
ELECE 0T8+PEY

I'm a Product Manager at IBM in Austin TX, after completing a Masters in Industrial Engineering and an MBA.

1T4

SARMAD ABDI
MENG 1T4

I originally came from Ottawa to pursue my Master's Degree in Engineering. Since I've always been a small townner, I knew there would be new and exciting challenges to overcome in the metropolitan city of Toronto and that was a very exciting adventure for me. On my first day of my graduate studies, I met many grad students who were from all over the world and had great passion for higher education. I stayed at the Grad House where I got to interact with aspiring lawyers, engineers, doctors, and opera singers! Sooner than later, Toronto started feeling more like home. The Grad Lounge on the 5th floor (still on the 5th?) in Bahen Centre had some very comfy couches where you could usually find me catching a midnight power nap! It has been two and a half years since I've graduated. I've moved back to Ottawa now working in a software company. I will always cherish the friends I made, the Professors I worked with, and the good ol' M.Eng Study Room that didn't have

many power outlets and of course, the abundant amount of empty coffee cups from Second Cup beautifully stacked in the shape of a pyramid in the study rooms. One important takeaway that U of T's ECE has taught me is, why reach for the sky when there are footprints on the moon?

HOUMAN RASTEGARFAR
PHD 1T4

After graduation from U of T, I have been spending time as a postdoc in the United States and Sweden which has been an amazing journey. I miss my days on campus.

1T5

XIA QU
MENG 1T5

I am going to join IBM Canada in late October of 2016!

PEIXI ZHAO
COMPE 1T5+PEY

I just graduated this June and joined RBC Capital Markets right after. It was a rewarding four years for me at U of T Engineering (1 year EngSci + 3 years ECE). I am proud to be a Skule graduate.

1T6

YINGJIAN LIU
MENG 1T6

Working at Pure Technologies Ltd. as a System Engineer.

YUHAN ZHOU
PHD 1T6

In January 2016, Yuhan joined Qualcomm Technologies Inc. in San Diego, CA. He is now working as a Senior System Engineer on LTE/WiFi modem design.

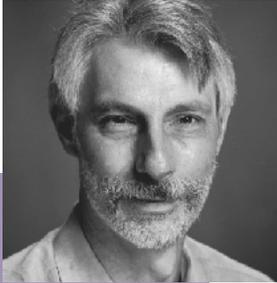
Jot a class note!
Email your news to
eceinquiry@utoronto.ca



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

Distinguished Lectures Series

2016 - 2017



Don Towsley
University of Massachusetts

Oct. 6, 2016
3:00 p.m.

Randomness, Everlasting Security, and Undetectability



Roland Ryf
Nokia Bell Labs

Nov. 10, 2016
3:00 p.m.

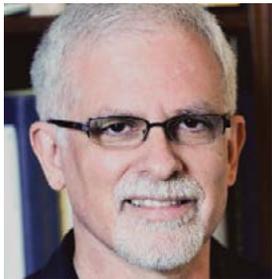
Novel High Capacity Fibers for Optical Communication



Yoshua Bengio
Université de Montréal

Nov. 24, 2016
3:00 p.m.

From Deep Learning to AI



Michael Roukes
California Institute of Technology

Jan. 19, 2017
4:00 p.m.

Nanotechnology for Massively-Parallel, Multi-Physical Interrogation of Brain Activity



Ian Hiskens
University of Michigan, Ann Arbor

Jan. 26, 2017
4:00 p.m.

Modelling and Control of Load Ensembles



T. Paul Chow
Rensselaer Polytechnic Institute

March 16, 2017
4:00 p.m.

Smart Power Devices and ICs with Wide and Extreme Bandgap Semiconductors

All lectures take place in room 1105 of the Sandford Fleming Building.

uoft.me/ece-dls



Marin Soljacic
Massachusetts Institute of Technology

March 30, 2017
4:00 p.m.

Nanophotonics in Material-Systems of Large Sizes



Asu Ozdaglar
Massachusetts Institute of Technology

April 13, 2017
4:00 p.m.

Incremental Methods for Additive Convex Cost Optimization

DIRECTORY

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



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

DIRECTORY



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

ECE RESEARCH DIRECTORY



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



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			•			•
SHEIKHOESLAMI, ALI	•	•	•	•		
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 underlying 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 *hiCUDA*, 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 currently extending 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 currently 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 disperse 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

■ ◆ **Coarse-Grained Reconfigurable Architectures**

Coarse-grained reconfigurable arrays (CGRAs) are a style of programmable logic device situated between FPGAs and custom ASICs on the spectrum of programmability, performance, power and cost. CGRAs contain large coarse-grained blocks and datapath interconnect, possessing less flexibility than FPGAs. The reduced flexibility implies less overhead for programmability and better performance, power and cost than FPGAs for certain applications, namely those applications for which the computational and communication patterns are well aligned with the underlying CGRA capabilities. CGRAs have been proposed by both academia and industry; however, prior works have mainly been “point solutions” without broad architectural exploration or comparisons with competing CGRAs. In this project, we are researching CGRA architectures, and in particular, are developing a software modelling and simulation framework, CGRA-Sim, that will enable the scientific exploration and evaluation of a wide range of CGRA architectures.

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

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

▲ ◆ **Neuroengineering of the Brain**

The main themes of the research fall within the general field of neural engineering and, in particular, the electrical rhythmic activities of the brain. The purpose is to (1) characterize both normal and pathological brain states, and (2) anticipate, detect then abolish, pathological brain states, such as epileptic seizures. The approach is to characterize the spatiotemporal relations of the electrical rhythmic activities in neuronal populations and use cognitive rhythm generators as adaptive devices to classify and mimic the dynamical features of the biological neural networks in the brain. Implantable neuro-mimetic cognitive devices will be developed as low-power hardware 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 important, 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 are also making FPGAs more power-efficient by finding ways to run them at lower voltages. 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; our recent work in this area has focused on faster light propagation physics simulation for medical purposes.

▲ ◆ **Simulation and Optimization of Photodynamic Cancer Therapy**

We are seeking to make a new form of light-activated cancer treatment, photodynamic therapy or PDT, more effective by simulating the outcome of this treatment for a range of possible light source placements and types, and by automatically suggesting optimized locations for the light sources. This requires advanced computation in which we simulate the paths of millions of photons in complex human tissue to determine an accurate assessment of where light will be absorbed and hence tissue will be destroyed. By

combining our fast, hardware-accelerated light physics simulator with advanced optimization algorithms we seek to determine the best arrangement of fibre-optic light sources 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.EECC.TORONTO.EDU/~BROWN

■ Computer Acceleration Using 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 Intel Programmable Solutions Group (formerly Altera Toronto Technology Centre), where I provide direction for the University Program for Intel FPGAs. By combining my involvement in both the University of Toronto and Intel, 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

● Circuits and Systems for Broadband and Spectrally-Efficient Communication

This project will explore circuits and systems to reduce the cost and power consumption of new and emerging spectrally-efficient broadband communications. Low-cost flexible transceivers achieving low power consumption are sought. Circuits and systems amenable to implementation in CMOS will permit their integration alongside other digital signal processing further reducing power and cost. Prototyping the circuits and systems in advanced (28nm and below) nanoscale CMOS provides unique training for graduate students, while also enhancing the work's commercial relevance and impact.

● 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.

● Integrated Circuits and Systems

Research in the Integrated Systems Laboratory spans practically all aspects of mixed-signal (analog-digital) integrated circuit and system design. Prototyping in advanced integrated circuit technologies is a key aspect of the research. Many projects are conducted in close collaboration with industry.

