Quest

A Catalogue of Research in Electrical and Computer Engineering at the University of Toronto **2011**

The Edward S. Rogers Sr. Department
of Electrical & Computer Engineering
UNIVERSITY OF TORONTO
www.ece.utoronto.ca

Inside:
- Articles on leading-edge research
- A complete listing of professors and their exciting research initiatives

TABLE OF CONTENTS

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

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

Telephone: 416-978-1801 General inquiries: susan.grant@utoronto.ca

www.ece.utoronto.ca

"Founded in 1909, The ECE Department has a long and proud history of education, research, and service."

As Department Chair, I am delighted to share with you our first Catalogue of Research for The Edward S. Rogers Sr. Department of Electrical & Computer Engineering at the University of Toronto.

Founded in 1909, our department has a long and proud history of education, research, and service. Over the past century, it has evolved to meet the changing needs of society and the dynamics of a large research university, with a mission that includes not only training but also research, innovation and knowledge creation. Today, this is the largest ECE department in the country, with an operating budget of about $21M, over 70 faculty members (nearly 100, including our Professors Emeritus), and over 1,300 undergraduate students.

Research is an integral part of our mission and activity. Our focus is to provide state-of-the-art training for our 450 graduate students in our MEng, MASc, and PhD programs. With an annual research budget of about $16M, 21 named research chairs, and approximately 36 PhDs awarded annually, the department literally buzzes with innovative research ideas and projects. This creative energy, in turn, benefits our undergraduate program enormously, enriching the course content and providing valuable research experiences.

As a result, ECE is engaged in the pursuit and dissemination of knowledge across a wide range of areas of engineering and applied science – from the fabrication of atomic level structures with special optical properties, through to the technologies that have revolutionized our world, such as microelectronics, computer systems, software, and networks.

In this publication, which we hope will be an annual event, you will find articles outlining some of the leading-edge research currently underway in our department, as well as a list of all our professors and their research projects.

I hope you find this to be a useful snap-shot of our research activities, and I welcome your feedback and comments – you may reach me at chair@ece.utoronto.ca.

Farid Najm

Professor and Chair
The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
University of Toronto

The intent of this publication is twofold. First, we wish to communicate our pride in the achievements of our vibrant community of researchers. Second, we want to announce to our partners – past, present and prospective – that we are open for business and committed to making an even greater impact on industry and society through our research.

In 2009, The *Times Higher Education World University Rankings* placed engineering at U of T eighth worldwide, and fifth in North America alongside MIT, Berkeley, Stanford, and Caltech. Internationally, our peers are Cambridge, Tokyo, and Imperial College. Most noteworthy, the key factors in these rankings include the impact of our work on other researchers, and our reputation among research leaders internationally.

Research for us is a major enterprise. Over 70 professors and 450 graduate students strive passionately to advance their fields within and beyond electrical and computer engineering. Each year we invest over $16M toward solving important problems in biomedical engineering, communications, computer engineering, electromagnetics, electronics, energy systems, photonics, and systems control.

"...we are open for business and committed to making an even greater impact– through our research on industry and society."

In addition to our valued corporate partnerships in Canada, and on a global scale, we are extremely proud of the over two-dozen companies incubated right here in engineering at U of T in the last decade alone. All of these companies, both large and small, are integral to our success.

Going forward, we will continue to strive. While we have enjoyed our many accomplishments, ECE at U of T is committed to being number one. We are determined to provide the most stimulating, industry-connected training environment for the graduate students we attract from around the world. Our promise is to deliver value to industry in the form of breakthrough innovation successfully transferred to our partners. We intend to transform Canada and the world through meaningful discoveries in fields as diverse as mobile computing, biology and medicine, and sustainability, that improve quality of life.

Please enjoy our Catalogue of Research, and let us know your impressions.
I invite you to reach me at: ted.sargent@utoronto.ca.

Ted Sargent

Professor, Canada Research Chair in Nanotechnology
Associate Chair–Research, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
University of Toronto

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

Professor Willy Wong

Associate Chair Graduate Studies, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering and Institute of Biomaterials and Biomedical Engineering, University of Toronto.

A Combinatorial

CROSS-DISCIPLINARY APPROACH BRIDGES GAPS TO NEW FRONTIERS

In the quest for cures and assistive technology for neurological diseases and impairments, Professor Willy Wong and his team of biomedical engineering graduate students are conducting breakthrough interdisciplinary research with far reaching implications for enormous impact on society.

One of the most profound outcomes of their work is the development of engineering tools to study speech that lead to their discovery of a non-invasive method for early detection of Parkinson's disease. Until now, no unambiguous test for the disease has existed.

"I have an ongoing collaboration with Professor Pascal van Lieshout at the Toronto Rehabilitation Institute on this project," said Wong. "When he suggested that our methods might be useful in the area of Parkinson's we felt confident, but it was unknown as to how it would work until we took the data, measured it and then applied our particular algorithms," explains Wong.

