YEAR IN REVIEW

ANNUM 2013

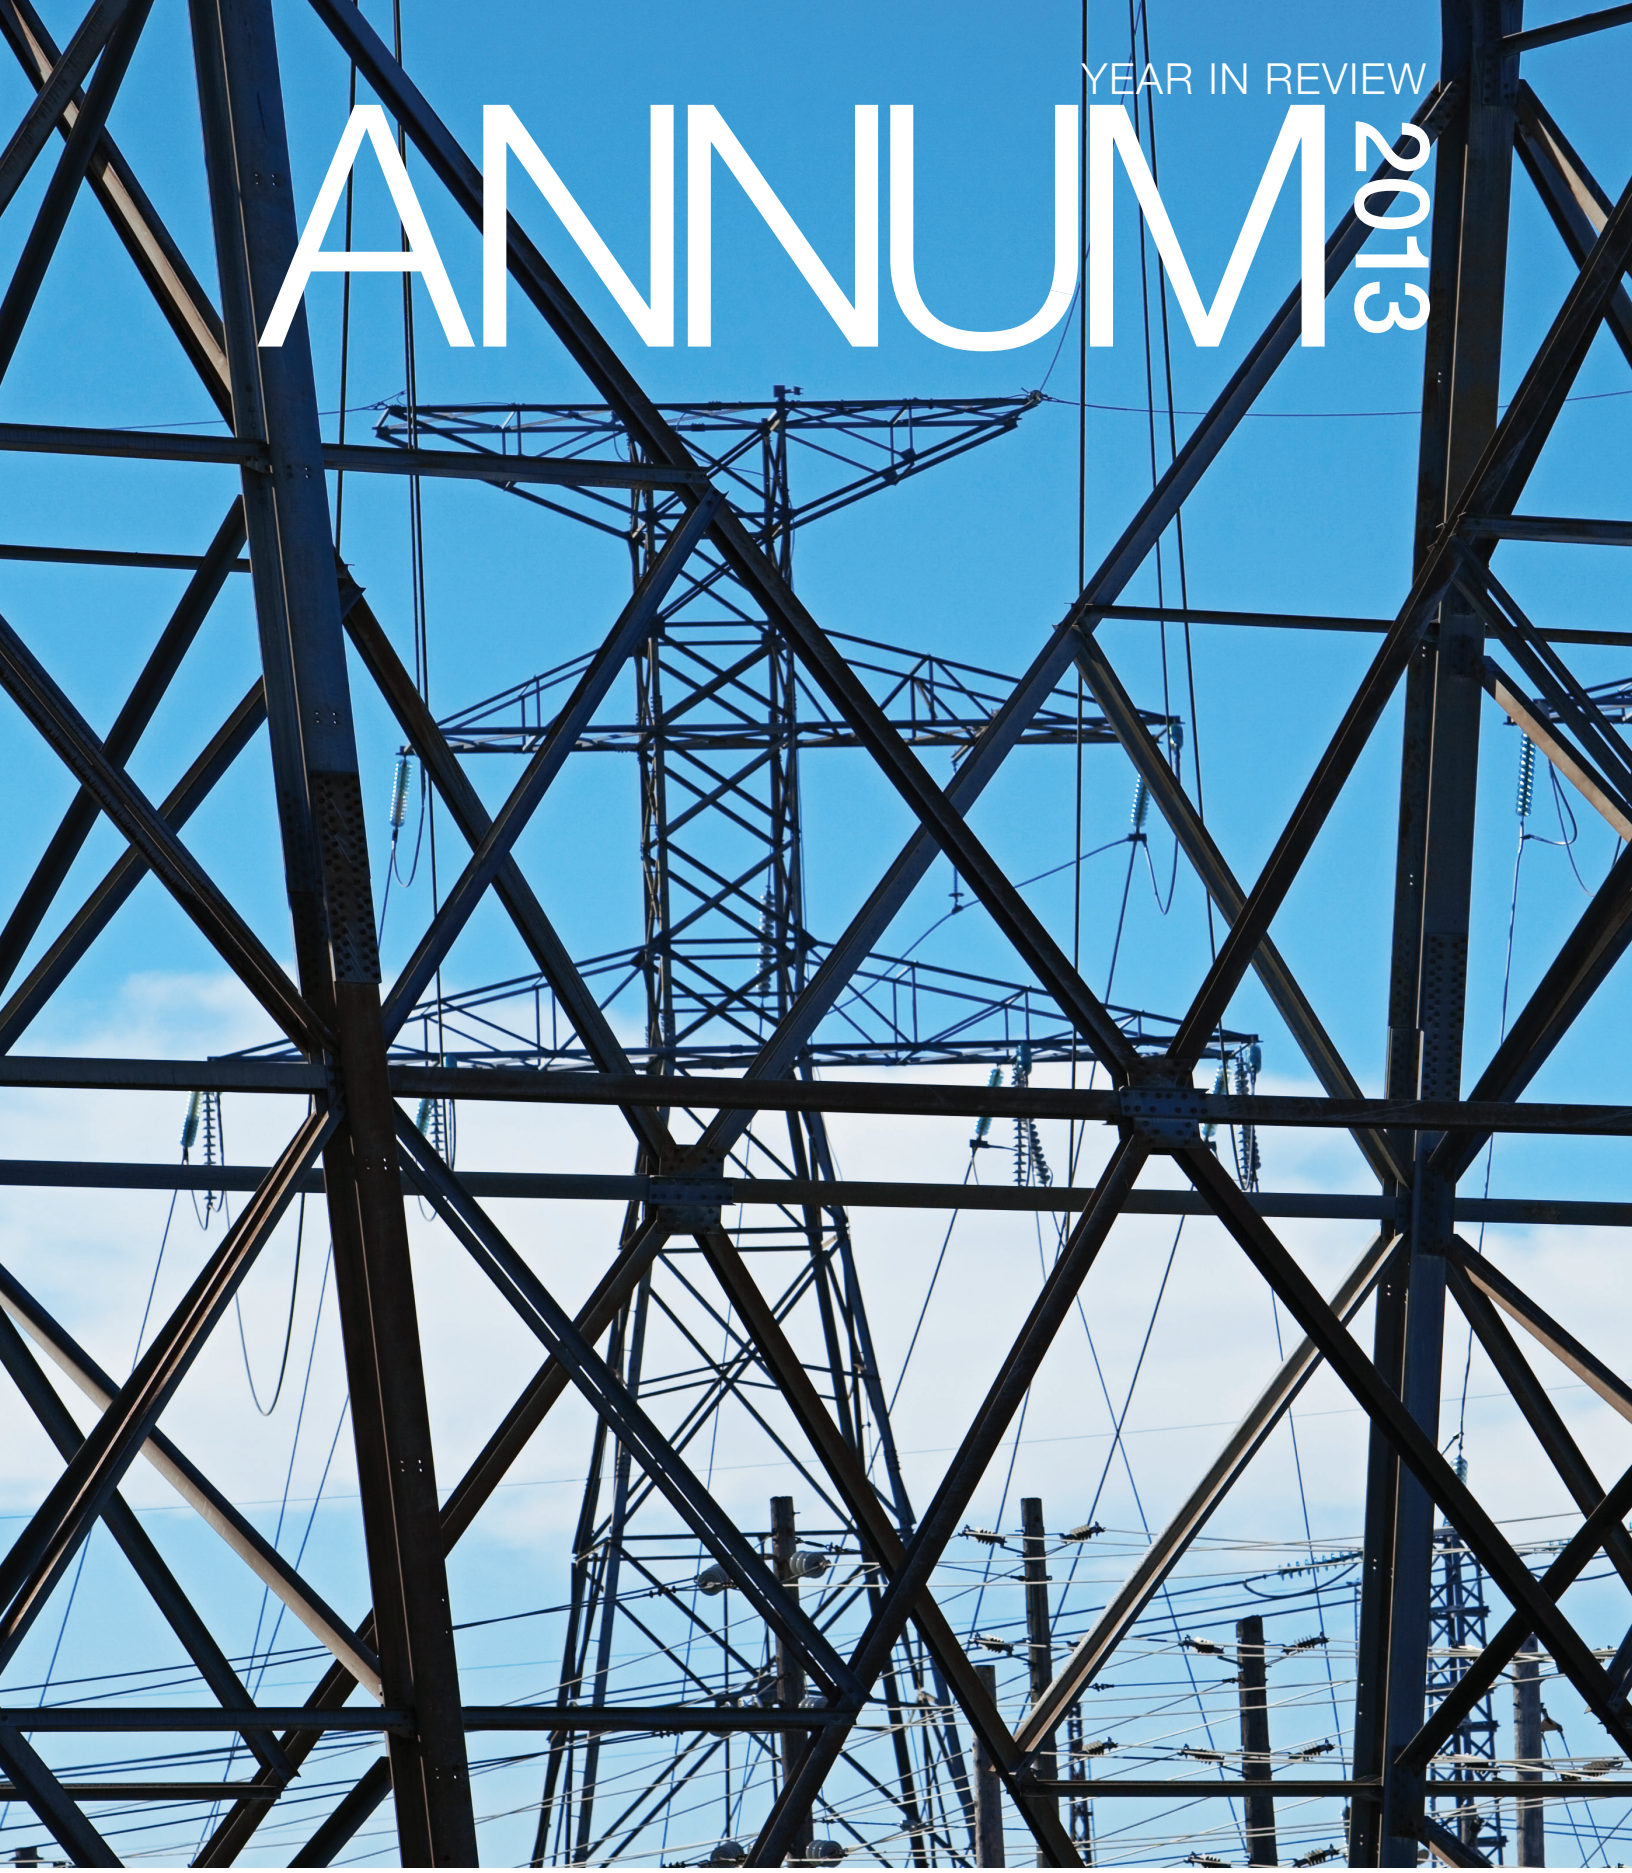

The Edward S. Rogers Sr. Department
of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

Table of Contents

Our Story

1st

The ECE Department at the University of Toronto ranks first in Canada according to **QS World University Ranking 2013.**

The past year has been one of growth, change and discovery in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) at the University of Toronto. In this year's annual report we have collected key statistics and stories that reflect the state of our department over 2012–2013.

With 78 research faculty, more than 1,200 undergraduate and 500 graduate students walking our halls, ECE hums with energy. We have much progress to report: over the past 12 months we have undergone two external reviews, scaled up our offerings and enrolment in the Masters of Engineering (MEng) program, assembled an Alumni Board of Advisors, and invested in our students and infrastructure.

In May 2013, we presented the results of an exhaustive self-study to an external review committee composed of faculty members from leading Canadian and American universities. The group evaluated our department on its academic programs, research, organizational and financial structure, long-range planning and internal and external relationships. The results were released in September 2013, and the review committee praised the outstanding quality of our faculty, research, students and programs. Also this past academic year, our undergraduate program underwent a thorough review by the Canadian Engineering Accreditation Board to ensure its compliance with professional engineering standards. The Board renewed our accreditation and gave ECE a glowing review.

As part of our ongoing focus on strengthening ties with industry and practicing engineers, we continue to expand our MEng curriculum and admissions. We have now rolled out five courses designed specifically and exclusively for MEng students, in electronics, computers, power systems, and two in communications. We also saw MEng enrolment reach the 150 mark including both full- and part-time students.

As chair, I am delighted to announce the formation of ECE's Alumni Board of Advisors. This diverse group of graduates and industry leaders meets throughout the year to consult on and assist with the department's alumni and advancement activities and strategic direction. You can find profiles of our board members in this issue's Community section (page 26).

Finally, the University of Toronto's Boundless Campaign is fully underway and ECE is dedicated to raising support for more graduate student scholarships, an ambitious renovation of our Energy Systems Lab, and the new Centre for Engineering Innovation & Entrepreneurship, which is slated to break ground in 2014.

The stories in this magazine highlight the successes of our world-class students, faculty, partners and alumni. I hope you find the issue both informative and enjoyable, and I hope it makes you proud of this department, as I am. I welcome your comments and feedback on this publication and our direction—you may reach me at **chair@ece.utoronto.ca**.

Our History

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering (ECE) at the University of Toronto was launched in 1909 as an offshoot of the Department of Mechanical Engineering. Up until the early 1960s only two such departments existed in Ontario — the other was at Queen's University in Kingston.

With the return of First World War veterans in the fall of 1919, the department saw enrolment swell from 241 to 819 students. It offered a practical undergraduate curriculum similar to that of the mechanical colleges in the United States. It was this North American model under which the department developed, as engineering was not considered a discipline for academic study in Britain or Western Europe at the time.

By the mid-1920s, the department expanded its applied science component and instituted the Master of Applied Science degree. The first PhD in Electrical Engineering was awarded in 1951.

Early research groups included those based in classical electromagnetic theory, followed by control engineering (dubbed systems disciplines) in the 1950s. Shortly thereafter came communications engineering, and electronic device and circuit engineering. In 1962 the department formed the Institute of Biomedical Electronics, now called the Institute of Biomaterials & Biomedical Engineering, and was joined by photonics. Computer engineering found its place in the department around 1965.

In June 2000, Ted Rogers Jr. made a substantial donation in honour of his father, who was a student of the department from 1919 to 1921. The department is now known as The Edward S. Rogers Sr. Department of Electrical & Computer Engineering.

4,400

Square metres of research labs in ECE.

2,700

Square metres of lab space for undergraduate teaching in the ECE department.

6

The number of buildings that comprise the ECE infrastructure.

Iron Ring

Many Canadian-trained engineers wear an iron ring on the pinky finger of their working hand. This tradition originated with University of Toronto mining engineering professor Herbert Haultain after the collapse of the Quebec Bridge in 1907, a construction disaster that killed 75 construction workers and was caused by poor planning and design, and inadequate diligence by the project's engineers. Professor Haultain wished to create stronger bonds between engineers and remind them of the obligations and ethics inherent to their profession. He enlisted the help of English writer Rudyard Kipling (author of The Jungle Book) to create a dignified obligation and ceremony known as The Ritual of the Calling of an Engineer, where new engineers undertake the obligation and receive their rings. The iron ring is worn as a symbol of the engineer's responsibility to society.

pg16 LegUp

UNDERGRADS AT A GLANCE:

12th

University of Toronto ranks 12th in the world for Engineering/Technology and Computer Science, according to the **Academic Ranking of World Universities 2013.**

The Undergrad Experience

Seated at the heart of most technical advances made today, electrical and computer engineering is the engine that powers 21st century technology. An undergraduate degree in electrical and computer engineering offers the widest range of career opportunities.

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering at the University of Toronto offers the broadest curriculum in Canada, making it the school of choice for students who prefer flexibility in their course selection and seek to learn from some of the world's most sought-after industry leaders. These professors, renowned experts in their fields, not only lecture but also supervise vibrant on-campus research laboratories, facilitating the transfer of knowledge from the lab to the classroom.

Undergrad Enrolment 2004–2005 to 2013–2014

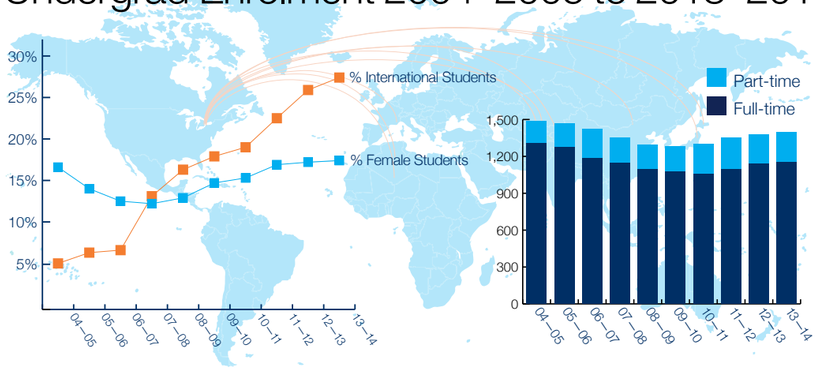

In ECE, the first two years of study provide essential background in basic science and mathematics and introduce students to important concepts in electrical and computer engineering such as circuits, digital systems, electronics, and communication systems. These two years of study are identical for both the electrical engineering and computer engineering programs.

In third year, students choose an area of specialization. In both third and fourth year, students may choose from six areas of study, depending on their individual strengths and interests. The options include: Photonics and Semiconductor Physics; Electromagnetics; Analog and Digital Electronics; Communications, Systems Control and Biomedical Engineering; Computer Hardware; and Computer Software.

Curriculum streams have been designed to help guide students; however, students are free to create their own unique path of study. Learn more about the eight curriculum streams at **uoft.me/ececurriculum.**

Undergraduate Awards

Academic Awards

Each year ECE recognizes the top three students in both Electrical and Computer Engineering with ECE Outstanding Student Awards.

The Sedra Medal is awarded annually to two students in the graduating class who have earned the highest cumulative grade point average in Electrical and Computer Engineering. For 2012–2013 the recipients were Zhi Li (CompE) and Weijian Zhou (ElecE).

Design Awards

Engineering is about more than marks—teamwork, creativity and execution are essential to success. Each year ECE awards three prizes for teams who have demonstrated excellence in their final year design projects. For 2012–2013:

The CNIB Hochhausen Prize for Excellence in Accessible Design in Engineering for People who are Blind or Partially Sighted was awarded to Wasif Iqbal, Ahmad Diab Marzouk and W. A. Kavindu Gayanath Amarasingha for their design project titled "Navigation System for the Visually Impaired."

The Gordon R. Slemon Design Award was awarded to Miad Fard, Richard Medal and Mehrad Mashayekhi for their project titled "Digitally Configurable Lab Platform for Design-Oriented Teaching of Analog Electronics."

The Aloha Design Award was awarded to Leon Chan, Guo Qin Low and Chia Chen Tan for their project titled "FPGA Video Processing – Cartoonizer."

Professional Experience Year

The Professional Experience Year (PEY) internship program allows students to apply their engineering knowledge in a 12- to 16-month project-based professional internship. The duration of the placement offers enough time for students to get involved in large-scale projects, build relationships with employers and achieve professional accomplishments. The PEY internship program has been running for more than 20 years and has earned a terrific reputation in both academic and industry circles. PEY offers students an outstanding education, a range of eligible engineering career paths to choose from and strong established industry partnerships. This year, PEY placed students in more than 150 companies located around the world: Canada, the United States, Belgium, Germany, Finland, China and South Korea.

ECE PEY Placements, 2004–2005 to 2013–2014

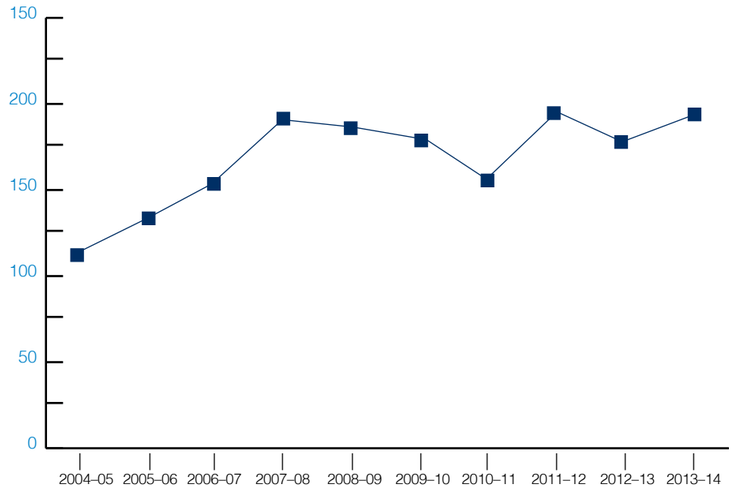

Blue Sky Solar Racing cracks top 10 at World Solar Challenge

U of T's Blue Sky Solar Racing team recorded an impressive eighth-place finish at the 2013 Bridgestone World Solar Challenge, a gruelling 3,021-kilometre race for solar-powered cars across the Australian outback.

The Blue Sky team achieved its goal of finishing in the top 10 among 23 teams from around the world. Blue Sky's car, B-7, completed the race in 45 hours and 38 minutes, achieving an average speed of 65.71 km/hour. The team recorded its journey online and in video, and live-tracked the vehicle throughout the race.

Built and driven entirely by students, B-7 features the latest advancements in photovoltaic technologies, a brand new aerodynamic design and improved vehicle dynamics systems. It weighs about 20 per cent less than its predecessor, Azure, which placed 24th out of 37 teams in the 2011 World Solar Challenge.

Blue Sky Solar Racing has been building and racing solar cars since 1996, when a small group of Mechanical and Electrical & Computer engineering students formed the Advanced Solar Electrical Vehicle Program. The 21-member team unites students from across Engineering departments, and is supervised by ECE Professor Olivier Trescases and Mechanical & Industrial Engineering Professor Kamran Behdinan.

ECE Student Among The Next 36

The Next 36 is a prestigious entrepreneurship leadership initiative helping to launch business careers for promising Canadian undergraduates.

It aims to transform participants into Canada's top entrepreneurs through a mix of teamwork, business experience and intensive entrepreneurship instruction.

U of T is a founding partner in The Next 36, launched by a group of business leaders, academics and entrepreneurs.

During the eight-month program, the students work in multidisciplinary teams to develop, launch and sell a product or service for the mobile or tablet market. The students also receive entrepreneurship training and mentorship from world-class business leaders.

The program provides up to $80,000 from top venture capitalists, a comprehensive offering of in-kind resources and instruction from some of the world's top faculty.

ECE student **Nikita Tarakanov** was selected among the 2013 Cohort. His venture, Sesame IO, is a cloud-based work environment geared towards saving teachers' time.

255

Undergraduate degrees awarded for 2012–2013

IEEE

The Institute of Electrical and Electronics Engineers (IEEE, pronounced eye-triple-ee) is the world's largest professional association, dedicated to advancing technological innovation and excellence for the benefit of humanity. With more than 425,000 members in more than 160 countries, IEEE's objective is to build a network of professionals and students in the electronics field and to promote the latest technological developments through conferences and published literature.

University of Toronto's IEEE Student Branch is the largest such branch in Canada and hosts numerous conferences, tours, dinners and networking opportunities throughout the year, as well as technical seminars and academic sessions aimed at preparing students for the professional and academic world.

This year two ECE students were honoured with IEEE scholarships worth $2,000 each:

Yue Lu received the IEEE Canada Toronto Section, Bruno N. Di Stefano scholarship

Chuanwei Li received the IEEE Canada Toronto Section scholarship.

Purple Power Lighting Event

Undergrads lit the CN Tower in purple to celebrate National Engineering Month in March.

WISE

Women in Science and Engineering, or WISE, is a co-ed student organization open to all University of Toronto students, staff and alumni. WISE offers activities and outreach programs designed to facilitate networking opportunities between students and professionals from a variety of industries. The group provides a welcoming, supportive environment through social events and a mentorship program that aims to foster positive relationships during university and after graduation. Established in 1999, WISE arose as a forum for peers to share their experiences. With support from the University and student body, this chapter of a national association has developed into a recognized campus organization.

GLEE

This spring, U of T's Faculty of Applied Science & Engineering welcomed more than 90 newly admitted female students with the Girls' Leadership in Engineering Experience weekend—also known as GLEE. Attendees got the chance to learn more about student life and seek guidance from current engineering students, faculty and alumni.

Now in its second year, GLEE is an initiative that aims to inspire graduating female high-school students to join the U of T Engineering community. The young women are joining a community where women are leaders in their fields, making a positive impact as educators, researchers and entrepreneurs.

pg18
Jeff Cassidy

GRADUATE STUDENTS AT A GLANCE

The Graduate Experience

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering consistently ranks among the top ECE departments in the world.

Throughout its history, ECE has been witness to groundbreaking discoveries and developments in almost every area of electrical and computer engineering, at the core of which is the outstanding research conducted by graduate students and faculty members.

Supervised by our faculty of 78 professors, many of whom are leading experts in their fields, graduate students may choose from a wide range of research areas including biomedical engineering, communications, computer engineering, electromagnetics, electronics, energy systems, photonics, and systems control.

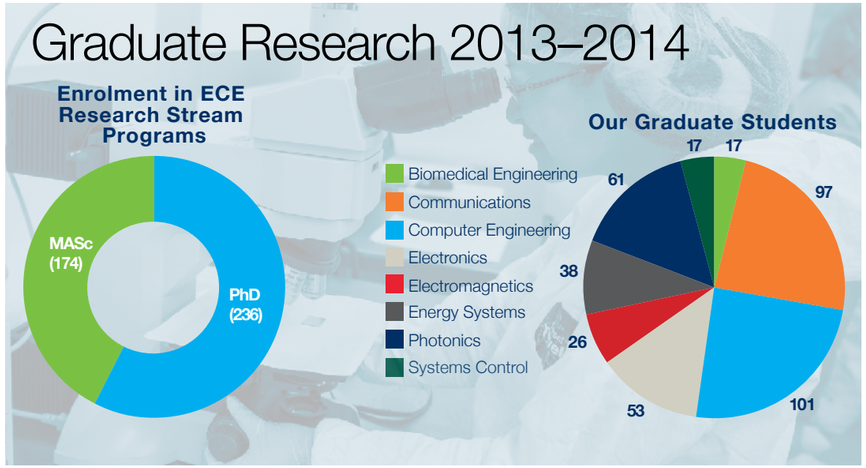

ECE offers three degrees: Master of Applied Science (MASc), Master of Engineering (MEng), and Doctor of Philosophy (PhD).

