24 ELECTRIC CHARGE
An industry partnership is leading the charge towards developing the vehicles of the future

30 SECTION **UNDERGRADUATE**
PROJECT PROJECTION
Undergraduates create an augmented reality system

38 SECTION **RESEARCH**
WHERE PROGRAMMER MEETS MACHINE
Machine learning hardware acceleration

45 SECTION **ALUMNI**
JUST LIKE COMING HOME
Alumna receives a Doctor of Laws, honoris causa

2017
A YEAR IN REVIEW

ANNUM

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

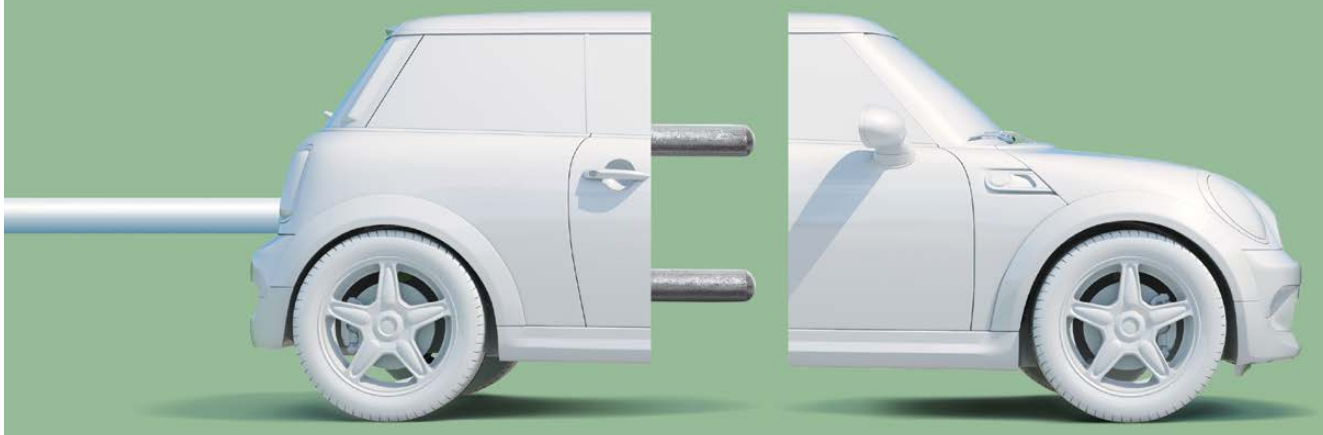

The Edward S. Rogers Sr. Department
of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

CONNECT

Whether you are an alumnus, current undergraduate student or graduate student, **ECE CONNECT** is a gateway to enriching your professional and personal network. By joining this new online community, you will have the opportunity to:

- Access a comprehensive directory of industry contacts
- Discuss or share relevant news, events and Skule™ memories
- Learn from or mentor fellow graduates
- Search the platform for new employees
- Post and search job opportunities
- Stay up to date on your department

SIGN UP TODAY
UOFTECECONNECT.CA

PHOTO BY RAINA+WILSON

Electrical and computer engineers are changing the world (as we know it)

Some of engineering's greatest achievements have been accomplished by creating new and better ways to move people and goods — as well as data and information — and here in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE), we are doing just that.

The driving force, so to speak, of the next revolution of transportation is electrification — moving people and goods by electric motor instead of by in-vehicle combustion. Just a decade ago, the widespread adoption of electric cars seemed a long way away — there were only hundreds of them on roads and highways around the globe. Today, there are more than two million electric vehicles worldwide, and growth continues to accelerate. Electric vehicles are not just poised to change our commute; they have the potential to significantly impact both our environment and our economy. The cover story of this issue, Electric Charge on page 24, highlights the early stages of a fast-moving and forward-looking industry collaboration, led by researchers in ECE that aims to bring electric vehicles to the mainstream, and to your driveway.

On the topic of acceleration, the place of machine learning in business and society is progressing at breakneck speed, as researchers and businesses are finding new and better ways to enable artificial intelligence in ways that are changing our daily lives. One ECE alumna has made it her mission to nurture — and keep — great minds in machine learning here in Canada. Dr. Foteini Agrafioti's journey has taken her from academia to entrepreneurship; read about how she's leveraging these experiences to help preserve and promote the machine learning ecosystem in Canada on page 12.

This magazine is a snapshot of the students, faculty, alumni and community who have helped us maintain our position as the top ranked ECE program in Canada and among the best in the world. Our collaborations with our industry partners help us keep our teaching relevant and our research impactful. This, in turn, helps us attract stellar undergraduate and graduate students who become members of our strong alumni community and make remarkable contributions to industry and society. Whether you are a prospective student, an industry partner or an alumnus, I welcome your feedback and the opportunity to connect — or reconnect — you may reach me directly at **chair@ece.utoronto.ca**.

FARID N. NAJM,
PROFESSOR & CHAIR

ANNUM 2017

FEATURES

EDITOR
Jessica MacInnis

ART DIRECTOR
Katina Constantinou,
Sugar Design

PHOTOGRAPHY & ILLUSTRATION
Raina+Wilson
Alvaro Dominguez

COPY EDITORS
Laura Brown
Avivah Wargon

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

Visit us online:
www.ece.utoronto.ca

Join the conversation:

All correspondence and undeliverable copies:

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

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

Publication Agreement No:
42887022

ISSN 2368-7037

Printed in Canada by
Flash Reproductions.

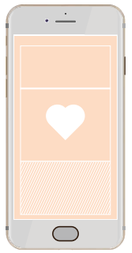

8

In the Hands of Patients

As a graduate student, Narges Norouzi published papers, developed a mobile app, and helped to create an online course, but she's not finished teaching — or learning — just yet

12

On the Pulse

First electrocardiograms, then entrepreneurship — alumna Foteini Agrafioti's next big challenge is cultivating Canada's machine learning talent

16

Thought Processor

To err is human, but Professor Natalie Enright Jerger is improving processor performance by deliberately introducing errors

20

The Sky's the Limit

Despite a challenging race against time and stormy weather, the future is bright for the Blue Sky Solar Racing team, led by ECE student Frank Gu

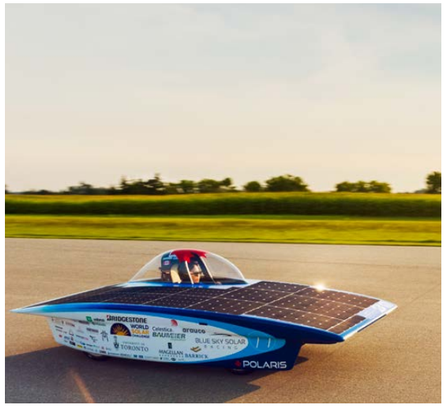

24

Electric Charge

Researchers in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering are leading a university-industry partnership to develop technologies for future electric vehicles

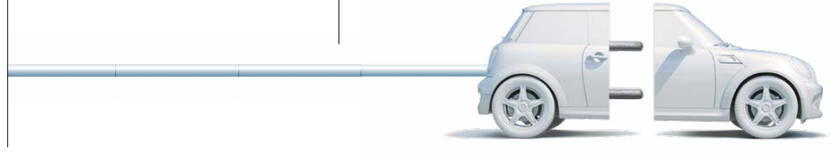

SECTIONS

30

UNDERGRADUATE

Project Projection: Augmented Reality

ANNUM sat down with the winners of the 2017 Gordon R. Slemon Design Award, which honours the year's most imaginative design project that exhibits both good design and execution

33

GRADUATE

Reza Rafie Receives Vanier Scholarship

Faster, cheaper, better: Reza Rafie's doctoral research is looking at how we can increase data rates to improve communication speed on fibre networks

> I really saw how it was not just the components of the car that had to work together, but the components of the team.
>
> p. 22

38

RESEARCH

Where Programmer Meets Machine

Andreas Moshovos is named director of the newly formed Natural Sciences and Engineering Research Council of Canada (NSERC) strategic partnership network on machine learning hardware acceleration

41

INDUSTRY

Fujitsu Laboratories to Open R&D Centre at U of T

An independent centre for research collaboration will be established, led by Ali Sheikholeslami

45

ALUMNI

Catherine Lacavera Receives Doctor of Laws, Honoris Causa

Triple alumna returns to U of T to cross the stage at Convocation Hall once again

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, 612 graduate students and 1,487 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, robotics and machine learning. Our graduates go on to prominent careers in engineering, finance, medicine, law and more — their options are truly unlimited.

Among our faculty are 30 Fellows of the IEEE, 10 Fellows of the Royal Society of Canada, 14 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 155 inventions disclosed and 97 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 2017

#1 *in* Canada
Electrical & Electronic Engineering

#13 *in* North America
Electrical & Electronic Engineering

#30 *in* the World
Electrical & Electronic Engineering

#1 *in* Canada
Computer Science & Information Systems

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

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

1,487 Undergraduate students

612 Graduate students

100 Professors, including Emeritus

73 Post-docs

53 Admin & tech staff

17 Research associates

9 Visiting professors

155 Inventions out of ECE since 2012

97 Patent applications out of ECE since 2012

14 Startups out of ECE since 2012

89 Percentage of ECE invention disclosures that include a student inventor

TIMES HIGHER EDUCATION WORLD UNIVERSITY RANKINGS 2017–2018

Engineering & Technology Universities

#1 IN CANADA
#15 IN NORTH AMERICA
#29 IN THE WORLD

ACADEMIC RANKING OF WORLD UNIVERSITIES 2017

Electrical & Electronic Engineering

#1 IN CANADA
#24 IN NORTH AMERICA
#46 IN THE WORLD

Computer Sciences & Engineering

#1 IN CANADA
#9 IN NORTH AMERICA
#13 IN THE WORLD

23 Chair titles held by ECE faculty in 2016–2017, including Canada Research Chairs, Endowed Research Chairs, Industrial Research Chairs and U *of* T Distinguished Professor Chairs

$53,143 Average salary earned by ECE students on PEY placements in 2017-2018

BY JESSICA MACINNIS

As a graduate student, Narges Norouzi published papers, developed a mobile app, and helped to create an online course, but she's not finished teaching — or learning — just yet

Emergency rooms are complex places. Physicians are often interacting with patients for the first time during an acute health issue. In this setting, health-care professionals can have limited access to aspects of their patient's health history that could provide clues for their diagnosis and treatment.

For example, patients going through alcohol withdrawal experience a type of hand tremor that is unique to that condition, but physicians had been relying on subjective indicators to figure out the optimal dosage of medication to counter the effects of alcohol withdrawal. If a physician over-estimates the severity of a tremor they could administer more medication than necessary, which could result in excessive sedation, among other issues. If a physician under-estimates the severity of a tremor, the patient may continue to experience symptoms of the withdrawal, such as seizures, irregular heartbeat and hallucinations. The medication used to treat the symptoms of alcohol withdrawal can also be misused by patients and have serious interactions when combined with alcohol or other medications.

When a group of physicians identified these and other challenges related to accurately assessing the severity of tremors in patients experiencing alcohol withdrawal, they looked to electrical and computer engineering researchers to help develop an objective way to measure these tremors so as to determine an ideal treatment plan. As an MASc student in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE), Narges Norouzi focused on multi-modal methods of measuring heart rates. Before wearable technology was mainstream, she worked on combining all the information from various sensors in a device to get a reliable heart rate measurement. But when her supervisor, Professor Parham Aarabi in ECE's Communications Group, was approached to help solve this common emergency room problem she pivoted

her research to applying signal processing techniques along with machine learning algorithms to measure the severity of alcohol withdrawal tremors. “When Professor Aarabi proposed this research topic to me I was already interested in applications of signal processing in the medical field,” says Norouzi. “So, it was a good fit for me to continue my research in this area as I transitioned into my PhD studies.”

Under Professor Aarabi’s supervision, Norouzi embarked on a collaboration with researchers at Mount Sinai Hospital in Toronto to find a more accurate way to diagnose patients who were experiencing alcohol withdrawal. “At the time, the diagnosis of the severity of alcohol withdrawal tremors was very subjective,” says Norouzi. “It depended on the experience level of the person who was assessing the patient — they would base their assessment on their past experience with patients experiencing alcohol withdrawal, using a scale to determine the dosage of medication they would administer.”

After collecting data, Norouzi and the team of researchers started working on a logarithmic model that would eventually be built into a mobile app, called AW Tremor. First, they had to determine how they could distinguish tremors caused by alcohol withdrawal from tremors caused by other health issues. Different health conditions result in different types of tremors: the frequency range of an alcohol withdrawal related tremor is in the range of five to 12 Hz, while tremors caused by Parkinson’s Disease lie in the range of four to seven Hz, in most cases.

Next, a panel of physicians at Mount Sinai Hospital who had years of experience in treating patients with alcohol withdrawal viewed videos of patients with suspected alcohol withdrawal related tremors. The panel rated the severity of these tremors on a seven-point scale to give Norouzi and the team a gold standard of training data for their algorithms. “Our goal was for the app to give similar results as this panel of experts,” says Norouzi. “What we found was that the result from our app was better than any other group of assessors — our estimation was more in line with the experts than any other group of physicians and nurses.”

Her PhD thesis resulted in eight peer-reviewed publications, an accomplishment for any graduate student, but the impact her work made in front-line health-care settings was the most rewarding for Norouzi. “Before she even finished her PhD, her work made a substantial impact in improving emergency room medicine,” says Professor Parham Aarabi.

What we found was that the result from our app was better than any other group of assessors — our estimation was more in line with the experts than any other group of physicians and nurses.

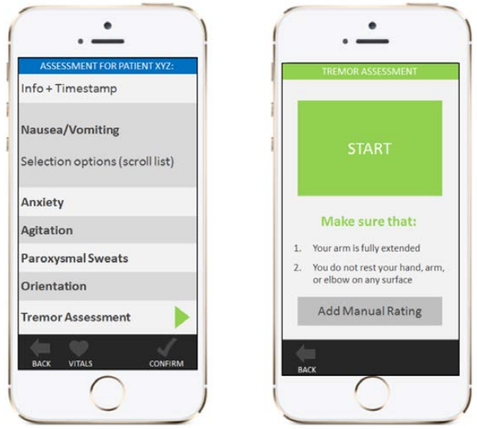

“In some cases, we saw the length of stay in emergency departments reduced by three and a half hours — all because of this powerful tool the app provided for health-care professionals.”

Building on the success with the app and her experience teaching undergraduate courses, Norouzi helped Professor Aarabi to develop a Massive Open Online Course (MOOC) on how to develop an iOS application. “I think we had 80,000 people enrolled in it; one of the first challenges was finding course material that was appropriate for so many students with such different backgrounds,” says Norouzi. “We took what we learned in developing that course and wrote a paper that was presented at the IEEE Frontiers in Education, a top-tier education conference.”

In late 2017, Norouzi joined the University of California, Santa Cruz, as a tenure-track teaching professor. “I’m teaching courses and doing engineering education research — helping to determine the best teaching practices in delivering engineering education in universities,” says Norouzi. “But, I am also continuing to pursue the biostatistics and machine learning research I was doing as a graduate student in ECE — I love the challenge of developing engineering solutions for health-care settings.”