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

Our capacity for digital communication continues to increase as we integrate more and more functionality into fewer electronic components. Integration enables lower-cost, smaller-size systems with lower power consumption. Unfortunately, our ability to integrate complete communication systems onto a single chip today remains limited by small but finite defect rates during chip fabrication, which limit the maximum number

of transistors that can be reliably integrated. In addition, the need to combine different fabrication technologies to realize, for example, a high-density memory alongside high-speed transmitter and receiver circuits, makes integration challenging. Emerging dense interconnect technologies may offer a path forward. These technologies place multiple chips in close proximity and connect them with densely packed wires that may be less than 1 cm long. Our long-term vision is to use these ultra-short-reach (USR) links to interconnect multiple chips so seamlessly that system performance can scale without bottlenecks. We are developing demonstration platforms for USR links, including transceiver circuits that are extremely small and consume very little power.

Cheng, Hai-Ling Margaret

WWW.IBBME.UTORONTO.CA/FACULTY/MEMBERS/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, deep-tissue 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.

■◆ Programming Models and Architectures for Large-Scale Reconfigurable 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.

■◆ 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.

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, piezoelectric and plasma devices in conjunction with the electrical source.

Draper, Stark

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

Anytime Coding for Distributed Computation

In this project we develop a novel coding scheme to speed up distributed computation through a form of approximate computing. It is known that task replication can greatly mitigate the “straggler effect” in cloud computing, wherein an overall computation can be significantly delayed by slowed processing nodes (or “stragglers”). It has also been demonstrated that, in certain contexts, ideas of error-correction coding can more efficiently deal with stragglers than pure replication. The approach proposed herein builds on these earlier observations through an “anytime” approach to approximate computing. In this paradigm, over time one can produce approximate solutions of increasing accuracy. To accomplish this we first decompose a computational job into tasks of various priorities. Next, we apply linear error correction coding to produce subtasks that are assigned to different processors. The decomposition used has a big effect on the type of anytime performance we attain. We study this scheme in a general framework in terms of the expected cost of the approximate solution. We further explore the approach in the context of vector-matrix multiplication. The proposed construction is numerically studied and, in comparison to previous work, demonstrates a significant improvement in the accuracy/latency trade-off.

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

In this project we re-examined 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 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. In addition to continued development of theoretical results, we have recently shown that the algorithms are compatible with hardware. We have completed a full-scale implementation in a field-programmable gate array (FPGA). This hardware implementation produced a number of interesting innovations needed to synthesize efficiently in hardware the computational primitives of ADMM-LP decoding.

Re-architecting Last-Level Caches for Low-Voltage Operation of Video Encoders

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. We are developing an application-specific implementation of this design in partnership with Advanced Micro Devices (AMD) customized to video encoding operations. The work with AMD is supported by NSERC.

● **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.EECCG.TORONTO.EDU/~ENRIGHT/

■ **Approximation Computing**

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 techniques 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. We have also proposed two novel cache architectures to reduce redundant data through approximation.

■ **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 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 multi-core 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.

■ 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 turn-around times.

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 μm 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 μm 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

WWW.COMM.TORONTO.EDU/~DIMITRIS/RESEARCH/BUSNET.PDF

▲◆◆ Biometrics User-Centric Sensor Networks (BUSNET)

We propose to develop an integrated security architecture to effectively and efficiently secure and protect sensitive information and data within the domain of a care enterprise, such as wireless health care and home care applications and services. Our proposal addresses the need for secure communication and authentication of personal information, which also requires enhanced privacy and confidentiality. The proposed security architecture, "Biometrics User-Centric Secure Networks (BUSNet)" will implement novel biometrics-based security solutions and technologies that can be effectively integrated into a plethora of wireless infrastructures. Specifically, this research initiative will be examining issues and developing solutions for processing of biometrics signals, biometrics registration and authentication and biometrics key generation and management, as well as biometrics-based data authentication. Implementations of the proposed architecture using specific realizations of suitable wireless Body Area Network (BAN) configurations will be also developed, examined and analyzed in collaboration with our industrial partners. The proposed security framework constitutes an effective mechanism that integrates a novel security architecture and a cost-effective networking configuration to develop a realistic, feasible and cost-effective solution for secure transmission of sensitive information, one of the most fundamental requirements of current and future health and home care services.

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 have fuelled 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 semi-conductor 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, CdTe 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

▲ Guided Light, Tightly Packed: Novel Concepts, Components and Applications

A multi-institutional and international (Canada-German) training program is proposed around advances in controlling the flow of light. Three collaborating Canadian institutions (U of T, INRS, ULaval) are partnering with the German institutes of Friedrich

Schiller University Jena (FSU) and Fraunhofer Institute of Applied Optics and Precision Engineering (IOF). The Canadian institutions represent the leading optical research-training centres in Canada, which in turn support a fledgling Canadian industry and service sector in optics. The German institutions define the central hub for optics education, training and industry in Germany, opening doors for our trainees to connect forefront science with manufacturing innovation. The training rationale seeks to lift PhD students out of research 'silos' that naturally are shaped by the limited resources and expertise found within individual labs, and enable higher impact research objectives that can harness the fuller spectrum of research expertise, advanced tools and resources provided by our multi-institutional team. Our CREATE philosophy is to drive this research with strong exposure to current industry practices, and participation in market-driven research activities through interaction with small and large Canadian industry and government partners. The key strategy with our German collaborators is to engage Canadian trainees early in their studies inside a world-leading place for optics education and manufacturing - Jena - The City of Lights. Trainees will participate in the highly acclaimed Fraunhofer model for technological innovation (IOF) and engage in German industry internships. The program will expose trainees to a much broader research experience, where research tools and resources will appear unlimited, where ideas can be tested and shaped amongst world leading experts with exceptional breadth in optics. The pathway ensures trainees will not only become future world leaders in optical research, but also be moulded for creativity, innovation, technology transfer and commercialization of new optical devices, advanced manufacturing processes, and more functional products with promising benefits to multiple facets of Canada's manufacturing, communication, health and security sectors.

▲ Laser-Fabricated 3D Optical Interconnects for Multi-core Fiber to Silicon Chip

High-density three-dimensional (3D) optical circuits will be designed, fabricated, integrated, packaged and tested at the University of Toronto (Prof. Herman Group) by combining silicon-photonics (Si-P) chips with 3D glass chip waveguides (WG) and multi-core fibre (MCF) input/output (I/O) ports. Femtosecond laser direct writing in optical glasses will enable formation of three-dimensional optical circuits as well as nano-structures to guide 3D chemical etching to micro-structure the glass block. Both the University of Toronto and Huawei will provide Si-P chips to test and verify various design approaches in chip-to-chip interconnection. The research will examine the compatibility of edge coupling (EC), grating coupling (GC) and adiabatic coupling (AC) for high-density, low-insertion packaging between high-index chips (Si, Si/InP, etc.) and MCFs.