The Master of Applied Science degree provides advanced study and research in a major field. This is a full-time program and requires the completion of courses, a thesis proposal, and a research thesis. The MASc is the recommended prerequisite for admission to a PhD program.

The Master of Engineering degree provides advanced training to individuals who wish to work or practice in the field of engineering. This course-based degree may be completed on a full- or part-time basis. It is not a research degree and is not a recommended prerequisite for admission to the PhD program.

The Doctor of Philosophy degree is intended for those who wish to pursue a career in fundamental or applied research. The PhD requires the completion of courses, a field-comprehensive examination, a thesis proposal, and a research thesis. Applicants to the PhD must hold the equivalent of an MASc degree with thesis. For admissions information visit **uoft.me/gradadmission.**

Graduate Enrolment, 2003–2004 to 2012–2013

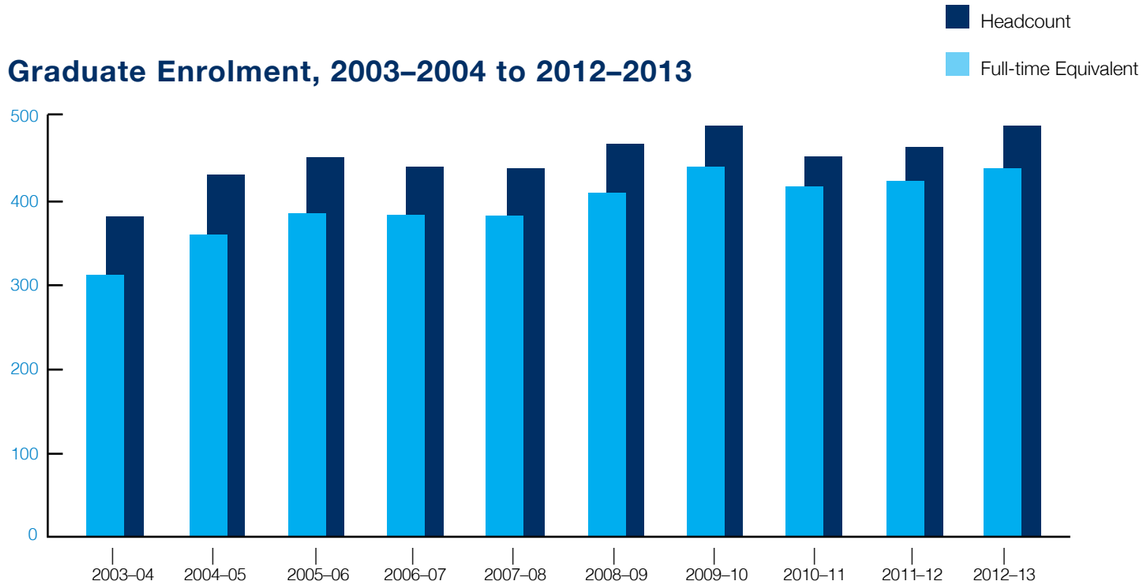

19th

ECE at the University of Toronto ranks 19th internationally according to **QS World University Rankings 2013.**

ECE is the top U of T department in the last five years (cumulative) in terms of new invention disclosures filed, new license and option agreements executed, and new start-up companies formed.

Degrees Awarded, 2003–2004 to 2012–2013

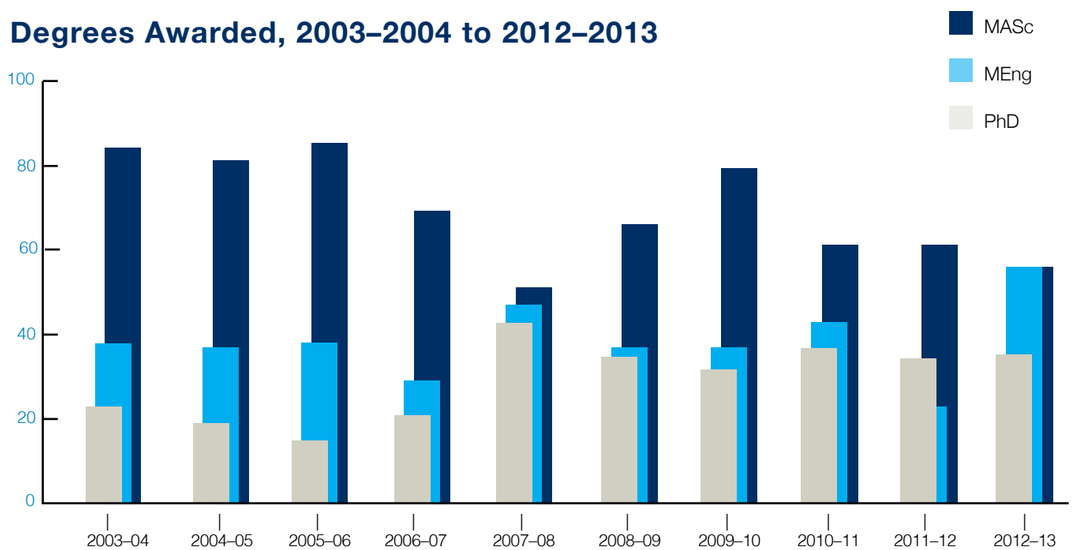

Vanier Scholars

Two ECE PhD candidates were selected as NSERC Vanier Canada Graduate Scholarship (Vanier CGS) recipients for 2013:

Mohamed Abdelfattah

Supervisor: Professor **Vaughn Betz**
Project title: "Communication-Centric Architectures and Design Styles for Next-Generation Programmable Systems-on-Chip"

Amit Deshwar

Co-supervisors: Professors **Brendan Frey** and **Quaid Morris**
Project title: "Computational Methods For Gene Expression Analysis of Heterogeneous Samples"

Worth $50,000 per year for three years, the Vanier CGS program is designed to attract and retain world-class doctoral students and help establish Canada as a global centre of excellence in research and higher learning.

PhD Enrolment, 2003–2004 to 2012–2013

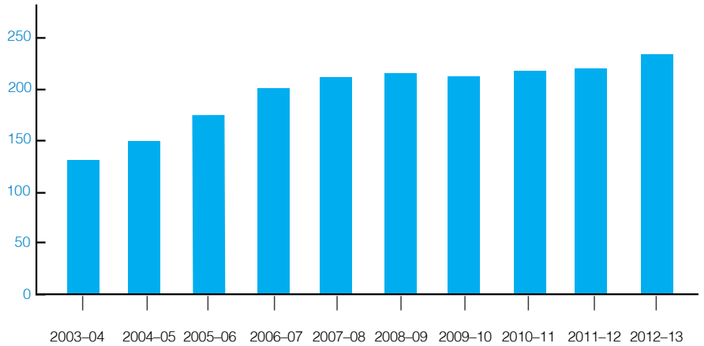

31

Spin-off companies sparked by ECE in the last 15 years

From 2009 to 2013, ECE inventors disclosed 163 inventions. In the same period, we filed 19 patent applications, signed 24 licenses and formed 16 start-ups.

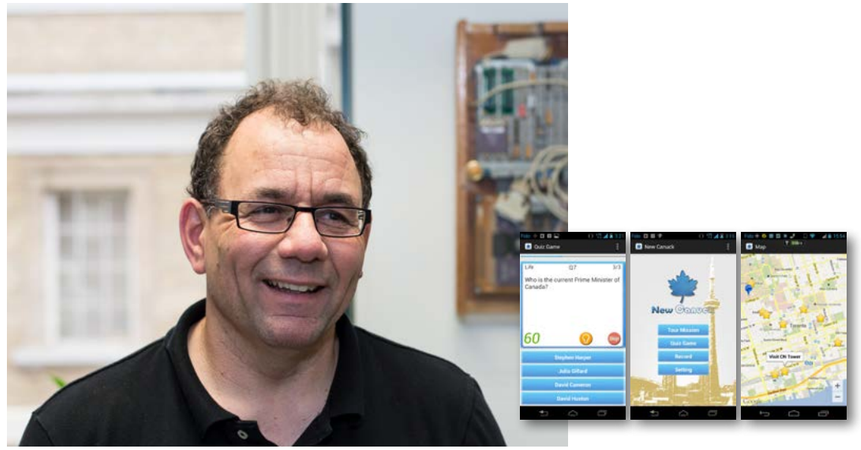

"That's what we can do uniquely here at U of T, being so interdisciplinary—once you work with these ideas, lots of exciting things happen."

–Professor Jonathan Rose

Interdisciplinary Problem-Solving: ECE 1778

You use your smartphone for calls, texts and email—but do you use it for treating addiction, diagnosing disease or teaching children? Smartphones are tiny computers with built in microphones, cameras, accelerometers, magnetometers and GPS—when approached imaginatively, their capabilities are virtually limitless. Professor Jonathan Rose encourages imagination in ECE 1778: Creative Applications for Mobile Devices, the interdisciplinary graduate course he launched. The course is open to graduate students from all fields across the University of Toronto. Students with non-programming backgrounds are grouped with programmers, typically from Computer Science or Electrical & Computer Engineering, and teams work together to realize the design and functionality goals of the non-programmer, or 'apper'.

The course ran for the third time in Winter 2013, and resulting applications have helped new Canadians acclimatize to a different culture, varsity athletes rehab from ankle injuries, and parents decide how much fever medicine is safe for their child. "We have a diverse set of apps, and that's because of people who aren't in the field jumping in," says Professor Rose. "That's what we can do uniquely here at U of T, being so interdisciplinary—once you work with these ideas, lots of exciting things happen." Professor Rose is launching a new research centre with a related focus: The Centre for Inter-disciplinary Mobile Software and Hardware, which seeks to put extensive effort into developing high-impact mobile applications in a variety of areas.

MEng Enrolment, 2003–2004 to 2012–2013

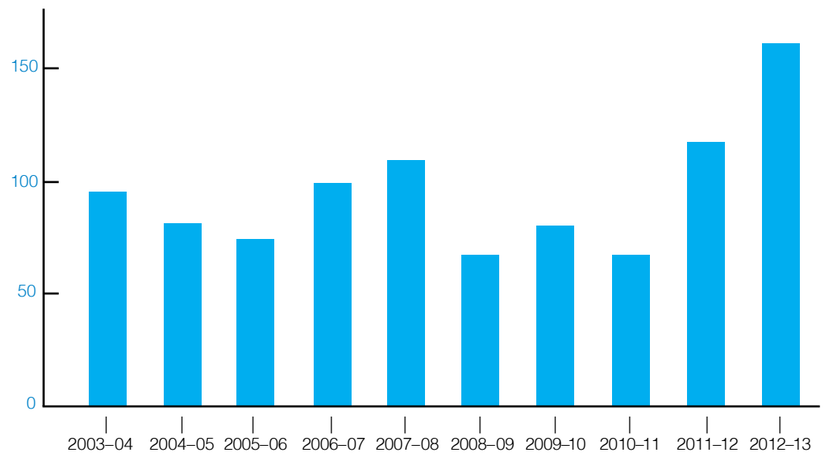

pg22
Deepa Kundur

pg24
Robotics

RESEARCH AT A GLANCE

The Research Perspective

Researchers at the University of Toronto strive to make discoveries that will influence industry and benefit society.

We do so by tackling challenging fundamental and applied problems; working across and beyond disciplinary boundaries; partnering with industry, government, and other academic institutions globally; and training the research leaders of tomorrow.

In The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, we are particularly proud of our ability to identify areas of strategic importance to global industry and society and to address these important problems. Initiatives include:

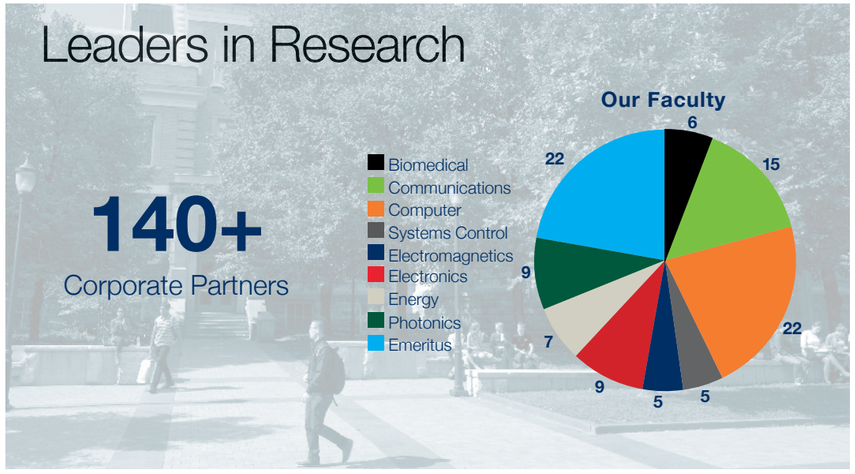

First, our focus on the Smart Grid, the electrical grid that seamlessly incorporates renewables such as wind and solar, and uses the latest technology to improve the efficiency, importance, reliability, economics, and sustainability of electricity services. This initiative leverages our strengths in energy systems, systems control, algorithms, optimization, security, communications, networking, electronics, and photonics;

Second, our emphasis on the technological, economic, and social dimensions of security and privacy in the mobile ecosystem and the emerging cloud;

Finally, our strategy to bring electrical and computer engineering expertise to important problems in biomedical engineering and medical devices, uniting expertise in sensors and stimuli, low-noise and low-power electronics, systems-on-chip, signal processing, systems control, optoelectronics and computing.

Our researchers partner with countless industry leaders worldwide to stimulate, enable, and translate our research into application. ECE continues to seek unique ways to make a global impact, benefiting Canada and the world through advances that improve quality of life.

ECE Invention Disclosures, 2009 to 2013

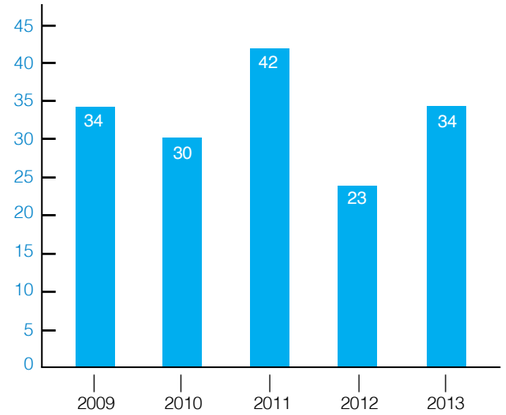

The University's Invention Policy requires inventors to completely disclose their inventions. In most cases, rights to their inventions are then jointly owned by the University and the inventor.

ECE is top five at U of T in new priority patent applications filed and new patents issued.

Canadian Peer Universities vs. University of Toronto Share of NSERC Funding
for Electrical and Electronic Engineering 2012 (Grant Year = April to March)

22nd

University of Toronto ranks 22nd in the world for Engineering and Technology according to the **Times Higher Education World University Rankings 2013.**

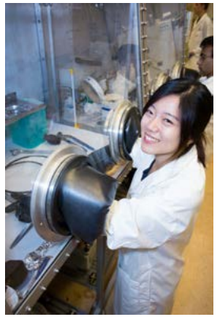

Professor Aleksandar Prodić
Professor Peter Lehn

Both named **U of T Inventors of the Year**

ECE faculty held 23 research chair titles in 2012–2013, including Canada Research Chairs, Endowed Research Chairs, Industrial Research Chairs and U of T Distinguished Professor Chairs.

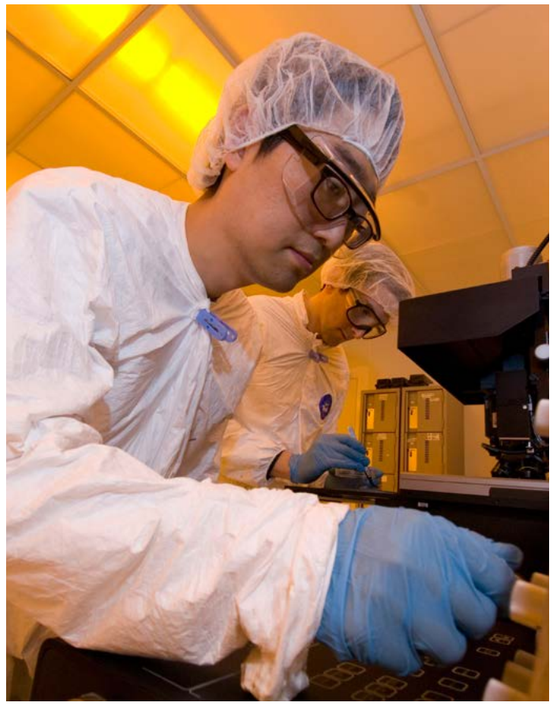

Early Career Teaching Award

Professor **Jason Anderson** received the Faculty of Applied Science & Engineering's 2013 Early Career Teaching Award, which recognizes an instructor in the early stages of his career who has demonstrated exceptional classroom instruction and teaching methods.

Sustained Excellence in Teaching Award

Presented to a faculty member who exhibits teaching excellence over a sustained period of time. This award is new this The first recipient is Professor **Tarek Abdelrahman** for 2012–2013.

Distinguished Professor in Application Platforms and Smart Infrastructure

Presented to Professor **Alberto Leon-Garcia.** He was also named a Fellow of the American Association for the Advancement of Science.

Fellows of the Royal Society of Canada

Professors **Frank Kschischang** and **Jonathan Rose**

Fellow of the Engineering Institute of Canada

Professor **Paul Chow**

$261K

Average research operating funding per ECE faculty member

ECE Research Funding from 2005–2006 to 2011–2012

	2005–06	2006–07	2007–08	2008–09	2009–10	2010–11	2011–12
Federal	$10,397,002	$8,743,709	$8,478,602	$8,004,336	$8,790,575	$9,327,932	$10,830,916
Provincial	$5,479,296	$3,703,586	$3,955,389	$4,862,786	$3,948,865	$3,251,806	$1,569,028
Industry	$2,266,096	$3,363,559	$2,940,516	$2,699.325	$1,471,958	$2,550,974	$2,959,090
Other	$575,955	$957,306	$863,281	$2,780,615	$2,735,435	$2,794,394	$2,599,683
Total	**$18,718,349**	**$16,768,160**	**$16,237,788**	**$18,347,062**	**$16,946,833**	**$17,925,106**	**$17,958,717**

A Year to Remember

Professor Joyce Poon was announced as the 2013 recipient of the University of Toronto's McCharles Prize for Early Career Research Distinction.

She is the third recipient of the McCharles Prize, which was established in 1907 by Aeneas McCharles and re-established in 2007 as an award for exceptional performance and distinction in early career research on the part of a pre-tenure member of the Faculty. Professor Poon is a Tier 2 Canada Research Chair in Integrated Photonic Devices.

Already considered a thought-leader in her field, her research focuses on the creation and exploration of novel photonic devices. Professor Poon received an Early Researcher Award in 2009. She garnered IBM Faculty Awards in 2010 and 2011 — the only Canadian to receive this prestigious award in either year. In 2012 Professor Poon was named one of the world's Top 35 Innovators Under 35 by *MIT Technology Review*.

TEDxToronto

ECE professors **Brendan Frey** and **Steve Mann** were among 12 expert speakers selected to present at TEDxToronto 2013. Canada's largest TED event, TEDxToronto is an independently organized program that builds on the central TED concept of sharing 'ideas worth spreading'.

Photo credits: Photo of Brendan Frey by Lana Khakham and Sara Stankiewicz - Parasol. Photo of Steve Mann and TEDxToronto by Andrew Williamson.

Publications by ECE Faculty, 2003 to 2012

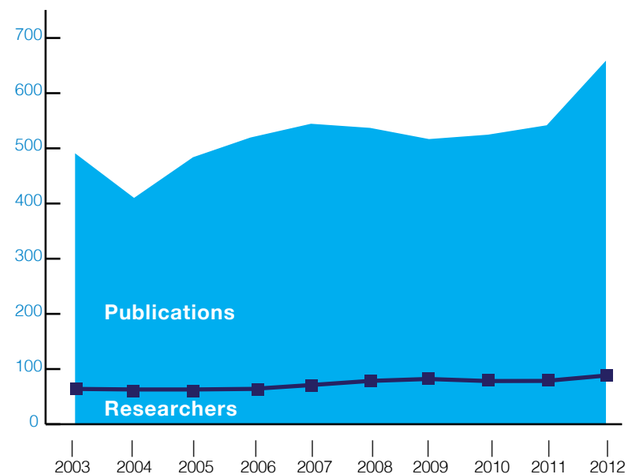

ECE is the top department in the Faculty of Applied Science & Engineering for all commercialization indicators in the past five years. ECE typically accounts for about 50 per cent of Faculty start-up companies and invention disclosures.

Discoveries are Made with Teamwork

From Software to Circuits with LegUp

This summer, while their peers were travelling, relaxing and working, four of ECE's top undergraduate students put their education into action.

William Cai, Emily Miao, Yolanda Wang and **Yvonne Zhang** got a real taste of research working on LegUp, a project jointly supervised by Professors **Jason Anderson** and **Stephen Brown**. LegUp is a high-level synthesis tool that lets software programmers benefit from the advantages of hardware — primarily speed and power efficiency — without struggling to learn complex hardware design languages. It's the first high-level synthesis tool that's completely free and open-source, and that is robust enough to handle realistic programs.

The end result is a reconfigurable field-programmable gate array, or FPGA, tailored to perform certain compute-heavy chunks of the program much more efficiently than software alone. Think of a program as an assembly line: in order to get the finished product, thousands of simple tasks must be performed in sequence. Software can be like running an assembly line with only one person on it — that worker can only complete one task at a time until the job is done. Incorporating parallelism, through FPGAs, is like adding many more people to the line. The additional workers can do many tasks simultaneously, and the finished product is produced more quickly.

LegUp is much more than just a great idea — it really works, right now. Launched in March 2011, the source code has been downloaded externally more than 600 times since its first release, and the first LegUp paper has already been cited more than 70 times. The group has presented for companies such as Intel, Samsung, Xilinx, Altera and NEC in Japan. The tool was developed and continues to evolve in consultation with Altera, the California-based FPGA manufacturer with a large R&D arm in Toronto.

"Coming out of second year, I knew all that theory but I didn't know how to apply it," says Cai. "In research you have to understand the theory to a deeper depth, and apply it to real-life scenarios."

In the summer everyone pulls their weight: Cai, Miao, Wang and Zhang each worked independently on solutions to various LegUp challenges under the guidance of PhD candidates **Andrew Canis, James Choi** and **Blair Fort**. "The undergraduate students that come to work with us are always top of their class," says Choi. "These guys are very smart, they know what they're doing."

And it helps that they're in expert hands — Canis and Choi joined Professors Anderson and Brown after completing undergraduate degrees at the University of Waterloo and have worked for several years to make LegUp the widely adopted tool it is today. Canis, along with alumnus **Mark Aldham**, wrote the original code for the program in 2009 after starting his master's degree. Now in the third year of his PhD, Canis is focused on improving the performance of LegUp's loop-pipelining component — in other words, the way it adds more people to the assembly line and assigns them tasks. He's designed many of the self-test features of the tool that make it robust, easy to use and modify. "I feel like it's important to contribute back and release something," says Canis.

Choi knew he wanted to concentrate on high-level synthesis for his graduate work. After completing his MASc with Professors Brown and Anderson, he applied to several American schools for his PhD but "nothing was as interesting as what we were doing here, so I stayed," he recalls. The direct connection to industry was a big plus, too. "It's motivating for us because we know that people are using it, we know that what we're doing now will make a difference," says Choi. "It helps us do a proper job because people are going to see our code." Choi's not the only student compelled to stick with the project: recently **Lanny Lian**, a veteran from the summer of 2012, rejoined the team to begin a master's degree. Fort just started his PhD full time, while maintaining a part-time role at Altera.

The group sits down three times a week to bounce ideas off each other, get input from Professors Anderson and Brown, and hear from **Tomasz Czajkowski**, a design engineer at Altera and former grad student of Brown's who consults on the project. "The undergrads are full-fledged members of the team, contributing to future releases of the tool and helping to produce publishable research results," says Professor Anderson. "Both the grad students and I are truly impressed by their capabilities."

Emily Miao grew up in Richmond, B.C., but moved to Delaware to attend a highly competitive high school known for excellence in math and science. "I've always loved math and science, even when I was little," she says. She chose U of T for the strength of its ECE program, and knew as early as the end of her first year that she wanted to see what academic research was like first-hand. "I wanted to get a feel for what it's like in this environment, kind of like you're going to grad school," says Miao. "This project is unique because you really get exposed to this industrial network as well as the academic network."