For Norouzi, there is still much work to be done on the app: “Initially assessing the severity of a patient’s alcohol withdrawal tremor is one thing, but I think there’s a lot more that we can do to help health-care professionals give the best care to their patients,” says Norouzi. “For example, how much would the severity of the tremor decrease by giving the patient medication — could we use this information to help physicians determine not only the severity of the tremor but the appropriate dosage of medication?”

The app is now used in nearly a dozen hospitals, including world renowned addiction treatment centres. “The Betty Ford Clinic — now known as the Hazelden Betty Ford Foundation — is going to collect data for us and help illuminate other areas that we might investigate,” says Norouzi. “They will help us explore things like the decay of the severity of the tremor, the best location to acquire the tremor signal and how we can improve the accuracy of our model using different kinds of sensors along with the accelerometer.”

With an increase in data, such as the information collected by the Hazelden Betty Ford Foundation, Norouzi believes that they can increase the accuracy of the model and standardize the accelerometer recordings to collect data from different devices. “Once we have that, we will launch the app in the App Store, making it available to help physicians provide their patients with the best care around the world,” says Norouzi. “It will be exciting to see this app in the hands of patients — literally — around the world.”

BY JESSICA MACINNIS
PHOTOGRAPHY BY RAINA+WILSON

ON THE PULSE

First electrocardiograms, then entrepreneurship — alumna Foteini Agrafioti's next big challenge is cultivating Canada's machine learning talent

Alumna Foteini Agrafioti (ECE MASc 0T8, PhD 1T1) has always been drawn to opportunities just beyond her comfort zone. As an undergraduate student at the Technical University of Crete, Agrafioti's thesis supervisor Professor Michael Zervakis (ECE MASc 8T5, PhD 9T0) urged her to consider pursuing graduate studies at the University of Toronto, where he himself completed his master's and doctoral degrees in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE). "I remember him showing me his alumni magazines from U of T," says Agrafioti. "And I would go through them and kind of see how life might be abroad in a completely different environment at a completely different school."

At his urging, Agrafioti applied to ECE and joined Professor Dimitrios Hatzinakos's research group as a master's student. "As an international student, there is a period of adjustment — learning a new language, meeting new people and navigating a new city," says Agrafioti. "But I was also trying to figure out in what direction my research would take me." Agrafioti had joined a group that had traditionally been working on biometric security; much of the group's work had been done in face and fingerprint recognition, some of the traditional biometric modalities. "I remember feeling that I had landed in a place surrounded by such smart people — everyone around me was extraordinary," says Agrafioti. "It felt like a place where there was so much opportunity to learn and be inspired by others; there was a big challenge ahead but I was excited."

At the time, Professor Hatzinakos's group had begun to explore heartbeat recognition: using the human electrocardiogram for biometric recognition. "I came across reports from cardiologists who were reporting that the human electrocardiogram tends to look quite different from one person to the next," says Agrafioti. "This made it difficult for cardiologists to identify disease and these doctors challenged biomedical engineers to build hardware and systems that would detect disease despite how differently each patient looked."

But where cardiologists saw a problem, Agrafioti and her colleagues saw potential. “It was really interesting: could a person’s heartbeat be different enough to use it to authenticate people?” says Agrafioti. “Fingerprints, face recognition and iris scans come with a host of issues: cost and reliability are just two — but if we could identify people by their heartbeat, every second their heart could produce another wave, a time dependant biometric signature, that could be used to authenticate and re-authenticate someone.” Agrafioti and her colleagues began to publish results of their research that showed that neither exercise nor illness affect the actual shape of the pulse of a person’s heartbeat; electrocardiograms were a viable means of identifying a person. These findings began to open doors in the then nascent field of medical biometrics: “Researchers began to look at other medical signals, like brainwaves, blood pressure and oxygen levels,” says Agrafioti. “But in the beginning, it was the heart.” Stemming from this research, Agrafioti and her colleagues invented HeartID: the first biometric technology to authenticate users based on their unique cardiac rhythms.

As Agrafioti was finishing her doctoral degree in Professor Hatzinakos’s group, she began to think about her next steps. “I really enjoyed the freedom and flexibility I had as a graduate student and I wasn’t sure if working at a large technology company would be a good fit for me,” says Agrafioti. “In talking with U of T’s Innovations and Partnerships office, they explained to us that we could either license the rights of our patent to a company or we could start a company ourselves and take our technology to market on our own.”

With that, Agrafioti set off on her next great adventure: entrepreneurship.

Agrafioti and her co-founder, started a company called Bionym, now Nymi, to commercialize their technology. “Our initial strategy was to sell our algorithms to companies that manufacture wearable devices; we didn’t want to be a hardware company because we thought that there was so much opportunity for biometric recognition that we shouldn’t limit it to one device,” says Agrafioti. “Like most entrepreneurs, our focus pivoted and narrowed when we eventually realized we needed to control our destiny and go to market with our own wearable device.” By the time the company had grown to about 60 people, Agrafioti had decided she was up for the next challenge. “I had spent six or seven years on one specific technology, and I was ready for a change,” says Agrafioti. “At that time, the product was complete and I had the opportunity to join another company so I jumped at the chance.”

By the time that the buzz around Canadian researchers in machine learning had gone mainstream, Agrafioti felt that she was in a position to leverage her experience of research rigor and successful entrepreneurship to help nurture this growing field. She joined RBC as Chief Science Officer in 2016 and was tasked with steering the bank’s innovation strategy. The company’s focus on machine learning theory and applications through open academic collaborations with research centres in artificial intelligence held a lot of promise for Agrafioti. “Initially we started off thinking we would be building machine learning capabilities within the banks, but machine learning has changed so much in the last few years,” says Agrafioti. “It started to become obvious to me, and many others at RBC, that we had the responsibility to create a place where people can do fundamental research in artificial intelligence after they graduate from school.”

IT STARTED TO BECOME OBVIOUS TO ME, AND MANY OTHERS AT RBC, THAT WE HAD THE RESPONSIBILITY TO CREATE A PLACE WHERE PEOPLE CAN DO FUNDAMENTAL RESEARCH IN ARTIFICIAL INTELLIGENCE AFTER THEY GRADUATE FROM SCHOOL.

FOTEINI AGRAFIOTI

Headquartered in Toronto, with another lab in Edmonton, RBC Research Institute has two teams: “The fundamental research team tackles the same problems as any academic or industrial research team around the world would — for example, unsupervised learning is an open machine learning problem, so we have a team working on this area, publishing papers and participating in conferences,” says Agrafioti. “The applied research team takes these innovations coming out of our fundamental research team, like algorithms, and applies them to banking.” Her teams are establishing research projects in areas like deep learning, high performance computing, cybersecurity and more.

Identifying money laundering, fraud detection and algorithmic trading are obvious areas where machine learning technologies can be applied to the banking sector. But Agrafioti points out that there are many applications that they are working on that are not just of interest to banks: “We work a lot on natural language processing and one of the challenges we are trying to solve is understanding news,” says Agrafioti. “News is extremely important to our business; we want to know very quickly what is happening around the world before it becomes mainstream.”

But identifying, quantifying and evaluating a news item is not just important for market analysis — identifying fake news and analyzing future impact of issues is of interest to the financial sector and beyond. “We’re trying to use machine learning to understand the contextual relationships of things that happen around that world and which have the potential to affect Canadian markets,” says Agrafioti. “Applications like this are not just for the financial industry — if we are able to really understand in real-time what is going on around the world, it can have enormous impact in other fields too. This is why RBC was the right fit: we have a clear goal of helping communities prosper.”

I remember feeling that I had landed in a place surrounded by such smart people — everyone around me was extraordinary.

Enabling projects that explore how machine learning theory and applications can be applied to the banking sector — and beyond — is a key goal of RBC Research. But the overarching strategy of RBC Research is to help retain and attract Canadian talent in the field of machine learning and to foster open academic collaborations with academic institutions. It’s this priority that attracted Agrafioti to the position in the first place. “It’s exciting to talk to some of our team and hear that they had job offers in the U.S. and U.K., but they chose to stay here in Canada,” says Agrafioti. “I’m extremely proud that we were able to keep these incredibly smart people here and they continue to do research for a Canadian company producing immense value for not only RBC but for the entire country — I’m excited to see where it will take us.”

BY JESSICA MACINNIS
PHOTOGRAPHY BY RAINA+WILSON

THOUGHT_PROCESSO/R

To err is human, but Professor Natalie Enright Jerger is improving processor performance by deliberately introducing errors

Moore's Law, the prediction that processing power will double every two years, has driven unprecedented advances in computer technology. For decades, each next-generation chip has had twice as many transistors on it than its predecessor and as a result, consumers have become accustomed to having access to faster, smarter, smaller and often more inexpensive computers every couple of years.

As the physical limits of the materials used for integrated circuits are reached, the debate about whether Moore's Law will slow down or even end rages on; but for Professor Natalie Enright Jerger in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) the march to improve processor performance and energy efficiency is ongoing. Enright Jerger's research in computer architecture — the design of processors and memory systems — has explored various avenues to improve processors along several dimensions. "Broadly, we've been looking to improve processor performance," says Enright Jerger. "But, we've also been thinking through how to improve processor energy efficiency."

In the field of computer architecture, there are many objectives in design and these objectives vary depending on the application context — performance, cost, reliability and power consumption are all balls that need to be kept in the air. Sure, maximum performance is best, but if it comes at the cost of higher power consumption, then it may not be ideal. Of course, cheaper is better, but if performance is impacted then it may not be cost-effective. Yes, energy efficiency is important, but if it negatively affects reliability the advantage is negated.

The juggle is real.

Interconnection networks — how processors communicate with each other — have been much of the focus of Enright Jerger's work, stemming from her dissertation research, and this area remains one of the main thrusts of her research group. Improving how processors within a single chip communicate with each other can reveal significant cost savings. Moving data from your memory to your processor and within your processor is expensive work: "The cost in terms of energy of moving that data, getting it from the memory to the processor, and then moving it within the processor exceeds the cost of the actual computation," says Enright Jerger. "So, if I am just doing an addition, the cost of getting the data for the addition is more expensive than completing the actual addition."

We're looking at interconnection networks as a way to efficiently move the data.

Looking at ways to move the data more efficiently with interconnection networks, Enright Jerger's group seeks to dynamically match network resources to demand. Here's where the juggling act comes in: providing more resources than required wastes power but providing fewer resources than required leads to a shortfall in performance. Scaling the voltage and frequency to match communications demands eliminates power overages and reduces power loss: "We're looking at interconnection networks as a way to efficiently move the data." says Enright Jerger.

In recent years, Enright Jerger's research has expanded to mitigate power consumption challenges through approximate computing. If you think that close only counts in horseshoes — as the saying goes — then perhaps you've never heard of approximate computing. From search engine results to video processing, there are a number of applications that are inherently error-tolerant: either there is no one correct solution, the computations involve noisy inputs (and therefore, outputs) or the applications can incorporate minor errors without being noticed by the human eye or ear.

Enright Jerger is exploring ways to build hardware that purposefully allows errors to happen — or even introduces errors — where these errors would result in a substantial reduction in energy or a substantial increase in performance. "Typically, processor design has operated on an energy and performance trade-off: if I want high performance, I have to spend a lot of energy and if I want to reduce energy consumption, I have to give up performance," says Enright Jerger. "But what if there was a third option where we introduce a bit of error in order to get a higher performance or lower energy expenditure?"

Some computations, search engine rankings for example, have no single correct solution. Each search engine has a different algorithm to rank the results of your query: one search engine may rank the results of an identical search differently than another search engine. "If I did something to my processor, for example, that caused it to switch the ranking of the first and second links that your search engine returns to you, it would be technically incorrect based on the original search algorithm used," says Enright Jerger. "But, from a user perspective, it's an equally valid result — for the user, there are multiple right answers."

Enright Jerger is exploring ways to build hardware that purposefully allows errors to happen — or even introduces errors — where these errors would result in a substantial reduction in energy or a substantial increase in performance.

Other computations, like video processing, have levels of error tolerance for humans that are below perfect. If a video is playing on your laptop or tablet, there may be millions of pixels on the screen — it would be nearly impossible for a human to tell if a handful of those pixels were incorrect. "When a video starts to pixelate or stutter, it becomes obvious," says Enright Jerger. "But we're talking about much smaller magnitudes of errors, errors that would be indiscernible to a viewer."

A third category of applications that are error-tolerant are those that have noisy inputs. "An example would be the GPS on your smartphone: in any given map application, it will show you as a dot with a wider radius around it," says Enright Jerger. "What it's saying is that it thinks you might be here, but you could also be anywhere within this radius." If you were using an application that was taking your GPS coordinates to calculate the speed or distance you are travelling, the results would never be completely accurate — because your starting point wasn't completely accurate to begin with. Leveraging these three instances of error tolerance can result in both energy savings and performance improvement — the less you access and move data the more efficient an application can be.

Another vein of Enright Jerger's research incorporates value similarity as a means to avoid these costly accesses to memory. If a certain set of data can be identified as similar to another set of data, can we limit data movement by substituting data already retrieved from memory? Enright Jerger thinks so: "If you looked at a photo of the Toronto skyline, a lot of values in the photograph are going to be the same — for example, the pixels that comprise the sky will all be a similar blue value," says Enright Jerger. "We're looking at ways to identify that a block of data is very similar to another block of data so we can reduce the amount of data that we bring in from memory."

Even though it may not apply to all applications, there are a large number of consumer applications where Enright Jerger feels that the error for performance and efficiency trade-off will bring significant value — especially in that she envisions its flexibility. "In our projects, we don't ever disable the correct option entirely, if you want your processor to be 100 per cent accurate then that is still an option" says Enright Jerger. "What we are looking at is a user configurable option, where the end user can decide on the parameters — and I think this could positively impact the trajectory of how technology continues to advance in the future."

Moore's Law might be slowing down, but Enright Jerger isn't.

THE SKY'S THE LIMIT

BY JESSICA MACINNIS
PHOTOGRAPHY BY RAINA + WILSON

Despite a challenging race against time and stormy weather, the future is bright for the Blue Sky Solar Racing team, led by ECE student Frank Gu

It was a seven day, 3,021-kilometre journey across the Australian outback powered by ingenuity, determination and the sun. But for Yu Cheng (Frank) Gu, a fourth-year student in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE), even though the journey finished in Australia, it started three years earlier at a U of T Engineering student club fair. Looking for a challenge, Gu signed up to join the Blue Sky Solar Racing team. To Gu, joining the team presented an interesting opportunity: the chance to apply technical knowledge from his program while learning new skills and meeting students in other years and disciplines. Besides, with all the buzz around electric vehicles, the opportunity to work on one while still an undergraduate student was too good to pass up.