▲ Quantized Structuring of Transparent Film and Plates with Ultrafast Laser Interference and Filamentation

Nano-structuring of materials is the basis of rapidly expanding application areas where ultrashort lasers are becoming a tool of choice with benefits of non-contact processing that precisely shape and pattern with very little residual damage. Short-pulsed lasers further underpin a unique capability for 3D structuring inside transparent glasses. In this direction, a new research project is proposed around our recent discovery

of a novel laser interferometric interaction discovered in silicon-nitride film that underlies a new means for laser material nanostructuring. Here, optical interference fringes are formed parallel with the surface that for the first time can create thin laser plasma disks and form lateral modification zones, nanocavities, thin blisters and quantized ejection from interaction zones substantially narrower than is possible even when focusing with high resolution lenses. This new means of control has not been previously anticipated, and is especially attractive in high temperature dielectric media like optical films that find widespread use in processing microelectronic, optical, lab-on-chip, photonic, MEMs, and photovoltaic devices. The research program will improve our fundamental understanding of how the interference and other nonlinear processes such as self-channelling can play together in transparent films and plates and serve as a basis for developing new nano-optical devices and forming novel nanostructured glasses. Reproducible manufacturing methods are pursued with two industry partners in directions that include bio-inspired concepts to strengthen windows for low-weight sunroofs, invent new types of anti-reflection surfaces and compact microlens arrays, and provide three-dimensional volume texturing for improving adhesion, wetting, or biosensing. Alternatively, ultrathin membranes offer new ways to package materials or sense mechanical properties. Lastly, quantum laser interaction offers the opportunity for generating complex multi-layered nanofluidic networks inside thin film that may permit flexible lab-in-film devices to be integrated with smart phones, cameras or microelectronic chips.

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.

● Efficient Numerical Analysis Techniques for Advanced Electromagnetic Surfaces

Next-generation antennas are being realized from novel technologies based on advanced electromagnetic surfaces, such as reflectarrays, meta-surfaces, and frequency-selective surfaces. These surfaces are very large and yet contain very fine features, making them challenging to design and analyze in conventional Computer-Aided Design (CAD) tools. This project is investigating computationally efficient yet accurate techniques for analyzing such surfaces. The approach is based on a macro-modelling-based method for encapsulating fine features with a coarse-grained model, to yield major savings in computation time and resources.

● 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 a single 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 next generation broadband telecommunication satellites.

● Multi-band Frequency Selective Surfaces for Manipulating Radio Wave Propagation in Buildings

Frequency selective surfaces are electromagnetic surfaces (FSSs) that act as filters. They are transparent to some frequencies and reflective to others. This project is exploring the application of FSSs for realizing building materials in order to control the propagation of radio waves around buildings. In this way, radio spectrum can be more effectively utilized by confining wireless networks to certain areas within buildings while enhancing coverage in others. The proposed surfaces will assist greatly in managing spectrum in future radio systems.

● Reconfigurable Spatially Fed Arrays

This project is developing 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 (particularly satcom-on-the-move), radar, and remote-sensing systems.

Iizuka, Keigo

WWW.KEIGO-IIZUKA.COM

▲◆ Noble Laparoscopes

Laparoscopic surgery is a less intrusive procedure than open surgery. The post-operative pain is less, the infection rate is lower, recovery is quicker and cosmetic outcome is better. As a natural extension of the earlier invention of Divcam (distance mapping camera), two applications of Divcam have been extended into the field of laparoscopy. (1) Application of Divcam to profilom. Images from ordinary laparoscopes and endoscopes are two dimensional (2D), meaning the surgeon's depth perception is hindered. The proposed method supplements the 2D image with an image of the depth profile of the surface. The depth profile is obtained in real time without surface contact. The profilometer uses the same principle for acquiring distance information as Divcam. The profilometer was added to an ordinary laparoscope with minimal increase in the weight and diameter of the shaft. With the profilometer added, there was a significant improvement in the ability to detect minute protrusions. It can detect a protrusion as small as 0.5 mm high that would not be recognized by an

ordinary laparoscope. K. Iizuka, *Applied Optics*, vol. 52, no. 19, pp. 4663–4671, July 2013. (2) Application of Divcam to omni-focus laparoscope. The depth of focus of a laparoscope, even the latest model, is limited to 10 cm. The proposed omni-focus laparoscope has stretched its depth of focus to 160 cm. Omni-focus means that everywhere, not just designated spot in the video is in focus. This property is significant in the case of a laparoscope. The entire scene, not just a specific spot, must be focused during the surgery. The proposed omni-focus laparoscope is equipped with an array of colour video cameras, each focused at a different distance. The distance information from the laparoscopic profilometer is used to select individual pixels from the ensemble of outputs from this array of colour video cameras and generates the final single “omni-focused” laparoscopic video image. The experimental results clearly demonstrate how significantly the quality of the image of the omni-focus laparoscope differs from that of a conventional laparoscope. K. Iizuka, *Applied Optics* vol. 52, no. 33, pp 7904-7911, 2013.

▲◆ 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) utility-grade urban, rural, and remote AC microgrids subject to the high depth of penetration of renewable and alternative energy resources, and (2) DC microgrids.

■ Electric Power Grid Modernization

The legacy AC power system exhibits major deficiencies in terms of technical, economical and environmental requirements for the 21st century due to (1) changing mix of types and characteristics of generation, (2) the need for a higher degree of resiliency and reliability and (3) the increasing demand for stakeholders’ participation in the electricity market. This research program proposes grid modernization based on utilization of the main and supplementary controls of overlay HVDC network, in coordination with the controllers of the host AC system and wind/solar power plants, and relies on pervasive use of ICTs and wide area measurements to achieve the objectives.

■◆ 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

■◆ 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.

● ■ ◆ Middleware Systems Research

We design state of the art middleware systems to simplify the design, development and management of complex distributed applications. We conduct research on middleware and distributed systems at the University of Toronto. Our current research is mainly focused on publish/subscribe, content-based routing, and event processing. We see middleware as the services and abstractions that facilitate the design, development, integration, and deployment of distributed applications and systems in heterogeneous networked environments.

● ■ ◆ 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.ECG.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

WWW.IBBME.UTORONTO.CA/FACULTY/EMERITI/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 voltages on the skin and relates them to possible conductivity distributions in the body. Unfortunately EIT is a very ill-posed problem and consequentially has poor resolution at depth. This severely limits its effectiveness. The novelty of CDI and CDII is that MRI can be used to measure the magnetic fields arising from internal electric current density and 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

▲ ◆ Hyperspectral Sensing and Imaging

The ability to control visible and infrared light vis-à-vis guidance and concentrated light localization below the diffraction limit can usher in remarkable sensing and imaging modalities. The objective of this project is to undertake in-depth research and development of novel nanoscale gratings that enable hyperspectral light-trapping within graded gratings using adiabatically coupled plasmonic waveguide resonators. Specifically, the goal is to explore various nanoscale designs and integrations that will permit unprecedented sensing capabilities at the molecular scale. Studies to date indicate

that this technology can create a powerful adjunct for biomedical applications wherein an optical device with micron-scale footprint would be capable of mapping the chemical composition of given samples with sub-cellular resolution at extremely high speeds.

▲ Novel Nanoparticle Material Systems

The ability to synthesize binary, ternary and quaternary nanoparticles using a non-solution based approach opens the potential of developing novel material systems with unique properties. The objective of this project is to undertake the synthesis and characterization of novel nanoparticles developed using the Advanced Materials-Integration Synthesis (AIMS) Facility. The AIMS facility permits the combination of both physical and chemical deposition techniques where the former includes size selective synthesis of nanoparticles.