Cai had already had a taste of summer research after his first year when he worked for a professor in Materials Science & Engineering, but wanted to see what life in ECE was like. "Jason made his research seem interesting — like he really enjoyed it. And the idea of making hardware out of software was pretty cool," says Cai. In his spare time Cai is teaching himself Java, and he enrolled in a singing competition. "It's definitely more of a hobby for me," says Cai.

Cai got a chance to put another one of his hobbies — a love of games — to good use toward the end of the summer. In collaboration with the team, Cai designed an algorithm for playing the Tetris-like game Blokus — an algorithm to be synthesized into FPGA hardware by LegUp. Teams from around the world submitted FPGA-based solutions to compete in a massive Blokus tournament at the International Conference on Field-Programmable Technology this winter in Kyoto, Japan. The tournament would be the first chance for LegUp-synthesized FPGA hardware to go head-to-head against human-designed hardware, still the industry gold standard, in a competitive gaming scenario. It's not the summer project Cai planned to take on, but he's thrilled it came his way. Says Cai, "Everyone loves games, right?.

"The undergrads are full-fledged members of the team, contributing to future releases of the tool and helping to produce publishable research results."

Discoveries are Made Across Disciplines

Jeff Cassidy: Engineering New Cancer Treatments

It's well known that an education in electrical and computer engineering can launch careers in technology development finance or medical research — but not many students have the skills and savvy to explore all three before their 30th birthday.

Jeff Cassidy is all about exploration — both intellectual and physical. The Hamilton, Ontario native has just completed his master's thesis, co-supervised by Professor **Vaughn Betz** of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering and Professor **Lothar Lilge** of the Ontario Cancer Institute and U of T's Department of Medical Biophysics. In his spare time he teaches sailing, goes heli-skiing, canoes, and plays squash once a week with the guys in his lab. This fall Cassidy began work toward his PhD, driving his project's focus toward developing hardware for faster and more power-efficient probabilistic Monte Carlo modelling of light propagation.

Why study light propagation? Photodynamic therapy, or PDT, is a potentially powerful method of treating cancer. First the patient is injected with a drug, called a photosensitizing agent, which is sensitive to light. The cancerous region is then irradiated with light of a specific wavelength, activating the drug and killing cells. The trick is to know how far the light will travel and where. Currently, PDT is in trials for treating prostate cancer, a well-understood and relatively simple part of the body to model. Cassidy and Professor Lilge hope that with fast access to better simulations, PDT will be useful for treating cancer of more complex and sensitive anatomy, such as the head and neck.

That's why Cassidy is working on building reconfigurable hardware, called field-programmable gate arrays, or FPGAs, to run compute-heavy algorithms to predict how light will disperse through different areas of the body. Running these algorithms with software can take many hours, but with FPGAs that time is reduced to minutes. Cassidy and Professor Betz predict you could optimize an entire treatment plan, by trying hundreds of simulations to find the best option, within a few hours.

To do this, they're teaming up with IBM's high-performance computing research consortium: the Southern Ontario Smart Computing Innovation Platform, or SOSCIP for short. SOSCIP is a consortium that offers researchers access to high performance computing systems, including an energy-efficient, reconfigurable computing platform built to perform supercomputer-worthy tasks with 20 to 40-times less power. Cassidy will program a part of SOSCIP to run his photon-predicting application, and tailor hardware to the new system. His is one of a few projects to have access to the SOSCIP's agile computing platform.

Cassidy is currently designing the FPGA bitstream that will program the SOSCIP machine to simulate light propagation. With close collaboration with researchers at Princess Margaret Hospital, a leading cancer research institution located in Toronto, it may be possible to see PDT for head and neck cancers begin clinical trials in a matter of just a few years. "The goal is to advance this towards pre-clinical and clinical trials within the timeframe of my PhD," says Cassidy. "That was one of the things that kept me on here, having something that may treat real people within four years."

It's a goal that means a lot to him personally. "Like anyone, I've had a lot of family and friends affected by cancer. My Uncle Andy had a recurrence of throat cancer and ultimately died. So I like that with this work, there's a sense that you're making a material difference in someone's quality of life."

He never planned to have a career in medicine. Always talented at math and science, Cassidy earned his undergraduate degree in Electrical Engineering from McGill University in 2006. From there he took a job at WESCAM in Burlington, Ontario where he'd completed internships while at McGill, and developed a digital video processing unit for their airborne military surveillance cameras. He then made the jump to finance, spending a year and a half at Scotiabank and then another year with Manulife. While working in finance, he studied and passed all his Chartered Financial Analyst exams. "Engineers get hired into finance because an engineering education provides both programming and mathematical skills," says Cassidy. "Quantitative modelling skills open doors everywhere."

But after a period learning how financial institutions work from the inside out, Cassidy started to feel his explorer's itch acting up. "I missed the challenge of designing large, complex systems," he recalls. In December 2010 he began scouting for grad schools and learned that Professor Betz was joining ECE's faculty and starting up his own FPGA group — an alluring prospect. This year Cassidy won a Banting & Best Doctoral Research Award from the Canadian Institutes of Health Research worth $35,000 each year for three years. The rest is history — albeit with an unwritten ending.

"I'm not a doctor or a biologist, but I'm working on treating cancer. I've done financial modelling, image processing," says Cassidy. "What field of endeavour or business doesn't rely on computing? For example, mapping the genome, searching for the Higgs Boson, commerce, data mining — all this stuff is stuff you can get involved in as a computer engineer. And not just one of those — you can hop between them because you have this portable skillset."

Today Cassidy is completely focused on driving his project toward clinical trials, one step at a time. But with his engineering training, countless unexplored paths lie in his future.

"What field of endeavour or business doesn't rely on computing?"

Discoveries are Made in Unexpected Places

Professor Deepa Kundur: Imagination in Engineering

A flock of birds sits on the steep slope of a church roof. Sparked by some invisible signal, the birds take off in unison, rising into the sky, banking a sharp left, diving and soaring together.

As they approach the church spire, the group separates to fly around it, seamlessly joining back together on the other side before landing on the roof without incident. How did so many individual creatures coordinate their elaborate movement, maintain formation and never collide even while avoiding an obstacle?

Professor **Deepa Kundur** thinks birds have a lot to teach us about increasing stability in the electric grid. Kundur, a professor in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, studies security of cyber-physical systems — networks we have always relied upon, that are becoming increasingly cyber-enabled or 'smart'. Cyber-enablement of systems is not new. Historical examples have included commerce, with the rapid adoption of online banking and shopping; entertainment, where musicians, artists, movies and TV series push a digital-first approach; and even friendship, as much social interaction now takes place online through tools such as Facebook, Instagram, Twitter and LinkedIn. The Smart Grid is another example — our old energy infrastructure is going digital, adding a new layer of communication and sensor connectivity on top of the existing physical network of generators, substations and transmission lines.

"In all of these systems, you cyber-enable to facilitate greater functionality, efficiency, and improved capability, but often security issues arise," says Professor Kundur. "Like with social networking, you don't really know if the person you're talking to is necessarily the person they say they are. Same with digital entertainment — it's easier to pirate material. So when you move into the digital world, there are clearly many advantages but security becomes an issue. And you'll see the same thing with the Smart Grid."

Professor Kundur's job is to model interactions between the existing physical grid and the new cyber layer, and to envision where emerging vulnerabilities might lie. There are three components to vulnerability: an inherent weakness in the system, an ability to access that weakness, and an ability to exploit that weakness. "When you cyber-enable a system you increase its complexity, which typically raises the number of weaknesses. You also increase its connectivity, which improves access to vulnerabilities, and you increase collaboration, which can allow additional parties to exploit weaknesses. Therefore, it's imperative that we study the impact of integrating advanced communications and computing to power system security," says Professor Kundur. A system is only as secure as its weakest link and it takes a little bit of cunning on Professor Kundur's part to think, "If I were a villain, what would I really want to do?"

But what does all that have to do with birds? Professor Kundur and her lab have borrowed from the biological world to model robust energy grids, capable of elegantly handling failure or chaos. "It's exciting when you see analogies," she says. Imagine a collection of synchronous generators contributing energy to the grid. If a tree falls across a set of high-voltage transmission lines, a major route for energy distribution is cut off, and generators can fall out of sync. All of a sudden one part of the grid is not getting enough power, while other sections get too much. This has the potential to significantly disrupt power flow, leaving it acutely vulnerable to another failure. Think of the individual generators as birds — Professor Kundur wants to apply birds' robust flocking behaviour to generators, letting the Smart Grid store and redistribute energy around a metaphorical obstacle resulting from the tree fault, much like birds avoid the church spire. The generators can then re-sync themselves and stabilize the system, skipping the brownouts or blackouts that typically result from a system under duress.

Born and raised in Toronto, Professor Kundur is herself a three-time ECE graduate, earning a BASc in 1993, an MASc in 1995 and a PhD in 1999. After completing her doctorate she joined the department as an assistant professor before moving to College Station, Texas, in 2003 to join the faculty at Texas A&M University. She returned to U of T in August 2012, bringing three adventurous graduate students and a postdoctoral researcher with her.

It's not easy to find models that accurately describe both the physical world and virtual world while elegantly describing their complex interactions. "The Holy Grail in this area is if you could develop a modelling framework that encompassed the many aspects of cyber-physical interaction, and you often can't without sacrificing performance or tractability. Our strategy is to develop models that elegantly and accurately represent select but salient aspects of cyber-physical interactions," says Professor Kundur. She believes this approach has allowed her group to discover a new strategy for attacking Smart Grids that involves remote corruption of circuit breakers. "Under certain circumstances, if you switch the corrupted breaker in a specific coordinated sequence, within seconds, you can create instability," says Professor Kundur. They are in the process of developing a comprehensive vulnerability analysis framework to help power utilities identify weaknesses and prioritize protection strategies.

Bridging the gap between the visceral present and the visionary future is what Professor Kundur does best, using imagination, open-mindedness and creative design. "What we aim to do is find elegant analogies that highlight the critical aspects of practical systems," she says. "So when we can find a metaphor that very beautifully meshes with a real-life situation and it inspires a new way of thinking, that's just a wonderful thrill."

"In all of these systems, you cyber-enable to facilitate greater functionality, efficiency, and improved capability, but often security issues arise."

Discoveries are Made through Collaboration

Systems Control Drives Robotics Revolution

Twenty years from now, we'll say the revolution started with a vacuum cleaner.

A Roomba, to be precise. That innocuous home appliance that takes an irritating chore off your list. But there's more to that little black disc than meets the eye. It's not the hardware that makes it special, it's what's inside: it can scan a room, map out a route, and follow that route meticulously without duplicating its effort or retracing its path. If the room configuration changes — if someone moves the couch or bumps the coffee table — the Roomba is running algorithms that allow it to make a decision and adjust accordingly. In other words, it has smarts.

It's the first smart robot to infiltrate our homes, but it won't be the last. "We're on the cusp of a revolution — we're going to be seeing more and more of these devices," says **Mireille Broucke**, a professor in the Systems Control Group of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, and a member of the steering committee of the Faculty of Applied Science & Engineering's new Institute for Robotics and Mechatronics. "People cannot yet conceive of how life with be in 20 or 30 years — robots will be in your home, doing tasks you can't even dream of."

To get a sense of what we're facing, think back to the early 1990s when personal computers started appearing in regular homes. They were handy for word processing and some games, but the average user couldn't do much else. Fast-forward 20 years, and most people can't imagine how they'd organize photos, choose a restaurant or complete their banking without a laptop, tablet or smartphone. We're about to see the same explosive proliferation and adoption of smart robotic devices, predicts Professor Broucke. And just like the early days of computers, we can't yet conceive of all the things we'll need them to do.

One thing we do know: control theory will be at the heart of the robotics revolution — or more accurately, at its head. "We control theorists design the brain of the system," says Professor **Mandfredi Maggiore**. Most human-made devices need controllers, from the antilock-brakes in your car, to the central heating in your house, to the high-tech Canadarm on the International Space Station. "Control systems are what make devices smart, to the extent that they need to be," says Professor **Bruce Francis**. Francis began working on theoretical robotics problems around 2002 after **Tim Barfoot**, now a professor in the University of Toronto Institute for Aerospace Studies (UTIAS), took one of Professor Francis's graduate courses and invited him to visit the Institute.

There Professor Francis saw penguin-like robots shuffling around on the floor and thought, "That's interesting, I could work on that."

One of his earlier projects, a collaboration with Professors Maggiore and Broucke, asked how robots, mathematically modelled as points or unicycles in the plane, could move to find one another, or rendezvous. The robots can see a few of the others but not the whole group. One of their solutions is called cyclic pursuit — each robot will pursue the next at a fixed speed, until they eventually arrive at a common point. More recently Professor Francis has worked with Professor Barfoot on control algorithms for wheeled vehicles in convoy, getting unmanned pursuit vehicles to faithfully follow a route set out by a human-driven lead vehicle. It's one of many flourishing collaborations between ECE and UTIAS.

Historically, control theory hasn't been the flashiest discipline as most of its action goes on under the hood. But students can expect to see more of it as this robotic renaissance advances. The Institute for Robotics and Mechatronics was established in 2010 to advance the art and science of robotics and mechatronics through collaborative research projects and innovative education programs. And just this past year ECE launched its first robotics minor option, offered in collaboration with the Department of Mechanical & Industrial Engineering, Department of Civil Engineering, the Institute for Aerospace Studies, and the Institute of Biomaterials & Biomedical Engineering. ECE470, Robot Modelling and Control, is the department's flagship robotics course, and a new 300-level Intro to Robotics course is slated to roll out in 2014. "The only way that a robot can possibly have impact on society is to give it some brains, and students come here to learn how to give it that," says Professor Broucke.

Professor Broucke herself is working on a project in collaboration with MDA Corporation to enhance the brains behind the Canadarm: she's writing an algorithm to help the robotic arm transition more seamlessly between its 'reaching' and 'contact' programs. Think of a child taking a glass of milk from her parent. The girl intuits when her mother is letting go of the milk and grasps it safely, making an effortless transition from stretching out her arm in a reaching function, then making contact and grabbing the glass. Current practice dictates that the Canadarm move really slowly when in this reaching-contact autopilot mode, but Professor Broucke's algorithm could enable much smoother and more efficient movement.

Professor Maggiore is also interested in smooth moves — he's currently working to design algorithms that mimic the locomotion of snakes. He has teamed up with a research group at the Norwegian University of Science and Technology that has built a snake robot specifically designed to navigate challenging environments. Between their hardware and Professor Maggiore's code, these future snake-bots could do dirty work in inhospitable spots, such as checking on deep sea oil pipes, slithering through rubble in disaster zones, or even climbing vertically into trees or up poles — places no wheeled vehicle could reach. "Theory and experiment together is a very good combination," says Professor Maggiore.

On the eve of the revolution, ECE students will be the ones to shape the brains of the dynamic machines that will soon live alongside us. "Decisions are being made on the fly — that requires more sophisticated algorithmic methods than the old style of robotics," says Professor Broucke. "ECEs are uniquely positioned to work in this frontier because they can use their math-based training to crack abstract problems." Math will be key — and so will imagination. "ECE is an exciting discipline for people that are interested in visionary work," says Professor Maggiore.

"We control theorists design the brain of the system."

BOUNDLESS ENGAGEMENT

"The world needs even more well-rounded engineers now — that evolution is happening in ECE and I want to be part of it."

Sohayla Praysner
ElecE 8T6

Sohayla Praysner chairs the Skule Society Committee at the Faculty of Applied Science & Engineering, a community of donors whose leadership helps our Faculty maintain its position as a global leader in training engineers for the future. Getting involved in Skule Society is just one way to support our Boundless campaign.

"I like being part of a team of educators and researchers enabling the next wave of transformation in communications, healthcare, business productivity, everything. The students we're training are going to make that happen."

Professor Vaughn Betz
ECE PhD 9T8

As a PhD student, Vaughn Betz co-founded a start-up that was acquired by a large chip-manufacturing company. After 13 years working in industry he returned to ECE as a professor — now he teaches the next generation of engineers how to translate their knowledge to the business world. Your support helps our faculty members bring the best of their innovative research into the classroom.

To find out more, contact
joannaf@ecf.utoronto.ca,
416-978-7270 or
donate.utoronto.ca/engineering

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

BOUNDLESS OPPORTUNITY

Our Community

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering at the University of Toronto is fortunate to have valuable, long-standing partnerships with visionary corporations, foundations, alumni and friends that guide and support the advancement of research and teaching in the country's largest ECE department. We invite collaborations with both individuals and industry to help ensure that our research remains relevant, and to create avenues for significant economic and social contributions to Canada and throughout the world.

Our industry partnerships are multi-faceted and include more than 60 industry funders and collaborators. As you'll read in this section, close partnerships with companies like TELUS let our researchers apply their work to solve real-world issues. Corporations such as Altera and IBM provide valuable opportunities for our students through work-experience programs such as the Professional Experience Year (PEY), and support the educational and research missions of the department through research grants and philanthropic support.

ECE is fortunate to have more than 10,800 alumni located around the world — from Toronto to Taipei, San Jose to Singapore. Our alumni have risen to prominence in a wide variety of fields, including business and academia — read more in our alumni profile of brothers Alex and Anthony Grbic, later in this issue. We are committed to staying in touch with this group through events such as our Alumni Lecture & Networking Reception, Fourth-Year Reception, which unites graduating students with members of the alumni community, and Spring Reunion, an opportunity for ECE grads from all classes to return to campus to socialize, tour our laboratories and share memories.

We regularly share alumni news on the front page of the ECE website, **www.ece.utoronto.ca** as well as through our social media channels on Facebook, Twitter and LinkedIn. If you have a piece of news you'd like to share with us, we'd love to hear from you — please send a note to Marit Mitchell, our senior communications officer, at **marit.mitchell@utoronto.ca.**

Our industry partners enjoy an effective multiplier of up to 5x their spending on joint research projects with ECE. Cash and in-kind contributions by our industry partners can be used to leverage additional cash investments from the provincial and federal governments, and from the University.

Board of Advisors

This year ECE formed its Alumni Board of Advisors, a diverse and hugely accomplished group of graduates and industry leaders. The Board meets throughout the year to advise the department on its advancement activities and strategic direction. Our Board members are:

Alan Boyce

Alan Boyce is a senior management consultant, executive and entrepreneur with more than 35 years of experience in consulting, technology, management and training. He is a registered Professional Engineer in Ontario, a certified Project Management Professional, holds a Master of Business Administration degree, and is a Certified Management Consultant. Mr. Boyce is the Chief Executive Officer of SOMOS Consulting Group, a growing consulting and training company with practices in project management, information technology, CRM, supply chain management, and benefits management systems. Alan built SOMOS from its conception in 1991 to the present day and oversaw the sale of the company in 2013 to SEB Inc. (TSXV:SEB). An alumnus of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, Mr. Boyce graduated in the class of 7T8.

John East

John East retired from Actel Corporation (a public company) in November 2010 in conjunction with the transaction in which Actel was purchased by Microsemi Corporation. He had served as the CEO of Actel for 22 years at the time of his retirement. Previously, he was a senior vice president of AMD, where he was responsible for the Logic Products Group. Prior to that, Mr. East held various engineering, marketing, and management positions at Raytheon Semiconductor and Fairchild Semiconductor. In the past he has served on the boards of directors of Adaptec and Zehntel (both public companies), and MCC and Single Chip Systems (both private companies). He currently serves on the board of directors of Alacritech — a private high-tech company involved in the data storage market — and Pericom — a public fabless semiconductor company. He holds a BS in Electrical Engineering and an MBA from the University of California, Berkeley. He has lived in Saratoga with his wife Pam for 40 years.

Dr. Alex Grbic

Dr. Alex Grbic is a Product Marketing Director at Altera, responsible for Software and IP products. At Altera, Dr. Grbic has held management positions in the Software and IP R&D organization, where he headed up external memory interfaces IP development, performance analysis of Altera's products and work on visualization/debug tools. More recently, he was Director of Applications Engineering, where he led escalated customer issue support, initiatives for early adoption of new products, and technical collateral. Dr. Grbic holds a PhD in Computer Engineering from the University of Toronto.

Catherine Lacavera

Catherine Lacavera joined Google in 2005. As Director of Litigation, she manages a team of over 20 intellectual property attorneys and technical advisors. She oversees a docket of more than 200 pending patent and other intellectual property litigation matters both within and outside of the United States, including the Viacom and other copyright litigation against YouTube, and the Apple, Microsoft and Oracle copyright and patent litigation directed at Android. Ms. Lacavera also advises on complex licenses and acquisitions, both within and outside the litigation context, including the acquisition of Motorola Mobility for $12.5 billion. In 2013, she was named one of Fortune's 40 Under 40, one of the most innovative in-house counsel by the Daily Journal, and a Rising Star in the Best Bay Area Corporate Counsel Awards. In 2012, she was recognized by the Silicon Valley Business Journal as one of the top 40 under 40 for significant contributions to the community. Prior to joining Google, Ms. Lacavera practiced patent litigation in New York City at a large global law firm. She graduated from Computer Engineering in the class of 9T7 and holds an MBA and law degree, all from University of Toronto.

Dr. Alex Shubat

Dr. Alex Shubat has more than 25 years of experience in design, management and executive roles in the technology industry. He previously co-founded Virage Logic in 1996, and served as the company's President and CEO until September 2010, when the company was acquired by Synopsys. During his tenure as CEO, the company doubled in revenue, grew profitability, gained market share through both organic and inorganic product expansion, completed three acquisitions, and expanded its global footprint. Prior to co-founding Virage Logic, he served as Director of Engineering at WaferScale Integration, a provider of programmable microcontroller peripherals (subsequently acquired by ST). During his 10-year career at WaferScale Integration, he managed various groups including the application-specific integrated circuit and high-speed memory groups. He holds 27 patents and has contributed to more than 25 publications. Dr. Shubat earned an Executive MBA from Stanford University, a BASc and an MASc in Electrical Engineering from the University of Toronto, and a PhD in Electrical Engineering from Santa Clara University. Dr. Shubat currently serves as the CEO of Goji Food Solutions, Inc.

Arbor Awards

University of Toronto celebrated three dedicated ECE alumni volunteers with Arbor Awards in September 2013. The Arbor Awards recognize volunteers for their outstanding personal service to the University.

K.C. Smith (ECE 5T4, 5T6, 6T0 and department chair from 1976-81) is one of the most active and dedicated alumni in the whole of the Faculty of Applied Science & Engineering. He has worked tirelessly for both ECE and the Division of Engineering Science, serving on EngSci's Board of Advisors, hosting alumni gatherings in California on behalf of ECE, and securing fundraising support for numerous scholarships.

Anahita Panthaky (ECE 1T0) began volunteering with ECE as a student, and since graduating in 2010 continues to serve at alumni and recruitment events. Her active and enthusiastic support has directly contributed to the department's efforts to recruit the most promising candidates from Ontario, Canada and around the world.

Warren Brown (ECE 5T3) has been serving ECE for more than 60 years. Brown helped organize and operate the department's Lunch & Learn Speaker Series, a gathering of classmates and experts that has been running for nearly 40 years.

CAPE & eCAMION Make Grid a Little Smarter

ECE's Centre for Applied Power Electronics (CAPE) has joined with Toronto Hydro and industry partners eCAMION and Korea-based Dow Kokam to install the world's first community-based power storage unit. The unit is tied into the Toronto Hydro grid and consists of a 250-kW battery system that can power an entire street for one hour if the electrical system fails.

A team led by Professor **Reza Iravani** developed the algorithms and software that integrates the power unit into the grid and provides control and operational needs of the battery system according to Toronto Hydro's requirements.

Pictured above: Professor **Farid Najm** and Professor **Reza Iravani** with **Bal Gosal**, federal Minister of State (Sport), who announced that Professor Iravani's Centre for Applied Power Electronics would receive a $560,000 grant for circuit breaker technology for fast protection or isolation of battery storage systems.

Young Engineer Award

Michael Branch, (CompE 0T3), received a Young Engineer Award from the Ontario Society of Professional Engineers (OSPE) and Professional Engineers Ontario (PEO). These Ontario Professional Engineers Awards recognize engineering excellence and community service. Branch is the founder and CEO of Inovex Inc., a web and mobile software firm specializing in products and services for the healthcare and energy sectors, and creators of Maps BI — a new internationally awarded cloud-based mapping and business intelligence platform. He has been instrumental in developing Inovex's lineup of secure healthcare data collection and decision-making tools for physicians and public health policy drivers in Ontario. A committed volunteer within and outside the engineering community, Branch served as President of the University of Toronto Engineering Alumni Association. He also serves as a Board Member of Streetwise Actors and a member of the Haltech Regional Innovation Centre.