Over the course of its 22-year history, the Blue Sky Solar Racing team has raced nine generations of solar powered electric vehicles built by U of T Engineering students. Students across all of the disciplines work together on not only the design aspects of the car — everything from aerodynamics to electrical design to mechanical manufacturing — but also the business aspects, like sponsorship and public relations. Throughout each successive build, improvements are made to the internal design, the shape of the car and its components — but the overall goal remains the same: to build the most efficient electrical vehicle powered by the rays of the sun.

In his first year on the team, Gu worked on the fabrication of the solar car. "I was learning about, and working on the composite materials — I thought as an ECE student I'd be working on something electrical," says Gu. "But the next thing I knew, I was working on all of these carbon fibre parts." Through this work he was establishing a baseline understanding of the car, but he was also beginning to understand the value of team dynamics.

By the time his second year rolled around, Gu was hooked: he loved the camaraderie of the team and the intellectual challenge of working within the confines of the team's resources and the limits of solar powered technology. He also had a new role and as a core member of the electrical team he earned his first trip to the Australian outback. "The Bridgestone World Solar Challenge is a 3,000-kilometre race — actually 3,021-kilometres — those numbers stick in your head once you've done the race," says Gu. "The competition is comprised of teams from engineering schools around the world, but there are also corporate teams, so the calibre of the teams is really top notch." The team took home 12th place but Gu took home much more: "I really saw how it was not just the components of the car that had to work together, but the components of the team," says Gu. "I thought I might be able to contribute in a leadership role, so while I was in Australia I made the decision to run for the managing director position." When the team arrived back in Toronto, Gu ran for the position and was unanimously elected by his peers on the team for the role — but that was the easy part.

I really saw how it was not just the components of the car that had to work together, but the components of the team.

In 2016, the team took their car, Horizon, to the American Solar Challenge, another 3,200-kilometre race, from Ohio to South Dakota over eight days. They made modifications to the car they had taken to Australia the year before: replacing the electrical system and modifying the mechanical system while maintaining the overall design. Those alterations paid off — the team placed third and identified a number of ways they could improve on the design, processes and technology for the next iteration of Blue Sky Solar Racing's car. "It was the right time to take everything we had learned racing Horizon and build a new car for our next race in Australia," says Gu. "So, when we got back to Toronto we had a team of about 100 students who helped to design, build, test and race our newest car: Polaris."

Named after the North Star, the brightest star in the constellation of Ursa Minor, Polaris was completely built from scratch. The team incorporated new lithium-ion batteries and installed a DC electric motor that allowed the team to adapt to different conditions. "Whether the car is cruising on a flat stretch of road or going up a hill, the motor can acclimate, which makes the car more nimble and adaptive," says Gu. "We also added cutting edge silicon solar cells — a top-of-the-line product that you wouldn't likely see in regular commercial use."

The team identified a need for large-scale, high performance computing so they ran a small-scale supercomputer to run simulations and analyses for various components of the car, from aerodynamics to the electrical systems. "We partnered with AMD to incorporate some of their latest compute units and their professional graphics design technology — and we also worked with the latest embedded computer systems and used a real-time operating system; it was really fun to take what we were learning in class during the day and directly apply it to the car on evenings and weekends," says Gu. "As engineering students, we are taught to identify gaps — make things and processes more efficient and that's exactly what the team did when designing and building Polaris."

The team unveiled the new car in the summer of 2017 and by early autumn they had shipped the vehicle to Australia and the travelling members of the team began their journey to unload the car and start the race. Polaris set out from Darwin in October, waved off by crowds of cheering spectators. While the race started smoothly for Polaris, suddenly all of the challenges of managing team members, finances, sponsorships and technical aspects seemed easy compared to their greatest challenge yet: the weather. "By day three we were in central Australia where there was supposed to be no possibility of rain and we hit not one, but two thunderstorms." says Gu.

A rainy drive is unpleasant in any car, but for a solar electric vehicle, a rainstorm is more than an inconvenience, it's a setback. "The team rolled through the storms at just 40 kilometres per hour, but soon those dark clouds turned to — well — blue sky," says Gu. "By the afternoon, we had finally passed both storms and made it out of that leg of the race."

Despite these challenges, the team persevered and upped their 2015 placing, finishing 11th in the world. "It was an incredible experience and the team was phenomenal," says Gu. "I'm proud of the car and our race but I am most proud of my team's unwavering dedication to this project that challenged what we thought was possible." As he prepares to leave his position when his two-year term winds to a close, he's excited to see where the next managing director will take the team and the memories new team members will create when they build the next car. "The team has been around for more than 20 years, so students have cycled through the Blue Sky Solar Racing team for decades now," says Gu. "I know that when people move on and graduate they leave the team — but what I've learned is that the team never really leaves them."

For Gu and the Blue Sky Solar Racing team, no matter the weather, it's blue skies ahead.

ELECTRIC CHARGE

BY JESSICA MACINNIS
ILLUSTRATION BY ALVARO DOMINGUEZ

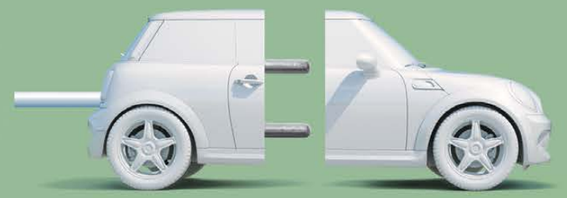

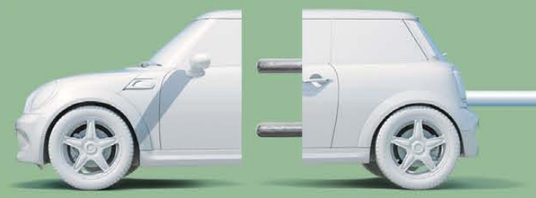

Researchers in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering are leading a university-industry partnership to develop technologies for future electric vehicles

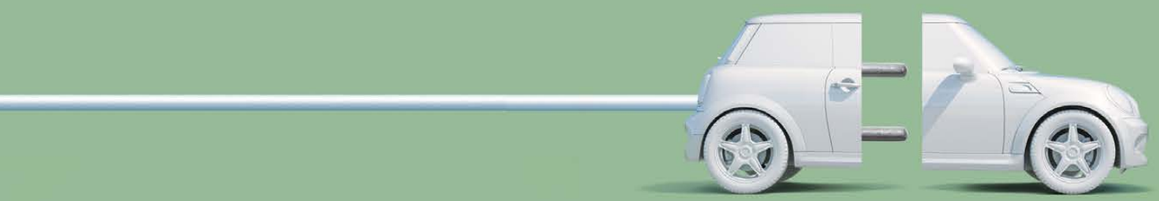

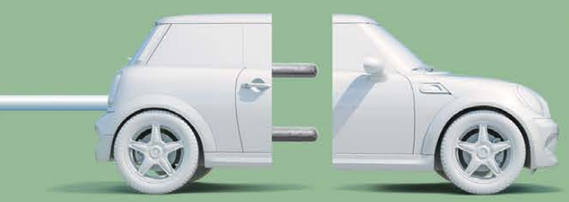

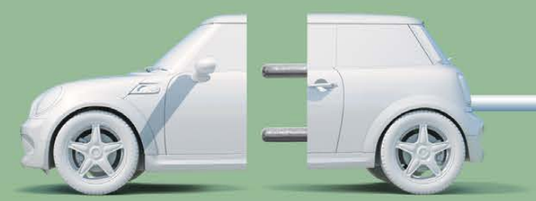

It was a routine trip to China.

Professor Olivier Trescases, of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE), had taken countless trips to Beijing, Hong Kong and Shanghai over the course of his career — to conferences, to meet with industry partners and to connect with research collaborators. The goal for this trip was simple: Ontario was embarking on a business mission to China and he was invited to highlight some of the research on power electronics and energy systems that his group was working on in ECE.

That flight turned out to be the first leg of a journey towards one of the largest industrial partnerships in the department, just over one year later.

While alumnus Tony Han was an undergraduate student in ECE, Professor Trescases had been working with Steve Dallas of Toronto Electric on what they referred to as the little yellow car — the first electric car to be completely designed and built in Canada. Trescases and his group built power converters and did experiments on the car, and it was working with this little yellow car that piqued Han's interest in electric vehicles (EVs). Han went on to pursue a master's degree under the supervision of Professor Peter Lehn, also in ECE's Energy Systems Group, working on research on bidirectional multi-port direct-current to direct-current converters — another area of great interest to the electric vehicle industry. So, it was not surprising that Han, with academic connections in Toronto and business connections in China, would reconnect with Trescases during this business mission overseas.

Han had previously met with Lehn to discuss how he might start a business focused on EVs — he had a good working relationship with his former supervisor and both Lehn and Trescases had solid connections in the industry: "I was interested in breaking into the Canadian market and it was important to me to work with people I trust, so Toronto was uniquely positioned," says Han.

"With the innovation stream coming from the University coupled with the proximity to the supplier network across southern Ontario — it made Toronto a great place to start Havelaar Canada."

Havelaar, a global company with operations in China, the Netherlands — and now Canada — set up shop in North York in 2016. With Han as CEO, they sought to create an environment in which Havelaar Canada's engineers and researchers at the University of Toronto were able to work closely together. Less than a year after the Ontario business mission to China, and with Havelaar as a founding partner, Trescases and Lehn established the University of Toronto Electric Vehicle (UTEV) research centre, a university-industry partnership focused on the advancement of electric vehicle technologies.

Led by Trescases and Lehn, UTEV is a multidisciplinary centre with professors from ECE, the Department of Mechanical and Industrial Engineering (MIE) and the U of T Institute for Aerospace Studies (UTIAS). "UTEV is focused on energy systems, battery management, drivetrain systems, inductors and energy efficiency," says Trescases. "But we are also working with other faculty members, for example, Dean Cristina Amon in MIE on mechanical aspects and liquid cooling, Professor Sean Hum in ECE on electromagnetics and Professor Tim Barfoot from UTIAS on autonomous driving."

UTEV came together quickly: just months after Trescases and Han reconnected in China, Havelaar Canada and ECE had signed a partnership agreement and initiated plans to renovate some ground level space on campus to house testing facilities for electric vehicles. "This is expensive work, so Havelaar Canada's investment gives us the opportunity get proper equipment to do really high impact work," says Trescases. "It allows us to take it beyond the small-scale prototype, to build a full scale electric vehicle battery pack, for example — to have impact in this space you need to have resources and this is something that Havelaar recognizes." But like most partnerships, the benefits work both ways.

ECE's strong Energy Systems Group was certainly a draw for Havelaar, but UTEV is also leveraging the newly renovated Energy Systems Lab: "It's truly one of the best power labs in North America," says Lehn. "So UTEV can leverage the millions of dollars of investment made in the renovation of the lab — and that, combined with the calibre of researchers here in the Energy Systems Group, made U of T a good fit for Havelaar Canada to partner with."

Broadly, UTEV's goal is to pursue research thrusts that advance power electronics for EV drivetrain systems and improve charging infrastructure and energy storage for EVs. In doing so, UTEV hopes to encourage consumer adoption of electric vehicles by helping to solve some of the main challenges preventing widespread adoption of EVs. "We're aiming to make EVs more efficient by increasing driving range and we are working on improving energy utilization to reduce the high costs associated with an EV's battery system," says Lehn. "We're really fortunate here to have the resources to be able to do large scale projects of this sort."

In a fast-moving research area with major industry players working in the space, being able to delve broadly into the research is beneficial to both the faculty and researchers involved as well as to Havelaar Canada. "Tony had this vision — which I think is very unique — of setting up a company to fully leverage the research outcomes that are generated here in the university," says Trescases. "So instead of saying 'We have this problem, can you solve it?' he set up the company to directly leverage all of the expertise here by saying 'What do you think are the interesting technological directions we should be looking at?' so, it has been really refreshing."

Expanding on the investments of the early stages of the partnership, Trescases and Lehn were each awarded Natural Sciences & Engineering Research Council (NSERC)'s Collaborative Research and Development Grants (CRD) — incorporating major industry funding by Havelaar Canada — totalling more than $9 million over four years for their projects. These grants will support multidisciplinary projects that aim to reduce the cost and improve efficiency of both electric vehicles and infrastructure. Projects awarded CRDs by NSERC are expected to lead to industrial or economic benefit to the country; Trescases and Lehn predict a strong electric vehicle industry in Canada will do both.

In May of 2017, Havelaar Canada launched what Han calls the world's first all-electric powered pickup truck, the Havelaar Bison, at an electric vehicle trade show in Markham. Han plans to first release these trucks to fleet operators such as municipalities and companies seeking to reduce their transportation emissions. That way, Han explains, Havelaar can roll out the Bison in a controlled way and monitor how they handle Canadian road conditions before making them available for the public to purchase in 2019. "After months of work, it was pretty thrilling to see the prototype model revealed," says Han. "I think we will see vehicles like the Bison all over our roads in the very near future — when the truck was unveiled at the trade show, it really was a first step in a very exciting journey."

And as for the colour of the Havelaar Bison unveiled in May? "Electric blue, of course."

With the innovation stream coming from the University coupled with the proximity to the supplier network across southern Ontario — it made Toronto a great place to start Havelaar Canada.

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**.

17:1
Undergraduate student-to-faculty ratio in ECE

Undergraduate degrees awarded, 2016-2017

ElecE *125*
CompE *213*
Total *338*

ECE'S SANDRO YOUNG NAMED U OF T'S TOP STUDENT

Sandro Young (CompE 1T6+PEY) always knew he was interested in computers, circuits and electronics.

What started as a middle-school passion — Young was an active participant in science fairs in grades seven and eight — eventually led him to study in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) where he graduated as U of T's top student in the spring of 2017.

"I never set out to be the top graduating student," says Young. "But I've pretty much always been interested in science, math and technology — especially programming, software and computer architecture, so ECE was a great fit for me, and I enjoyed all of the courses and projects along the way."

Young graduated with a number of scholarships, awards and medals in recognition of his academic achievements, including one of three Governor General's Silver Medals bestowed to U of T's most academically outstanding undergraduate students who graduated in June 2017. He is also the recipient of the John Black Aird Scholarship, given to the top student at U of T as well as the W.S. Wilson Medal, the Ontario Professional Engineers Foundation for Education Gold Medal, the Henry G. Acres Medal and the Adel S. Sedra Gold Medal.

PROJECT PROJECTION: **AUGMENTED REALITY**

Each year, fourth-year ECE students enrol in a full-year, team-based design course that allows students to integrate the knowledge acquired over the first three years of study in the planning and execution of a project proposed either by the students or by a supervisor.

ANNUM sat down with the winners of the 2017 Gordon R. Slemon Design Award, which honours the year's most imaginative design project that exhibits both good design and execution — Joey Bose, Lyndon Chan, Da Hoon Pyun and Ayush Shrestha — to hear about their project.

Q: WHAT WAS YOUR PROJECT CALLED AND BRIEFLY DESCRIBE IT?