▲ Photonic Materials for Next-Generation Optical Coatings

Novel photonic materials and integrations provide the potential of developing next generation optical coatings for a variety of energy efficiency applications. As an example, the buildings sector consumes approximately 40% of the total energy consumed in the U.S. (followed by the industry and transportation sectors). Of this 70% of the energy consumption in buildings is for HVAC (heating ventilation and cooling) and lighting at approximately 50% and 20%, respectively. In this context, windows have long been recognized as a major cause of radiative heat loss and solar optical energy gain in buildings. Integration of economically effective measures necessitates advanced research on novel, high-efficiency, next-generation photonically selective energy coatings. The aim of this project is to develop novel photonic materials that embody the requirements of high performance, sustainable material system, and economic viability. Specifically, three optical applications are targeted. One, next generation spectrally selective coatings which are inherently stable (optically and thermally), tunable vis-à-vis visible, near infrared and mid infrared, and provide integrated active heating functionality; two, radiative cooling photonic devices that avail the large transparent mid infrared bandwidth of the atmosphere in relation to outer space; and three, novel strain-tunable active/adaptive photonic devices for a variety of applications.

▲ Ultra-Thin Silicon Micro-Opto-electronic Devices: Using Versatile Low-Temperature Fabrication Techniques

The ubiquity of silicon extends into both electronic and optical devices. Within the framework of energy generation, silicon solar cells continue to dominate the field. While confronted by new emerging solar photovoltaic materials and devices, silicon continues to rise to the challenge with scientific advances and technological enhancements. Perhaps the next holy grail is ultra-thin silicon micro-opto-electronic devices with versatility in fabrication, particularly at low temperatures, while preserving or promoting the highest level of performance. From an energy generation perspective, silicon PV would benefit from advances in ultra-thin (≤ 10 microns) high-efficiency silicon PV—rendering the benefits of pliable, lightweight, durable, economical and ubiquitous photovoltaic electricity. Further, it is conceivable that

micro-opto-electronic circuits could also benefit from this platform with an eye to “more than Moore.” The objective of this project is to research, develop and integrate a set of thin-film and allied versatile and low-temperature fabrication techniques that will lead to the demonstration of high-performance micro-opto-electronic devices based on ultra-thin silicon foils.

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 non-coherent 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

● ■ ■ **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 Professor Berj Bardakjian of ECE and IBBME at the University of Toronto.

● ■ ■ **A Real-Time Federated Co-simulator for Cyber Security Analysis of Microgrid Systems**

Effective modeling 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 fragility of the Canadian grid to incidental and natural disruptions; given the increasing dependency 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. They 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 Engage Plus project builds upon an existing Engage-based collaboration amongst Professor Deepa Kundur, University of Toronto; the Institut de recherche d’Hydro-Québec (IREQ); and Opal-RT Technologies, Inc. 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 microgrid and cyber-security projects within IREQ by providing a framework to test communication and control strategies. Furthermore, the real-time 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 Security and Privacy Framework for a Cloud-Based Green Button Standard**

The smart-grid design mantra can be described, in part, as “knowledge is power.” The application of information technology to the power grid aims to provide more efficient and reliable operation and facilitate new consumer-centric applications through greater situational awareness. Advanced metering infrastructure (AMI) provides unprecedented opportunities for informed decision-making on energy consumption and home energy management. Cloud-based services provide a convenient, cost-effective and scalable platform upon which to host emerging AMI applications, but

practical considerations arise from this greater dependence on information, namely those related to security and privacy. To address these challenges, many utilities are adopting Green Button, an ongoing North America-wide initiative to provide utility customers with secure access to their energy usage information while protecting privacy. This Engage proposal, through a three-way collaboration amongst the PI's research group at the University of Toronto, London Hydro and First North, Inc., aims to address security and privacy of consumer AMI data as they relate to the emerging Green Button standard in the context of a cloud-based environment. The following tasks are addressed: (1) identify critical privacy challenges in a utility sector multi-tenant public cloud-based environment; (2) assess the Green Button standard to determine necessary extensions for multi-tenant public cloud privacy; (3) analyze a specific use case focused on secure federated aggregation of de-identified AMI data sets; and (4) develop a cloud-based Green Button App to test the use case in the context of social benchmarking. This work has the potential to have significant impact on the evolution of the Green Button standard, enhancing Canadian leadership in this international effort. Moreover, this proposal is strongly aligned with Canada's Cyber Security Strategy and the National Strategy and Action Plan for Critical Infrastructure. The health and happiness of a society is strongly coupled to its national security. The proposed project supports a strengthened society in which Canadian academia and industry can collaboratively participate in its defense and security.

●▲■ 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 defence-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/PEOPLE/PROFS/KWONG

▲ Security and Fault Tolerance of Cyber-Physical Systems

Critical infrastructures such as power grids, water supply systems, and traffic networks are referred to as cyber-physical systems. One key requirement for such systems is security from malicious attacks, often in the form of cyber-attacks aimed at compromising supervisory control and data acquisition systems. The second key requirement is fault tolerance: the ability of the control system to continue functioning despite component failures. Component failures can occur as a consequence of malicious actions by an attacker, but can also occur as a result of hardware breakdown or unforeseen operating conditions. Most approaches to combat cyber-attacks focus on prevention through firewalls. Ubiquitous connectivity and stealth technologies have rendered preventing access and detecting intrusion increasingly difficult. Traditional approaches to fault-tolerant control design are also inadequate as they often rely on first carrying out accurate fault diagnosis, followed by control reconfiguration. The system may be damaged or irrecoverable while waiting for diagnostic information. The proposed research aims to improve the security and fault-tolerance of control systems central to cyber-physical systems. The two main objectives of the proposed research program are to develop a novel methodology to detect malicious intruders and prevent them from inflicting damage to the control system, and to maintain functionality and safe operation of control systems when components fail. The first objective focuses on security at the system level. Our approach will make use of the theory of supervisory control and diagnosis for discrete event systems to perform intrusion detection and to counteract malicious behaviour. The second objective focuses on resilience of the control system at the physical level when actuators and/or sensors fail. The novelty of our approach is that we focus on fault-tolerant control design, which does not require accurate diagnostic knowledge. We propose to study fault-tolerant control as an integrated problem of diagnosis and control reconfiguration, emphasizing the interaction of these two modules and exploiting available analytical redundancies. Protecting cyber-physical systems is of great national interest. The proposed research can provide significant improvements on the security and resilience of cyber-physical systems.

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
PORTAL.CVST.CA

●▲◆■ Connected Vehicles and Smart Transportation

The 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 is enhanced by analytics to enable users and applications 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 portal.cvst.ca, and one showing New York City and San Francisco traffic at <http://usportal.cvstproject.com/>.

●■ 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 datacenters 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.

WWW.SAVINETWORK.CA

●▲◆■ Smart Applications on Virtual Infrastructures (SAVI)

The SAVI project on Smart Applications on Virtual Infrastructures (SAVI), is the outgrowth of a partnership between Canadian industry, academia, government, education research networks and high-performance computing centres. SAVI has developed a virtualized converged computing and communications infrastructure to 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 a multitier cloud that leverages cloud computing and software-defined networking. 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 and an additional node operates in Korea.

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
IQUA.ECE.TORONTO.EDU/SPOTLIGHTS/DATACENTER.HTML

■ Bandwidth Allocation in Datacenter Networks

Web service providers like Google and Facebook have built large scale datacenters to host many computationally intensive applications, ranging from PageRank to machine learning. In order to efficiently process a large volume of data, these applications typically embrace data parallel frameworks, such as MapReduce. It is commonly accepted that bandwidth in datacenters should be shared in a fair manner. We have proposed the notion of performance-centric fairness that is specifically customized for bandwidth allocation among data parallel applications in private datacenters.