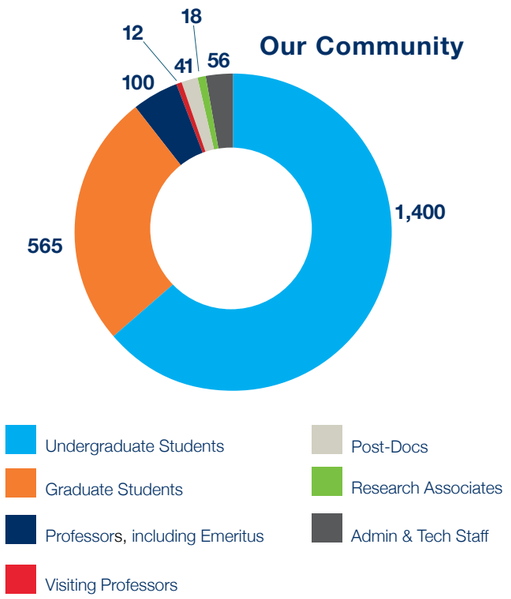

ECE Revenue 2013: $30 million

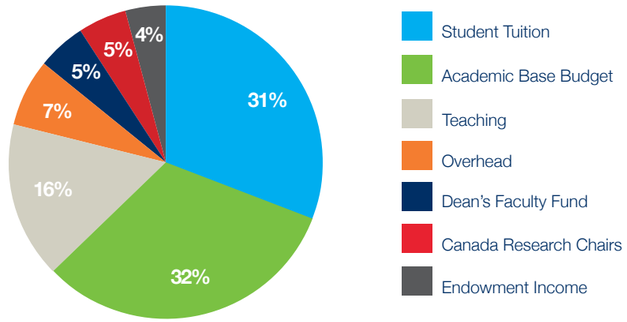

ECE Expenses 2013: $30 million

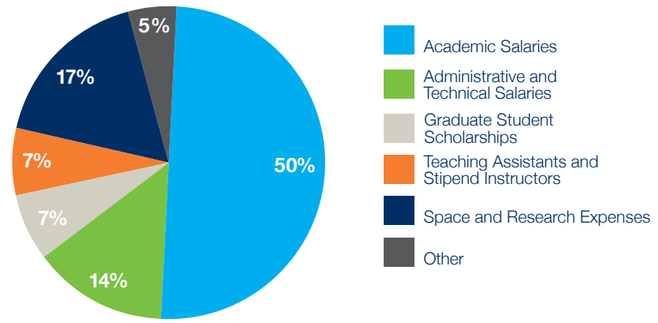

Get in touch! Stay current!

@eceuoft

Electrical & Computer Engineering at U of T

Electrical & Computer Engineering Alumni, U of T

10,805
Number of ECE alumni worldwide

ECE's Boundless Potential

Our individual and corporate partnerships are of vital importance to us and enable today's scholars and tomorrow's engineering leaders to make groundbreaking discoveries and advance knowledge in every area of electrical and computer engineering. Ongoing engagement, support and collaboration ensures ECE's ability to maintain its leadership position among the world's top authorities on electrical and computer engineering, and to build a global society of boundless innovation, creativity and economic development.

Individual donations from alumni and friends ensure that engineering scholars can focus their time and efforts on developing new technological breakthroughs, rather than worrying about financial constraints.

There are many different ways to support ECE projects and initiatives. Skule Society is a special community of alumni and friends who demonstrate leadership with annual gifts of $1,000 or more to the department. The creation of endowed chairs and professorships, graduate fellowships, and undergraduate scholarships are also important elements to Engineering's $200-million component of Boundless: The Campaign for the University of Toronto.

For further information on how you can support The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, please contact Joanna Forbes at **joannaf@ecf.utoronto.ca**.

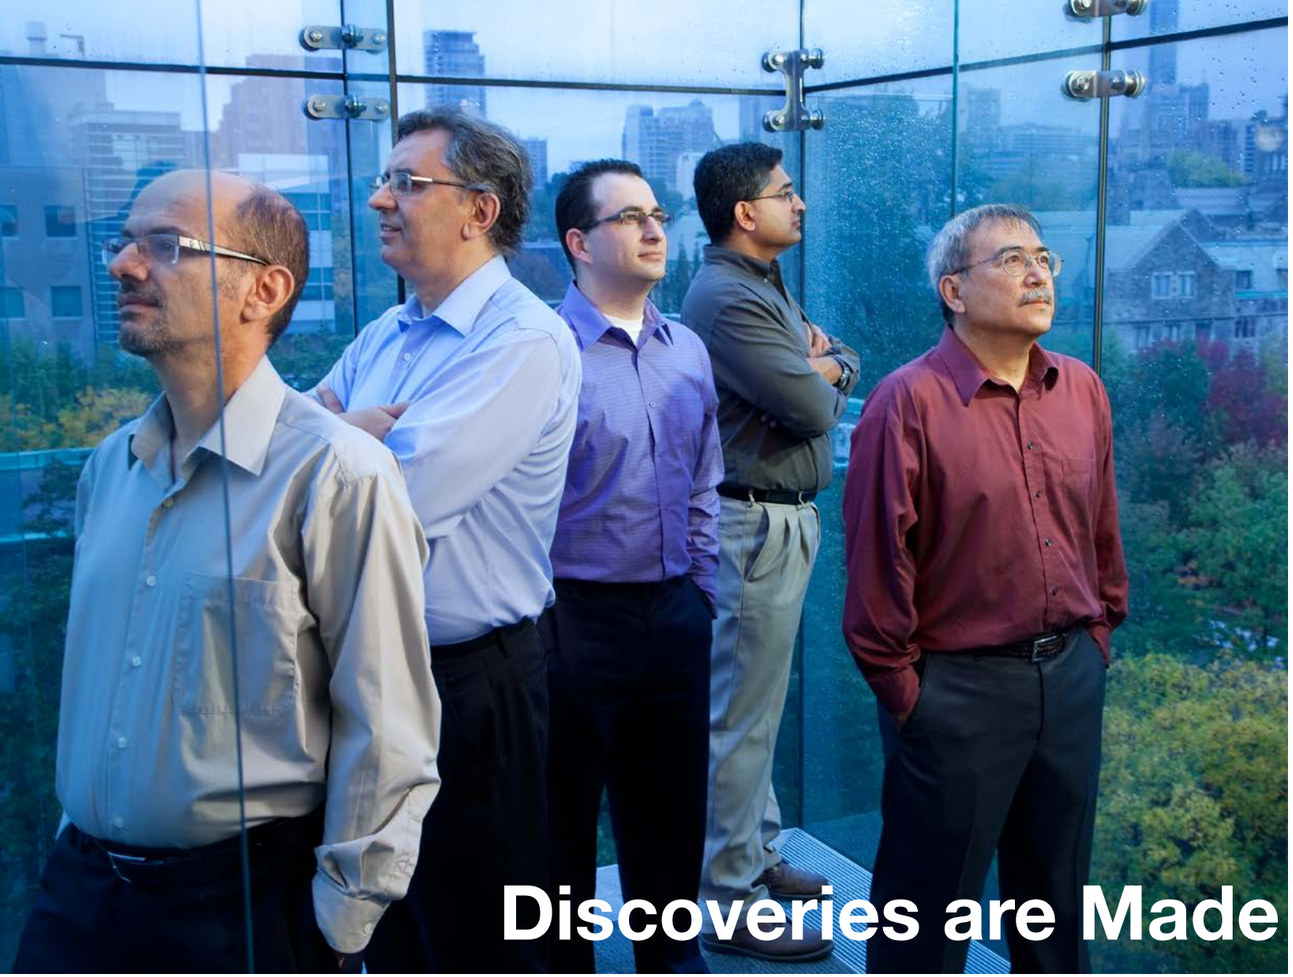

Discoveries are Made while Looking to the Future

TELUS & ECE: Industry Partnership Fuels Innovation

It's a Wednesday morning and you're finishing your breakfast cereal while streaming Futurama on Netflix.

You take a last sip of your coffee, grab your tablet, and rush out the door with a chuckle. As you leave your house the show keeps rolling, your device making a seamless transition from your home WiFi connection to your service provider's LTE. Down into the subway, then up 60 stories to your office in a glass skyscraper, you arrive at your desk just in time to catch the last punch line.

Sound too good to be true? Not for long — that kind of super data-rate mobile service may be reality soon, thanks in part to an ongoing collaboration between TELUS and three researchers in The Edward S. Rogers Sr. Department of Electrical & Computer Engineering. Professors **Alberto Leon-Garcia, Elvino Sousa and Raviraj Adve** are being backed by the telecommunications service provider to explore and design critical components for the network of the future — one equipped to handle skyrocketing demand for data over limited spectrum bandwidth, while maintaining high quality of service.

This industry-academic partnership was spearheaded by Dr. **Ibrahim Gedeon**, TELUS' Chief Technology Officer, who is a big believer in forging stronger ties with the academic community. Dr. Gedeon holds firm to his philosophy that a company should cultivate a healthy ecosystem by pursuing harmonized growth and development through collaboration with academia, business, government and industry. As demand for higher data-rate services pushes TELUS and its competitors to the limits of their carrying capacity, it's more important than ever to invent new strategies to maximize performance out of a system that is confined by limited spectrum bandwidth. Gedeon "is the driving force behind TELUS' funding for university research," says Dr. **Adam Tenenbaum**, a Radio Access Network (RAN) Design Specialist for TELUS and one of the collaboration's coordinators. "It's about developing good engineering resources to support the Canadian market."

Think of the wireless world like a highway — perhaps Toronto's heavily trafficked Highway 401. Lanes of the highway correspond to sections of spectrum bandwidth, dedicated to different services — some for traditional television transmission, some for AM radio, and some for FM radio that you listen to in your car. The lanes dedicated to voice and data transmission for mobile service providers are further subdivided, allowing TELUS and its competitors a fixed number of lanes each. When you sign up with TELUS, you add to the traffic in its lanes. Like the 401, the system was set up long ago with no expectation of the intense level of traffic it now carries. With mobile data demand exploding, the lanes are now clogged with users wanting high-speed all the time, and traffic jams ensue.

But it's not so easy to add more lanes. "We're talking about billions of dollars to buy bandwidth," says Professor Adve. That's why TELUS has turned to ECE at U of T — in hopes of finding ways to get more people on the wireless highway without congestion or conflict.

Enter TELUS' Dr. Adam Tenenbaum, Senior RAN Design Specialist and Dr. **Ivo Maljević**, both ECE alumni themselves. It's their job to scour the country for potential research and development solutions to TELUS' real-world challenges, and to support projects likely to have future applications for their company. "We've selected professors that are at the cutting edge of research," says Dr. Maljević, a Senior RAN Design Specialist. "We strongly believe that the work they do benefits both TELUS and Canada."

The association dates back to a partnership between TELUS and Professor Sousa's group in 2004, and in 2011 when TELUS joined the Smart Applications on Virtual Infrastructure (SAVI) network, an NSERC Strategic Network led by Professor Leon-Garcia. Professor Leon-Garcia then began working with TELUS on a wireline network quality of service project, and Maljević and Tenenbaum expanded into wireless research with Professors Adve and Sousa. Industry investment has been the key to the partnership's success — the funds invested by TELUS are augmented by an NSERC Collaborative Research and Development grant, designed specifically to compound the value of connections between universities and Canadian industry.

"Everyone knows when you use Internet on your desktop computer, it's much faster than using it on your mobile device," says Professor Sousa. "Our goal is to come up with really new ways to close the gap between using your mobile device or desktop computer." The solution involves rejigging the existing network infrastructure to fill service holes between big long-range cell phone towers, called base stations, by incorporating thousands of local relay stations, called small cells, to handle calls and data demand across many small geographic zones.

This mix of base stations and small cells is known as a heterogeneous network, or HetNet, and Professor Sousa's group has been studying these new architectures for over a decade. It's his vision to design the next generation of autonomous infrastructure wireless networks. "We come up with methodologies to deploy small cells in such a way as to manage interference," says Professor Sousa. He likens base stations to people at a party speaking very loudly — it's hard to carry on multiple conversations over their noise. Sousa works on arranging the party-goers so the loud-talking base stations and soft-spoken small cells can be heard simultaneously. He also tries to find the best ways to phase out old technologies, designed primarily for voice calls, to make room for data-intensive 4G and LTE service, without causing major disruptions.

> "It's about developing good engineering resources to support the Canadian market."

Professor Adve's project runs jointly with Professor Andrew Eckford of York University. Their research focuses on resource allocation in HetNets — determining how much bandwidth and how much power each small cell needs to maximize service while minimizing consumption. The closer small cells are to the user, the more effectively they can recycle bandwidth as that person moves from one to the next, and the handoff should be invisible. Professor Adve is looking toward a future of streaming live video in the subway, or Facetime chatting at the top of an office tower. "You should be able to use anything you want, wherever you want. I think that would be a long-term goal," says Professor Adve. "And not just you — hundreds of people."

Professor Adve is excited about the possibility of incorporating real traffic data into his models, further closing the gap between academic theory and practical application. "They're encouraging us to look at very forward-thinking questions," he says. "We've already had quite a few productive discussions on this." And both sides look forward to many more.

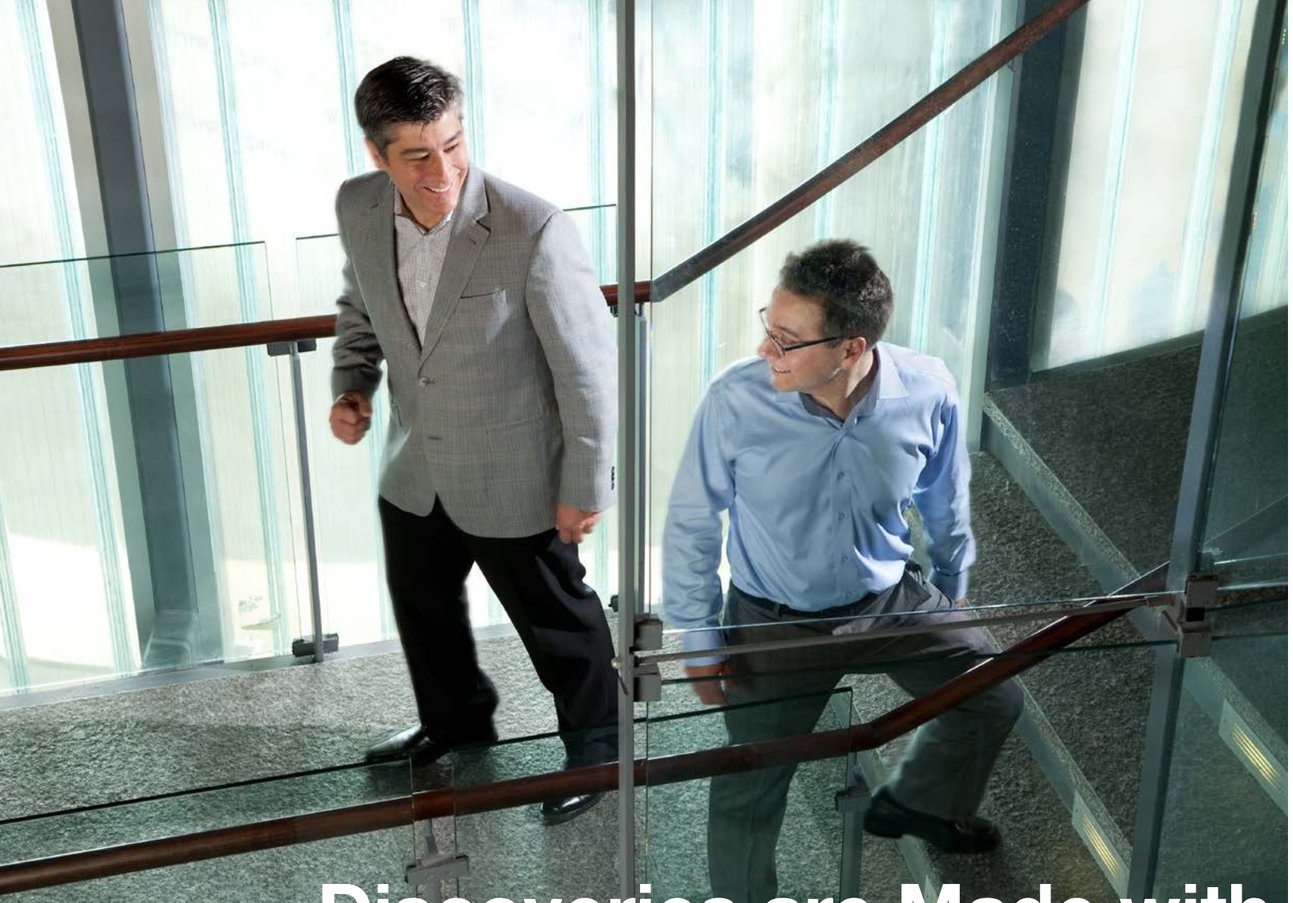

Discoveries are Made with Shared Commitment

Alex and Anthony Grbic: ECE Alumni on Top

If you want to know about Alex Grbic's career highlights, ask his brother Anthony. And if you'd like to hear about Tony's accomplishments, Alex will fill you in.

Each too humble to mention his own successes, the Grbic brothers have good reason to be proud of one another — Alex is currently director of product marketing for Altera and the newest member of ECE's Board of Advisors, and Tony is an associate professor at University of Michigan. Both Alex and Tony are triple graduates of The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, each earning undergraduate, master's and PhD degrees at University of Toronto. But look past the parallel credentials and it's clear that each has carved a career path as unique as it is extraordinary — one to the top of industry, the other to the top of academia.

"We're both ECE grads, but we were in pretty much opposite disciplines within the department," says Alex. "I'm into the computer hardware and digital world, and he's into electromagnetics, math and all that 'fun stuff'. So we're in the same department, but attracted to opposite ends of the spectrum."

"Alex wasn't my guidance counsellor in picking courses!" jokes Tony. "I took some advice from him in undergrad — he convinced me to take an operating systems course in my senior year and it nearly killed me."

Born and raised in Brampton, Ontario, Alex and Tony showed strong interest in technology and aptitude for science and math. When it came time for Alex, the elder brother, to pick a school, "University of Toronto was clearly a favourite," he says. Four years later, Tony considered the program's ranking against Waterloo's and McMaster's, and followed suit. (Skule truly runs in the family blood — their younger sister, Mary Vanda Grbic, graduated from Chemical Engineering in 2007.)

Now widely recognized as a leading researcher in the fields of antennas, microwave circuits and metamaterials, Tony's focus revealed itself gradually. "I wasn't sure in undergrad what I wanted to study until I took the Fields and Waves course taught by Professor **Sergi Dmitrevsky**," he remembers. "That sparked my interest in electromagnetics."

Many years and thousands of students later, Professor Dmitrevsky still remembers teaching Tony. "His dominant feature was a rich imagination, and an ability to accept new concepts," he says. New concepts were flying when Tony launched into his master's in applied electromagnetics with Professor **George Eleftheriades**. "Tony was an exceptional student," says Professor Eleftheriades. "He has a very positive personality, and he's very determined. I was really blessed to work with Tony." Metamaterials was then an emerging area, generating a lot of excitement. Their group published some of the first papers in the field and Tony was starting to enjoy academic research. But one particular evening still stands out to him.

Sir John Pendry, one of the world's leading theoretical physicists and the man largely responsible for inventing the field of metamaterials, was visiting U of T and Professor Eleftheriades invited Tony to join the pair for dinner. "I was ecstatic, I'll never forget that dinner for the rest of my life," says Tony. "The simple gesture of inviting me to dinner with such a prominent researcher kind of meant the world to me and inspired me to commit myself further to this field." After completing his PhD with Professor Eleftheriades, he and his wife Ana made the move to Ann Arbor where Tony joined the faculty at University of Michigan in early 2006.

Pendry wasn't the only high-profile figure Tony would meet — in July 2009 he was named one of just 100 recipients of the Presidential Early Career Awards for Scientists and Engineers, the highest honour bestowed by the United States government on young professionals in the early stages of their independent research careers. In January 2010 he and the other winners received their awards and met U.S. President Barack Obama in the East Room of the White House. Obama offered each researcher a handshake or fist-bump — your choice. (Tony opted for the handshake.)

Both as a student and professional, Alex has a rare ability to combine his deep understanding of technology with a keen business acumen and outgoing personality. "I've always liked working in groups," he says. "If I think back to my ECE days, I was involved in ECE club, and participated in the Engineering Society activities. I enjoyed the personal interaction quite a bit."

In his fourth year as an undergraduate, Alex got involved in a multiprocessor project called NUMAchine, a joint effort between four professors and their grad students, led by Professor **Zvonko Vranesic**. "I really liked the prospect of building something and trying it out, so I hung around for grad school," he recalls. After the FPGA-manufacturing company Altera established its Toronto R&D facility and acquired ECE start-up Right Track CAD, Professor **Stephen Brown**, an Altera Toronto founder, brought Alex in as one of its first hires. "Alex's passion lies at the intersection of business and technology," says Professor **Vaughn Betz**, former senior director of software engineering at Altera and now a professor in the ECE department. "He's a skilled manager, and a gifted communicator — someone who's able to combine the technology perspective and the end-user perspective, and join that seam," Betz says. "First in R&D, then as a manager, senior manager and most recently as a director, Alex stands out as someone who makes the people around him better."

> ...each has carved a career path as unique as it is extraordinary— one to the top of industry, the other to the top of academia.

So when the director of applications engineering role opened up four years ago, Professor Betz recommended Alex. He got the job, and within six weeks he and his wife, Gordana, and their four children had packed up and moved to San Jose, California. "Both of us had spent many years at U of T, so taking that experience outside of U of T and really putting it to the test I think is important," says Alex.

Does he miss his hometown? "Alex needs to get some snow time in," says Tony. "You kind of miss the weather and the fun you can have in the cold," Alex agrees. "But I don't think I've convinced my wife that she should miss it yet!"

Though they're flourishing in California and Michigan, the brothers' roots are still in Toronto. "Everywhere you go, people ask you 'Where did you go to school?'" says Tony. "I definitely keep an eye out for Toronto graduates."

Keep an eye out for Alex and Tony Grbic, because they're going nowhere but up.

Corporate Partners

ABB AB
Actel Corporation
AMD Inc.
AEG Power Solutions Inc.
Agile Systems Inc.
Akai Electric Co., Ltd.
Alcatel Canada Inc.
AlliedSignal Canada Inc.
Altera Corporation
AMP of Canada Ltd.
ARISE Technologies Corporation
Asahi Kasei Microsystems Co. Ltd.
ATI Technologies Inc.
Bell Canada
Bell Laboratories
Bell-Northern Research Ltd.
eCamion Incorporated
ELCAN Optical Technologies
Electro Scientific Industries
EMS Technologies Canada Limited
Exar Corp.
Fiber Optic Systems Technology Inc.
Firan Technology Group Corp.
Fuji Electric Co., Ltd.
Fujitsu Laboratories Ltd.
General Electric
Genia Photonics Inc.
GM Canada
Gennum Corporation
GO Lighting Technologies Inc.
H2Green Energy Corporation
Hewlett-Packard Company
Macroblock Inc.
Magna E-Car Systems
Mark IV Industries Corp.
MaRS Innovation
Mattson Technology Canada
Maxim Integrated Products, Inc.
MDA
Metalink Inc.
Microcore Technology Corporation
Microsoft Corporation
Mircom Technologies
Mitsubishi Electric Research Lab.
Morgan Solar
Motorola Inc.
MPB Technologies
National Semiconductor
Robert Bosch Corporation
Rockwell International of Canada Ltd.
Rotoflex International Inc.
Saft Power Systems Inc.
Samsung Thales Co., Ltd.
Semiconductor Research Corporation
Sendyne Corp.
Siemens
Silicon Mitus Inc
SINTEF Energi AS
Sipex Corporation
Siradel Inc.
Solana Networks
Solantro Semiconductor Corp.
Sony Corporation
Spar Aerospace Ltd.

Name: **Yoley Li**
PEY Placement: Ontario Power Generation
Title: Engineering intern - Plant Computers section, Computer and Control Design department
Experience: "Working in an industry where safety is the number one goal has changed my attitude towards delivering work. It makes me realize that even as an intern, my work could potentially impact many other people and there is no room for mistakes. It is rewarding to see my name appearing on the milestone completion form. Sometimes the staff forget than I'm an intern because I am working on projects just like any other full-time employee."