A: Our project was called DARI — Depth-variable Augmented Reality Interface — and it was an augmented reality (AR) system that could dynamically project images at a range of focal depths. Augmented reality "augments" a view of the real world by projecting virtual images on top of what you're seeing, but usually this means the image is projected at a fixed depth. There are a lot of AR systems out there but, unlike DARI, none of them project dynamic virtual content at variable depths, and allow users to interact with the virtual content in real space with hand gestures and facial tracking.

Q: WHAT DO YOU MEAN BY VARIABLE DEPTH?

A: By variable depth, we mean that DARI is able to project virtual graphics which appear to be at changeable 3D positions. Many other AR systems only project virtual graphics at a single depth, so even as the user interacts with the virtual image in real space, it appears to be at the same 3D position — this makes the user experience unintuitive and unconvincing. By supporting variable-depth virtual images, DARI enables real-space user interaction: as the user interacts in real-space with DARI, a camera captures their hand gestures and facial position and vertically positions a sliding display placed behind a one-way mirror so that the image appears to the user at the correct depth, thus eliminating the perceived discrepancy between user input and virtual graphic.

Q: WAS AR SOMETHING YOU WERE INTERESTED IN FROM THE BEGINNING? HOW WAS YOUR GROUP FORMED?

A: We formed a team and were looking for a professor to guide us — we weren't entirely sure what we wanted to do but knew we wanted to work in a space that was exciting to us, like machine learning, AR and computer vision. So, we proposed this project to Professor Aarabi who was open to our idea of doing an AR interface. It was a great fit because he has a lot of experience in computer vision and augmented reality and he was able to guide us along the way while making really helpful suggestions — like incorporating a mirror component, which we hadn't originally thought of.

Q: WHAT WOULD BE SOME REAL-WORLD APPLICATIONS OF YOUR PROJECT?

A: The retail sector is definitely an application for this type of technology. Imagine being able to "try on" augmented reality glasses and move around in front of a mirror and the glasses follow your face at whatever distance you are from the mirror. Or, it could be used for educational purposes — a person could hold an image projected onto their hand and rotate it right in front of them.

Q: WHAT'S NEXT FOR DARI?

A: Professor Aarabi showed a video of DARI at the Augmented World Expo in the US — it was really exciting to have it seen at one of the largest symposiums on augmented and virtual reality. With his help, we were also able to file a provisional patent on DARI. We designed the system so that people could always add on top of it. We think of it more as giving future developers an API that they can use to build custom applications: we've created the technology for people to use in whatever purposes they see fit. We're looking forward to seeing where it goes from here.

PEY: MORE THAN A CO-OP

The Professional Experience Year (PEY) paid 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 2017–2018, 241 ECE students are working in PEY placements at 84 companies located across Canada and around the world, earning an average salary of $53,143. To learn more, visit **uoft.me/ece-pey**.

$89,895
Top salary earned by an ECE student on PEY placement in 2017-2018

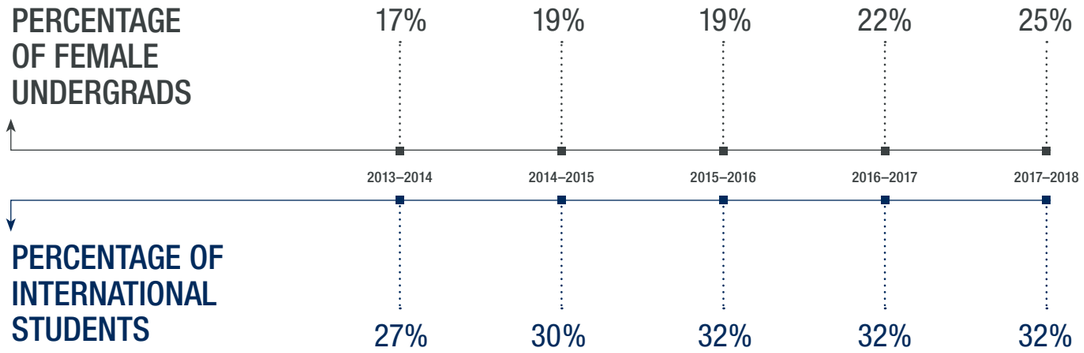

UNDERGRADUATE ENROLMENT, 2013–2014 TO 2017–2018

	PART TIME	FULL TIME	TOTAL
2013-2014	259	1133	1392
2014-2015	272	1208	1480
2015-2016	313	1199	1512
2016-2017	283	1230	1513
2017-2018	289	1198	1487

ECE PEY PLACEMENTS, 2013–2014 TO 2017–2018

	ElecE	CompE	TOTAL
2013-2014	133	77	210
2014-2015	156	70	226
2015-2016	160	92	252
2016-2017	135	91	226
2017-2018	121	120	241

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, computer 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.**

$2.7M
in graduate student scholarships awarded by ECE for 2017-2018

OUR GRADUATE STUDENTS BY FIELD OF STUDY

	MASC	PHD	TOTAL
BIOMEDICAL ENGINEERING	8	3	**11**
COMMUNICATIONS	39	50	**89**
COMPUTER	43	45	**88**
ELECTROMAGNETICS	11	16	**27**
ELECTRONICS	29	39	**68**
ENERGY SYSTEMS	28	20	**48**
PHOTONICS	23	32	**55**
SYSTEMS CONTROL	10	7	**17**

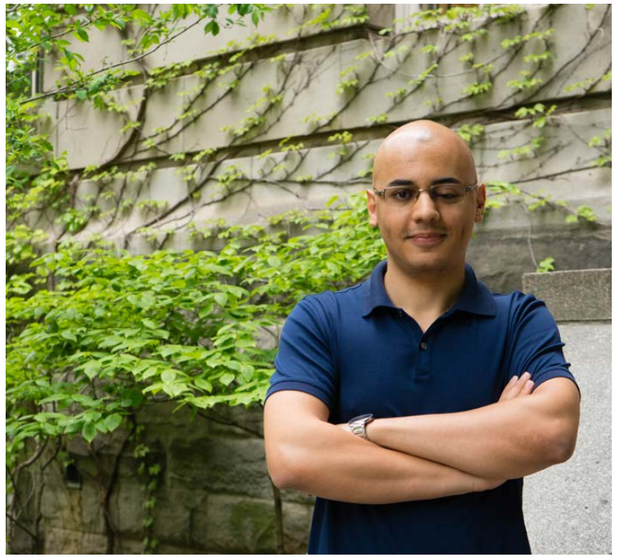

Reza Rafie receives Vanier Scholarship

Faster, cheaper, better: Reza Rafie's doctoral research is looking at how we can increase data rates to improve communication speed on fibre networks. Rafie's research, under the supervision of Professor Frank Kschischang, in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE), involves exploring the information-theoretical aspects of fibre networks and devising mechanisms that can exploit their full capacity.

Rafie received a 2017 Vanier Scholarship for his ongoing research in fibre-optic communications. Worth $50,000 per year for three years, Vanier Scholarships are awarded by the Government of Canada to doctoral students at Canadian universities who demonstrate excellence in three areas: academics, leadership and research impact. "Professor Kschischang is a renowned leader in information theory, I'm very fortunate to be a student of his," says Rafie. "Being named a Vanier Scholar is an honour, particularly in light of the very competitive nature of this award, and it is a great motivation for me to do high quality research during the rest of my PhD studies."

MASTER OF ENGINEERING SPOTLIGHT

Advanced technical courses and the opportunity to choose from a variety of specializations take Master of Engineering (MEng) students well beyond what they learned as undergraduates and equip them to tackle industrial and societal challenges of the 21st century

Dr. William Chisholm, an internationally renowned expert on the effects of adverse weather on electric power system reliability, has more than forty years of utility experience. But for students taking his MEng courses in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) his decades of experience translate into coursework with both an industrial focus and a solid technological base.

Dr. Chisholm is just one adjunct professor in ECE's MEng graduate program — a course-based degree program that provides individuals with advanced training to work or practice in the fields of electrical or computer engineering. MEng students enrol in one of eight fields of specialization — from biomedical engineering to systems control — and take courses on a full-time, part-time or extended full-time basis to complete their degree requirements. "The depth and breadth of their educational experience in the MEng program is an asset to the students," says Chisholm. "And it's also an asset to their current or future employers."

From Smart Grid Case Studies to Electrical Insulation Design and Coordination to Grounding and Bounding, Dr. Chisholm has developed and delivered a number of courses to students in the MEng program. "The classes are often a mix of full-time and part-time students — which makes both the projects and the discussions quite robust," says Chisholm. "For the students already working full-time in the field, their perspective can be valuable to other students because they can bring things that they're doing at work into the presentations and into the experiences that are shared."

While the MEng courses provide students with specialized technical knowledge, they also equip them with the skills required to apply this knowledge in the workplace. "The number one value add for students is to give them new experiences, for example picking a favourite technical paper from the four million found online in the library," says Chisholm. "We spend a lot of time on industrial standards in my courses: knowing how to find, follow, understand and make effective use of standards is another example of the value that my MEng students can bring back to employers."

Enrolment in Research-stream Programs 2017–2018

MASc 191

PhD 212

Total 403

TOTAL GRADUATE ENROLMENT, 2013–2014 TO 2017–2018 (HEADCOUNT)

	FALL 2013	FALL 2014	FALL 2015	FALL 2016	FALL 2017
MASC	167	173	145	158	191
PHD	249	253	245	242	212
MENG	140	169	249	249	209
TOTAL	556	595	639	649	612

GRADUATE DEGREES AWARDED, 2012–2013 TO 2016–2017

	2012–2013	2013–2014	2014–2015	2015–2016	2016–2017
MASC	56	70	71	70	53
PHD	37	27	35	47	38
MENG	56	85	76	97	144
TOTAL	149	182	182	214	235

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**.

$331K – AVERAGE ANNUAL RESEARCH 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

SELECTED AWARDS WON BY ECE FACULTY, 2016-2017

ROYAL SOCIETY OF CANADA, FELLOW
Professor Reza Iravani

ROYAL SOCIETY OF CANADA, COLLEGE OF NEW SCHOLARS, ARTISTS AND SCIENTISTS, MEMBER
Professor Wei Yu

INSTITUTE OF ELECTRICAL & ELECTRONICS ENGINEERS, FELLOWS
Professors Raviraj Adve and Sorin Voinigescu

ASSOCIATION FOR COMPUTING MACHINERY, FELLOW
Professor Andreas Moshovos

SPIE (THE INTERNATIONAL SOCIETY FOR OPTICAL ENGINEERING), FELLOW
Professor Peter Herman

OPTICAL SOCIETY (OSA), FELLOWS
Professors Mo Mojahedi and Joyce Poon

CANADIAN ACADEMY OF ENGINEERING, FELLOW
Professor Wei Yu

TECHNICAL UNIVERSITY OF MUNICH, TUM AMBASSADOR
Professor Frank Kschischang

ONTARIO PROFESSIONAL ENGINEERS AWARDS, ENGINEERING EXCELLENCE MEDAL
Professor Vaughn Betz

ONTARIO PROFESSIONAL ENGINEERS AWARDS, RESEARCH AND DEVELOPMENT MEDAL
Professor Stewart Aitchison

CONNAUGHT INNOVATION AWARD
Professor Glenn Gulak

CONNAUGHT GLOBAL CHALLENGE AWARD
Professor David Lie

VAUGHN BETZ RENEWED AS INDUSTRIAL RESEARCH CHAIR

Industrial Research Chairs are jointly funded by NSERC and by industry

Moore's Law — the observation that the number of transistors on an integrated circuit doubles approximately every two years — might not be dead, but it is slowing. A key challenge for computer engineers is finding new ways to increase compute power and reduce energy consumption to handle the ongoing demand for faster and cheaper computing technologies.

Professor Vaughn Betz of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) is working on this challenge with reconfigurable chips called Field Programmable Gate Arrays (FPGAs), and he has been renewed as the NSERC/Intel Industrial Research Chair in Programmable Silicon.

FPGAs are computer chips that give the end user the flexibility of reprogramming them after manufacturing, unlike other chips like application-specific integrated circuits which are customized during manufacturing for a specific use. During this term as research chair Betz and his team will work on making FPGAs more efficient, develop software that makes FPGAs easier to design and improve the power efficiency of these reprogrammable chips.

The funds associated with the research chair will support research activities and equipment to develop new FPGA hardware along with software tools to implement designs that make FPGAs more energy-efficient and versatile, potentially improving areas like environmental impact and healthcare. "Data centres consume a tremendous amount of power — about two per cent of U.S. electricity demand," says Betz. "If we can employ FPGAs in data centres, which today mainly use conventional processors, we could reduce the environmental impact while enabling new and better technologies."

Where programmer meets machine

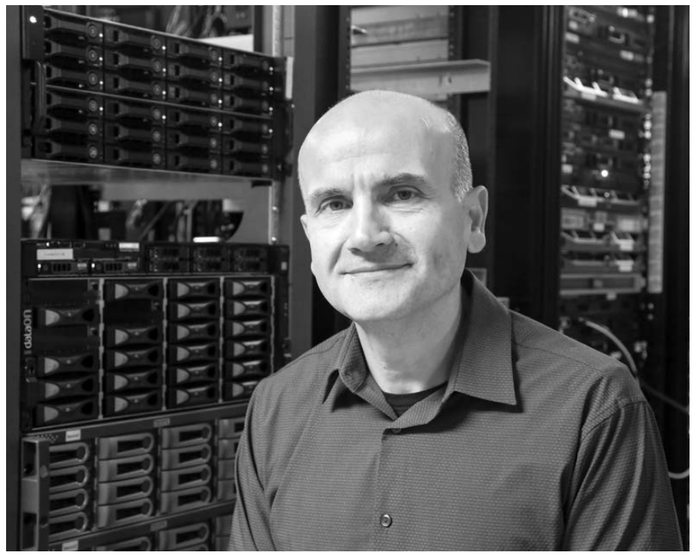

PHOTO BY TYLER IRVING

Machine learning has the potential to revolutionize nearly everything around us — from law to retail, from medicine to transportation, from banking to media. But a key question for computer engineers is: how do we provide sufficient computational power for these technologies?

Professor Andreas Moshovos of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) has been named director of the newly formed Natural Sciences and Engineering Research Council of Canada (NSERC) strategic partnership network on machine learning hardware acceleration. Named the Computer Hardware for Emerging Sensory Applications (COHESA), this network includes 19 researchers across seven universities and eight industrial partners and hopes to solve some of these challenges collaboratively.

"The goal of COHESA is to look at machine learning algorithms and see how we can deliver performance improvement," says Moshovos. "But this network is not just comprised of hardware designers; we're lucky to have some of the very best minds in machine learning working with us." The goal of the network is to bring together researchers from academia and industry to co-design machine learning algorithms simultaneously with the hardware to deliver next-generation deep learning engines.

"For decades, application designers and hardware designers were working in silos; we knew that every two years we would get machines that were twice as fast," says Moshovos. "This is no longer the case, so we are working with top minds in areas spanning computer vision to circuits to help continue the improvement in performance and realize the promise of machine learning."