IQUA.ECE.TORONTO.EDU/SPOTLIGHTS/SCHEDULING.HTML

■ Fair Scheduling in Cloud Datacenters with Multiple Resource Types

In the age of big data, it has been the norm for cloud datacenters to run data analytic applications on a large scale. Yet, as multiple applications share resources in these datacenters, it is important to design scheduling disciplines so datacenter resources are shared in a fair and efficient manner. We have designed a new class of scheduling disciplines specifically for sharing multiple resource types in cloud datacenters. This new design allocates resources to applications by scheduling their computing tasks onto datacenter nodes. I will then focus on the problem that an increasing number of datacenter jobs specify placement constraints and can only run on a particular class of machines that meet specific hardware/software requirements (e.g., GPUs). Our recent research shows that directly extending existing policies to constrained jobs either compromises isolation guarantees or allows users to gain more resources by deceiving the scheduler. It remains unclear how multi-resource fair sharing is to be defined and achieved in the presence of placement constraints. We have designed a new sharing policy, called task share fairness (TSF), to provide provable isolation guarantees and to be strategy-proof against gaming the allocation policy. Challenges remain in real-world implementations.

■ On Data Parallelism of Erasure Coding in Distributed Storage Systems

Employed in various distributed storage systems, erasure coding has demonstrated its advantages of low storage overhead and high failure tolerance. Typically in an erasure-coded distributed storage system, systematic MDS codes are chosen since the optimal storage overhead can be achieved and meanwhile data can be read directly without decoding operations. However, data parallelism of existing MDS codes is limited, because we can only read data from some specific servers in parallel without decoding operations. In this project, we wish to propose new coding techniques that are designed to allow data to be read from an arbitrary number of servers in parallel without decoding, while preserving the optimal storage overhead of MDS codes.

■ Privacy-Preserving Inference in Crowdsourcing Systems

Machine learning has widely been used in crowdsourcing systems to analyze the behaviour of their mobile users. However, it naturally raises privacy concerns, as personal data need to be collected and analyzed in the cloud, and results need to be sent back to the users to improve their local estimates. In this project, we focus on the use of a specific type of learning algorithm, called maximum a posteriori (MAP) inference, in crowdsourcing systems, and use a crowdsourced localization system as an example. With MAP inference, the accuracy of private estimates of each user in a crowdsourcing system may be improved by analyzing estimates of other users. Naturally, the privacy of such estimates from each user needs to be preserved. Within the general framework of differential privacy, we show how private user data can be perturbed while preserving the distance between such data, so that data privacy can be preserved while producing statistically accurate learning results. Our proposed work

is not limited to localization, and can be applied to other learning algorithms that use Euclidean distances between data points as their metrics.

■ Scheduling Jobs across Geographically Distributed Datacenters

It has become routine for large volumes of data to be generated, stored, and processed across geographically distributed datacenters. To run a single data analytic job on such geographically distributed data, recent research proposed to distribute its tasks across datacenters, considering both data locality and network bandwidth across datacenters. However, a problem remains where multiple data analytic jobs need to fairly share the resources at these geo-distributed datacenters. In this project, we focus on the problem of assigning tasks belonging to multiple jobs across datacenters, with the specific objective of achieving fairness across jobs sharing these datacenters, in terms of their job completion times.

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 enables mobile terminals to dynamically and autonomously adapt to changing environmental conditions

and reduces their energy consumption. There are many examples in nature of dynamic fitting of performance to changing boundary conditions, since it represents the best way to achieve maximum efficiency in highly complex systems. Almost all ecosystems are based on this principle, which allows them to evolve while minimizing energy dissipation.

▲ Ultra-Low-Power Transceivers for Wireless Sensor Networks

Wireless communication represents one of the most important revolutions of the last century. Although initially based only on star-mesh networks (e.g., cellular), at the end of the 1990s some wireless systems started to also adopt peer-to-peer (P2P) architectures, Wireless Sensor Networks (WSNs) being a prime example. These systems do not require base stations since they are formed by autonomous short-range wireless nodes. All these nodes monitor and control the environment defining the working area by their spatial distribution. Since the high density of units makes the system more flexible and relaxes the sensitivity of the single receiver, in ZigBee network performance is exchanged with the possibility of enabling long-lasting and cheap devices. Unfortunately the target of a large-scale diffusion of WSNs was partially missed due to difficulty in realizing both long-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 μ W 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.

▲ CMOS Circuits and Systems for Broadband Wireless Communication

This project will explore circuits and systems to reduce the cost and power consumption of wireless communication equipment. Low-cost flexible RF transceivers achieving low average power consumption with few, if any, external components are sought. Specifically, wireless receiver electronics that are resilient to unwanted interfering signals, and transmitters with high spectral purity will be researched. Circuits and systems amenable to implementation in CMOS will permit their integration alongside other digital signal processing with attendant reductions in cost and power consumption.

Lo, Hoi-Kwong

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, in 2012,

we (Lo, Curty and Qi) 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. MDI-QKD has attracted a lot of attention in the scientific and engineering community. We will be building a quantum communication network based on our MDI-QKD proposal and demonstrating 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. We are also investigating the applications of quantum communication in Smart Grids and also reconfigurable free-space networks. With the recent launch of a Chinese quantum satellite and the construction of quantum networks all over the world (Japan, China and the USA), our work is timely and important.

● 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 communication 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) aircraft, 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

METAVISION.COM

■ ◆ Augmented Reality Eyeglass

The Digital Eye Glass provides real-time digitization of everything you see and experience, and real-time computational mediation, augmentation, or the like. See publications in <http://www.eyetap.org/publications/> as well as the website <http://metavision.com> for some examples. Our metavision eyeglass is also used for teaching and research purposes. If you would like to join this project, please do one of the instructables on augmented reality: <http://www.instructables.com/member/SteveMann/>

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/

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.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. HDR (High Dynamic Range) overlays augment dark regions of the scene and diminish bright areas, and makes overlaid content easier to read. See www.eyetap.org

WWW.WEARCAM.ORG/COMPARAM

◆ Comparametric Equations and High Dynamic Range (HDR) Imaging

High Dynamic Range (HDR) imaging was invented in our lab, and is now widely used, having 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. Please get our attention by doing the Instructable, uploading an image, and clicking "I made it" at: <http://www.instructables.com/id/HDR-Eye-Glass-From-Cyborg-Welding-Helmets-to-Wearab/>

GLOGGER.MOBI

◆ Lifelogging: 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 lifelogging 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.

WWW.SPLASHTONES.COM

▲ ● 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 See also our website, <http://splashtones.com>

WEARCAM.ORG/PAR/

■ ● Phenomenological Augmented Reality

This is the original wearable AR (Augmented Reality) invented by S. Mann in 1978, nearly 42 years ago, to see radio waves and sound waves. Presently we are implementing this through miniaturized wearable computer systems that overlay physical-world sensory content on reality. This can be used for a wide range of applications from teaching and research to everyday life. For more on our AR eyeglass, see also <http://metavision.com>

WEARCAM.ORG/MANNFIT.PDF

▲ ● Physical Fitness through Integral Kinesiology

Traditional physical fitness is based on "kinematics," which is the study of distance and its time derivatives. We have a new invention, integral kinesiology, which looks also at distance and its time integrals. So rather than only considering distance, speed, acceleration, jerk, jounce, etc., we also consider absement. <https://en.wikipedia.org/wiki/Absement> This represents a new breakthrough in physical fitness. Absement-based training devices and apps develop the capacity to have simultaneous strength and dexterity.

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 hydraulophononic 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 electro-magnetic 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.EECS.TORONTO.EDU/~MOSHOVOS

Machine Learning Accelerators

Deep learning is the state-of-the-art approach for many tasks such as automatically identifying image content or interpreting speech. We are developing hardware engines that offer best-in-class performance and energy efficiency for such machine learning algorithms targeted at mobile and datacentre applications.