Bioscrypt Inc.
Black River Systems Company, Inc.
BlackBerry
BLiNQ Networks Inc.
Brammo Inc.
Broadcom Corporation
CANAMET Inc
Carinthian Tech Research Institute
Celestica
Chameleon Systems Inc.
Chip Express
CICLON Semiconductor Device Corp
Ciena Canada Inc.
Cisco Systems
Corning Inc.
Cybermation Inc.
Digital Predictive Systems Inc.
DPL Science
Dynex Semiconductor Ltd.
Honeywell ASCa, Inc.
HRL Laboratories, LLC
Huawei Technologies Co., Ltd
Hydro One
IBM
IMAX Corporation
Inometrix Inc.
Intel Corporation
Invisage Technologies (Canada) Inc.
IPPLEX Holdings Corporation
IVG Fiber Ltd.
Jazz Semiconductor
JDS Fitel Inc.
Kapik Inc.
Kinectrics Inc.
KT Micro Inc.
LG Electronics
Litton Systems Canada Ltd.
Lockheed Martin Canada
Northrop Grumman ISS
NVIDIA Corporation
NXP Semiconductors
Netherlands B.V.
ON Semiconductor
OneChip Photonics Inc.
Ontario Lottery and Gaming Corp.
Ontario Power Authority
Ontario Power Generation Inc.
Opticom ASA
Panasonic Semiconductor Singapore
Peraso Technologies Inc.
Philips Electronics North America C
Qualcomm
Quanser Inc.
Rambus Inc.
Raytheon Canada Limited
Redline Communications Inc.
Ricoh
STMicroelectronics Inc.
Sun Microsystems of Canada Inc.
Systems for Research
Taiwan Semiconductor Manufacturing
TELUS
Telus Mobility
Texas Instruments
Thales Rail Signalling Solutions
Toronto Electric Ltd.
Toronto Hydro-Electric System Ltd
Toshiba Corporation
Total American Services Inc.
UBM TechInsights
Ultra Electronics
USHIO Inc.
ViXS Systems Inc.
Xilinx Inc.
Xogen Technologies Inc.
Zentrum Mikroelektronik Dresden AG

Name: **Yu (Jack) Luo**
PEY Placement: IBM
Title: Continuous Engineering - DB2 Kernel
Experience: "This position features flexible work hours,good benefits and competitive pay. But most importantly, it confirmed my future careerpath of being a computer engineer."

2013–2014

ECE PEY students hired by:

360incentives.com
Adexa, Inc.
Aercoustics Engineering Limited
Altera Corporation
AMD Inc.
Analog Devices Inc.
Apple
Autoliv Electronics Canada Inc.
Bell
Bombardier Aerospace
Cancer Care Ontario
CAST Software Inc.
Celestica Inc.
CommunityLend Finance
Eastern Power Ltd.
Enersource
Envision Mobile
Esna Technologies
General Motors of Canada Ltd.
HP Canada
Hydro One
IBM
Imperial Oil
Intel Canada
Intelliware Development Inc.
ITG
Magna Closures
Marin Software
Marvell Technology Group
Microsemi Corporation
Motorola Solutions/Psion
National Instruments
Ontario Financing Authority (OFA)
Ontario Ministry of Health
& Long-Term Care
Ontario Power Authority
Ontario Power Generation
Ontario Teachers Pension Plan
Qualcomm
QuickPlay Media
SciCan
Semtech
Smith and Andersen
Consulting Engineering
Sunwell Technologies Inc.
Symantec Canada Ltd.
Thales Rail Signalling Solutions
Toronto Hydro
Trapeze Group
TXIO
Tyco Electronics Canada Ltd
University Health Network
Wattpad
Xtreme Labs Inc.

Name: **Ahsan Sardar** (Left)
PEY Placement: Qualcomm
Title: Linux Software Engineer
Experience: "I chose Qualcomm for work-life balance, and the fact that Qualcomm is a leading tech company with a bright future."

Name: **Piyush Gupta** (Right)
Title: Multimedia Software Engineer
Experience: "One of the most innovative and prestigious companies to work for, Qualcomm offers myriad opportunities to turn my ideas into reality."

Directory

Colour-Coded Legend

Each research category has a corresponding colour. Search by colour to locate the lead researcher and project details.

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

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

Quick-Search by Colour-Coded Listing

Communications

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

Quick-Search by Colour-Coded Listing

Computers

• Computer Communications • Computer Software/Hardware

Engineering/Sciences

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

Quick-Search by Colour-Coded Listing

Information Technology

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

Information Technology

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

Energy

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

Life Sciences

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

Aarabi, Parham — Internet Video, Audio and Image Processing

www.apl.utoronto.ca

In the past few years, we have seen an exponential increase in the number of videos and images that have been recorded and placed on the internet. Smart mobile phones (Blackberry, iPhone, etc.) now enable seamless recording, transmission and sharing of videos in near–real time. Whereas just two decades ago there were a few video broadcasters and publishers, today there are millions, if not billions, of video broadcasters and online publishers. With all this visual content, how do we find what we want? How do we categorize the content? How do we develop search engines that bring order to visual content just as text-based search engines (Google, Bing, etc.) brought order to the textual web?

With videos and images, there are of course certain tags manually entered by users that define and categorize the video. However, the 10–15 words that usually accompany a video or image can hardly describe the entire content of the video and at best help to generally categorize the video or mention 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 copyright detection, among other applications). It can also help identify the important parts of a video segment or the most content-filled section of an image (which is useful for compressing videos and intelligently resizing images for mobile devices). The more information we can extract automatically from images and videos, the more we can address a range of practical problems including better search, better (and more contextual) monetization, copyright detection, video compression, intelligent image resizing, as well as a broad range of other important applications.

Although different approaches for visual information understanding have been explored in the past, one of the most promising directions is that of utilizing Extremely Large Datasets (eld). ELDs allow for greater accuracy in extracting information from images and videos, but in return require a substantially greater number of computations for each image processed. Examples of successful ELD systems include the Tiny Images image categorization system, which used a database of 80 million tiny images for image classification, or the Tiny Videos framework ,which utilized a large library of videos for video classification. In both cases, the images and videos were resized to a “tiny” representation in order to minimize the rather large computational load.

We aim to extend these research directions by exploring new hardware and software solutions that enable real time image and video searching using large databases. Our goal can be described in the following two phases: (1) algorithm research and development (including finding ways to utilize ELDs for better image and video understanding and improved visual classification accuracy); and (2) hardware acceleration of the developed algorithms in order to enable accurate real time searching of images and videos using ELDs. To summarize, using currently available images and videos that are either tagged or partially tagged, it is possible to develop highly accurate (and computationally demanding) systems that use this information for understanding and classifying vast amounts of untagged images and videos. In turn, the computational load can be addressed through FPGA-based hardware acceleration, which would enable the classification of an image or video to be performed in rea time.

Abdelrahman, Tarek — Architectural Support for Parallel Programming

www.eecg.utoronto.ca/~tsa

The last few years have witnessed the introduction and then dominance of multicore processors as the mainstream approach to improving performance. Multicores have more than one processor on a single chip. Today’s multicores have four processors on a single chip and industry experts predict that by the year 2020 the number of cores on a chip will exceed 100. However, to exploit the performance potential of multicores, software must be developed to execute on multiple processors. This process, referred to as parallel programming, is a difficult one and is currently one of the major challenges in the field. This problem of parallel programming is exacerbated by the fact that multicores are designed by replicating the hardware of single-core processors. We believe that novel multicore architectures are needed to ease the tasks of parallel programming.

The goal of this project is to explore architectural support for parallel programming. One example of such support is our recent work on the memory versioning scheme. In this scheme, hardware is added to each processor to monitor the accesses the processor makes to shared data. The hardware is designed in such a way that processors access shared data in exactly the same order as in a single processor execution. The result is a parallel programming model that requires little effort on the part of the programmer and delivers good performance. This is achieved at the expense of hardware that is dedicated to versioning. We continue to explore efficient implementation of this hardware as well as other novel approaches to supporting parallel programming at the architecture level.

Compiler Support for GPU Programming

Graphics Processing Units (GPUs) have emerged in the last few years as a promising platform for cost-effective and energy-efficient computing. These highly parallel systems are the core that processes data for graphics-intensive applications, such as computer games on almost every desktop computer. Recently, vendors such as AMD and NVIDIA have provided standard high-level programming interfaces that allow the use of GPUs for more general-purpose non-graphics computation. However, one of the major obstacles still facing GPUs today is the considerable effort needed to program them. This effort is needed to extract and package computations in the form of parallel kernels, to partition computations among threads, to explicitly manage the complex memory hierarchy that exists within the GPU and finally to deal with the various resource constraints that exist inside a GPU.

The goal of this project is to develop compiler support that eases the programming difficulty of GPUs. We have defined a directive-based language that provides users with directives that allow them to perform common tasks such as kernel definition, computation partitioning, data movement, local memory allocation, etc., directly to the sequential code base. Such an approach will make GPU programs easier to develop, to debug and to maintain since the sequential code base is maintained. We are exploring common optimizations patterns performed by GPU programmers and are developing compiler optimization to automate the applications of these patterns.

Dynamic Acceleration of Soft Processors

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

Adve, Raviraj

www.comm.utoronto.ca/~rsadve

Adaptive Signal Processing for Wireless Communications and Radar Systems

Adaptive processing in the physical layer in communication systems: exploiting the spatial and temporal dimensions to improve the quality, capacity and reliability of wireless communication systems; enabling co-operation for energy savings in wireless sensor and data reliability and/or access-point networks.

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

Nano-Photonics 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 nano-structured surfaces for biology, sensing and photonics.
Optical frequency conversion, based on second or third order nonlinearities, provides a mechanism of 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 like nonlinear properties of the III-V semiconductor AlGaAs to develop efficient wavelength conversion devices based on second order nonlinearities (difference frequency generation) and third order effects (four wave mixing).The ability to engineer the dispersion and field profile in a nanowire waveguide has applications in optical sensing. By narrowing the waveguide and incorporating a photonic crystal, or defect state it is possible to control the overlap of the optical field with the sensing material. Typically sensing can be done through a change in refractive index, or absorption of an intermediate material which is sensitive to the substance to be measured, for example, platinum for hydrogen detection. Using this approach it is possible to develop a single chip, with multiple sensors which could detect multiple gases, temperature, humidity and pressure.

Amza, Cristiana

www.eecg.toronto.edu/~amza

Automated Self-Management in Cloud Environments

The economy has been transformed by the delivery of web services over the Internet in the past three decades. Behind the scenes of web-based service delivery technologies are large-scale, complex information systems. These information systems consist of thousands of servers, which store client data in a consistent and persistent manner and are multiplexed among several applications.

Many aspects of content delivery still currently depend on manual fine-tuning and troubleshooting by humans. As a result, the service provider is strapped, with huge human administrator costs for performance debugging, resource allocation to applications and infrastructure configuration adjustments. This is a major impediment to the efficiency of this industry, limiting reinvestment, research and development. To address this problem, we investigate innovative self-diagnosis and adaptive reconfiguration techniques for scalable and available information systems.

We have designed and developed novel online performance modelling and anomaly detection algorithms and tools that form the basis for self-configuring, self-tuning and self-healing servers. We use these techniques in our data centre laboratory towards automatically providing quality of service for a range of dynamic content services such as e-commerce, online bidding and massively multi-player games.

System Support for Parallel and Distributed Software Transactional Memory

Because of the increase in complexity and ubiquity of large-scale parallel and distributed hardware environments, simpler parallel programming paradigms become key. Transactional Memory is an emerging parallel programming paradigm for generic applications that promises to facilitate more efficient, programmer-friendly use of the plentiful parallelism available in chip multiprocessors and on cluster farms.

We developed and optimized libTM, a Transactional Memory library that can be used in connection with C or C++ programs to facilitate more efficient, programmer-friendly use of the plentiful parallelism available in chip multiprocessors and on cluster farms. libTM implements Transactional Memory (TM) for generic applications, 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

www.eecg.toronto.edu/~janders

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

High current density in sub–100 nm ICs has created a power wall, limiting the rate of clock speed scaling in general purpose microprocessors. Attaining higher speed performance and improved energy efficiency motivates the need to develop processors that are customized to specific applications. Performing computations in custom hardware can deliver orders-of-magnitude improvement in energy efficiency and throughput. However, custom processor design, as with any hardware design, is difficult and time-consuming compared to software design. Moreover, hardware design in VHDL or Verilog lies outside the skill set of most software programmers. A further challenge is in identifying what custom hardware would be of benefit to the speed and/or power of an application, as this depends on the run-time and power profile of the application under specific datasets. There is a need then for the rapid and automated design of processors tailored for particular applications.

We are building a self-accelerating adaptive processor by modifying the architecture of a standard processor to create the ability to profile the execution of its own code. Using this profiling ability, our unique processor will be able to identify sections of its code that require optimization. Specifically, the profiling results will drive the selection of program code segments to be re-targeted to custom hardware from their original high-level language implementation. C-to-RTL synthesis will be used, with the RTL subsequently compiled by standard back-end tools. Once the hardware “compute accelerators” are available, the program binary will be modified to access the accelerators accordingly. Programmable logic devices, such as field-programmable gate arrays (FPGAs), are an ideal implementation platform for such adaptive processors, as FPGAs can be configured in milliseconds to implement any digital circuit. The reconfigurability of FPGAs also permits functionality to evolve over time, based on application needs.

Computer Hardware: Applications, Tools, Architecture, Circuits for Programmable Logic

Field-programmable gate arrays (FPGAs) are programmable semiconductor chips that are part of a revolution and poised to be present in every piece of electronic equipment within 10 years. The rapid growth of FPGAs stems from technology scaling that today allows billions of transistors to fit onto a single chip. Each advance in technology is tied to rapidly escalating complexity, such that building a custom chip

now costs tens of millions of dollars and is out of reach for all but a few large companies. FPGAs provide inexpensive access to advanced semiconductor technology, allowing innovation to happen across the spectrum, from small start-ups to large industry to academia.

As they are programmable chips, FPGAs incorporate additional circuitry relative to chips that implement a single fixed function. A consequence is that FPGAs use more power than fixed-function chips and the advantages of FPGAs cannot be realized in the mobile electronics world. Our research aims to drastically improve the energy efficiency of FPGAs to enable their use in the battery-powered electronic devices pervasive in society today.

Today, it can take hours or even days to implement a design in an FPGA and our research seeks to reduce this time to minutes or seconds. FPGA chips are regular — the same circuits repeat in a tiled fashion on an FPGA. This regularity can be used to create patterns for how circuits should be implemented on the FPGA. Our approach is akin to taking a complex problem and breaking it into smaller problems whose solutions can be rapidly looked up in a library and stitched together to form a good solution to the complex problem.

Balmain, Keith — RF Phenomena in Magnetized Plasmas

www.waves.utoronto.ca/prof/Balmain/balmain.html

Magnetized plasmas that are encountered both in space and in the laboratory are highly anisotropic media which have a large impact on RF antennas operated in such media. One effect of a transmitting antenna in a magnetized plasma is the occurrence of luminous discharges such as those observed during the 1995 OEDIPUS-C sounding rocket experiment as well as in subsequent experiments carried out in our laboratory. In the OEDIPUS-C experiment, luminous discharges were observed close to the transmitting antenna early in the flight while Argon thrusters driving the separation of two subpayloads were running. The formation of these RF-discharges as well as the unique shape of the resulting glow patterns are studied in order to better understand the processes that led to their occurrence. In the laboratory experiment, intended as a scaled-down version of the OEDIPUS-C RF antenna, similar RF-discharges were observed that extended as glowing band-shaped beams directed along the ambient magnetic field and extending all along the laboratory plasma column. Research consists of the evaluation of data from the OEDIPUS-C experiment and computational simulations of both the rocket experiment and the laboratory experiment, taking advantage of the vastly improved computer hardware and software now existing, which wasn't available in 1995. Eventually, improved laboratory experiments with emphasis on relating the computed and measured phenomena will be designed and carried out.

Simulation of RF Tag Interrogator Units Underneath Urban Rail Trains

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

Bardakjian, Berj — Bioengineering of the Brain

heart.ibme.utoronto.ca/~berj/berj.html

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

Betz, Vaughn — Improved FPGA Architecture and CAD

www.eecg.utoronto.ca/~vaughn

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). Field-Programmable Gate Arrays are a type of computer chip that can be reprogrammed to perform any function. As the cost of creating chips with billions of transistors has risen to $100 million, most applications cannot justify a custom-fabricated chip and instead are best served by a reprogrammable chip.

Our research seeks to find the best "architectures" for FPGAs — what function blocks should they include and perhaps even more importantly, how can we best programably 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 interconnect with packet-switched networks on chip. This fundamentally raises the level of abstraction of communication on the chip, but requires new CAD tools, which we are also developing, to automate this new and different design flow.

We also seek to find new algorithms and Computer-Aided Design tools to allow FPGA designs to be completed more quickly and to run at higher speeds and make more efficient use of the chip. In particular, we are looking at how to make highly scalable placement and routing algorithms that can handle the latest chips, which contain billions of transistors, in a reasonable run time.

Finally, we are also investigating new areas in which FPGAs can accelerate computation. A key application in this area we are currently researching is the modelling of photodynamic cancer therapy. By simulating the paths of millions of photons in complex human tissue, we seek to determine the best arrangement of fibre optic probes to destroy a tumour (and minimize the exposure of healthy tissue) by activating a photosensitive catalyst only in a local area of the patient's body.

Broucke, Mireille — Control for Complex Specifications

www.control.utoronto.ca/~broucke

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

Patterned Linear Systems

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

Brown, Stephen — CAD and Architecture for FPGAs

www.eecg.toronto.edu/~brown

My research is focused on many different aspects of field-programmable gate array technology, including the design of the chip architectures and the algorithms that are used to implement circuits in these devices, as well as applications of FPGAs. In addition to my faculty position at the University of Toronto, I maintain an active involvement in the Altera Toronto Technology Centre, where I provide direction for the University Program that is offered by Altera. By combining my involvement in both the University of Toronto and Altera, it has been possible to develop research results that are both interesting from the academic point of view and 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, Anthony

Highly Integrated Optical Transceivers

www.isl.utoronto.ca

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

Integrated Circuits for Ultra-High-Speed Digital Signal Processing

This project strives to enable the use of robust and reconfigurable digital signal processing (DSP) for our highest-speed communication links. Presently, analog circuitry performs a lot of of the key signal processing in multi-Gb/s communication links. Technology trends favour an increasing use of digital circuits and software, but ironically the transition will require significant advances in analog circuits such as analog-to-digital and digital-to-analog converters, as well as signal processing algorithms to automatically calibrate and adapt the links to the conditions of real-world use. Moreover, the power of the circuits must be lowered, in some cases by two orders of magnitude. We develop integrated circuit prototypes of our solutions, demonstrating their potential impact on real-world applications.

Ultra-Short-Reach Chip-to-Chip Communication

Our capacity for digital communication continues to increase by integrating more and more functionality into fewer electronic components. Integration enables lower cost, smaller-size systems with lower power consumption. Unfortunately, our ability to integrate complete communication systems onto a single chip today remains limited by small but finite defect rates during chip fabrication, which limit the maximum number of transistors that can be reliably integrated. In addition, the need to combine different fabrication technologies to realize, for example, a high-density memory alongside high-speed transmitter and receiver circuits, makes integration challenging. Emerging dense interconnect technologies may offer a path forward. These technologies place multiple chips in close proximity and connect them with densely-packed wires that may be less than 1 cm long. Our long-term vision is to use these ultra-short-reach (USR) links to interconnect multiple chips so seamlessly that system performance can scale without bottlenecks. We are developing demonstration platforms for USR links, including transceiver circuits that are extremely small and consume very low power.

Chow, Paul

Programming Models and Architectures for Reconfigurable and Heterogeneous Computing Systems

www.eecg.toronto.edu/~pc

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

Internet-Scale Memory Systems

With the vast amount of data accessed and stored using the Internet, new memory architectures are required to host the data that can provide low latency access, low power dissipation and a compact form factor. Current systems use the collected main memory of a cluster of high-end servers for an application that does not need the computation power of such systems. This project explores the use of Field-Programmable Gate Arrays and novel architectures for building Internet-scale "Big Data" memory systems.

Davison, Edward

Control of Large Scale Decentralized Systems

www.control.utoronto.ca/people/profs/ted/ted.html

Our research is focused on the control of large scale systems, where only limited information of 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. Problem areas that immediately arise from large scale systems are current areas of research: decentralized control, intelligent control, fault-tolerant control and the control of unknown systems. A direct application of this research is presently being applied to the control of large flexible space structures, earthquake-resistive building structures, electric power systems with particular focus on micro-grid systems and spinal cord injury patients.

Dawson, Francis

Improving Energy Efficiency of Energy Conversion Processes

www.ele.utoronto.ca/~dawson

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

Draper, Stark

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

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

Large Scale Linear Programming Decoding Via the Alternative Direction Method of Multipliers

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

Re-Architecting Last Level Caches For Low-Voltage Operation

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

Eleftheriades, George — Artificial Materials (Metamaterials) from Microwave to Optical Frequencies

www.waves.utoronto.ca/prof/gelefth/main.html

We are developing paradigm-shift metamaterial devices and subsystems and related technologies at RF/microwave and optical frequencies.

Metamaterials are artificially structured media with unusual electromagnetic properties. Such properties include negative refraction, enhanced evanescent waves through resonant amplification and sometimes a negative group velocity. Our vision is to develop metamaterials that can manipulate and control electromagnetic waves, much as conducting wires manipulate the flow of electrons.

Both three-dimensional volumetric and surfaces (metasurfaces) metamaterials are being developed. A recent effort concerns the development of ultrathin metasurfaces for wavefront manipulation, such as refraction (bending of incident plane waves or Gaussian beams), lensing and controlled beam formation.

Application areas include super-resolution microwave and optical microscopy, detection and sensing, advanced hardware for wireless communications, wireless power transfer, reduction of interference, space technology, radars, defence, solar-cell concentrators, thermophotovoltaics, infrared focal-plane arrays and many more. Examples of devices include small antennas, multi-functional RF/microwave components (including active devices), sub-diffraction imaging lenses and probes (even operating in the far field), ultrathin lenses, invisibility cloaks and related “transformation optics” lenses, plasmonic optical circuits, plasmonic waveguides and nano antennas.

Research includes both experimental work and fundamental theory. Our research is supported by several industrial partners, government agencies and laboratories. Graduates from our group are quite successful in securing faculty positions in academia (e.g., U. Michigan, U. Alberta, McGill, U. Toronto and UBC) and industry (e.g., Apple, AMD, Blackberry, Freescale and Motorola).

Enright Jerger, Natalie — Interconnection Networks for Heterogenous Multicore Systems

www.eecg.toronto.edu/~enright

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

Semantically Rich Networks for Many-Core Architectures

Parallel architectures are rapidly becoming ubiquitous. To leverage the computational power of these multiple cores, communication between cores or devices is essential. This project looks at streamlining the communication between cores via on-chip network innovations to increase its efficiency. The most efficient network design — the one that provides the greatest performance at the lowest cost (area and power) — will be one that most closely matches the functionality required by an application. However, customized, application-specific networks are not appropriate for general-purpose many-core architectures since they run many diverse applications. Therefore, we focus on communication behaviours that are evident across a range of workloads, specifically behaviour exhibited by cache coherence protocols. Cache coherence protocols introduce communication overhead and can substantially impact performance, as many of these operations lie on the critical path. Examining several coherence protocols, we note the presence of coherence primitives that use multicast and reduction operations. Our reduction routing combines redundant messages (such as acknowledgements) during their network traversal to reduce network load. This insight and router architecture can be more broadly applied to any many-to-one communication pattern.

In addition to many-to-one and one-to-many messages that are common to many coherence protocols, coherence protocols exhibit additional exploitable behaviour. For example, these protocols are characterized by having a mix of long and short messages; cache lines (data messages) represent long messages while coherence requests (control messages) are typically only a few bytes long. Existing theories for deadlock freedom in fully adaptive routing require a conservative virtual channel allocation scheme. We show that this restriction severely limits performance for short coherence control messages. We proposed a novel flow control technique, whole packet forwarding that accelerates the handling of short coherence messages in the network to improve performance and efficiency.

An important feature of these solutions is the low hardware overhead they incur. Small hardware modifications and modest additional logic are required to support each of the above-mentioned designs. These optimizations yield significant throughput and latency improvements for a variety of workloads.

Simulation Methodologies for On-Chip Networks

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

Improving Parallel Application: Focus on Communication

Our recent work explores opportunities to ease programming of large parallel systems. First, we explore improved software broadcasting algorithms on the Intel SCC. We propose several novel broadcasting techniques implemented that leverage the message passing hardware on the SCC. We demonstrate that broadcasting can be critical to application performance on this system. We also demonstrate that by improving the communication substrate, we can ease the burden on the programmer by allowing them to use broadcasting rather than devoting extra time to carefully partitioning their algorithm and orchestrating point-to-point communication. Further enhancement to the hardware and software of these types of many-core systems will reduce the burden placed on the programmer. Next, we propose a runtime system that effectively distributes a single OpenCL kernel across a cluster of GPUs. Distributing a task across a cluster involves partitioning the work into smaller tasks, scheduling these tasks, partitioning memory and tracking and transferring memory that is written and read by various tasks. DistCL abstracts many of these challenges associated with partitioning the kernel to run on multiple GPUs and managing memory from the programmer. DistCL allows efficient distribution of unmodified OpenCL kernels across a cluster of GPUs; as a result, significant speed-ups can be achieved.