ECE INVENTIONS DISCLOSED OVER THE PAST FIVE YEARS

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

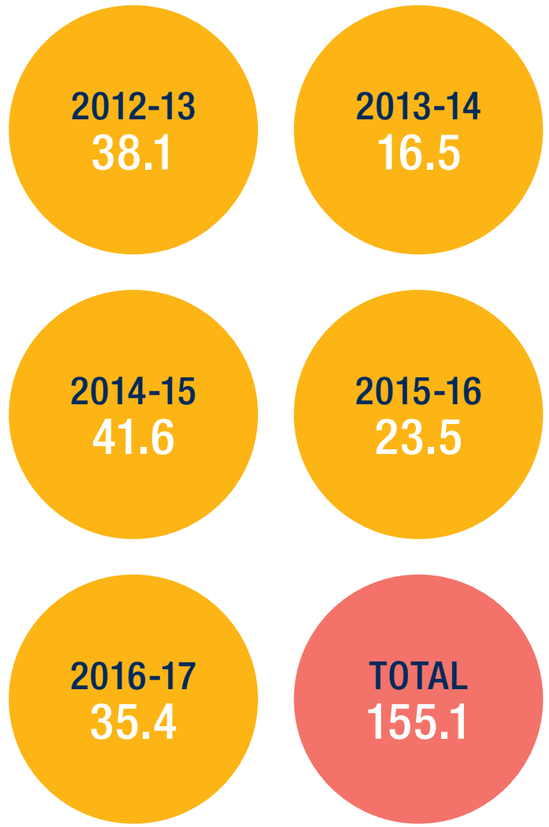

ECE is the top U*of*T department over the past five years for new invention disclosures, new patent applications filed, and new license and option agreements executed.

ECE RESEARCH FUNDING 2011–2012 TO 2015–2016

	FEDERAL	PROVINCIAL	INDUSTRY	OTHER	TOTAL $
2011-2012	10,168,061	1,416,798	2,891,235	3,582,326	**18,058,420**
2012-2013	9,861,957	2,300,498	2,305,784	3,647,052	**18,115,291**
2013-2014	11,668,423	3,310,688	3,001,004	4,557,823	**22,537,938**
2014-2015	14,659,734	4,906,558	3,176,267	2,119,945	**24,862,504**
2015-2016	14,360,193	4,042,251	3,088,408	2,344,957	**23,835,809**

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.

This year marked the expansion of existing industry partnerships as well as the cultivation of several new ones. In October, Huawei CEO Ren Zhengfei visited U of T and met with researchers in ECE to hear about their ongoing research collaborations which are part of a $3 million research partnership agreement Huawei signed with U of T in the spring of 2016. Learn more about upcoming events at **uoft.me/eceevents.**

ECE accounts for more than 40% of all U *of* T Engineering invention disclosures

FUJITSU LABORATORIES TO OPEN R&D CENTRE AT U OF T

Last year, *ANNUM* sat down with Professor Ali Sheikholeslami to discuss the research collaboration he has maintained with Fujitsu Laboratories for nearly two decades. This was an exciting year for this research collaboration so *ANNUM* caught up with Professor Sheikholeslami to hear how the partnership is moving forward.

Q. HOW HAS THIS COLLABORATION ADVANCED OVER THE PAST 12 MONTHS?
A. In the fall, Fujitsu Laboratories Ltd. signed a memorandum of understanding with the University of Toronto to establish a new research and development centre here in Toronto. This is really exciting in the evolution of this partnership.

Q. YOU'VE WORKED WITH YOUR LONG-TIME COLLABORATOR AT FUJITSU, HIROTAKA TAMURA, ON A NUMBER OF PROJECTS OVER THE YEARS, WHAT ARE SOME HIGHLIGHTS?
A. We have worked together on three major areas: non-volatile memory, spin electronics, high-speed chip-to-chip communications, and most recently on a digital annealer. We have published over 30 research papers together — it's been a successful partnership.

Q. TELL US MORE ABOUT THE DIGITAL ANNEALER YOU HAVE CREATED WITH FUJITSU.
A. We use quantum-inspired computing methods to solve combinatorial optimization problems without dealing with the complexities of a quantum computer. This allows the user to respond to a variety of combinatorial optimization problems more quickly. The hardware accelerator can find the optimal solution from a large set of feasible solutions. Tamura-san and I are keen to see this applied to real-world challenges — like optimizing radiation treatment for cancer patients.

Q. DID YOU EVER IMAGINE TWENTY YEARS AGO THAT YOU'D BE WORKING ON SUCH A LONG-TERM AND SUCCESSFUL COLLABORATION WITH FUJITSU?
A. When I was a graduate student, I did an internship at Fujitsu Laboratories in Japan. I had no idea at the time that I would maintain such a collaborative and productive research partnership with them all these years later. I'm looking forward to continuing this research and to the benefits that come with working near my collaborators — it brings a whole new meaning to putting our heads together!

WHERE THE MONEY GOES:

OUR 2017-2018 FINANCIAL PICTURE AT A GLANCE ($ MILLION)

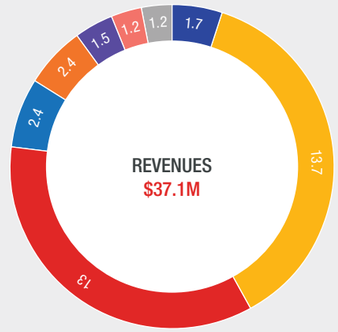

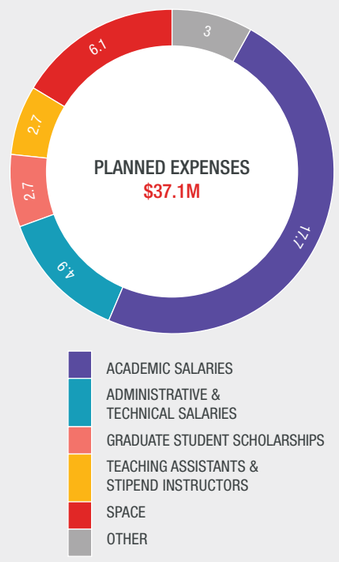

199 Employers Hiring ECE Students on PEY Over the Past Five Years:

2mpower
360incentives.com
407 ETR Concession Company Ltd.
AB SciEx (Formerly MDS Sciex)
Acuity Ads
Adexa Inc.
Aercoustics
AeroFS
AGFA Graphics
Altera
Amazon
AMD
Analog Devices Inc.
Apple
Arista Networks
Arup
AssetMine Inc.
Autoliv Electronics Canada Inc.
Bayer Healthcare
Bell
BELL Mobility
Bibliocommons Inc.
Big Viking Games
BlackBerry Limited (formerly RIM)
BlueCat Networks
BMW Group Canada
Bombardier Aerospace
Buro Happold
Caine Health Inc.
Cancer Care Ontario
CAST Software Inc.
Celestica Inc.
Ceridian Canada Ltd.
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
Deloitte
DEMONWARE
Desire2Learn
Doublethink Inc. (myBlueprint)
Eastern Power Ltd.
eCAMION
Enersource
Environment and Climate Change Canada
Envision Mobile
Epson
Ericsson Canada Inc.
Esna Technologies
EventMobi
Evertz Microsystems Ltd.
FinanceIt Canada Inc.
Finastra (formerly D+H)
Fio Corporation
Fixmo
Flipp Corporation
Ford of Canada
Forma AI Inc.
Fu Tin Business Co. Ltd.
General Electric Canada
General Motors of Canada Ltd.
Genesys Canada Laboratories Inc.
GEO Semiconductor Inc.
Google
Guomao Group Co. Ltd.
Healthcare of Ontario Pension Plan
HH Angus
Honda
Honeywell
HP Canada
Husky Injection Molding Systems Ltd.
Hydro One
IBI Group
IBM
IESO (formerly Ontario Power Authority)
Imagine Communications
Imperial
Indigo Books and Music
Infinera
Intact Financial Corporation
Intel Canada
Intelliware Development Inc.
iNTERFACEWARE Inc.
Intronix Technologies Corp.
Intuit Canada
Isowater Corporation
ITG
Just Energy
Kapsch TrafficCom Canada Inc.
Kijiji
Knowroaming
Labatt Breweries of Canada
Loblaw Companies Ltd.
M&E Engineering
Magna
Marin Software
Marvell Technology Group
MAXXIAN
MediaAdx
Mercatus Technologies Inc.
Microsemi Corporation
Microsoft
Ministry of Children and Youth Services
Ministry of Transportation
ModiFace
Moneris
Morgan Solar
Morgan Stanley
National Bank Financial Markets
National Instruments
National Research Council Canada
NexJ Systems Inc.
NVIDIA
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.
Patheon Inc.
Peraso Technologies Inc.
Periscope Capital Inc.
Perspecsys
Philips (China) Investment Co., Ltd.
Pivotal
PrecisionHawk
Proofpoint Inc.
Psion
QNX, Blackberry
Qualcomm Canada Inc.
Quanta Technology
QuickPlay Media
QuickTapSurvey
Rambus
Rapyuta Robotics Co., Ltd.
Red Hat Canada Ltd.
Rockwell Automation
Royal Bank of Canada
Rubikloud Technologies Inc.
Sandvine
SAP Canada Inc.
SAP Labs, LLC
SciCan
Scotiabank
Security Compass
Semtech
Siemens
Smith and Andersen Consulting Engineering
SOTI Inc.
Southpaw Technology Inc.
Statflo
Strategic Mapping Inc.
Suncor Energy
Sunwell Technologies Inc.
SurfEasy
Symantec
Synaptive Medical
Synopsys
Teck Resources Ltd.
Telus
Temenos
Thales Group
The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
The Hospital for Sick Children
The Independent Electricity System Operator (IESO)
TheRedPin.com Realty Inc.
Think Research
TimePlay
Tofun Marine Supply Inc.
Toronto Hydro
Toronto Transit Commission (TTC)
TradeRev. (formerly NthGen Software)
TransCanada
Trapeze Group
TunnelBear Inc.
TXIO
Tyco Electronics Canada ULC
Uber Technologies, Inc.
Ultimate Software
University Health Network
University of Toronto
Upstream Works Software
Veeva Systems
Verity Studios
Wattpad
Wealthfront
WinMagic Data Security
Xagenic Inc.
Xilinx Inc.
XLV Diagnostics Inc.
Zebra Technologies

ECE CORPORATE PARTNERS, 2013-2017

A.U.G. Signals Ltd.
ABB AB
Advanced Micro Devices Inc.
Alcatel Canada Inc.
Altera Corporation
Analog Devices, Inc.
Asahi Glass Co., Ltd.
Avertus Epilepsy Technologies Inc.
Bell Canada
BLiNQ Networks Inc.
Brammo Inc.
Broadcom Corporation
Candura Instruments
Christie Digital Systems Canada Inc.
Ciena Canada Inc.
Commissariat à l'energie atomique
Dana Canada Corp.
Diros Technology Inc.
E. I. du Pont Canada Company
Ericsson Canada Inc.
Finisar Corporation
Ford Motor Company of Canada Ltd.
Fuji Electric Co., Ltd.
Fujitsu Laboratories Ltd.
Fujitsu Labs of America, Inc.
Futurebound Corp.
Gener8 Inc.
Genia Photonics Inc.
Gennum Corporation
Hatch Ltd.
Havelaar Canada
Hewlett-Packard Company
Huawei Technologies Co., Ltd.
Huron Digital Pathology
Hydro One Networks Inc.
IBM Canada Ltd.
Inphi Corporation
Intel Corporation
International Business Machines
Kapik Inc.
Keysight Technologies Canada Inc.
Kinectrics Inc.
Lattice Semiconductor Ltd.
Maxim Integrated Products, Inc.
NXP Semiconductors Netherlands B.V.
OneChip Photonics Inc.
OPAL-RT Technologies, Inc.
Opus One Solutions Energy Corp.
Polaris Industries
QD Solar Inc.
Qualcomm Canada Inc.
Qualcomm Technologies Inc.
Resertrac Inc.
Robert Bosch Corporation
Royal Bank of Canada
Samsung Advanced Institute of Technology
Semiconductor Research Corporation
Sendyne Corp.
Silicon Mitus Inc.
SINTEF Energi AS
Solantro Semiconductor Corp.
Solar Ship Inc.
Tactual Labs Co.
Taiwan Semiconductor Manufacturing
TELUS
Telus Mobility
Texas Instruments
Thales Canada Inc.
THERALASE Inc.
Toshiba Corporation
Total American Services Inc.
Ultra Electronics
Unisearch Associates
Varilume Lighting Inc.
Wurth Elektronik eiSos GmbH&Co. KG
Xilinx Inc.
Zentrum Mikroelektronik Dresden AG

$16.1 million
Corporate funding to ECE over the past five years

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, FELLOW, PLATFORM ARCHITECTURE // AMD
SAI-KIT ENG, DIRECTOR OF PROGRAM MANAGEMENT // QUALCOMM
DR. RON HO, SENIOR DIRECTOR OF ENGINEERING, PROGRAMMABLE SOLUTIONS GROUP // INTEL
WALTER KINIO, VP, RESEARCH AND INNOVATION // THALES
ALLEN LALONDE, SENIOR INNOVATION EXECUTIVE // IBM CANADA
SONG ZHANG, DIRECTOR, TECHNOLOGY PLANNING & PARTNERSHIP // HUAWEI CANADA

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

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.

INDUSTRY FUNDING
$52,500

+

FEDERAL FUNDING
$75,000
NSERC PROGRAMS

+

PROVINCIAL FUNDING
$75,000
ONTARIO CENTRES OF EXCELLENCE PROGRAMS

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

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

Alumni

STRENGTH OF THE ECE ALUMNI NETWORK

The ECE alumni community is a powerful global network of leaders in a variety of fields ranging from entrepreneurship to sustainable energy, aerospace to artificial intelligence. ECE alumni are extremely generous with their time, and many volunteer to help current students — by meeting fourth-year students on the brink of graduation at our spring Fourth-Year & Alumni Reception, by sharing stories of startup successes and failures as part of our Engineering Entrepreneurship Speaker Series or through mentoring students in ECE CONNECT, our online alumni engagement platform. 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. Stay in touch with your fellow alumni at **uoftececonnect.ca**. A

More than 12 thousand active alumni worldwide

2017 ECE ALUMNI BOARD OF ADVISORS:

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

ALEX GRBIC
CompE 9T4, MASc 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, MASc 8T5,
PhD Santa Clara,
MBA Stanford
Co-founder and CEO of Espresa

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, four 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 Peter Denyer (ElecE 7T2), Alex Grbic (CompE 9T4, MASc 9T6, PhD 0T3), Anthony Lacavera (CompE 9T7) and Michael Salamon (ElecE 8T6).

(PHOTO BY ENGINEERING STRATEGIC COMMUNICATIONS)

ALUMNA CATHERINE LACAVERA RECEIVES DOCTOR OF LAWS, HONORIS CAUSA

This spring, alumna Catherine Lacavera (CompE 9T7) returned to the University of Toronto where she was recognized for her career accomplishments and for her contributions to her alma mater with a Doctor of Laws, honoris causa.