Co-Designing Hardware and Software for Analytics

We are developing systems that can deliver real-time transactions and data analytics. By rethinking software and hardware we are exposing bottlenecks in conventional implementations, boosting performance and energy efficiency and for that utility and application sophistication. This is a joint project with Professor Nick Koudas, of the Department of Computer Science.

Digital Imaging Accelerators

We are characterizing digital imaging applications and developing specialized yet programmable hardware that boosts performance and energy efficiency.

Nachman, Adrian

WWW.CURRENTDENSITYIMAGING.ORG

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 Multi-contact 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.EECS.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. We also have ongoing collaborative projects with our industrial research partners to develop discrete and integrated power devices, including power MOSFETs, IGBTs, and GaN power HEMTs. In recent years, we have focused on the design and implementation of VLSI power management circuits with special focus on integrated DC-DC converters with digital control. We have demonstrated the world's first integrated DC-DC converter with dynamically adjustable power transistor sizing to optimize the power conversion efficiency at ISPSD'06. In 2010, we demonstrated a superjunction power FinFET at IEDM 2010 for the first time. This is exciting work toward enabling the next-generation FINFET CMOS fabrication technology to be compatible with the implementation of smart power ICs. Recently, we have also demonstrated an all-digital on-chip temperature sensors for thermal management applications, gate driver circuits with dynamically adjustable driving strength for EMI suppression and efficiency enhancement, and dead-time control circuits. Currently, our group is working on the integration of the controller and gate driver circuits with current sensing for various power output stages. We also have ongoing work on the design and fabrication of GaN power transistors as well as their applications in power converters and class-D power amplifiers. Finally, we have just started an exciting project with Dana Canada on liquid cooled IGBT modules for electric vehicle applications.

Pavel, Lacra

WWW.CONTROL.TORONTO.EDU/~PAVEL/

▲ 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 readjusting 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.

WWW.COMM.UTORONTO.CA/~KOSTAS

◆ Digital Pathology: Fast and Reliable Image Analysis Software

In collaboration with Huron Digital Pathology Inc., located in Waterloo, Ontario, the University of Toronto team aims to develop an innovative image processing algorithmic framework for digital pathology (DP) scanners in tissue imaging applications. Using tools from signal and image processing, computer vision, and machine learning, we applicants plan to research, develop, and test new, cutting-edge algorithms to acquire, analyze, and store histo-pathological images. The proposed framework will greatly decrease the time and computational complexity of the image acquisition and scanning process while improving accuracy and reliability, and enhancing the end-user experience.

WWW.COMM.UTORONTO.CA/~KOSTAS

◆ eDREAMs: enhancing Driver inteRaction with digital mEdiA through cognitive Monitoring

Recent advances in digital media systems, have resulted in a growing interest in the use of in-vehicle digital interactive systems over the past decade, which in turn has led to serious concerns regarding drivers' safety. The E-DREAM project aims to develop smart interactive devices that take into account the driver's cognitive state as well as the vehicle's conditions to (i) quantify and monitor the driver's workload and arousal level, and (ii) provide the driver with adaptive feedbacks that will mitigate the risks associated with the particular level of workload the driver is experiencing. E-DREAM ultimately aims to pave the way for developing an advanced human–computer interaction system that (a) provides proactive and seamless support in daily life situations, (b) utilizes smart sensors, multi-modality, and data aggregation to enable services that enhance the way to use technology, and (c) models user behaviour, detects changes in context and adapts accordingly to achieve pro-activeness.

WWW.COMM.UTORONTO.CA/~KOSTAS

● ● **Opportunistic Mesh Networks for Smart Home Applications**

This project contributes to the advancement of the state of the art in smart home infrastructure by addressing challenges often hampering the successful implementations of ambient intelligence solutions. The primary focus of this initiative is to provide seamless integration of sensors and actuators by utilizing an OMesh-based wireless mesh network. The project aims to develop a distributed information system which will collect data from low cost sensors, utilize artificial intelligence and multimedia algorithms to extract information, and drive a servicing robotic assistant. The output of this project will be a smart home prototype configuration, and a software module which, along with the hardware set up, will allow for replication and expansion of the results.

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.

▲ ● **Photonics for Neuroscience and Brain Health**

We are designing photonic devices and optical systems for brain activity mapping and neurosurgery guidance. 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 optical 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.

● **Nonlinear Optical Devices: Ultrafast Switching and Frequency Conversion**

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

● **Quantum 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/RE-SEARCH-PROJECTS.HTML

■ 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.

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

■ 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.

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

■ Smartphone-Based Diagnosis of Social Anxiety

Social anxiety disorder is one of the most prevalent mental health disorders, yet it is treatable. The goal of this research project is to develop a reliable and objective measure of social anxiety using data passively collected from a patient's smartphone. The standard approach for subjects seeking treatment for social anxiety is to receive the care of a clinician. The clinician extracts knowledge of the patient's activities and emotions, among many other attributes, from the patient themselves, either verbally or through written self-report scales and questionnaires. Subjects can be unwilling or unable to give accurate accounts of their life in between visits to clinicians, and questionnaire fatigue may also arise after completing a long series of self-report questionnaires. A measure of social anxiety that is objective, consistent, and requires no input from the patient would address this problem, and that is the goal of this research project.

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 CO₂ 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 CO₂ 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**

Wireless power transfer (WPT) is an area where significant research and development efforts are being undertaken, for applications including the wireless charging of electric vehicles, biomedical and communications devices. Recently, WPT systems with multiple transmitters were investigated, both theoretically and experimentally. Multiple transmitters provide more degrees of freedom of the primary field or current distribution and, therefore, promise the possibility of enhanced transfer efficiency as compared to single transmitter systems. We have shown that WPT systems with multiple transmitters can be designed and optimized semi-analytically, using convex optimization. Convex optimization leads to closed-form expressions on maximum achievable efficiency and optimal loading of multi-transmitter WPT systems, illuminating their physics and providing valuable insights into their optimal operation. Moreover, the robustness of the optimal loading conditions can be evaluated with regards to the tolerances that the loading elements are subject to. With this design framework in place, one can envision wireless power transfer networks, in analogy to wireless communication networks, whereby multiple distributed transmitters would be coordinated to charge multiple devices in a cooperative fashion. We investigate the power transfer efficiency bounds of such fully multiple-input multiple-output (MIMO)-WPT systems and their practical realization.

▲ **An Optimization-Driven Methodology for the Design of Optically Transparent Sensors and Antennas**

Touch sensors are indispensable components of the user interface of personal electronic devices such as smartphones, tablet computers and wearable health monitors. These sensors are ordinary circuits, whose usual copper wires are replaced by optically transparent (practically invisible) conductors such as Indium Tin Oxide (ITO). The use of optically transparent conductors, which suffer from higher losses compared to copper, makes the design of these circuits a challenging task. Circuit designers have by and large dealt with this challenge on a trial-and-error basis, using a combination of circuit models and electromagnetic simulations to derive topologies that trade performance for optical transparency. With the market for personal and wearable electronic devices as competitive as ever, new touch sensors are expected to support high-rate sensing of multi-tasking gestures with low latency. These stricter specifications test the limits of current design tools, motivating further research on the optimal performance bounds of transparent circuits and the topologies that meet those bounds under given fabrication constraints. Leveraging advances in convex optimization of electromagnetic structures and uncertainty quantification techniques, we are deriving a dictionary of building blocks for transparent circuits made with standard (ITO) and emerging (silver nanowires, graphene) materials, just as one/two-dimensional microstrip lines are for copper-based ones. In addition to touch sensors, applications of the EOTC concept to the design of transparent antennas and multi-functional touch-input/antenna surfaces are being explored.