Francis, Bruce — Control Theory and Applications

www.sites.google.com/site/brucefranciscontact

1. Distributed robotics theory: We study the mathematical theory of robot formations. The robots are typically modelled as unicycles with only onboard sensors and no leaders. The objective is to design local motion strategies so that a team of robots performs a coordinated task, such as forming a circle. (with Professors Mireille Broucke, Manfredi Maggiore and Luca Scardovi)

2. Applied robotics: Starting in 2007, we conducted an application in collaboration with Defence Research and Development Canada (DRDC), Suffield, Alberta. Motivating this research is a military situation in which a manned vehicle convoy traverses hostile territory to deliver supplies. We designed and tested a vehicle-following system to allow a convoy of full-sized autonomous vehicles with large inter-vehicle spacing to follow a manually-driven lead vehicle’s trajectory without cutting corners on turns. Our testing was done on MultiAgent Tactical Sentry (MATS) vehicles that were provided by DRDC. Since there are no inter-vehicle communications to relay the lead vehicle’s position, the goal of an autonomous follower is to track the trajectory of its immediate leader. (with Professor Tim Barfoot)

3. Infinite lattices of dynamical systems: In studying the formation of a very large number of vehicles, one approach is instead to model an infinite number of vehicles .The question then arises as to what mathematical framework to take so that the latter model correctly describes the behaviour of the former. This leads to the subject of infinite chains or lattices of dynamical systems. (with Professor Avraham Feintuch)

Frey, Brendan — Algorithms for Inference and Machine Learning

www.genes.utoronto.ca

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

Data Analysis and the Affinity Propagation Algorithm

Summarizing data by identifying a subset of representative examples is important for scientific data analysis and in engineered systems. Such exemplars can be found by randomly choosing an initial subset of data points and then iteratively refining it, but this only works well if that initial choice is close to a good solution. Dr. Frey's group developed a new method called affinity propagation, which takes as input measures of similarity between pairs of data points. Real-valued messages are exchanged between data points until a high-quality set of exemplars and corresponding clusters gradually emerges (Frey and Dueck, Science 2007).

Because of its simplicity, general applicability and performance, the affinity propagation algorithm is widely used in science and engineering. In the past year, an on-line web tool developed by Dr. Frey's group was accessed over 100,000 times by over 3000 users (unique IP addresses), 600 of which were from Canada. Google returns over 10,000 hits for the search term affinity propagation . Dr. Frey's method has been applied to solve problems in biology, genetics, genomics, medicine, physics, chemistry, telecommunications, electronics, archeology, economics and social networks.

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

Portable, Wearable and Implantable Sensory Biomedical Electronics

www.eecg.utoronto.ca/~roman

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

Binary Instrumentation of Operating Systems

www.eecg.toronto.edu/~ashvin

A binary instrumentation system enables monitoring and manipulating 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 runtime 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

VLSI for Digital Communications

www.eecg.toronto.edu/~gulak

In the area of digital communications, we have continued to develop several practical ways to improve the performance and implementation of wireless systems that use multiple antennas (MIMO) for improved diversity and capacity. One of the key elements that we have investigated is the subsystem in the baseband known as the detector, which is responsible for data detection. A key contribution is the creation of an innovation that we call an on-demand K-best algorithm (a breadth-first search technique) whose complexity scales linearly with constellation size. This innovation is key to supporting higher order modulation schemes such as 64-QAM and 256-QAM systems that will appear in next-generation communication standards, necessary for Gbps performance. We have implemented and tested our algorithm in 0.13um CMOS and have generated the best-known results published in the literature to date, with respect to data rate, power efficiency and area. Our results have been extended to soft detection and tested with CMOS prototypes for use with iterative FEC decoding schemes.

We have also made recent contributions to an important channel pre-processing block found in all MIMO systems, namely that of QR decomposition, a function needed for decomposing the channel matrix. Our key contribution in this area is the development of both algorithms and a 0.13um CMOS implementation that demonstrates the world's lowest (best) processing latency.

Another area of recent accomplishment is in a channel preprocessing element known as Lattice Reduction, which can be used to mitigate scattering and antenna correlations that exist in practical MIMO systems. Lattice Reduction is a baseband signal processing algorithm to re-orthogonalize the signal space with the objective of improving BER performance. We have developed several algorithmic innovations and the world's first CMOS prototypes for Lattice Reduction; the concepts developed will be particularly attractive for low-power implementations. Future work focuses on next-generation wireless OFDM baseband signal processing algorithms and their high-performance, low-power CMOS realization.

Hatzinakos, Dimitrios

Biometrics User-Centric Sensor Networks (BUSNET)

www.comm.toronto.edu/~dimitris/research/busnet.pdf

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

Efficient Resource Allocation Strategies for Wireless Multimedia Communications

www.comm.utoronto.ca/~dimitris/research/multisignproc

One of the major objectives of future-generation communication networks is to provide high-quality multimedia content to users. This demand necessitates more efficient utilization of limited resources such as power and spectrum. While it is essential to minimize consumption of limited resources, the conservation should not come at the cost of inferior quality of service (QoS). As a result, power- and spectrum-efficient strategies, that can also guarantee some level of QoS, are highly desirable. We propose a generalized framework of resource allocation, which enables efficient integration of various adaptation methods and strategies: efficient use of available bandwidth and power; adaptive modulation and coding coping with highly varying wireless channel conditions; and an integrated and multilayer design for overall performance gain. With dynamic resource allocation, improved flexibility and robustness can be obtained in the hostile wireless channel environments. Various channel distortions can be mitigated efficiently, successfully accommodating various user needs in a wide range of scenarios. Depending on the application, a preselected level of QoS can be guaranteed while keeping resource consumption to a minimum. Together, these strategies offer an attractive communication framework of increased power and spectral efficiency, which will enable high-data-rate wireless multimedia communication to be an affordable and practical reality.

Medical Biometrics

www.comm.utoronto.ca/~biometrics/medical/index.html

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 wave-form 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 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.

Self-Powered Sensor Networks

www.comm.toronto.edu/~spsn

The University of Toronto, AD Telecom and SRADEL are partners in developing compelling materials, communication architectures, software and other critical technologies necessary to create self-powered, ubiquitous and wireless ad hoc sensor networks. Substantial benefits will be realized by the citizens of Ontario and by Canadian society in general with the commercialization of a family of products that take advantage of these sensor networks, along with the novel energy harvesting and power generation technologies used to support them. The panoply of envisioned applications include effective, responsible and sustainable monitoring and governance in structural health, disaster relief, 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 current state-of-the-art extensive infrastructure for collecting massive amounts of sensor data in order to provide critical functionality for (i) management of inconsistent and uncertain data; (ii) light-weight data integration; (iii) data cleaning and social network analysis; and (iv) various enhanced security functions for device authentication and data protection under a wide range of attack scenarios.
(3) Creation of innovative energy conserving, capture and storage technologies that use novel nanoscale materials, energy harvesting methods and renewable energy resources to supply consistent power to sustain autonomous sensor networks. The research on self-powered sensor energy systems will focus on five major areas: (i) power conditioning and conservation; (ii) electromagnetic energy harvesting; (iii) solar energy harvesting; (iv) vibrational-thermal energy harvesting; and (v) energy storage. The ultimate target is low-cost, miniaturized, readily integrable, 24/7 energy generation systems that can sustain on the average a few hundred mW pulses, as well as support continuous current draw at mA scale. Several, if not all, of the energy generation techniques will undoubtedly be extended to the next-generation, large-scale wireless technologies to further reduce dependence on fossil fuels and other environment-taxing resources.

Helmy, Amr S.

Infrared and THz Semiconductor Laser Chips

Bragg reflection waveguide lasers 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 epi-layers. The core is a layer of the low refractive index material and the device operates in Bragg reflection waveguide (BRW) mode, not the conventional total interface reflection mode.

Bragg reflection waveguide lasers have the potential to enable the realization of high power single mode lasers and amplifiers with larger mode volumes, high gain coefficient and strong mode discrimination. Moreover, this class of novel lasers has also shown potential in applications related to nonlinear frequency conversion, monolithically integrated optoelectronic integrated circuits.

We have demonstrated the first edge-emitting BRW laser-diode with low threshold current and high characteristic temperature. This group's research focuses on using this class of lasers to develop high performance single-mode lasers, monolithically electrically injected optical parametric oscillators. These chip-based sources can provide continuous coverage of spectral regions, which are not accessible by other technologies including quantum cascade lasers. Examples of niche applications served by this unique platform include sources for environmental and biomedical sensing elements in the 1–4 μm window and chip-based THz spectroscopy sources.

Nanophotonic Devices and Networks

photonics.light.utoronto.ca/helmy/nanophotonics

Our work in the field of nanophotonics focuses on the design of functional devices based on plasmonic slot waveguides (PSWs).
With subwavelength footprint, versatile functionality and low parasitics, PSW is a promising platform for creating high-speed optoelectronic devices with low power consumption and can potentially help alleviate the latency and power dissipation bottlenecks in current VLSI technology. Specifically, our research addresses four components of PSW device design: excitation mechanism, functionality, modelling and loss.

For efficient PSW excitation, we have developed a broadband, non-resonant excitation mechanism by interfacing conventional dielectric waveguides and PSWs in a direct, orthogonal junction to achieve phase matching. By minimizing the momentum mismatch between the two waveguides, efficient energy transfer across a large spectrum can be obtained instantaneously at the junction interface. These devices exhibit record broadband performance and occupy the smallest footprint possible, both of which are ideal for temperature-insensitive hybrid optoelectronic applications.

To create PSW devices suitable for on-chip integration, our research explores reconfigurable PSW networks, which consist of 2D networks of intersecting PSWs. Utilizing the ability of bound surface waves to propagate over sharp bends, the simultaneous power distribution and interference within a network of PSW junctions can enable the design of multi-input multi-output optical devices with diverse functionalities.

In order to reduce the computational cost associated with designing PSW devices through numerical techniques, our research also involves creating analytical model for PSW mesh structures. By approximating PSWs as microwave transmission lines, a scattering matrix model based on characteristic waveguide impedance has been formulated to model the dispersion of localized plasmonic waves within mesh structures. The transmission responses of the mesh are encapsulated into generic closed-form expressions that can handle arbitrary combinations of junctions without requiring numerically-extracted parameters. Thus, the model serves as a scalable, generalized framework that can be adapted to model specific mesh configurations as well as perform device design, optimization and sensitivity analysis.

Photonic Integrated Circuits for Ultrafast All Optical Signal Processing

Parametric processes based on second-order optical nonlinearities in III-V semiconductors are an ideal platform for the development of novel parametric devices for all optical signal processing. The efficiency of parametric processes chiefly relies on the phase-matching technique employed. Because of the lack of natural birefringence in compound semiconductors, phase-matching can be challenging in these materials. Our group has proposed and successfully demonstrated an exact phase-matching technique using Bragg reflection waveguides (BRWs) in the AlxGa1-xAs material system. Current research focuses on improving the conversion efficiency of the nonlinear processes by investigating advanced transverse waveguide geometries as well as extending the technique to other second-order nonlinear processes including sum- and difference-frequency generation that benefit other domains of all optical signal processing in photonic integrated circuits.

Bragg reflection waveguides are also used to achieve phase-matching for spontaneous parametric down-conversion in monolithic AlGaAs waveguides. Through the dispersion control afforded by this technique, bandwidth tunability between 1 nm and 450 nm could be achieved using the same vertical wafer structure. This tuning can be achieved through the lithographic process used to define the waveguides. It can also be achieved by utilizing both type-I and type-II phase-matching conditions. This technology offers a promising route for realization of electrically pumped, monolithic photon-pair sources on a chip with versatile tunable characteristics.

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 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 main stream 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 breakthroughs by our group, fully integrated, portable and robust entangled photon sources have been made possible using the mainstream semiconductor technology. We successfully demonstrated the generation of entangled photons from a semiconductor chip. This chip is specially engineered, which not only increases the photon generation efficiency from the bulk crystal counterparts, but also makes the integration with other optical components possible. Our technique could lead to the world's first fully integrated, room temperature entangled photon source in the foreseeable future. Meanwhile, our group has been working on engineering the generated photon properties on the same platform. Ultimately, our techniques will allow the entire photon generation and manipulation processes on the same chip, which could be a big step towards a practical, commercial quantum computer and other quantum information processing systems.

Sensing Liquid-Phase Nano- and Bio-Materials in Optofluidics Using Raman Spectroscopy

Conducting Raman spectroscopy in hollow core photonic crystal fibres (HCPCFs) results in significant Raman intensity enhancements (approx. two orders of magnitude) compared to direct sampling scheme 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 HCPCF for optical sensing originates from the increased light-matter interaction volume and efficient accumulation of the Raman scattering along the extended length of the HCPCF. The well-confined excitation interacts directly with the sample molecules while propagating along the length of the HCPCF and Raman scattering can be efficiently excited along the fibre's entire length.

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. By employing the HCPCF platform for light-matter interaction, both the pump laser and the QD solution can be confined within the central core of the HCPCF. Subsequently, the confined laser power within the core induces a strong interaction with the solution that is filled inside which allows an enormous amount of Raman signals to be induced and collected compared to the conventional Raman scheme. With this novel technique, Raman signals from aqueous or diluted solutions can be enhanced by two to three orders of magnitude without the use of additional metallic nanoparticles. This enhancement 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, CdS0.7Te0.3 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 bio-conjugation abilities.

Herman, Peter

photonics.light.utoronto.ca/laserphotonics

3D Laser Fabrication: Enabling Nano-optics for the Nanosciences

Nanoscience and nanotechnology define significant trends today that seek to exploit the new physical laws encountered as the structures and devices we make become smaller and smaller. Most attention is on the quantum effects that dramatically alter the electrical, magnetic, optical, mechanical – virtually all – properties of materials in surprising but exceedingly useful ways. Optical physics is playing a significant role in both these trends. Near-field, multi-photon and confocal microscopy and laser tweezers are opening the frontiers of nanoscience by probing and manipulating individual nanostructures, while laser lithography is a nanotechnology used in high-volume manufacturing of nano-size transistor gates in microelectronic chips.

As this race to shrink the world goes forward, optical physics is also evolving with its own set of challenges to understand and to harness the new optical phenomena in nanostructures much smaller than the wavelength of light. Nano-optics has emerged as the new discipline that promises new optical materials (photonic bandgap crystals, metamaterials, plasmatronics) to guide light at dimensions below conventional diffraction limits or probe the electron wavefunction of protein molecules with powerfully enhanced optical resolution.

To this end, the proposed NSERC program seeks to invent a new means of laser optical beam delivery that will facilitate the fabrication of 3D nano-optical systems. Near-field and phase-shifting techniques will be exploited in multi-level diffractive optical elements to design "intensity defects" within 3D periodic interfering laser patterns. Photosensitive optical materials exposed to these modified laser patterns will see nano-optic devices precisely embedded at the critical points of a 3D periodic lattice to enable the nanofabrication of compact 3D photonic crystal circuits, 3D optical-domain metamaterials and nanofluidic chromatography sensors for cell proteomics. This significant extension of laser holography promises a powerful advance in nanooptics and defines a new paradigm for high-volume manufacturing — contactless 3D nanomolding — of significance to Canada's optics, biophotonics and nanotechnology industry.

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

The symphony of colours and light flashes generated during laser machining attest to the dramatic undulating physics evolving rapidly as material is transformed through various states of matter, heated to sun-like temperatures, shocked to explosive pressure and finally ejected at supersonic velocity. Understanding and controlling such complex phenomena is a major science challenge and, regretfully, too poorly understood to benefit today's industry as it attempts to steer wanton laser processes into reliable nanoscale manufacturing methods for a new generation of smart medical devices, bio-sensors and Telecommunication products.

The proposed program aims to improve the fundamental understanding of laser interactions at the forefront of "burst" ultrafast laser processing and self-focusing "filamentation" machining — effects first discovered by our group. Our aim is to turn this understanding into "intelligent" laser control methods that can manage the highly nonlinear light interactions in transparent materials and possibly open a new direction for three-dimensional manufacturing. A novel burst generator provides tailored laser profiles to build up heat accumulation and annealing effects to counter shock and other collateral damaging effects. State-of-the-art delivery systems with self-learning algorithms for spatio-temporal beam shaping are to be developed for dynamic focusing into transient plasma and defect centres that promise to control the size, position, stress and morphology in various laser machining directions. Powerful "5-D" spectroscopic and phase-contrast microscopy tools will uniquely harvest the rich optical signature of the laser physics to offer real-time monitoring as optical and microfluidic devices take shape. Femtosecond laser filamentation is a new opportunity for deep penetration machining and stress-induced scribing of transparent media like flat-panel display, silicon wafers and lab-on-a-chip devices. The program aims to deliver new laser diagnostic and control systems for 3D manufacturing of Telecom optical circuits, optical fiber assemblies, smart medical catheters, optical sensors and other high value photonic systems for our Canadian partners.

Hum, Sean

www.waves.utoronto.ca/prof/svhum

Antenna Arrays for Wide-Angle Beam Scanning

This project is exploring a range of technologies for improving the beam-scanning range of electronically-scanned antenna arrays, including that of phased array, reflectarray, and array lens architectures. Currently, these arrays cannot scan too far from the broadside direction of the array, which is constraining in many applications. Using unique electromagnetic materials and antenna designs, this constraint can be removed, allowing these systems to scan over a wide angular range. For example, for satellite tracking, such antenna arrays can scan from horizon-to-horizon, a critical feature in such systems.

Reconfigurable Antennas for MIMO and Compact Handsets

This project is exploring the development of reconfigurable antennas for use in compact terminals such as handsets. Such antennas can significantly improve signal diversity which is crucial for improving the performance of multi-input multi-output (MIMO) systems. They can also benefit handsets operating in highly dynamic environments where antennas with agile characteristics can be used to effectively deal with changing channel conditions. This type of technology improves the capacity and reliability of wireless networks while maintaining a low-cost form of adaptive antenna diversity on wireless terminals.

Reconfigurable Wideband Spatially Fed Arrays

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

Transparent Reflectarrays for satellites

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

Ultra-Wideband Antenna Technolgy for Wireless Localization

This project is exploring the capabilities of ultra-wideband (UWB) technology for wireless localization (the determination of an object's position using wireless technology). UWB can provide very high localization accuracy, especially when coupled with UWB antenna arrays or advanced antenna designs. This project is developing such technology and exploring applications for vehicular applications such as rail signalling systems.

Iravani, Reza

www.ele.utoronto.ca/prof/iravani/iravanimain.html

Real-Time Simulation, Control and Protection of Integrated AC-DC Power Systems

This work includes research and development of analytical and time-domain simulation tools and control and protection strategies/algorithms for: (1) interconnected AC power systems that imbed overlay High-Voltage Direct-Current (HVDC) grids, mainly for large-scale integration of wind and solar power; and (2) microgrids with high-depth of penetration of distributed generation and storage units.

Iizuka, Keigo

www.keigo-iizuka.com

Omni Focus Video Camera

Our major achievements during the recent past were the invention of two novel types of distance mapping video cameras. The first invention, called the Axi-Vision Camera, is a distance mapping camera that is based on the combined principles of time of flight and modulated light illumination. Television programs produced by using the Axi-Vision Camera have been broadcast from NHK, Japan.

In a contest sponsored annually by *Optics & Photonics News* of the Optical Society of America, the paper on the Axi-Vision Camera was selected as one of the most significant scientific accomplishments described in a refereed journal in 2002. The Axi-Vision Camera was commercialized by NHK Enterprises, Japan and the first unit was sold for $400,000. We received the 2003 Fujio Frontier Award in recognition of our leading-edge research and development of the Axi-Vision Camera.

The second invention, called the Divcam (short for Divergence Ratio Axi-Vision Camera), is a distance mapping camera that utilizes the universal decay rate of the illuminating light with distance. The Divcam is light weight, 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 Omni-focus Video Camera" highlighted this achievement: the image appeared on the cover of the February, 2012 issue of the journal *Applied Optics*.

Jacobsen, Hans-Arno — ACC - AspeCt-oriented C

www.msrg.org

AspeCt-oriented C (www.AspeCtC.net) implements an aspect-oriented extension to C and offers one possible language design for an aspect-oriented C language. AspeCtC is open-source, released under GPL.

AspeCt-oriented C is a research project conducted by the Middleware Systems Research Group at the University of Toronto. ACC enables aspect-oriented software development with the C programming language. AspeCt-oriented C consists of a compiler that translates code written in AspeCt-oriented C into ANSI-C code. This code can be compiled by any ANSI-C compliant compiler, like for example gcc.

AspeCt-oriented C is a proposed language design and compiler. ACC serves as one viable AspectC language design. AspeCt-oriented C ships with a set of Compiler Tools that help to use ACC as part of larger development projects, either to integrate aspects and ACC compiler into larger builds or to organize new software development builds with aspects in mind.

The objective of the AspeCt-oriented C project is to build a robust compiler to support aspect-oriented programming with C. ACC achieves this by building on proven aspect-oriented language designs for other languages, most notably the AspectJ language for aspect-oriented programming with Java.

eQoSystem: Towards Declarative Distributed Applications

The eQoSystem project (eqosystem.msrg.org) seeks to simplify the development and management of business processes deployed on a distributed Service Oriented Architecture (SOA).

The target architecture is an enterprise system with distributed services coordinated by application workflows or business processes. Declarative goals, specified in Service Level Agreements (SLA), 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.

The PADRES ESB - Events and Services Bus

PADRES (padres.msrg.org) is an open-source, enterprise-grade event management infrastructure that is designed for large-scale event management applications. Ongoing research seeks to add and improve enterprise-grade qualities of the middleware.

The PADRES system is a distributed content-based publish/subscribe middleware with features built with enterprise applications in mind. These features include: (i) Intelligent and scalable rule-based routing protocol and matching algorithm; (ii) Powerful correlation of future and historic events; (iii) Failure detection, recovery and dynamic load balancing, (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, (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 — Advanced Interface Circuits for MEMS Technology

www.eecg.toronto.edu/~johns

Micro-ElectroMechanical Systems (MEMS) refer to tiny devices that combine micrometre-scale mechanical devices with micro-or nanoscale electronic circuits to sense physical quantities. Some recent examples of commercial applications for MEMS are pressure sensors used as microphones in devices such as cellphones and hearing aids, inertial sensors used in airbag deployment as well as positional control in handheld games and cellphones and gyroscopes used for image stabilization for cameras as well as angular velocity measurement in handheld games.

This research program investigates new circuits and architectures that will significantly improve MEMS power dissipation as well as improve accuracy performance. With improved accuracy, new applications can be developed that are not otherwise possible. For example, a highly accurate inertial sensor can be used to track position by integrating acceleration to obtain velocity and then integrating velocity to determine distance travelled.

Joy, Mike — Current Density and Conductivity Imaging with MRI

www.currentdensityimaging.org/

This is a very active research area. Every decade a new medical imaging technology appears or matures. The most recent of these is magnetic resonance imaging (MRI). We are investigating the use of magnetic resonance imaging for tissue characterization. The goal here is to physically characterize not just the shape of an organ or neoplasm but also the tissues within it.

We are also studying electric current density imaging with potential application in electrical safety, therapeutic electrical stimulation, management of electrical burns, impedance tomography and solution of the inverse problem of the measurement of bioelectric sources.

Our analytical tools include Fourier transforms, singular value decomposition and phase unwrapping. We use magnetic resonance imagers to image electric current. We have access to a clinical 1.5 T magnetic resonance imager.

Our research is closely related to impedance tomography (the electrical measurement of tissue conductivity) and biomagnetic encephalography (the measurement of the biomagnetic fields produced by neural currents). The patient group benefiting from our research has not yet been identified.

Kherani, Nazir — High-Efficiency Silicon Photovoltaics

www.ecf.utoronto.ca/~kherani

The objective of this project is to research, develop and integrate a set of thin-film technologies that will lead to prototype demonstration of high-efficiency silicon photovoltaic solar cells. The novelty of the research lies in the development and integration of unique, production-worthy technology elements which will ultimately make it possible to attain the lowest cost per watt peak (Wp) of silicon-based photovoltaic solar cells.

Silicon offers one of the highest photovoltaic energy conversion efficiencies. This property, along with its stability, abundance, environmental compatibility and technological maturity, make silicon a prime material for photovoltaics. However, the challenge today is cost. Cost can be reduced by decreasing production and material costs and by increasing energy conversion efficiency.

The unique technology elements comprising this project are (1) development of high-efficiency solar cell concepts with the objective of producing the greatest quantity of solar electricity per gram of silicon (i.e., thin silicon); (2) use of low-temperature, high-quality thin-film synthesis techniques with the objective of implementing low-thermal-budget and high-production-rate processing; and (3) integrated development of PV cell concepts, photon harvesting techniques and production processes compatible with the drive to continually reduce the silicon absorber thickness.

Micro-Power Sources and Sensors

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

NICE Composite Materials

Nano-Integrated Carbon-Enveloped (NICE) composites are being developed (1) as smart coatings for energy conservation in buildings' applications and (2) as novel materials for photonic applications.