"My connection to the University of Toronto is so special to me," said Lacavera. "It's been 20 years since I graduated from ECE and this is where my career started: U of T Engineering was absolutely the foundation of everything I've done since."

An internationally recognized leader in the field of intellectual property law and the Director of IP, Litigation and Employment at Google, Lacavera and her team have handled more than 1000 patent lawsuits — along with all types of other litigation involving Google — around the world.

A triple alumna of the University of Toronto, Lacavera received her Bachelor of Applied Science in Computer Engineering, Juris Doctorate and Masters of Business Administration at U of T. "I remember thinking that it was such an exciting time to be an engineer," said Lacavera about her own Convocation. "And the same is true now: U of T Engineering students are graduating with so much power to shape and influence the world."

Fourth-Year and Alumni Celebration

PHOTOS BY BEN OUYANG

Alumnus Somen Mondal (CompE 0T2) delivered an inspiring keynote address to graduating fourth-year students at the end of their busy Design Fair week. Students were joined by alumni and professors to celebrate their accomplishments and to welcome them to the ECE alumni community.

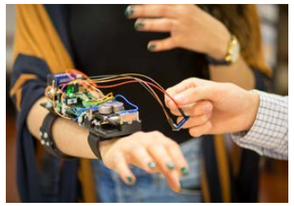

Fall Alumni Networking Lecture & Reception

PHOTOS BY JONATHAN SABENIANO

Alumni of all ages returned to Skule™ to network with their classmates and to hear Professor Milos Popovic discuss the key role electrical and computer engineers play in tackling some of medicine's most difficult neurological challenges.

Spring Reunion

PHOTOS BY LAURA BROWN

ECE alumni returned to campus for our annual lunch and lab tours. Attendees sat down for lunch with their classmates and exchanged Skule™ memories before they toured the Systems Control Lab and the Blue Sky Solar Racing team's workshop. Join us for the next Spring Reunion on June 2, 2018!

Convocation

PHOTOS BY CAITLIN FREE PHOTOGRAPHY

At our Convocation ceremony, graduating ECE students heard a moving speech from alumna Catherine Lacavera (CompE 9T7) who received a Doctor of Laws, honoris causa from the University of Toronto. Our newest alumni were then joined by friends, family, faculty and staff as we celebrated the graduating class of 1T7 and 1T6+PEY at a lunch hosted by ECE.

CLASS NOTES

For this edition of Class Notes, we caught up with some of the members of the ECE class of 9T2 at this year's Spring Reunion where they celebrated the 25th anniversary of their graduation. They've stayed in touch over the years and across vast distances to remain great friends.

On how they met:

SUJATHA SIVARAJAH
ELECE 9T2

I met most of my ElecE friends during frosh week — and actually met at least four of my friends on the very first day of orientation! It must have been the camaraderie derived from finding ourselves travelling back home on the TTC with our purple arms and yellow hard hats. I met the rest of the gang during the first semester of my first year and our friendship just grew stronger over the course of four special years that we spent together — we just gelled and got along really well.

On why they stayed in touch:

NOOR AL-SHAIKH
ELECE 9T2+PEY

We graduated and then we all started our careers in different paths and sometimes different countries. We got married, we had children, we came across personal challenges but we were always in touch, thought of each other and celebrated life in holidays, birthdays, major milestones and we continue to do so.

On memorable moments as an ECE undergraduate student:

DHAMAY KANTHAN
ELECE 9T2

My most memorable times were during FROSH. I loved my FROSH week and it was so much fun that I signed up to be on the team for the years afterwards — I still have my spray painted blue overalls that I wore when I led my new FROSH groups to do the same silly things year after year.

On career trajectories after graduation:

BRANIMIR TASIC
ELECE 9T2, ECE MASC 9T5

I went on to complete my MASc in ECE in 9T5 and then worked at Celestica/IBM in Toronto and Poughkeepsie until 2000. I moved to California to join a startup as a hardware design engineer and that startup was acquired by Cisco in 2001. I stayed at Cisco until 2013 when I joined Infoblox, where I'm currently the Director of Hardware and Platform Engineering.

On how they stay connected:

SUJATHA SIVAMOORTHY
ELECE 9T2

As a group, we've stayed committed to maintaining our friendships through the years, thanks in large part to Noor and Laura Lau (IndE 9T2). They've been instrumental in keeping us all connected — even before we all had email and smartphones.

Advice for ECE's class of 1T7:

MAI MUNSHI
ELECE 9T2

The connections you make at school will be your most precious — try to maintain these friendships you've developed through your undergraduate years. The connections you've made in undergrad are the basis of your most valuable professional — and personal — network.

Are you interested in reconnecting with your classmates? Join our online alumni engagement platform, CONNECT, where you can search the directory of thousands of active U of T Engineering alumni, mentor younger alumni and stay on top of the news from your home department. Register at uoftececonnect.ca.

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

TODAY'S BREAKTHROUGH RESEARCH & TOMORROW'S ENGINEERING LEADERS

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) is the top-ranked department of its kind in Canada and a world leader in research and education. We are generating new knowledge and innovation, creating informed citizens and helping build a more prosperous, inclusive and globally engaged society. Our commitment to excellence in teaching, research and innovation is only possible through generous private support.

By giving to the **ECE Chair's Fund**, you are helping to provide our students with the best possible education and to further our research mission.

The ECE Chair's Fund supports initiatives such as:

Scholarships and Fellowships for Students
As Canada's leading public university, U of T has pledged that no qualified undergraduate student will have to decline admission or withdraw from studies due to financial circumstances. Philanthropic support from our alumni and friends makes this remarkable commitment possible.

High-Impact Research
Donor support enables ECE researchers to create innovative technologies and produce high-impact research in a range of fields, including smart grid technology, mobile application development and cloud computing, next-generation networks, and in emerging areas such as machine learning and robotics.

Lab Renovations
The DC Microgrid Project is a unique opportunity to provide commercial-grade infrastructure and research and teaching facilities while utilizing solar energy as the backbone of the lighting systems at the ECE Energy Systems Lab.

Since 1909, ECE has been at the forefront of the development of the most advanced technology humankind has seen. Building on our success to date, we will seize on new opportunities to strengthen our role in civic and global society, provide a transformational undergraduate learning experience, pave the way for game-changing startups and new industries, drive breakthroughs in health care and create the knowledge that illuminates humanity. Together, we can reimagine the world.

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

If you would like to learn more about how you can support these initiatives, please contact Caitie Spears at caitie.spears@ecf.utoronto.ca or 416 946 0372.

DISTINGUISHED LECTURES SERIES 2017-2018

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

Stephen W. Keckler
NVIDIA

September 28, 2017

Consumer-focused High Performance Computing Architectures

Daniël De Zutter
University of Ghent

October 5, 2017

Broadband Electromagnetic Modelling and Stochastic Signal Analysis of Multiconductor Interconnections

Jelena Kovačević
Carnegie Mellon University

November 2, 2017

From Biomedical Imaging to Online Blogs: Graph Signal Processing

Francesco Bullo
University of California, Santa Barbara

November 30, 2017

On the Dynamics of Influence and Appraisal Networks

Bogdan Kasztenny
Schweitzer Engineering Laboratories Inc.

January 18, 2018

Breaking the Speed Barrier of Today's Line Protective Relays

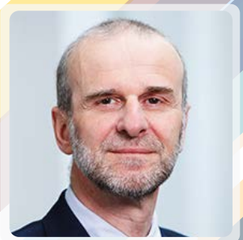

Manfred Morari
University of Pennsylvania

February 1, 2018

The Past, Present, and Future of Control

Rinaldo Castello
University of Pavia

February 8, 2018

Wireless Transceivers: Architectural and Circuit Evolution toward a Complete Integrated Solution

David Reitze
University of Florida

April 5, 2018

Black Holes Last Tango: LIGO and the Dawn of Gravitational-wave Astronomy

All lectures take place in room 1105 of the Sandford Fleming Building at 4 p.m.

uoft.me/ece-dls

DIRECTORY

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

DIRECTORY

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

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

COMMUNICATIONS

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

COMPUTERS

Computer Communications
Computer Software/Hardware

ENGINEERING/SCIENCES

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

INFORMATION TECHNOLOGY

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

ENERGY

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

LIFE SCIENCES

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

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

	●	■	▲	⬢	▰	⬟
LIE, DAVID		•				
LIEBEHERR, JORG	•					
LISCIDINI, ANTONIO	•		•	•		
LO, HOI-KWONG	•					
MAGGIORE, MANFREDI	•		•			
MANN, STEVE		•	•	•		•
MOJAHEDI, MO			•			
MOSHOVOS, ANDREAS		•		•		
NACHMAN, ADRIAN			•			•
NAJM, FARID		•				
NG, WAI TUNG			•	•	•	
PAVEL, LACRA			•		•	
PHANG, KHOMAN				•		
PLATANIOTIS, KONSTANTINOS N. (KOSTAS)	•	•		•		•
POON, JOYCE	•		•	•		•
PRODIC, ALEKSANDAR			•		•	
QIAN, LI	•		•	•		
ROSE, JONATHAN		•				
SARGENT, EDWARD			•		•	•
SARRIS, COSTAS	•		•	•		
SCARDOVI, LUCA			•			•
SHEIKHOLESLAMI, ALI	•	•	•	•		
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, 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 *hi*CUDA, a directive-based language for GPUs. The language facilitates programming GPUs through simple directives added to the sequential code while maintaining the well-adopted CUDA/OpenCL programming models, and does so with no penalty to performance. We developed a prototype *hi*CUDA 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 current exploring compiler-based solutions for automatic extraction of DFGs of applications; determining the best overlay instance for a given application; extending the design of the overlay to multiple FPGA devices; and exploring a just-in-time compilation framework for dynamically and transparently translating binary code into overlay circuits.

Adve, Raviraj

WWW.COMM.UTORONTO.CA/~RSADVE

Adaptive Signal Processing for Wireless Communications and Radar Systems

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

Aitchison, Stewart

PHOTONICS.LIGHT.UTORONTO.CA/AITCHISON

Nanophotonics for Optical Signal Processing and Sensing

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

Amza, Cristiana

WWW.EECG.TORONTO.EDU/~AMZA

Automated Self-Management in Cloud Environments

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

System Support for Parallel and Distributed Software Transactional Memory

Because of the increase in complexity and ubiquity of large-scale parallel and distributed hardware environments, simpler parallel programming paradigms become key. Transactional Memory (TM) 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 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, particularly for specific application domains, such as machine learning.

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-ME, 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

Computer Models of Electrical Rhythms of the Brain

Meso- and micro-level computer models of the electrical rhythmic activities of the brain are developed to (1) elucidate the biophysical basis of cross-frequency coupling in the electrical activities in the neuroglial networks of the brain, and (2) provide platforms for testing neuromodulation strategies for dealing with disorders associated with brain pathologies (in particular epilepsy). Meso-level models consist of coupled nonlinear oscillators based on cognitive rhythm generation dynamics (as developed by our team), whereas micro-level models consist of cellular models based on Hodgkin-Huxley dynamics.

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 neuromimetic 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. Another focus is the analysis of scalp EEG for (1) epileptogenic source localization using cross-frequency coupling features of the brain’s electrical rhythms, and (2) prediction and detection of epileptic seizures using machine learning.

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 where 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 activatinga 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.

Model of the Human Cerebellum

The aim of this project is to develop a mathematical model of the human cerebellum, the part of the brain responsible for motor control. A complete understanding of the cerebellum is one of the great open problems of neurobiology today. As a first step, we focus on experimental findings from cognitive psychology regarding a visuomotor experiment in which a human subject makes repeated fast reaches to a target in the presence of an unknown visual disturbance. This experiment has been investigated in psychology for at least 20 years. Our methodology contrasts with all other approaches taken by psychologists today. Rather than deriving a model based on analysis of experimental data, we first formulate the fast-reach problem as a mathematical problem of control theory. Then we synthesize a solution using control design methods. Finally, we use our synthesized solution as the starting point to fit to the observed data of the experiment. Using eight different qualitative behaviours, we argue that we have the correct model of the experiment. We believe this model is a breakthrough in the understanding of the cerebellum. As a final and significant step we will investigate the underlying neurobiology that gives rise to our derived mathematical model.

Patterned Linear Systems

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

Brown, Stephen

WWW.EECG.TORONTO.EDU/~BROWN

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.

Intelligent Sensing Hardware

We target the development of optimized computing hardware for intelligent sensing applications. Specifically, we are interested in fully customized and dedicated architectures and circuits that provide significant acceleration and low power consumption for a wide variety of applications such as object recognition, speech recognition and natural language processing. As part of this work we are developing and testing prototype integrated circuit demonstrations in advanced CMOS technologies.

Next-Generation Chip-to-Chip Interfaces

This project is exploring circuits and systems to reduce the cost and power consumption of digital interfaces between chips inside of our networking and computing infrastructure equipment. These interfaces are increasingly the bottleneck in overall system performance. Targeting serial data rates above 100 Gb/s, this project combines research on application-specific signal processing and high-speed analog circuit design. Prototype interfaces are being developed in nanoscale (below 20 nm) CMOS technologies to permit their integration within the digital processors that drive infrastructure networking and computing. Close collaboration with industry provides unique opportunities for impact and training.

Cheng, Hai-Ling Margaret

WWW.IBBME.UTORONTO.CA/FACULTY/MEMBERS/CHENG/

3D Bioprinting the Mammalian Heart

The goal of this project is to build a mammalian heart using principles of developmental biology and tissue engineering.

Early Detection of Fibrosis

Fibrosis, or scarring, is a pathological process that occurs in many life-threatening diseases, such as heart failure, kidney failure, and liver disease. To date, there is no specific diagnostic tool that is capable of detecting fibrosis, especially in the early stages when organ function has not deteriorated and is still amenable to treatments that halt progression. Our program in fibrosis imaging seeks to develop a magnetic resonance imaging platform that allows direct detection of fibrosis even at the early stage of scar development. The end goal is to enable early detection in the various organs in the body susceptible to fibrosis.

Magnetic Resonance Imaging of Microvascular Dysfunction

The microvasculature, or the nutritive blood vessels in our body, is essential to maintaining tissue health. Loss of proper microvascular function underlies a broad spectrum of conditions, including neurological disease, diabetes, and cardiovascular disease. This is why being able to evaluate the health of these small blood vessels is important for early 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.

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, scaffold-based and cell-based therapy.

Non-invasive Monitoring of Stem Cell Therapy by Magnetic Resonance Imaging

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.

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 computing elements for many types of computation and network processing. In a cloud-based platform, the hardware of the computing resources is abstracted from the user in the form of a virtual machine that looks uniform regardless of the physical hardware platform. This project explores (1) how to place FPGAs into this virtualized environment when using Open Stack; and (2) how to scale up heterogeneous programming environments to work in a large-scale heterogeneous environment. The prototyping platform is the SAVI Networks platform (savinetwork.ca) where real applications are being built to drive the development of the programming environment as well as to characterize the behaviour and performance of large-scale systems.