▲ 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

●▲◆ 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 acts like a transmission line and as such exhibits 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 crosstalk and timing jitter so as to reduce the bit error rate (BER) of the communication link, while using less than a few mW per Gb/s operation. In the past few years, we have been able to contribute to this research through the design of ADC-based receivers that allow for extensive signal equalization in the digital domain. We have also contributed to the design of non-data-aided equalization techniques and 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.

●■ New Machines for Machine Learning

Machine learning has received much research attention in recent years, but the focus has always been on the “learning” part, with not much attention given to the “machine” side. So far, most of the machine learning algorithms are written for the typical von-Neumann-based computers. However, to be able to solve more complex optimization problems, these computers cannot offer the required speed, nor are they expected to in future due to the end of Moore’s law. In this project, we strive to look at other, non-conventional, computer architectures, and their implementations in CMOS, in order to attain speedups a few orders of magnitude larger than those available today.

Sousa, Elvino

◆ Autonomous Infrastructure Wireless Networks

The research focuses on our vision for 5G 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. This research complements our Two-Tier vision for 5G.

◆ Data Mining from Point Processes

This research focuses on the analysis of certain point processes with applications in data mining in the context of IoT and biomedical engineering.

◆ Internet of Things in the Smart Home and Healthcare Environments

This research focuses on the development of sensors and analytics for application in smart home environments and healthcare.

◆ Two-Tier 5G Wireless Networks

This project addresses a novel architecture for public cellular wireless networks. This architecture constitutes a key vision for fifth-generation wireless networks (5G). The goal of this architecture is to provide network capacities that are significantly higher than the 4G capacities and at the same time are amenable to the expected increases in traffic, not only from a higher density of users of smartphones with greater traffic demands than the current levels, but also capacity increases resulting from emerging applications such as the Internet of Things (IoT). The proposed architecture consists of a Two-Tier system structure where the network consists of traditional base stations and intermediary, or secondary, nodes that act as traffic aggregators. The intermediary nodes are deployed strategically in locations where traffic concentration is required, including homes, offices, vehicles, etc. Such traffic aggregators may aggregate traffic from various user terminals that are closely located or devices in machine-to-machine or IoT applications. The secondary node is designed using antennas with a relatively large number of elements and the transmission from the base stations to the secondary nodes constitutes a link that we may refer to as using the emerging massive MIMO technology. The proposed architecture will result in the links from the base stations to the secondary nodes having a very large capacity, due to the massive MIMO link. In order to design such links, novel channel estimation methodologies are required. We foresee an approach where these channel estimation methodologies make use of a database consisting of channel propagation information for different areas in a city relative to the system of base stations. The design of such a database will depend not only on topography information but also on time—and location-dependent traffic demand. In addition, other factors such as the mobility of people will come into play. Such a database will form a component of the future Smart City.

Stumm, Michael

WWW.ECG.TORONTO.EDU/~STUMM

■ Software Engineering: Continuous Deployment

A small number of companies have started to implement a new way to manage the software life-cycle with continuous deployment: developers produce software updates in small increments and push the updates out to production as soon as the developers think it is ready. The end result: continuous

deployment process substantially increases agility and responsiveness, and at the same time keep developer productivity and software quality high. We did an extensive study of software development and deployment over a period of eight years at two companies: OANDA and Facebook. For example, FB deploys software out to production thousands of times a day, yet the quality of the software has only been increasing. More recently, we have been studying the development and deployment of mobile software, and were able to draw similar conclusions.

■ System Software Performance Optimizations

Our primary research objective is to improve the performance of large, complex computer systems (e.g., distributed systems or multiprocessor/multicore systems). Our research approach is usually experimental: we implement systems exploiting novel technology, and we evaluate them experimentally using realistic workloads. Recently, for example, we developed a new way of doing operating system calls without requiring exceptions/traps that would allow software servers more efficiently. We were able to show that with this optimization, Linux Apache's throughput increased by over 100% without any modification to Apache, and MySQL's, Nginx's, and MemcacheD's throughput increased by 40%, 120%, and 35%, respectively. On two other projects, we worked extensively with Facebook on some of their back-end systems. Specifically, we developed SocialHash, which performance improves the assignment of objects to resources (e.g., assignment of HTTP requests to servers and storage sharding), and we helped optimize the performance of RocksDB, a FB-designed key-value store. SocialHash, resulted in a 25% cache miss rate reduction for servicing FB HTTP requests and the optimization to storage sharding resulted in a reduction in storage access response times by 50%. RocksDB optimizes for space amplification, which reduces the amount of disk space required by 50% while increasing DB read latencies by only a marginal amount.

Tate, Joseph (Zeb)

WWW.ECE.UTORONTO.CA/PEOPLE/TATE-J/

■ Energy Storage Control for Regulation and Frequency Support

With a higher penetration of solar and wind generation, there is a greater chance of large frequency excursions along with mismatches in the supply and demand of electricity. The proposed storage controller combines online system identification with model predictive control to arrest and recover from frequency excursions during major grid events (e.g., unexpected generator outages).

■ 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.ECE.UTORONTO.CA/PEOPLE/TAYLOR-J/

■ 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 like 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 at 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 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 molecular factors that select preferred biosensor conformations. Last, to design synthetic protein networks or model larger existing networks, I will develop a computational tool for simulating the spatial and temporal kinetics of protein signalling networks. The sum of this work will yield insights into protein sequences and their networks that will ultimately aid in developing therapies for human illnesses.

◆ Live Cell Imaging and Control of Caspase Kinetics Using Engineered Proteins

Over the past decade, members of the caspase family of proteases have been extensively studied for their critical role in apoptosis. The caspase family displays rich spatial and temporal kinetics in living cells, such as cascading activation and differential subcellular expression. While such characteristics confound many biosensor designs, they accentuate the strengths of fluorescent protein biosensors. By employing the principle of fluorescence resonance energy transfer (FRET), protein biosensors can be created to image the kinetics of caspase activation in living cells. Furthermore, we can control the exact moment that caspase activation occurs within the cell using an inhibitory protein of caspase that is engineered to be switchable on $[Ca^{2+}]$. 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.

Valaei, 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 converters for next-generation optoelectronic transceivers with multilevel modulation formats. The circuits will be fabricated in the world's most advanced 55-nm SiGe BiCMOS technology. 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 100 GHz and 40-GHz input bandwidth.

▲ Atomic-Level and 2D Crystal Electronic Devices

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

▲ 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

◆ Auditory Gap Detection and Clinical Implications

Auditory processing disorder affects nearly 5% of school-aged children as well as 70% of older adult population. Such individuals will have healthy peripheral hearing but deficits in the processing ability of auditory information in the central nervous system. One reason is due to the low temporal resolution of the auditory system which leads to difficulty in perceiving speech under noisy circumstances. Hence diagnosis is extremely important as it may affect a child's ability to develop normally, and may allow schools to better tailor their teaching to the students. For older adults, an early diagnosis may provide a means for a better quality of life. We are developing new electrophysiological methods that will allow for rapid detection and diagnosis of auditory processing disorder in collaboration with Vivosonic Inc., a world leading Canadian developer of medical devices for audiometric solutions.