NICE composites, based on diamond-like carbon film, which is a silicon-compatible material, have thermal, mechanical, optical and electrical properties that can be tailored over extremely wide ranges, yielding a versatile material for photonic, optoelectronic and micro-electro-mechanical systems (MEMS) applications.

The objective of this project is to demonstrate NICE composites as a viable platform material for the development of smart coatings for building energy applications and to explore their viability as passive and active rare-earth base photonic materials.

Photonic Crystal–photovoltaics

The aim of this project is to investigate photonic crysta–photovoltaic integrations with the aim of creating high-efficiency, economic, third-generation solar cells. The novelty of the research lies in innovative integrations of nanomaterials and thin film semiconductors.

As thin-film crystalline or nanocrystalline silicon solar cells are made thinner, light trapping at wavelengths near the absorption edge becomes increasingly important (e.g., absorption lengths are 10 mm and ~1 mm for wavelengths of 800 nm and 1100 nm, respectively). Upon applying perfect random scattering on an incident silicon surface with a lossless back reflector, a maximum path length enhancement of ~50 is expected, though in reality the actual value is closer to 10. However, much larger path length enhancement factors, on the order of 103 to 104, are required to effectively absorb the longer wavelength light. An alternative approach is light localization through the application of photonic crystals.

Photonic crystals (PC) are periodic dielectric structures that affect the behaviour of electromagnetic waves similar to periodic potentials in semiconductor lattices that affect the behaviour of electron waves.

Through the exploration of a range of nano-integrations, we have recently proposed a novel class of transparent conducting porous nanocomposite films amenable to a variety of device applications. One application involves the use of a selectively transparent and conducting photonic crystal as an intermediate reflector for efficiency enhancement of tandem thin-film silicon micromorph solar cells. Other application areas include bifacial PV, LEDs and catalytic process applications.

Khisti, Ashish — Low Delay Communication Systems for Streaming Media

www.comm.utoronto.ca/~akhisti

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 — Energy of Decoding

www.comm.utoronto.ca/frank

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

Fibre-Optic Communication Using the Nonlinear Fourier Transform

Fibre-optic transmission systems are evolving at a rapid pace towards achieving greater spectral efficiencies. Coherent detection is supplanting noncoherent detection and polarization multiplexing and advanced modulation schemes are being implemented. Today’s high-speed electronics enables 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.

Physical-Layer Network Coding

Nested-lattice-based physical-layer network coding is a type of compute-and-forward (C&F) relaying strategy that is emerging as a compelling information transmission scheme in Gaussian relay networks. While most wireless protocols try to avoid interference, C&F allows two (or more) users to transmit simultaneously. In this case, a (random) linear superposition of the signals is observed by the relay. Rather than discarding this information, or attempting to separate the messages, the new idea of C&F is that the relay can attempt to decode a linear combination of the messages being transmitted, i.e., a “linear equation.” Such linear equations can be forwarded to the next relay and so on. Exploiting the property that a linear combination of linear equations is again a linear equation, the eventual destination is presented with a linear system to solve to recover the original messages. There are many potential practical applications of this idea; in particular, it may be possible to design new Wi-Fi routers that can tolerate much more interference (many more active users) than existing ones while still providing high-quality network access.

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 — Cyber-Physical Protection of the Smart Grid

www.comm.utoronto.ca/~dkundur/

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 behaviors and new forms of vulnerabilities. However, opportunities may also exist through the coupling to improve system survivability to 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, re-develop 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.

A Cyber Security Impact Analysis Framework for the Electric Smart Grid

The scale and complexity of the smart grid, along with its increased connectivity and automation, make the task of cyber protection particularly challenging. Recently, smart grid researchers and standards bodies have developed technological requirements and potential solutions for protecting cyber infrastructure. However, grid protection remains daunting to asset owners because of resources limitations.

Important questions arise when identifying priorities for design and protection: Which cyber components, if compromised, can lead to significant power delivery disruption? What grid topologies are inherently robust to classes of cyber attack? Is the additional information available through advanced cyber infrastructure worth the increased security risk?

The goal of this project is to develop a framework to assess the impact of cyber attacks on the electric smart grid. Our approaches borrow from mathematical principles from control and communication theory to identify new vulnerabilities stemming from the use of cyber infrastructure and the relative physical impact of cyber attacks. One outcome is a vulnerability analysis tool that can be employed by smart grid stakeholders to identify critical cyber infrastructure that must be prioritized for system hardening.

Kwong, Raymond — 3D Conformal Thermal Therapy of Soft Tissues for the Treatment of Localized Cancer using MRI-Controlled Ultrasound Therapy

www.control.utoronto.ca/~kwong

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

Dependability and Security in Control and Multimedia Systems

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

Lehn, Peter — Power Electronics to Enable More Sustainable Electrical Energy Networks

www.ele.utoronto.ca/~lehn

Professor Lehn’s research lies in the area of medium- and high-power applications of power electronics to form more reliable, cost-effective and sustainable electrical energy systems. Of specific interest is the development of converter systems and network architectures for low cost, low-loss integration of wind, solar and energy storage resources, including plug-in hybrid/electric vehicles. Improving robustness and

power quality of the electrical grid via intelligent control of power-electronically interfaced sources and loads is a major focus. Research into high-power applications revolves around exploitation of established and emerging HVDC and FACTS technologies to improve utilization and stability of power transmission systems.

Leon-Garcia, Alberto

Application Platforms and Smart Infrastructure

www.met.utoronto.ca/alg/alg.htm

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

Autonomic Service Architecture

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

Connected Vehicles and Smart Transportation

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

Design of Converged Communications and Computing Infrastructure

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

Green Networking

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

Green Telco Cloud

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

NSERC Strategic Network on Smart Applications on Virtual Infrastructures

www.savinetwork.ca

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

Optical Networks for Ultrascale Datacentres

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

Levi, Ofer

Optical Biosensors and Biomedical Imaging Systems

www.biophotonics.utoronto.ca

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

Li, Baochun

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

iqua.ece.toronto.edu

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

GestureFlow: Streaming Multi-touch Gestures

To support collaboration among multiple users in real time, we propose that gestures are streamed in a broadcast fashion from one user to all participating users. Streaming gestures themselves, rather than application-specific data, makes it possible to optimize the design and implementation of a gesture broadcast protocol that can be reused by any gesture-intensive application that needs to support multi-party collaboration. We believe this is a more elegant and reusable solution that serves the needs of an entire category of gesture-intensive applications. Once received, a gesture stream can be rendered in real time by a live instance of the same application on a receiver. To take such broadcast of gestures a step further, we believe that multiple gesture broadcast sessions need to be supported concurrently, so that any participating user can be the source of a gesture stream.

In this work, we have designed GestureFlow, a new gesture broadcast protocol specifically designed for multiple concurrent broadcast sessions of user gestures. We point out that gesture streams typically incur low yet bursty bit rates, unlike traditional media streams. They do pose unique challenges, as gesture streams need to be received with the lowest possible delay and packet losses are not tolerable. In our design of GestureFlow, we use network coding and present a detailed design that takes advantage of inter-session network coding to support low latencies across multiple broadcast sessions.

Optimizing Datacentre Operations with Practical Complexity

The unprecedented growth of mega-datacentres, in which hundreds of thousands of machines are assembled to process a massive amount of data for Internet-scale services, has been driving the evolution of computing. Designing algorithms to optimize datacentre operations is thus imperative. At the same time, the scale of the infrastructure calls for novel approaches to reduce the complexity of the solutions in order to make them practical.

In this project, our research objective is to resolve the tussle between optimality and practicality in designing algorithms for datacentres. First, for a single datacentre we have designed Anchor, a resource management system that effectively allocates server resources to virtual machines. Instead of being optimal, Anchor is designed to be flexible and practical and uses a unified mechanism to support diverse allocation policies expressed by operators and tenants. It abstracts performance goals as preferences and uses a novel stable matching algorithm to solve the matching problem efficiently. We have also studied the problem of workload management for multiple centres distributed over a wide geographical area, where it is possible to go for both optimality and practicality. I propose to exploit the geographical diversity to reflect the electricity and bandwidth price difference at different locations and ISPs and develop a novel distributed algorithm to solve the large-scale optimization problem with faster convergence than that of traditional methods.

Liang, Ben

Broadband Multimedia Communication in the Mobile Environment

www.comm.utoronto.ca/~liang

Multimedia content is the single most influential factor driving the need for increased mobile network capacity and device capability. The proliferation of cloud-based content distribution services and video social networking applications will severely stress the existing mobile systems. The inherent heterogeneity of both wireless access technologies and mobile devices enables rich and ubiquitous multimedia services, but it also significantly complicates system design. We are interested in investigations into system optimization, resource management and

algorithm design to create innovative technologies for multimedia communication in the mobile environment. Promoting a synergistic approach, we work at the interface between mobile access, broadband communication and distributed-system technologies.

Fair Resource Scheduling in Large-Scale Networked Systems

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

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, networking algorithms and communication protocols for efficient allocation of spectrum, hardware and power in high-throughput wireless networking environments. Our investigations include stochastic optimization, distributed computing and co-operative communication.

Lie, David

www.eecg.toronto.edu/~lie

Computer Systems Security

Cloud computing offers a new exciting form of service to users in need of computer infrastructure. It provides users a pay-as-you go model and allows them to outsource costs such as management, power and cooling, procurement and provisioning. Unfortunately, cloud computing poses serious security concerns. Users want to ensure the security of their data and code while executing in the cloud. At the same time cloud providers want to protect their infrastructure from being abused. We are working on a new paradigm that will offer users unprecedented flexibility to control how and where their data is stored in the cloud, while at the same time giving them the benefits of mobility, durability, availability and elasticity that cloud computing has to offer.

We are exploring security and usability issues in modern mobile phones. Today's smartphones have more computing power than yesterday's workstations. However, they fit in your pocket, travel with you wherever you go and are involved in a multitude of daily activities. In addition, capabilities on the phones allow them to track your location, record conversations and maintain a list of people you interact with. While these are useful, mobile phones also pose a grave threat to privacy and security, larger than that of any other computing device. In this research, we explore the development of mobile phones to make them more secure improve their power utilization and overall usability.

Liebeherr, Jorg

www.comm.utoronto.ca/~jorg

Network Architectures and Services for a Mobile World

With their ability to create large scale self-organizing networks, on-the-fly peer-to-peer overlay networks have proven to be a disruptive technology, which has enabled new application services in support of content distribution, streaming and social networking. We believe that the role of self-organizing overlay networks can be much greater and that the full potential of this technology remains largely unexplored.

In our research, we explore the potential and fundamental limits of a network architecture for mobile users that is entirely based on the principles of self-organizing overlay networks. The architecture is characterized by the coexistence of virtually unlimited numbers of mostly mobile users in peer networks that can quickly grow to arbitrarily large sizes and adapt to changes in the number of peers and substrate networks.

Liscidini, Antonio

www.eecg.toronto.edu/~liscidin

Smart Power Optimization for Wireless Transceivers

Mobile Internet access has become very popular with the introduction of 3G and 4G networks that offer high-speed wireless connections. Until now, this mobile revolution has been driven by the possibility of having low-cost mobile terminals with Internet access, enabling ICT applications in education, health, government, banking, environment monitoring and business.

Although several "smarter" phones with multi-standard capabilities have been introduced, the path towards a universal mobile radio is far from smooth. Nowadays smartphones are still extremely expensive compared to simple phones and have a battery life limited to a couple of days. The main reasons for these limitations are the use of dedicated transceivers for each standard supported and the ever-increasing demand for better performance and thus faster communication. These two factors nullify all attempts to reduce power dissipation and the overall bill of material.

The main idea of this project is to apply the concept of reconfigurability, enabling mobile terminals to dynamically and autonomously adapt to changing environmental conditions and reducing their energy consumption. There are many examles in nature of dynamic fitting of performance to changing boundary conditions since, it represents the best way to achieve maximum efficiency in highly complex systems. Almost all ecosystems are based on this principle, which allows them to evolve while minimizing energy dissipation.

Ultra Low Power Transceivers for Wireless Sensor Networks

Wireless communication represents one of the most important revolutions of the last century. Although initially based only on star-mesh networks (e.g. cellular), at the end of the 1990s some wireless systems started to also adopt peer-two-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 simultaneously realizing 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 100uW to operate from harvested energy, sustaining an autonomous short-range communication to enable an ultra low power wireless sensor network. In this case low power consumption and low costs will be achieved by combining the functionality of several building blocks. Indeed, recycling bias current and devices is the prime strategy to minimize, area, power consumption and complexity of the transceiver. Furthermore, minimizing the overhead associated with each start-up-wake-up cycle will maximize the efficiency of the node.

Lo, Hoi-Kwong

www.comm.utoronto.ca/~hklo/QRNG/Quantoss.html

High-Speed Quantum Random Number Generator

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

We plan to develop the prototype further to make it compact and low cost. We also plan to develop the software for the classical post-processing.

Measurement-Device-Independent Quantum Key Distribution

Quantum cryptographic systems are, in principle, unconditionally secure. In practice, quantum hacking has emerged as a key challenge to their security. To foil quantum hacking, we have recently proposed an entirely new approach — measurement-device-independent quantum key distribution (MDI-QKD) — that can "short-circuit" all detector security loopholes. In other words, the system will be automatically immune to all detector side channel attacks. This is remarkable because it means that commercial QKD detection systems would no longer require any special security certifications and, in fact, they can even be manufactured by a malicious eavesdropper.

We will build a prototype MDI-QKD system and demonstrate its robustness against detector side channel attacks. We will also develop the theory of MDI-QKD and take into full account various imperfections in real-life devices. Our work will allow us to use our enemy in quantum cryptography.

Quantum Cryptography: From Theory to Practice

www.comm.utoronto.ca/~hklo/index.html

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

Maggiore, Manfredi

www.scg.utoronto.ca/~maggiore

Formation Control in Multi-vehicle Systems

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

Formation Control of Nanosatellites

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

Virtual Constraints: A New Paradigm for the Control of Motion

The traditional approach to making robots perform complex motions relies on a hierarchical decomposition of the control task: motion planning at the high level and reference tracking at the low level. This approach has proven to be inadequate in complex motion control problems such as locomotion in multi-legged robots or flight in bird-like 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 in the literature to induce stable walking in biped robots and we believe it can be used to emulate the flight of birds and insects and the swimming of fish and, more generally, to induce complex behaviours in robots.

Mann, Steve

www.eyetap.org/research/medr.html

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

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

www.InteraXon.ca

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

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

www.wearcam.org/comparam.htm

Comparametric Equations and High Dynamic Range (HDR) Imaging

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

www.eyetap.org

EyeTap Electric Eyeglasses, Personal Safety Devices and Systems

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

Lifeglogging: Lifelong Videocapture

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

www.wearcam.org/funtain.htm

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 thehydraulophone; take a look at this video and then imagine the possibilities when we miniaturize it and use air instead of water: vimeo.com/14018088.

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/absemen.

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

Nano-plasmonic and Nano-photonic Devices

Performance of computers is expected to eventually reach its fundamental limits in terms of speed, bandwidth, power consumption and electromagnetic interference. The problem lies partly in the degrading performance of electrical interconnects. Unlike transistors, in which functionality increases with miniaturization, the functionality of electrical interconnects degrades substantially with miniaturization. One suggestion is to replace the electrical interconnects with optical interconnects, which do not suffer from signal latency, limited bandwidth or high power consumption compared to their electrical counterparts. However, there is a major problem with optical interconnects and waveguides. The optical mode size and hence the device size, are approximately proportional to the operational wavelength. In other words, while transistors with dimensions of approximately 50 nm are common today, the micron size of optical devices makes their integration with electronics difficult. Surface plasmon polariton (SPP) — surface waves at the interface between a metal and dielectric — may provide a

solution. These plasmonic waveguides, like optical interconnects, have small latency and large bandwidth but, unlike the optical interconnects, they can easily be miniaturized. However, plasmonic waveguides have their own challenges. Chief among these are (1) large propagation losses, and (2) lack of various efficient and integrated plasmonic devices such as polarizers, directional couplers and bends, to name a few. In order to overcome the losses associated with SPP while maintaining a small device size, our group was among the first to propose a hybrid plasmonic waveguide (HPWG). The HPWG can be viewed as an optimized structure exhibiting a compromise between loss and mode size. Moreover, fabrication of our HPWG is compatible with the existing silicon technology. Our HPWG can be used as a building block for the next generation plasmonic devices such as TM- and TE-pass polarizers, polarization independent couplers and other novel components.

Moshovos, Andreas

www.eecg.toronto.edu/~moshovos

Bandwith Efficient DRAM Controllers in Non-coherent Systems

Embedded and mobile hand-held devices have been proliferating, enabling applications that were not impossible or cumbersome with the big iron machines of the past. Each new generation of these devices offers more capabilities enabling new applications: While early mobile devices were capable of simple tasks and low bandwidth communication, today's devices offer many more capabilities such as multimedia, navigation, digital photography, etc. As their capabilities are increasing, novel applications such as health monitoring will be possible. For these possibilities to materialize, mobile and embedded systems need to become more powerful while maintaining reasonable up-time.

A mobile system today contains several compute engines that are all supported by an external memory device. As the computation needs increase, more data needs to be fed to these engines. The link between these engines and the memory is the memory controller. The memory controller can greatly affect how much data the memory can provide, and at what energy cost. This work will develop memory controller technologies that will boost data feeding capabilities while taking power into account. The goal is to develop the memory controller technology that will be used in future-generation mobile devices in support of more demanding applications, while allowing the device to stay on for longer periods of time.

Exploiting Multi-Megabyte On-Chip Memory Hierarchies

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

The key questions addressed by this research are (1) How do we manage these multigigabyte caches; are the techniques currently used still adequate, or is there room or need to rethink these decisions? and (2) Can we exploit this tremendous wealth of on-chip storage to further optimize performance beyond what is possible by simply caching instructions and data? Accordingly, the proposed research comprises two thrusts: The first considers the use of coarse-grain tracking to achieve performance that is otherwise not possible with conventional cache management techniques. The second exploits the on-chip caches to store program metadata in addition to instructions and data. Program metadata is information collected at runtime about program behaviour that can be used to anticipate and optimize for future program demands.

FPGA-Friendly Processor Architectures for Irregular Applications

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

Power-Aware Cache-Based Structure Design

Computing devices comprise processing elements that process digital information and memory elements for storing digital information. Because of technological constraints, memory tends to be significantly slower than the processing elements it supports. Accordingly, virtually all modern computing devices employ caches, which are additional small and fast temporary memories that serve to accelerate most references to the otherwise slow memory elements. In recent years, power dissipation has emerged as an additional critical design constraint in computing device design; power limits performance for all devices and up-time for portable ones. Low power dissipation and performance are at odds: high performance typically comes at the price of high power dissipation. Caches account for a significant portion of total power dissipation (e.g., 25% to 45% in a modern processor) because of trends in semiconductor technology, their

power dissipation is expected to increase in relative terms. Accordingly, there is a need for techniques to reduce their power dissipation while maintaining performance and usability. In addition, virtualization is emerging as a key technology for future server systems. Caches will play an important role in virtualization, as they can accelerate access to memory from devices without intervention from the processor (this is necessary for adequate performance). Accordingly, there is a need to understand and develop caching mechanisms to support this aspect of virtualization. In addition, this project investigates the caching and communication architecture for fused CPU and graphics processor systems.

Smartphone and Mobile Platform Architiecture

Cellphones are changing the way we interact with each other, access information and do business. Just a few years ago, cellphones offered just voice calling and short text messaging. Today cellphone capabilities parallel those of recent, high-end desktop systems, while they include several additional capabilities to communicate and interact with the physical world, such as embedded cameras, and touch, position and acceleration sensors. New applications are continuously emerging including image-based searching, speech recognition and translation.

All of this is possible because smartphones incorporate significant power. However, providing this power is a very challenging task , because smartphones must operate using limited energy sources while maintaining a reasonable manufacturing cost and a relatively small physical size. A key mechanism for improving smartphone capabilities is computer architecture that studies how to build smartphones given the available manufacturing technologies, while taking into consideration the applications that these phones will use. Computer architecture faces continuous challenges for two reasons: (1) the properties of the underlying technology change significantly over time and (2) so do the applications. This work seeks to understand smartphone applications and to identify opportunities to improve smartphone architectures, leading to next-generation smartphones. The primary target of the work is applications that acquire, manipulate and use images and video in smartphones. Expected benefits include increased compute capability and functionality, novel imaging applications, improved energy efficiency and reduced cost for smartphone platforms.

Nachman, Adrian

www.eecg.utoronto.ca/~sorinv/mm_wave_lab.html

Millimetre-Wave Imaging System

This is a joint project with Professor Sorin Voinigescu's group. It seeks to integrate their breakthrough design of novel silicon systems on chip (capable of transmitting and receiving very high frequency electromagnetic waves) with novel inverse scattering and compressed sensing algorithms to produce a millimetre-wave imaging system.

MRI-based Impedance Imaging

www.currentdensityimaging.org

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.

Najm, Farid N.

www.eecg.utoronto.ca/~najm

Power Grid Verification

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

Ng, Wai Tung

www.vrg.utoronto.ca/~ngwt

Smart Power Integration and Semiconductor Devices

Our research group is focusing on the integration of power devices, smart power integrated circuits and power management systems. Our group has worked extensively in the development of CMOS-compatible HV fabrication processes for automotive and consumer applications in the 40–100V range. We also have ongoing collaborative projects with our industrial research partners to develop discrete and integrated power MOSFET and silicon- and SiGe-based

BiCMOS fabrication processes for smart power ICs and wireless applications respectively.

In recent years, we have focused on the design and implementation of VLSI power management circuits. This includes the demonstration of integrated soft-switching topology with predictive dead-time control and a practical DVS (Dynamic Voltage Scaling) system in 2004. We are also working on the integrated DC-DC converters with digital control. One of our aims is to prove that digital controller is a viable approach for portable power management. We were able to demonstrate an integrated DC-DC converter with dynamically adjustable power transistor size for power conversion efficiency optimization at ISPSD'06 for the first time. In 2010, we introduced a superjunction power FINFET at IEDM 2010. This is exciting work toward enabling the next-generation FINFET CMOS fabrication technology to be compatible with the implementation of smart power ICs.

Power management is currently an area of intense activity. Our group is focused on the integration of the controller with power output stages. We currently have activities in all-digital on-chip temperature sensors for thermal management applications, gate driver circuits with dynamically adjustable driving strength to suppress EMI and improve efficiency, and dead-time control circuits, as well as integrated class-D power amplifiers.

Pavel, Lacra — Decentralized Optimization and Game Theory

www.control.utoronto.ca/~pavel/index.html

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

Dynamic Optical Network Control and Self-Optimization

www.control.utoronto.ca/~pavel/LP_research2.htm

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

Phang, Khoman — Friends of Design

www.eecg.utoronto.ca/~kphang

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. Inquiriesare welcome from students, faculty and staff, as well as undergraduate students, outside faculty, academe and industry and alumni.

Plataniotis, Konstantinos N. (Kostas) — Affective Signal Processing: Unravelling The Mystery of Emotions

www.dsp.utoronto.ca

Emotion plays an important role in our daily activities and greatly influences many areas, such as learning, decision making and interaction with others. Our decisions and courses of action are adapted according to the emotional cues we receive while interacting with others. This allows the exchange of information to be much smoother and more effective. Integrating the emotional states of a user into a human-mobile interface will provide a user-centric experience that enables the interaction to be more intuitive, flexible and efficient. We are proposing an affective signal processing system that enables real time analysis, tagging and inference of cognitive-affective mental states from facial video and EEG recordings. This framework combines vision-based processing of the face (e.g., a frown or smile) with EEG predictions of mental-states (e.g., interest or confusion) to interpret the meaning underlying EEG and facial signals over time.

Privacy-Enhancing Face Recognition

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

Privacy-Protected Video Surveillance

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

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

We seek to support for development in the areas of (1) signal-image processing for 3D imaging applications in the frequency regimes of visible and infrared stereoscopic camera systems and biometrics sensors; (2) image and data fusion for these multisensor systems; (3) Implementing of the signal-image processing developments for biometrics sensors (i.e. face tracking features using stereoscopic cameras, vital signs from EEG and ECG sensors) and imaging aid systems for helicopter landing operations in visually degraded environments.

Poon, Joyce — Integrated Photonics for Communications and Computing

www.photon.utoronto.ca

We invent, design, fabricate and measure integrated photonic devices and circuits, such as transmitters, switches and receivers, for communications and computing. Our unique strength is the breadth of technologies we access. We partner with a rich variety of collaborators in industry, academia and institutes around the world to use the most sophisticated electronic-photonic integration platforms. Our photonic devices and circuits are implemented in the following material systems and platforms: silicon-on-insulator (SOI); indium phosphide on SOI; silicon nitride on SOI; indium phosphide; correlated electron materials (vanadium dioxide).