Internet-Scale Memory Systems

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

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

Davison, Edward

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

Control of Large-Scale Decentralized Systems

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

Dawson, Francis

WWW.ELE.UTORONTO.CA/~DAWSON

Improving Energy Efficiency of Energy Conversion Processes

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

Draper, Stark

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

Anytime Updating in Synchronized Stochastic Gradient Descent for Machine Learning

In modern deep neural network training problems the scale of data has outpaced the memory and processing capabilities of individual computers. For example, applications in image recognition can require ExaFLOPs of computation and petabytes of data. Many deep neural network training problems are solved using stochastic gradient descent (SGD) algorithm. However, in these large scale training problems, performing SGD in a single processor can be infeasible due to limited storage and computation power or to high wall-clock time. These facts together with the advent of high performance computing, GPU-accelerators, and computer clusters, have driven the development of parallelized variants of SGD. Synchronous SGD is a one method to perform SGD in parallel. In Synchronous SGD, workers in parallel all receive the latest parameter vector, compute their gradients, and send their updates to the master node to be combined. Synchronous SGD waits for all the parallel workers to finish the updates. Therefore, the time to update depends on the slowest worker. These slow workers, referred to as stragglers, can introduce significant delays in obtaining the final output. In general, stragglers cannot be completely removed from a distributed computation system. In this project, we exploit stragglers (of course, along with non-stragglers) rather than avoiding them. The central idea is to give all workers a fixed amount of time to work in each computational epoch. As such, the waiting time of the master node is deterministic (except for possible communication delays) and is no longer limited by variability in finishing times. Furthermore, we use the outputs of all workers, both fast and slow. Our method effectively introduces data redundancy to enhance robustness, but at the same time it uses that redundancy to effect faster convergence rather than producing wasteful (unused) computation. We perform extensive numerical evaluations in the Amazon EC2 cloud.

Distributed Coding of Multispectral Images

The acquisition and compression of multispectral images is often performed in an environment where resources such as computational power and memory are scarce. To that end, we propose a new extremely low-complexity encoding approach for compression of multispectral images, shifting the complexity to the decoding. Our method combines principles from compressed sensing and distributed source coding. Specifically, the encoder compressively measures blocks of the band of interest and uses syndrome coding to encode the bitplanes of the measurements. The decoder has access to side information, which is used to predict the bitplanes and decode them. The side information is also used to guide the reconstruction of the image from the decoded measurements. Our experimental results demonstrate significant improvement in the rate-distortion trade-off when compared to similar low-complexity coding schemes.

Feature Transform Based Semi-Supervised Learning

In this project, we develop a semi-supervised learning (SSL) approach based on feature transformation and the boundary tree algorithm. The recently proposed boundary tree algorithm allows efficient nearest-neighbour classification on labelled data. In general, nearest-neighbour methods work best when input data clusters are well separated and possess a simpler topology. However, most real-world datasets do not have these favourable qualities, therefore they give an inferior performance when tested against nearest-neighbour methods such as the boundary tree algorithm. The recently proposed differentiable boundary tree algorithm makes use of a neural network to transform real-world data to a simpler intermediate representation that works well on the boundary tree classifier. This joint classifier has proven to be producing better results than individual classifiers in terms of accuracy and representability of original data. In our project, we extend this idea of differentiable boundary trees to the SSL paradigm where a large amount of unlabeled data is available with a small number of labelled data. We reformulate the original problem to leverage unlabelled data to train a joint classifier. Performance testing is done on the standard MNIST handwritten image dataset.

Hardware-Aware Motion Estimation via Low-Resolution Motion Hints

Developments in video coding standards such as HEVC, and the earlier development of H.264/AVC, focused on the compression of camera-captured video sequences. However, screen content coding (SCC) has rapidly emerged

as an extension in today's video coding standards to enable compression targeted towards videos containing substantial amount of still or moving rendered graphics, text, and animation. Screen content typically exhibits a high number of identical blocks, large regions containing uniform motion, and no sensor noise. These properties, if properly leveraged, can provide significant performance improvements for SCC. On the other hand, power management, latency, and portability are high-priority design considerations when it comes to the characteristics of the devices to display screen content. We first propose a methodology for hardware-constrained video coding by taking advantage of the freedom given to designers to alter encoder input data at the encoder. We design an algorithm that exploits the characteristics of screen content to improve compression performance while prioritizing hardware design constraints. The algorithm determines a set of motion hints that are passed to the hardware-constrained encoder as side information to aid in the video coding process. We develop an application-specific implementation of this design in partnership with Advanced Micro Devices (AMD) customized to AMD's video encoding operations. The implementation shows that our proposed methodology and algorithms effect significant improvement for SCC under hardware constraints and can come within 1.5 dB of pure software video codec compression. The work with AMD is supported by NSERC.

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

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

Eleftheriades, George

WWW.WAVES.UTORONTO.CA/PROF/GELEFTH/MAIN.HTML

▲ Engineered Materials (Metamaterials) and Surfaces (Metasurfaces) 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 antenna 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, low-cost steerable antenna arrays, 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, Space X and Motorola).

Enright Jerger, Natalie

WWW.EECG.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.

■ Architectures for Energy Harvesting IoT Devices

This project combines research in approximate computing, power-efficient design and machine learning. We are broadly interested in designing next-generation energy-harvesting IoT (Internet of Things) devices with a focus on architectural enhancements to allow efficient computation given limited and unpredictable power budgets and applications such as machine learning which will become prevalent in edge devices.

■ 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.

■ 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

Fully Homomorphic Encryption

Fully Homomorphic Encryption (FHE) is a recently developed quantum-secure, public-key encryption technology that enables arbitrary computations, such as addition and multiplication, on encrypted data, without decrypting the data and without requiring knowledge of the secret key. Multiplication and addition operations are precisely the operations required to perform matching operations directly on ciphertext, thus enabling secure lookup and authentication operations. Homomorphic encryption is unique in its ability to protect encrypted data while at rest, during transport and during computation, thus simplifying key management and reducing attack surfaces. Beginning with Gentry’s mathematical breakthrough, presented in his 2009 PhD thesis, there has been rapid development in the theory and algorithmic implementation of homomorphic encryption schemes. However, to date, commercial homomorphic encryption applications are not yet deployed in the marketplace because of the computational requirements of the mathematical algorithms. We have developed, demonstrated and patented scalable parallel algorithms and hardware acceleration innovations that will enable lattice-based homomorphic encryption with large security parameters to be used at very high transaction rates for a wide variety of performance-demanding applications that are run in either private and/or cloud-based server environments. Future work is focused on improved parallel algorithms, new applications and next-generation acceleration techniques. References A. Khedr, P.G. Gulak, and V. Vaikuntanathan, “SHIELD: Scalable Homomorphic Implementation of Encrypted Data Classifiers,” *IEEE Transactions on Computers*, vol. 65, no. 9, pp. 2848–2858, Sept. 2016.

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 also be 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.

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

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

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.

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 fledging 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 in 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-photonic (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.

▲ Spatio-Temporal Polarization Control of Ultrafast Laser Interactions

Nanoscience and nanotechnology are evolving with the new physical laws encountered as the structures and devices we make become smaller. Quantum-level effects manifest in new ways to manipulate electrical, magnetic, optical, mechanical – virtually all – properties of materials in surprising but exceedingly useful ways. Optical physics and quantum optics have embodied these trends to push resolving power well below the classical diffraction limit in areas of near-field, confocal, plasmatronic and multi-photon microscopy, prodding a Nobel Chemistry Prize in 2014 for super-resolution optical microscopy. Such tools are opening the frontiers of nanoscience by enabling the formation, interrogation and manipulation of nanostructures down to the size of single molecules while laser projection lithography has come to define the nanotechnology forefront for high-volume manufacturing of microelectronic chips with transistor gate widths of 10 nm – just a few atoms thick. In this race to shrink the world, our program is seeking to understand and harness the new optical phenomena found in nanostructures much smaller than the wavelength of light. Beginning in the domain of linear optics, novel types of optical materials with unusual photonic bandgap, metamaterial, or plasmatronic properties have been developed that reshape how light can propagate or reach below conventional diffraction limits with powerfully enhanced optical resolution. The advent of high power lasers with extremely short duration pulses has accelerated the study of nonlinear optical interaction physics. High brightness light poses significant questions when propagating inside bulk transparent and nanostructured media, becoming 'self' manipulated by the nonlinear response of the medium, while also receiving resonance feedback in the proximity of nanostructured media. Such nonlinear interactions define a new unexplored opportunity for manipulating the phase and absorption response that we aim to study and control with new spatio-temporal polarization tools. Our program proposes various modes of amplitude, phase and polarization shifting to create novel beam shapes and patterns — non-diffracting (Bessel), vortex, 'self-accelerating' — and harness advanced real-time characterization tools such that the nonlinear absorption, Kerr-effect, plasma response, phase explosion and shock physics can be followed in transparent media and eventually be controlled to drive open new manufacturing methods for photonics, biology and medical devices.

Hum, Sean

WWW.WAVES.UTORONTO.CA/PROF/SVHUM

● 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.

● 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.

● 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.

● Reconfigurable Leaky-Wave Antennas

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

● Reconfigurable 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 the 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 Hybrid AC-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 Interconnected DC-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.EECG.TORONTO.EDU/~JOHNS

Advanced Interface Circuits for MEMS Technology

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

Joy, Mike

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 photonic spectrally 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

Coding and Modulation for Data-Centre Fibre-Optic Communications

We are investigating information-theoretic limits and coded-modulation schemes for Stokes-vector direct-detection receivers in short-range (<100km) high-speed optical communication systems useful for data-centre applications.

Coding for Resistive Memories

We are investigating efficient coding schemes for information storage in resistive memories organized in multi-layer cross-bar arrays. This involves combinatorial analysis (counting the number of essentially distinguishable patterns) for a given array size, as well as the design of appropriate encoding and decoding functions.

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.

Cyber Security Enhancement for Smart Grids Using an IT/OT Convergence Approach

The December 2015 cyber attack against Ukraine’s power grid revealed the developing capacity of energy infrastructure attackers to manipulate and to spread their actions across a variety of systems in both information technology (IT) and operational technology (OT) domains, such as business networks and supervisory control and data acquisition (SCADA) systems. Traditional security approaches relying on siloed IT security, OT security and physical security programs have become increasingly unable to deter such advanced cyber attacks. Moreover, despite ensuring conformity to different security standards and regulations, traditional security approaches potentially run the risk of becoming inefficient against evasive and dynamic threats that can bypass classic solutions. Hence, a new approach is needed in order to bridge the security gaps resulting from the divergence of IT and OT in energy infrastructures. This divergence stems from disparate technical realities (such as different evolution pace, lifecycles, and drivers) and organizational aspects (dictated, for instance, by the fact that IT and OT teams are often associated with different administrative entities and follow different regulations). As a result, no collaborative security approaches have been tailored to solve cyber security incidents by jointly involving IT, OT and physical security teams. The main objective of this project resides in developing a new converged security approach that not only leads to more collaborative solutions but also to more advanced ones using a broader variety of defences and network management tools.

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.

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 in 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 activities encompass on-board power electronics for electric vehicle drivetrains and charging, power converters for AC and DC microgrids along with their local and systems-level controls and utility-scale power electronic systems for energy storage, FACTS and HVDC transmission. Of specific interest is the development of network architectures and requisite power electronic conversion systems that provide lower-cost and lower-loss integration of renewable energy sources, stationary energy storage and plug-in electric vehicles. Improving robustness and power quality of the electrical grid via intelligent control of power-electronically interfaced sources and loads is an additional area of interest.

Leon-Garcia, Alberto

HTTP://PORTAL.CVST.CA

Connected Vehicles and Smart Transportation

The project on Connected Vehicles and Smart Transportation (CVST) 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. The CVST platform is currently being deployed in smart city projects.

Design of Next-Generation Smart Infrastructures and Service Platforms

The convergence of three technologies—cloud computing, software-defined networking, and Internet of Things (IoT)—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. IoT 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 datacentres 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 testbed (Smart Applications on Virtual Infrastructures) is the outgrowth of a partnership between Canadian industry, academia, government, education research networks and high-performance computing centres. SAVI was developed to prove the concept of 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. The current focus in SAVI is to incorporate Internet of Things at the edge to support experimentation in smart city research.

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/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.

HTTP://IQUA.ECE.TORONTO.EDU

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.

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.

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

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.

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

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.

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 uW to operate from harvested energy, sustaining an autonomous short-range communication to enable an ultra-low-power wireless sensor network. In this case low power consumption and low costs will be achieved by combining the functionality of several building blocks. Indeed, recycling bias current and devices is the prime strategy to minimize area, power consumption and complexity of the transceiver. Furthermore, minimizing the overhead associated with each start-up-wake-up cycle will maximize the efficiency of the node.

Lo, Hoi-Kwong

WWW.COMM.UTORONTO.CA/~HKLO

● 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, augmediation, 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/

GLOGGER.MOBI

⬢ Blockchain Lifeglogging: Lifelong Sensory (Audio/Video) capture

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

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. For more on our startup company see http://wearcam.org/interaxon/ and http://interaxon.ca.

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/

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.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, 39 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 an historical perspective, see http://wearcam.org/par.htm. For more on our AR eyeglass product, 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. See research paper in http://wearcam.org/mannfit.pdf

Physics-Based Modelling Using Presement and Absement

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

WWW.EYETAP.ORG/RESEARCH/MEDR.HTML

Virtual Reality, Augmented Reality, Mediated Reality, and, Most Importantly, Real 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 and works much better. We've developed various mediated-reality iPhone apps as well as eyeglass apps, etc. that help people see better and find their way better. This 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/.See http://wearcam.org/kineveillance.pdf and http://www.eyetap.org/papers/docs/HumanisticComputing_Mann1998_ProcIEEE.pdf and http://mannlab.com

Mojahedi, Mo

WWW.MOGROUP.UTORONTO.CA

Engineering the Electric and Magnetic Dispersive Responses of Artificial Media

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

Nanoplasmonic and Nanophotonic Devices

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

Moshovos, Andreas

WWW.EECG.TORONTO.EDU/~MOSHOVOS

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 this reason, utility and application sophistication. We are analyzing SPARK's implementation of various analytics algorithms. 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. Our first accelerator makes possible for the first time high-end denoising (Bm3D) for high-resolution images and in real-time for HD frames.

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 data center applications. We have developed accelerators that exploit value-based properties and that can be up to 9 times faster than state-of-the-art designs. Our designs do not require any changes to the underlying networks.