◆ Neuromodulation for Motor Movements

We are interested in studying how the different regions of the brain communicate together to coordinate complex motor movements like a reaching task. There are over 100 billion neurons in the brain, but how do the different networks of neurons communicate to complete a complex task? We tackle this problem both experimentally, in collaboration with neurologists and surgeons at Toronto Western Hospital through studies on deep brain stimulation, and we also collaborate with mathematicians to develop models of synchronous neural activity which mimic the experimentally observed behaviour. We use Kuramoto oscillator models and study phase transitions in order to develop understanding of the underlying mechanism. Through furthering basic knowledge, we hope to pioneer new neuromodulation techniques which help those with restricted mobility to restore their capability for movement.

◆ Theoretical Modelling of Sensory Processing

We develop theoretical models to understanding the functioning of the senses (seeing, hearing, etc.). These models are based on a new approach involving statistical physics and information theory. The end result is a system of equations with very few adjustable parameters that is able to predict, with incomparable generality and universality, the neural response of sensory systems in all types of organisms and sensory modalities. This theory is currently used to guide the development of sensory prostheses (see entry under artificial eye).

◆ Visual Prosthetic Devices

In collaboration with ophthalmologists in Japan, we are designing the next generation of visual prosthetic devices to restore sight in those people with retinitis pigmentosa (loss of photoreceptor cells). This new method involves implanting a photosensitive dye layer into the retinal space and using this passive, biologically compatible method to induce electrical responses in the eye, thereby triggering “vision” and allowing the person to see. The photosensitive dye is superior to other existing prosthetic devices because it does not require external cameras or arrays to be implanted into the eye. Using theory developed at the University of Toronto, we help guide the development to ensure that what the person sees is faithful to the actual scene he or she is viewing.

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-NEUROLAB.COM

▲ 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.

◆ Mechanisms of Obstructive Sleep Apnea

Although there are several effective treatment options for patients with obstructive sleep apnea (OSA)—such as continuous positive applied pressure (CPAP) and hypoglossal nerve stimulation—many patients are either not suitable for a permanent nerve stimulation device or not compliant with wearing a mask during sleep. The primary goal of this research project is to better understand the physiological mechanisms that contribute to airway collapse and thereby develop improved methods of treating OSA.

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

■ JVM Performance Overhead on Big Data Analytics Systems

We are seeing a recent trend: many widely used, latency sensitive, big data analytics systems, such as Hadoop, Cassandra, Spark, etc., are built on top of Java Virtual Machine (JVM), despite debate on the overhead of doing so. This project aims to analyze the extent and causes of JVM performance overhead in the above mentioned systems.

Surprisingly, we have found recently that the warm-up overhead, i.e., class loading and interpretation of bytecode, is frequently the bottleneck. For example, even an I/O intensive, 1GB read on HDFS spends 33% of its execution time in JVM warm-up, and Spark queries spend an average of 21 seconds in warm-up. The findings on JVM warm-up overhead reveal a contradiction between the principle of parallelization, i.e., speeding up long running jobs by parallelizing them into short tasks, and amortizing JVM warm-up overhead through long tasks. We have built HotTub, a new JVM that amortizes the warm-up overhead over the lifetime of a cluster node instead of over a single job, by reusing a pool of already warm JVMs across multiple applications, to solve this problem. The speed-up is significant. For example, using HotTub results in up to 2.38X speedups for Spark queries, despite not adhering to the JVM specification in edge cases. This result will appear in OSDI 2016, the best conference in systems software.

■ **Non-intrusive Failure Diagnosis for Distributed Software Systems**

Real software systems inevitably experience failures. When such failures occur, it is critical to diagnose them so as to bring the systems back online. However, the challenge is also daunting: the most widely used systems are built using distributed systems, where hundreds of thousands of servers work together to process user requests. When a failure occurs, one has to reconstruct the complex execution of the processing of this request across these servers to diagnose the problem. Professor Yuan pioneers the area of non-intrusive failure diagnosis. Unlike intrusive approaches, his does not require any modifications to the target systems to collect additional data, which face serious deployment hurdles in practice as vendors are often reluctant to alter their production systems. Instead, his technologies can reconstruct the (failed) software execution entirely from the system's existing log output. His research has been published in the proceedings of the best conference in systems software (OSDI) multiple times, and has been licensed by Huawei Technologies.

WWW.ECG.TORONTO.EDU/FAILUREANALYSIS/

■ **Simple Testing Can Prevent Most Critical Failures in Distributed Software Systems**

Large, production-quality distributed software systems, as exemplified by Google.com, Facebook.com, Amazon.com, etc., still fail periodically, sometimes catastrophically where most or all users experience an outage or data loss. Conventional wisdom has it that these failures can only manifest themselves on large production clusters and are extremely difficult to prevent a priori, because these systems are designed to be fault tolerant and are well-tested. This project asks a simple question: why do these systems still fail and what can be done about them? By investigating 198 user-reported failures that occurred on production-quality distributed systems, we found that almost all (92%) of the catastrophic system failures are the result of incorrect handling of non-fatal errors, and, surprisingly, many of them are caused by trivial mistakes such as error handlers that are empty or that contain expressions like "FIXME" or "TODO" in the comments. We therefore developed a simple static checker, Aspirator, capable of locating trivial bugs in error handlers; it found over 200 new bugs that have been fixed by the developers. After appearing in the best conference in systems software

field (OSDI'14), this work has had an immense impact on the software industry. Google and HBase developers will be using Aspirator. The work has been the subject of many blog entries and thousands of Twitter tweets, and has appeared twice on the front page of *Hacker News*. The paper has been disseminated in the developer's email list for HBase (by its founder) and ZooKeeper. It has been discussed at numerous industrial conferences, including the Cassandra Summit (keynote by Cassandra's founder), InfoQ's QCon, RICON 2014 (keynote), Surge 2014, and is included in "Papers We Love," a developer community interested in reading and discussing computer science papers, with over 10,000 members and chapters in over 30 cities around the world. Some universities, like the University of Illinois at Chicago, have made this paper mandatory reading for its undergraduate programming class.

Zhu, Jianwen

WWW.ECG.TORONTO.EDU/~JZHU

■ **Analysis and Characterization of Cloud Applications with Flash Memories**

In the past, datacentre applications often find poor utilization-of CPUs due to the idle wait for the slow disk-based storage systems. The rise of flash memory based storage brings the opposite problem: the CPUs cannot keep up with the storage. This project aims to quantify the new gap, and obtain new insights on the interaction between applications, operating systems, and storage.

■ **Channel Modelling of NAND Flash Memory**

The goal of this project is to investigate physical channel characteristics of NAND flash devices to extract soft information for soft-decision error control coding, in the hope of permitting flash memories to operate well beyond their ordinary service life.

■ **Erasur Coding Acceleration**

Erasur coding is a generalization of RAID that can enhance the reliability and availability of petabyte scale storage systems. One problem that prevents its adoption in high-performance storage systems is its computational complexity, that commands a large amount of CPU cycles. The project aims to find efficient acceleration hardware built on FPGAs.

■ **Performance Modelling of Garbage Collection Algorithms**

A central concern of flash memory is its limited number of cycles a flash memory cell can be erased/written as it directly relates to the service life of flash-based storage. Another nuance flash-based storage has to address is that there is no in-place update like in traditional disk drives. Garbage collection therefore has to be used. The key performance metric of GC is write amplification, or the amount of non-productive write. The goal of this project is to establish a parametric analytical model of realistic garbage collection algorithms, such that write amplification can be not only predicted based on design parameters and workload characteristics, but also provisioned to deliver the best result.