Our goal is to demonstrate integrated photonic-electronic devices and circuits that are ultra-low-power, high-speed and compact for communications at >1 terabit per second using as few wavelengths as possible.

Integrated Quantum Photonic Circuits

We are using silicon-on-insulator (SOI) platforms available through foundry services to design and implement photonic integrated circuits for quantum physics experiments and quantum information applications. The goal of the research is to integrate the typically large, tabletop sized quantum experiments into single chips.

We are presently designing and testing entangled photon pair sources based on spontaneous four wave mixing, quantum interference experiments, high-performance optical filters and polarization controllers and modulators for quantum cryptography.

Prodic, Aleksandar — Power Management and Integrated Switch-Mode Power Supplies

www.ele.utoronto.ca/~prodic

Low-to-medium switch-mode power supplies (SMPSs), used in cellphones, computer systems, communications, vehicles, medical devices and other applications consuming 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 SMPSss and is currently focusing on fully utilizing the digital control advantages as well as on the development of novel converter topologies.

Qian, Li — Fibre-Optic Sensing

www.ecf.utoronto.ca/~qianli

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; multi-point optic sensing for stress and/or temperature sensing (used in civil structures); multi-point chemical gas sensing for environmental monitoring as well as industrial monitoring in hazardous environments; and liquid level sensing in cryogenic environments (required in space applications). Virtual-reference interferometry (VRI), a technology developed by one of our graduate students, has been commercialized in a successful start-up company.

Nonlinear Optical Devices — Ultrafast Switching and Frequency Conversion

Many photonic devices rely on nonlinear optical properties of materials. For example, ultrafast switching devices that operate in the 100 GHz range and beyond, utilize the ultrafast nonlinearity of optical materials. They can be widely used in high-speed data communication and signal processing. We are developing a sophisticated model for nonlinear optical materials that possess ultrafast and resonant

optical nonlinearity and 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 Communication

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

Rose, Jonathan — Architecture of Field-Programmable Gate Arrays

www.eecg.toronto.edu/~jayar/research/architecture.html

Field-Programmable Gate Arrays (FPGAs) are pre-fabricated chips that can be programmed to perform any digital hardware function. They reduce the time it takes to manufacture an integrated circuit from months to seconds, and the cost of a prototype from millions of dollars to under$1,000. They play an essential role in the wireless, automotive, consumer and industrial markets, with total FPGA annual sales approaching $5-billion. They enable essentially all hardware development including the emulation of high-volume processors and ASICs and represent the key technology for medium-volume systems.

Our research explores FPGA architectures, focusing on heterogeneous architectures that mix the efficiency of full-fabrication silicon with the programmability of an FPGA. This effort will require new CAD algorithms and architectural description capabilities in our world-renowned FPGA architecture exploration software.

Computer-Aided Design for FPGAs

www.eecg.toronto.edu/~jayar/research/CAD.html

FPGAs present new problems in Computer-Aided Design that sometimes differ from those in other implementation media such as Mask-Programmed Gate Arrays, Standard Cells and full-custom design. We are currently engaged in a large-scale collaborative effort to enhance our world-leading FPGA architectural exploration softwareto become a complete flow from the Verilog Hardware Description Language through to placement and routing, and timing analysis. It is called the “Verilog to Routing” (VTR) project. The new CAD software has the ability to describe far more complex FPGA architectures, including FPGA logic blocks with arbitrary hierarchy, modes of operation and interconnection structures.

Creative Applications for Mobile Devices

www.eecg.utoronto.ca/~jayar/ece1778/

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, many new possibilities will occur to each of us. As new hardware sensors and other capabilities are added to the phones, ever more applications will become possible. The purpose of this research/graduate course is to build a collaborative environment of creativity for new applications of mobile devices. Graduate students from all disciplines at the University of Toronto are invited to take the course for credit. This course has been taught three times, and will be available for the fourth time beginning in January 2014. It is primarily a project-based course in which the goal was to produce a working app by the end of 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.

Sargent, Edward — A Biochip for Gene-Based Disease Detection

www.light.utoronto.ca

We are building integrated circuits for the detection of a panel of biomarkers that indicate the early onset of specific types of disease. We configure nanostructured electrodes on a conventional integrated circuit; functionalize these electrodes with a nucleic acid probe having a sequence complementary to the target molecules of interest; and sensitively detect hybridization when it occurs. We are applying the chip to the early detection of cancer and to the sensitive and rapid detection of “superbugs” such as MRSA at the point of need.

Low-cost High-Efficiency Photovoltaics

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

Sarris, Costas — Advanced Radio Propagation Modelling for Next Generation Rail Signalling Systems

www.waves.utoronto.ca/prof/sarris

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 of position of access points will maintain wireless connectivity for the trains. While lack of coverage in a cellular communication network may result in dropped calls or slow data speeds, a similar effect in a CBTC system may have much more serious consequences for passenger safety. As a result, a detailed radio survey, whereby wireless propagation measurements are carried out over the entire railway network, precedes the installation of CBTC systems. Conducting a radio survey requires significant resources (time put in by qualified personnel and funds), while the line remains out of service. Often, the pressure to complete the survey results in overestimating the necessary number of access points. This redundancy results in higher installation and maintenance costs that reduce the competitiveness of CBTC solutions.

This project is focused on the development of a powerful software package that can significantly accelerate the radio survey, using advanced propagation modelling techniques to optimize the distribution of access points for CBTC systems. The project will develop a comprehensive modelling framework for radio-wave propagation in complex railway environments, validated through measurements.

Multi-user Wireless Power Transfer

The principle of inductive coupling offers a route to energy transfer for very small distances. On the other hand, the engineering of coupled magnetic resonances between capacitively loaded loops has led to the recent demonstration of energy transfer over a distance of about two metres. Subsequently, several research groups have focused their efforts on the possibility of wireless charging of handheld devices from a medium distance. This research faces two particularly important challenges: first, the rapid degradation of energy transfer efficiency with distance and, second, the safety limitations on exposing the general public to magnetic fields.

The twin challenges of efficiency and safety are intimately connected, as higher efficiency implies that lower magnetic field amplitude is needed. In this project, a new direction in wireless energy transfer will be investigated, based on near-field antenna beamforming and the emerging technology of electromagnetic meta-materials. Also, new algorithms will be developed for the creation of multiple beams that can enable the concurrent charging of multiple devices.

Realisting Assessment of Small Antenna MIMO System Performance

Multiple-input multiple-output (MIMO) systems using multiple receiving and transmitting antennas can be used to increase channel capacity relative to a single-input single-output (SISO) system. The performance of a MIMO system mainly depends on the channel matrix of the propagation environment, which in turn depends on the type of antenna elements being used, among other things. The capacity of a MIMO system has been studied using measurement campaigns; however, these results apply only to the given propagation environment. Modelling efforts based on deterministic methods such as ray-tracing and vector parabolic equations, which can require far fewer resources and less time than measurements, have also been used. In this project, a 3D ray-tracing tool is employed to study the effect of using closely spaced antennas with low mutual coupling on the correlation coefficient at the receiving antenna elements, the channel matrix and hence the channel capacity of a 2x2 MIMO system. This study is complemented by an experiment using a channel characterization test-bed.

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.

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 significance of modeling uncertainty is 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 modeling 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

Analysis and Control of Complex Interconnected Systems

www.scg.utoronto.ca/~scardovi/

It is well recognized that control has proven to be an essential ingredient in almost every engineering system, ranging from power and automotive systems to space missions, and that feedback is a key element in many natural phenomena, ranging from molecular pathways in living organisms to ecological systems. Recent years have witnessed an increasing interest in systems that are composed of (possibly many) interconnected units. As a whole, those systems often exhibit one or more features that cannot be predicted from the properties of the individual parts. These properties (called emergent behaviour) are not an attribute of any single entity: they are irreducible and are generated by their interconnection. Emergent behaviour can lead to surprising and useful phenomena such as memory, intelligence and self-organization in cells, but can also have disastrous consequences. Examples include the spread of infectious diseases, neuronal synchronization disorders in the brain, collective motion in bacteria, and locust swarms. It is therefore of great interest to understand the principles behind the emergence of such properties and investigate solutions to control 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, amongst others.

Sheikholeslami, Ali

Circuits for Spin Electronics

www.eecg.utoronto.ca/~ali

Spin Electronics (or spintronics) is a new and emerging field of science and engineering that exploits the spin of electrons, in addition to their charge, for the purpose of information storage, transport and manipulation. The ultimate aim of research in spintronics is the discovery and invention of new devices, such as spin transistors and their integration into semiconductor technology to create better functionality and performance at lower cost and complexity. The purpose of this research is to explore circuit techniques for spin-based devices that are suitable for nonvolatile memory applications, replacing conventional memory technologies such as DRAM, SRAM, Flash and EEPROM.

The basic structure of a spin-based memory cell is a magnetic tunnel junction (MTJ) that consists of two ferromagnetic layers separated by a thin layer of insulating material. One of the two ferromagnetic layers is a thick layer whose magnetization is fixed. The other one is a thin layer, also called free layer, whose magnetization can be switched so that its direction either parallel or anti-parallel to that of the fixed layer. This corresponds to storing a digital 1 or 0 in the cell. Reading the stored bit is achieved by examining the resistance of the cell. This resistance is low for the parallel state and high for the anti-parallel state.

There are several challenges in the operation of the MTJ device that must be addressed before the spin-based memory can compete favourably against Flash. Our goal in this research is to devise circuit techniques to circumvent the device shortcomings and ease the requirements for the underlying technology.

High-Speed Wireline Signalling

This research targets circuit design for high-speed chip-to-chip signalling, backplane signalling and optical communication. This includes circuit designs for the transmitter (such as the design of MUX, equalization and driver) and the receiver (such as adaptive equalization and clock and data recovery).

At speeds beyond 1 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 cross-talk and timing jitter so as to reduce the bit error rate (BER) of the communication link while using less than a few mW per Gb/s operation.

In the past few years, we have been able to contribute to this research through the design of ADC-based receivers that allow for extensive signal equalization in the digital domain. Our latest work in this area was presented at ISSCC 2013 where we presented a 10 Gb/s blind baud-rate receiver using an ADC front end. In the past few years, we have also contributed to the design of non-data-aided equalization techniques and to burst-mode CDRs.

Moving forward, there are still many challenges in the area of high-speed signalling as demand for signalling speeds of 28 Gb/s and beyond grows. These data rates impose stringent requirements on both the channel equalization and the power budget for these links. We strive to address these challenges in the near future.

Smith, Peter W. E.

Ultrafast Photonics

www.ecf.utoronto.ca/~upl/

My teaching and research interests have centred on the study of ultrafast photonic and nonlinear optical effects in materials, and the development and characterization of ultrafast optical devices for all-optical switching and signal-processing.

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

Sousa, Elvino

Autonomous Infrastructure Wireless Networks

www.comm.toronto.edu/~sousa/sousa.html

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

Steffan, Gregory

Making Programming Multicores Easier

www.eecg.toronto.edu/~steffan

We are developing compiler and analysis tools for making modern multicore processors easier to program. In particular, we are investigating support for optimistic parallelism and transactional memory, with a current focus on parallelizing CAD applications.

Overlay Architectures for FPGAs

Field-Programmable Gate Arrays (FPGAs) are pre-fabricated integrated circuit "chips" that can be programmed to become any digital circuit. They are now widely used in all types of communications, computer and industrial hardware because they are often far more economical than fully fabricated chips. Our goal is to allow software programmers to more easily program FPGAs for high-performance applications, by developing new "overlay" architectures for FPGAs: structures programmed onto FPGAs that are themselves programmable.

Stumm, Michael

System Software Performance Optimizations

www.eecg.toronto.edu/~stumm

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

Tate, Joseph (Zeb)

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

www.ele.utoronto.ca/~zeb

The rapid increase in variable generation technologies such as wind and solar power throughout many nations' power grids has the potential to significantly reduce reliability. To ensure this does not happen, power companies must run a multitude of simulations that identify potential problems before they occur in the real system. The ability of these simulations to accurately inform decisions is only as good as

the models being used and the lack of confidence in dynamic models is one of the main problems associated with these new-generation technologies. For several reasons — such as model order reduction to make simulations tractable, the reluctance of generator manufacturers to release detailed models or parameter sets and the relatively high installation of wind generation — the accuracy of wind generator models is becoming increasingly important to planning engineers.
This project is looking at ways to use ambient wind power generator measurements (i.e., without introducing artificial stimuli to the system) to determine wind parameter models for use in simulation-based studies. Thus far, we have been exploring the performance of various nonlinear parameter estimation schemes, in particular the Extended and unscented Kalman Filters (EKF and UKF respectively), to select the most appropriate algorithm for this application. We have seen via simulation with high-bandwidth sampling that, for a relatively simple DFIG model, the UKF is generally superior to the EKF in both robustness and speed of convergence, confirming the UKF's superior performance when applied in other disciplines. Two extensions are currently being investigated: first, the impact of reduced bandwidth sampling (e.g., using 30–60 samples per second) and second, testing of the UKF estimator with acmeasurements.

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 generic lossless compressors.

Power System Simulation Using Programmable Graphics Processing Units

Modern power grids comprise millions of individual nodes, which are interconnected in a variety of ways (e.g., transmission lines, transformers, switches and power electronic converters.). At each node, there may be anywhere from zero to hundreds of sources and sinks of electric power, each with its own complex model. As an example, consider a modern household, which may include slowly varying electrical loads such as washing machines in parallel with rapidly changing loads such as CPUs. Because of the complex physical structure of the power grid and the components to which it is connected, simulation of power grid behaviour can be a challenging task. Presently, evaluation of system events is done on an ad hoc basis, in which planning engineers have to guess at both the likelihood of events occurring and the impact on the rest of the network. Reducing the solution time of power system simulations allows planners and operators to consider a wider variety of events and/or more detailed modelling of power system components and this has led to a renewed interest in algorithm design and implementation for power system simulations. This project focuses on the development of power system simulation software that targets a particular parallel computer architecture — programmable graphics processing units (GPUs). These processors pose several unique challenges because of the hundreds of cores on each chip and the unique game-driven memory access patterns; as a result, prior work in both serial and parallel power system simulation cannot be immediately adapted to these architectures. Thus far, we have developed a GPU-based linear solver designed to deal with the large, sparse, ill-conditioned systems that are typically solved during power system simulations. Preliminary results have shown that the GPU is a viable architecture for power system simulations and that utilizing polynomial preconditioners combined with Krylov-subspace-based solvers can offer significant gains over serial code. Future work will be focused on profiling and optimization of the already developed algorithms and the development of a complete power system simulation package based on GPU-accelerated computing.

Wind Impact Metrics for Short-Term Power Grid Operations

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

Taylor, Josh

Learning to Manage Electrical Loads

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

Gaming in Modern Electricity Markets

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

Trescases, Olivier

www.ele.utoronto.ca/~ot

Battery Management for Electric Vehicles

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

Making better use of the energy capacity by increasing the system efficiency is the key to reducing the overall size and cost of the EV battery. Regenerative braking (Regen) is often used in electric vehicles to capture kinetic energy that is otherwise wasted in the brake pads when the vehicle comes to a stop. Instead of simply applying the mechanical brakes during deceleration, an EV equipped with a Regen system uses the motor as a generator in order to transform mechanical energy into stored charge in the battery. Even the latest lithium-based batteries have a relatively poor ability to quickly absorb energy without affecting long-term performance. The maximum output power of modern Li-ion batteries is typically at least three times higher than the maximum input power. Repeatedly using Li-ion batteries to both absorb this large negative burst of power during Regen and provide 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 ten 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 SMPSs acts as the interface between the energy source, such as a battery and the load ICs. A typical SMPSs uses a combination of high-speed, low-resistance semiconductor switches, energy storage components, sensors and control circuits to regulate one or more output voltages in the presence of disturbances. State-of-the-art SMPSs have a power conversion efficiency above 90%. The resulting low heat dissipation allows multiple SMPSs to be integrated with their load circuits into a single IC. The clear trend in SMPSs 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 then ideal. High brightness (HB) LED modules are rapidly emerging as a promising candidate to replace CFLs in numerous lighting applications, since HB-LEDs are mercury free, scalable and can be easily dimmed. The main goal of the project is to develop smart digital switched-mode power supplies (SMPSs) to optimize the lighting efficiency and chromatic properties of HB-LED systems using closed loop thermal and optical feedback. The targeted controller can efficiently regulate the electrical-to-optical energy conversion process under a wide range of operating conditions.

Power Electronics for Photovoltaic Applications

Solar energy has long been recognized as one of the most abundant forms of clean renewable energy. Countless research efforts around the globe are contributing to the steady decline in the cost of photovoltaic (PV) power, with the promise of reaching grid parity in the near future. This is a complex target, as the prices of conventional energy sources are constantly in flux and heavily dependent on government subsidies. The penetration level of solar power is rapidly increasing in most developed countries because of government incentives and multidisciplinary technological advances. The exponential growth of PV technology presents tremendous opportunities for all companies in the semiconductor supply chain, ranging from discrete power devices to mixed-signal control ICs. Performing maximum power point tracking (MPPT) on a PV array is used to continuously optimize the total harvested power under time-varying temperature and illumination fluctuations. It has been demonstrated that performing distributed MPPT (DMPPT) on a per-panel or even per-cell basis, instead of using a single MPPT controller across the entire PV array, can substantially improve the total system efficiency under partial shading conditions. The main goals of this project are to quantify the benefits of DMPPT for different levels of granularity and to develop new high-efficiency power electronic converter topologies and control schemes for both mono-crystalline 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 circuitry and to the thermal simulation of 3D integrated circuits with liquid cooling.

Signal Integrity and Electromagnetic Compatibility Engineering

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

Truong, Kevin

apel.ibbme.utoronto.ca

Computational Tools for Protein Sequences, Structures and Networks

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

Live Cell Imaging and Control of Caspase Kinetics Using Engineered Proteins

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

Valaee, Shahrokh

www.comm.utoronto.ca/~valaee

Localization of Wireless Terminals in Indoor Environments

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

Wireless Communications in Vehicular Environment

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

Veneris Andreas

CAD for VLSI Verification, Debugging, Testing and Synthesis

www.eecg.utoronto.ca/~veneris/AndreasVeneris.htm

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

Venetsanopoulos, Anastasios

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

www.dsp.toronto.edu/

In digital signal/image processing our work is contributing to both fundamental and applied areas. New digital filters, such as linear, non-linear. multidimensional and multi-spectral, both fixed and adaptive, have been introduced.These filters have been applied to geophysics, imaging and financial data, and to biomedicine, including radiology, tomography, mammography and MRI.

In the area of telecommunications, we have dealt with the problems of sonar transmission through fading dispersive random channels and image compression, focussing on techniques called perceptually lossless techniques. We have also explored biometrics and multimedia.

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

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

Voinigescu, Sorin

56 GS/s 7-bit DAC and ADC

www.eecg.toronto.edu/~sorinv

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

Atomic-level and 2D Crystal Electronic Devices

This exploratory research focuses on the simulation, design and fabication of novel nanoscale 2-D crystal and metal nanowire transistors. The ultimate goal is the physical implementation of sub–10 nm transistors.

High Efficiency mm-Wave Transmitter Array

This research focuses on novel transmitter array architectures based on high-efficiency direct mm-wave digital modulators. The goal is to deliver 64 QAM modulated transmitters with over 65% power added efficiency and over 36, 30 and 26 dBm of output power at 45 GHz, 94 GHz and 138 GHz, respectively.

New circuit topologies, IQ correction techniques and ci and free space power combining and modulation schemes are being contemplated. Both SiGe BiCMOS and 45-nm SOI technologies are being investigated for the practical implementation.

Digitally Enhanced Analog Equalization Techniques for 50–110 Gb/s Wireline Applications

This research focuses on the development of new high-data rate wireline communication circuits and systems based on mm-wave DACs.

The goal is to take advantage of MOSFET layout segmentation in deeply scaled nanoscale CMOS technologies, in order to realize arbitrary waveforms for transmitter equalization directly at frequencies in the 50–110 GHz range. In receivers, oversampling with a 100+ GHz clock will be employed to recover the clock and data information and to recondition the signal before transmission. An 80-Gb/s equalizer capable of equalizing over 30 dB of cable loss has been demonstrated. A retimer capable of operating at clock frequencies up to 108 GHz has been fabricated and tested. Both circuits represent world-record performance for silicon technologies.

Large Swing DACs for 200 Gb/s Optical Transmitters With QAM and OFDM Modulation

As internet traffic continues to increase exponentially because of the explosion of mobile multimedia devices, there is renewed demand for electronic circuits and optoelectronic systems that can operate at serial data rates in excess of 50 Gb/s. Current OOK- and QPSK-modulated optical transmitters are limited to 50-Gb/s bit streams because of the limitations of optical filters and optical fibres. This is one of the reasons for the widespread adoption of QPSK modulation in the 110-Gb/s Ethernet standard. Rather than transmitting and receiving a single 110-Gb/s NRZ bit stream, four NRZ bit streams, each at 28 Gb/s, are employed to modulate two orthogonally polarized optical carriers in-phase and in-quadrature. The next generation of optical fibre standards will inevitably imply the deployment of both higher-level modulation schemes such as 16-QAM and OFDM and of higher serial bit streams at and beyond 112 Gb/s.

Optical modulators typically require 3–5V electronic signal swing for proper operation. The large voltage swing and the very broadband operation from DC to over 55 GHz are beyond the best performance reported for digital-to-analog converters in nanoscale CMOS and SiGe BiCMOS technologies. One option would be to use a low-voltage swing DAC followed by a very large voltage-swing, large gain, broadband linear amplifier fabricated in III-V technology. This is an expensive multi-chip solution. In addition, because of the large resolution (7–10 bit) required, it is almost imperative that the DAC directly drive the optical modulator to avoid signal distortion.

This proposal seeks to research and develop large swing (>3V differential), multi-bit (6–10 bit) BiCMOS-55-nm and 28-nm SOI CMOS DAC topologies for 16-QAM, 64-QAM and OFDM optical transmitters.

One such SiGe BICMOS DAC, developed in our group and operating at 50 GS/s with 6 Vpp differential swing was presented at IEEE IMS in 2012 and won the best student paper award.

Silicon SoCs in the 100–500 GHz Range

This research focuses on developing characterization, modelling and design methodologies, as well as circuit topologies and system architectures for future integrated systems operating in the virtually uncharted 100–500 GHz band.

We envisage applications in industrial sensors, automotive radar, active and passive imaging, remote sensing, environmental monitoring, radio astronomy and 0.5Tb/s wireless communication.

On-die noise source-pull test set-ups have been developed for the extraction of the noise parameters of 400-GHz SiGe HBTs in the 110–170 GHz band.

The first integrated systems that have already been demonstrated in the lab are (1) 120-GHz and 150-GHz high resolution position sensors with above IC and in-package antennas fabricated in a SiGe BiCMOS process, and which work over 2–5m; (2) a 160-GHz radio transceiver array with on-die antennas operating as a 4Gb/s radio and as a Doppler radar; (3) a 165-GHz passive imaging receiver with sub 0.3 K temperature resolution and consuming 82 mW; and (4) 200–300GHz low-noise signal sources and amplifiers.

Wong, Willy

Sensory Neuroengineering

www.individual.utoronto.ca/willy

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

Wonham, Murray

Supervisory Control of Discrete-Event Systems

www.control.utoronto.ca/~wonham

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, 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

Electrical Neuromodulation for Bladder Dysfunction

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

The Edward S. Rogers Sr. Department of **Electrical & Computer Engineering's**

Distinguished Lectures Series, 2013-2014

Ray Liu
University of Maryland
Green Wireless Communications: A Time-Reversal Paradigm
October 10

Harel Shouval
U Texas Medical
Reverse Engineering the Cellular Basis of Learning and Memory
October 3

Kameshwar Poolla
UC Berkeley
Modeling, Control, and Optimization: Critical Technologies in Semiconductor Manufacturing
November 14

Peter Asbeck
UC San Diego
Pathways to Power with Micro- and Millimeter Waves
January 9

Jason Nieh
Columbia University
Virtual Smartphone Computing
December 5

David Perreault
MIT
In Search of Powerful Circuits: Developments in Very High Frequency Power Conversion
January 23

Ness Shroff
Ohio State University
Low-Complexity Scheduling Policies for Achieving Throughput and Delay Optimality in Multi-Channel (OFDM) Downlink Systems
February 27

All talks are held at 3 pm in room 1105 of the Sandford Fleming Building, 10 King's College Road, University of Toronto.

www.ece.utoronto.ca/about/distinguished-lectures

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
University of Toronto
10 King's College Road
Toronto, Ontario, Canada
M5S 3G4

www.ece.utoronto.ca

General Inquiries: eceoffice@utoronto.ca
Telephone: 416-978-1801

Editor: Marit Mitchell

Art Direction: Laurie LaBelle
www.icon-art.ca

Photography: Robert Waymen
and
Engineering Strategic Communications