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.EECG.UTORONTO.CA/~NAJM

Power Grid Verification

With increased power dissipation and reduced supply voltage, modern large microprocessor chips draw over 150 amperes from the external supply! These levels of current are unprecedented in microelectronics and are a key challenge for design. Apart from the design issues of delivering a well-regulated low-voltage supply at such high current, a key problem for chip designers is to make sure that the increased voltage drop and/or rise (due to IR-drop and/or Ldi/dt drop) in the on-chip power/ground grid does not lead to functional failures. Another major problem is designing the grid so that the grid metal branches do not suffer from electromigration failures. We are aware of at least two industrial instances, (a DSP core and a large microprocessor) where the chip had to be redesigned because functional failures

on silicon were caused by current-induced noise on the power grid. However, checking the grid node voltages and branch currents is very time-consuming and expensive, so that it is often done incompletely or not at all. We are developing efficient techniques for verifying that the voltages and currents of the power/ground grid are safe and within user specifications and, if the grid is found to be unsafe, for redesigning and optimizing the grid to achieve safety.

Ng, Wai Tung

WWW.VRG.UTORONTO.CA/~NGWT

Design and Fabrication of GaN Power Transistors

Gallium nitride (GaN) power transistors promise to be the game changer for the next generation power converters. Traditional silicon based power MOSFETs are already reaching their material limits for power conversion. Aluminum gallium nitride (AlGaN)/GaN high-electron-mobility transistors (HEMTs) with high breakdown field, high-mobility 2-D electron gas (2DEG), high saturation velocity, and low intrinsic carrier density are emerging as the ideal candidate for the implementation of high frequency (10s to 100s MHz) and high efficiency (>90%) converters. The wide bandgap materials also allow GaN power devices to operate with a high junction temperature (>200°C), relaxing the need for expensive heat removal mechanisms. The availability of GaN on silicon substrates can further reduce the cost of fabrication by making good use of existing silicon-based manufacturing facilities. This project is focused on the development of two key enabling technologies to fully exploit the potential of integrated GaN power converters: enhancement mode metal insulator semiconductor field effect transistors (MISFETs) and intelligent gate driving techniques. Currently, the majority of the GaN based power transistors are depletion mode HEMTs. The normally “on” characteristics made them less favorable when robustness and fail-safe characteristics are critical. Monolithic integration of silicon-based devices and GaN devices is still not practical at this time (although GaN-on-Si technology will eventually allow the co-existence of CMOS circuits and GaN power HEMTs). As a result, there is a critical need to develop novel methods/structures to implement true enhancement mode GaN power devices with silicon processing compatibility. The development of enhancement mode GaN power transistors will involve the investigation of suitable gate insulators and device structures, passivation techniques and silicon compatible ohmic contacts. In order to take advantage of the inherent switching speed of the GaN power devices to implement integrated power converters with fast transient response, high efficiency, and compact form factor, dedicated gate driving techniques are essential. The rapid switching will produce ringing oscillation, leading to unwanted power losses (reduced efficiency) and electromagnetic interference (EMI). The second theme of this project encompasses the design of intelligent gate driver ICs to provide precision control of the gate voltage to suppress these issues.

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 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.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 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.

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.

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.

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.

Poon, Joyce

WWW.PHOTON.UTORONTO.CA

Integrated Neurophotonics

We are designing integrated neurophotonic probes 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.

●▲⬣ Integrated Photonics for Communications, Computing, and Sensing

We invent, design, fabricate, and measure integrated photonic devices and circuits, such as electro-optic transceivers and optical switches, for communications, computing, and sensing. We work with a wide range of photonic platforms, often in partnership with industry and research institutes from around the world. 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 power efficient, high-speed, and compact, to meet the needs of communication, computing, and sensory systems of the future.

Prodic, Aleksandar

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

▲▰ Power Management and Integrated Switch-Mode Power Supplies

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

Qian, Li

WWW.ECF.UTORONTO.CA/~QIANLI

▲⬣ Fibre-Optic Sensing

We utilize photonics technology to create instrumentation for fibre-optic sensing and metrology. Our frequency-shifted interferometry technique has been demonstrated to have a variety of applications, such as dispersion measurement; fibre length measurements; multipoint optic sensing for stress and/or temperature sensing (used in civil structures); multipoint chemical gas sensing for environmental monitoring as well as industrial monitoring in hazardous environments; and liquid level sensing in cryogenic environments (required in space applications). We are currently extending this technique to vibration sensing. Another technology we developed is called the virtual-reference interferometry (VRI). It was developed by one of our graduate students, and has been commercialized by a successful start-up company. It is used to characterize the dispersion of optical fibres, waveguides, and devices.

● 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/RESEARCH-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.

■ Mobile Apps and Chatbots to Help Smokers Quit

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) has two thrusts: first, building and validating 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. The second thrust is to build a 'chatbot' (a program that talks to people through texting apps, or ultimately speech) that helps smokers decide to quit, and to help them practice the actual act of quitting.

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 themself, 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 and Perovskites for Light Emission Applications

We develop new 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 and also perovskites engineered to provide narrow light emission in the green, red, blue, and IR.

▲■ Renewable Fuels

Our investigations focus on converting solar energy into stored chemical energy, mimicking the natural photosynthesis process. We develop highly active catalysts for the splitting of water to hydrogen and for the reduction of CO_2 to carbon-based products (CO, methane, methanol) through electrochemical processes using electricity from solar photovoltaic cells. The ultimate aims of our projects are to generate renewable fuels that help 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 the demand for signalling speeds of 56 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.EECG.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 to operate significantly 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).

Grid-Connected Energy Storage Scheduling and Control

This project focuses on the scheduling and control of grid-connected energy storage to achieve multiple objectives (e.g., load balancing, frequency regulation, and prevention of under-frequency load shedding). The control design includes online parameter estimation during initial power injection along with use of model-predictive control to recover energy without the introduction of additional, unwanted disturbances.

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 likes the Vickrey-Clarke-Groves mechanism enable us to make such tactical behaviours unattractive by imposing auxiliary payments like an upfront tax.

Learning to Manage Electrical Loads

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

Trescases, Olivier

WWW.ELE.UTORONTO.CA/~OT

Battery Management for Electric Vehicles

Despite numerous technological innovations, the proliferation of EVs in Canada is primarily limited by the range and cost of today's vehicles. Reducing the cost and extending the range of EVs are a major multidisciplinary challenge for the global automotive industry. Advances in lightweight materials, battery chemistry, battery management and power electronics are needed to meet future customer expectations and convert entire fleets from gasoline to EV technology. Another major hurdle in the widespread acceptance of EVs is the uncertainty about the lifetime and reliability of the battery pack, especially in the harsh Canadian climate. This has delayed the adoption of Lithium-Ion (Li-ion)–based battery technology until very recently, despite vastly superior energy density compared to the Ni–MH batteries used in the first generation of hybrid vehicles. Making better use of the energy capacity by increasing the system efficiency is the key to reducing the overall size and cost of the EV battery. Regenerative braking (Regen) is often used in electric vehicles to capture kinetic energy that is otherwise wasted in the brake pads when the vehicle comes to a stop. Instead of simply applying the mechanical brakes during deceleration, an EV equipped with a Regen system uses the motor as a generator in order to transform mechanical energy into stored charge in the battery. Even the latest lithium-based batteries have a relatively poor ability to quickly absorb energy without affecting long-term performance. The maximum output power of modern Li-ion batteries is typically at least three times

higher than the maximum input power. Repeatedly using Li-ion batteries to both absorb this large negative power burst of Regen and provide a large positive power burst during acceleration can significantly raise the pack temperature and accelerate aging. Automotive-grade Ultracapacitors (Ucaps) have recently been developed as an energy storage technology to complement batteries. Commercial Ucaps have input and output power densities on the order of 12 kW/kg, which is at least one order of magnitude higher than that of Li-ion batteries. On the other hand, the 6 Wh/kg specific energy of these Ucaps is at least 10 times worse than that of Li-ion batteries, leading to the concept of using a hybrid storage system consisting of a smaller Li-ion battery and a Ucap. Using this approach, the battery serves purely as an energy tank, while the Ucap is sized to meet the surge input and output power requirements. Effectively managing the energy flow between the Ucap, the battery and the motor requires new power-electronic topologies and advanced control schemes. The main goal of this project is to develop new models, control schemes and power-electronic converters to extract the maximum performance from modern EV energy storage systems.

High-Frequency Digitally Controlled DC-DC Converter ICs

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

Power Converters for High-Efficiency LED Lighting

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

Power Electronics for Photovoltaic Applications

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

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.

Non-Invasive Assessment of Coronary Artery Disease via Computational Fluid Dynamics

In Canada, about 1 person out of 7 dies because of coronary artery disease. Coronary arteries play a vital role, bringing oxygen-rich blood to the heart. However, lipid plaques can progressively occlude them, leading to heart failure. With graft surgery, new vessels are created to increase blood flow to the heart and prolong life. However, graft failure is quite common (for example, about 50% of vein grafts fail within 10 years). Unfortunately, we still do not know exactly why grafts fail and, consequently, how to prevent their failure. Using computational simulations guided by medical images, we are investigating how blood flows in coronary arteries after graft surgery, in order to gain new insights on the mechanisms that contribute to graft failure. Insights will then be translated into better treatment procedures, with the ultimate goal of improving patient care and reducing costs.

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 industry.

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.

Valaee, Shahrokh

WWW.COMM.UTORONTO.CA/~VALAEE

Localization of Wireless Terminals in Indoor Environments

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

●● Wireless Communications in Vehicular Environments

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

Veneris, Andreas

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

■ CAD for VLSI Verification, Debugging, Testing and Synthesis

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

Voinigescu, Sorin

WWW.EECG.TORONTO.EDU/~SORINV

▲ 56 GS/s 7-bit DAC and ADC

The research focuses on architectures and physical implementation of low-power 56-GS/s, 6-8 bit digital-to-analog and analog-to-digital convertors for next-generation optoelectronic transceivers with multilevel modulation formats. The circuits will be fabricated in the world's most advanced 55-nm SiGe BiCMOS 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. An 8-bit 70-GS/s DAC with over 32-GHz analog output bandwidth and 1.2-Vpp differential output swing has already been fabricated and is being tested. Two 6-bit SAR ADCs with 32-GHz analog input bandwidth and 40-60 GHz sampling clocks, based on a patent-pending architecture, have been designed and are now in the fab.

▲ 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, 28-nm and 22-nm SOI CMOS DAC and 60-GHz bandwidth 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 have been published. This year we are focusing on the implementation of the analog front end of ADCs with over 200 GHz sampling rates and 60-GHz analog bandwidth. Already, a record-breaking DC-60GHz non-overlapping quadrature clock signals generator with 25% duty cycle, based on a novel digital mm-wave architectures, is being measured in our lab.

▲ Atomic-Level and 2D Crystal Electronic Devices

This exploratory research focuses on the simulation, design and fabrication of novel nanoscale semiconductor and metal nanowire transistors and their natural evolution to coupled-quantum well qubits beyond the 2030 ITRS horizon. The ultimate goal is the physical implementation of sub-7-nm-gate-length transistors and qubits. Some of the results of this work have been published in two IEEE *Proceedings* papers published in June 2017. This is the final year of the project. We are setting up international collaborations with Intel, Global Foundries, and several European research groups to continue this work with focus on the development of silicon qubits in commercial CMOS foundry processes for room-temperature quantum computing.

▲ High-Voltage Differential Distributed Amplifier in Si Technology

Distributed linear large-swing driver topologies have been investigated for operation at 128 GBaud with 6-Vpp differential output swing as needed to drive optoelectronic modulators in future 1-Tb/s fibre-optic systems. Two distributed drivers were designed, and fabricated in an advanced 55-nm SiGe BiCMOS process through Ciena. Both circuits were tested after fabrication at the University of Toronto and achieved world-record performance with NRZ operation at 120 Gb/s with 6V differential output swing, significantly larger than the state of the art. Operation with 4-PAM and 8-PAM signals up to 64 GBaud and 56 GBaud, respectively, was also demonstrated. A paper was presented at the 2017 IEEE RFIC Symposium in June 2017. The paper was selected as a finalist for the Best Student Paper Award.

▲ 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. Three papers on these circuits have been presented at IEEE IMS 2017 and IEEE RFIC 2017. This year, the research focuses on developing a daisy-chain sensor in the 60-GHz band and also on exploring new low-power architectures and 160-GHz operation in 22-nm and 12-nm FDSOI CMOS technologies.

▲ Receiver and Transmitter Circuit Blocks for Si Photonics

Several silicon photonics linear driver topologies were investigated. Two of these drivers were co-designed with a Si photonics Mach-Zehnder Modulator in a hybrid electronics-photonics design kit and were implemented in a 55-nm SiGe BiCMOS technology with 350 GHz fT and fMAX. A new linear lumped circuit topology was invented, designed and fabricated in the same technology. The circuit was tested in the U of T lab with a new 92-GS/s arbitrary waveform generator capable of generating 64-GBaud 4-PAM and 8-PAM signals. The new modulator driver is the first to achieve 56-GBaud 8-PAM and 64-GBaud 4-PAM operation with output differential voltage swing of over 4.5 Vpp. A world record data rate of 156 Gb/s was achieved with 56-GBaud 8-PAM modulation. The results were presented at IEEE IMS 2017. The research continues with the investigations of several other driver topologies which have been fabricated and are now being tested, as well as with the design of a 130-GS/s 5-bit flash ADC with over 40-GHz input bandwidth and which is now in the fab.

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 the 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 Pelvic Dysfunction

The objective of this project is to develop therapeutic platforms for treating pelvic dysfunction in individuals with neurological or idiopathic disorders. This involves (1) the investigation of neural mechanisms underlying bladder-inhibitory reflexes evoked by peripheral nerve stimulation; (2) development of minimally invasive methods of electrically activating these neural pathways; and (3) the clinical translation of these technologies in patients. Our current areas of interest include chronic clinical problems of bladder and bowel function.

⬟ 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

● ⬣ Cooperative Communications for Wireless Cellular Networks

Prof. Wei Yu's research program aims to use novel cooperative 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.

● ⬣ Massive Device Connectivity

Massive connectivity is an important application scenario for future wireless networks in which millions of devices may be connected to the network, but with possibly sporadic traffic. The main goal of this research project is to understand the fundamental limits and to design coding strategies to support massively connected low-latency wireless random access in a cellular environment.

● Ultra-Reliable Communications

Prof. Wei Yu engages in research in ultra-reliable communications for next-generation cellular networks. These new communication protocols are motivated by emerging machine-type communication scenarios where reliable and low-latency transmissions are required for sensing, control, and actuating applications.

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.EECG.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 their undergraduate programming class.

Zhu, Jianwen

WWW.EECG.TORONTO.EDU/~JZHU

■ Analysis and Characterization of Cloud Applications with Flash Memories

In the past, data centre 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 projects 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.

■ Erasure Coding Acceleration

Erasure coding is a generalization of RAID that can enhance the reliability and availability of peta-byte scale storage systems. One problem that prevents its adopting in high-performance storage systems its computational complexity that commands 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 (GC) 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.

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

ECE RESEARCH
DIRECTORY

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