Although everyone has some degree of voice tremor, Wong discovered that the characteristics of tremor from Parkinson's have fundamental differences. "We very quickly realized we had outstripped any other method and were able to characterize the low frequency voice variations specific to Parkinson's. It was so obvious because it's a completely different signature than the run-of-the-mill voice tremor."

Wong tested a population size of 30, including age-matched 15 Parkinson's patients and 15 control patients. The result was an outstanding 95 percent prediction rate.

Applying the same type of thinking in other areas such as schizophrenia and physical disability, Wong and his group are on a continuous search for solutions to problems that affect quality of life for so many.

"Being a biomedical engineer means that I have a foot in engineering and a foot in life sciences," says Wong. "The challenge is to be able to master the medicine, and master the engineering. The upside is that the payoff is huge–more than worth it."

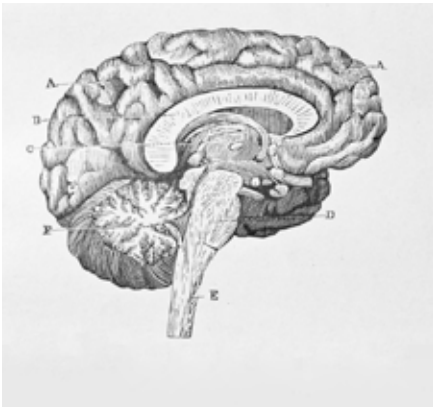

Explosion

"For me, all of this is quite promising," says Wong. "We are taking engineering which classically is physics, mathematics and chemistry, and using it to help with one of society's most pressing concerns–healthcare."

Wong's passion for biomedical engineering and the endless benefits it offers society is evident. "It's a wonderful field and we need to give these projects a chance to develop," says Wong. "Society as a whole needs to invest time, money, effort, and opportunity for growth to happen. It's a combinatorial explosion. If we can get more students into this area and encourage the cross-disciplinary research, the sky's the limit."

Biological signals and markers for the detection and diagnosis of neurological diseases.

Professor George V. Eleftheriades
Professor and Canada Research Chair/Velma M. Rogers Graham Chair in Engineering, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, University of Toronto.

Bending Beams

AND EXTENDING THE LAWS OF PHYSICS TO FOCUS ELECTROMAGNETIC WAVES WITH UNPARALLELED PRECISION

Big things really do come in small packages. This is evident in the huge impact that the University of Toronto's Professor George V. Eleftheriades is having on the scientific community using miniscule, quilted metallo-dielectric patterns called 'metamaterials'.

Since 1997, Eleftheriades has led graduate students in the areas of electromagnetic negative-refraction metamaterials, small antennas and components for broadband wireless communications, novel antenna beam-steering techniques, plasmonic and nanoscale optical components and electromagnetic design for high-speed digital circuits.

Credited with pioneering contributions to the science and technological applications of electromagnetic materials, Eleftheriades and his group produced the first experimental demonstration of "left-handed" or negative-refractive-index metamaterials (LHM) lenses.

"Negative refraction had previoiusly been observed by a group at the University of California San-Diego," says Eleftheriades. "I had read about this in IEEE Spectrum and got very excited because electromagnetics is my specialty. We originally started out with very primitive planar prototypes but based on a fresh concept for constructing metamaterials by making use of loaded transmission lines. Today we are working with three-dimensional structures."

Eleftheriades' metamaterials are composites of miniature, photolithographed-loaded wires (transmission lines) that form a tiny quilt-like pattern fashioned in three-dimensional structures in which each side could be several centimetres long. When a radio wave is beamed through such metamaterials, the wave refracts, or bends in the opposite direction towards the source.

This is contrary to conventional understanding that when a beam is shone through a slab of a material such as glass, the beam bends away from the source. It's the unique structure of the metamaterials that cause the electromagnetic waves to experience a negative index of refraction thus focusing on a precision point instead of spreading out in different directions.

"Metamaterials enable us to form flat lenses that allow us to see details that we couldn't see before," says Eleftheriades. "The amplitude of 'evanescent waves' that carry the object's smallest details increases inside the lens thus enhancing the resolution."

By using Eleftheriades' metamaterials and related next-generation 'meta-screens', it is now possible to focus waves on sub-wavelength spots with unprecedented precision and at increasingly longer distances from the lenses. Alternatively, related properties of metamaterials enable antennas and other radio-frequency components to be miniaturized, resulting in significantly smaller wireless devices while providing extensive bandwidth.

"In recent years we've invented new ways which are much simpler for making super resolution image focusing happen," said Eleftheriades. "My hope is to some day use this knowledge to significantly improve the imaging capability of MRI equipment or perhaps create a low-cost alternative that can be readily accessible by everyone who is in need. Imagine the benefit to being able to see inside the body in greater detail–to detect things like small cancerous tumours that we presently can't see with medical imaging."

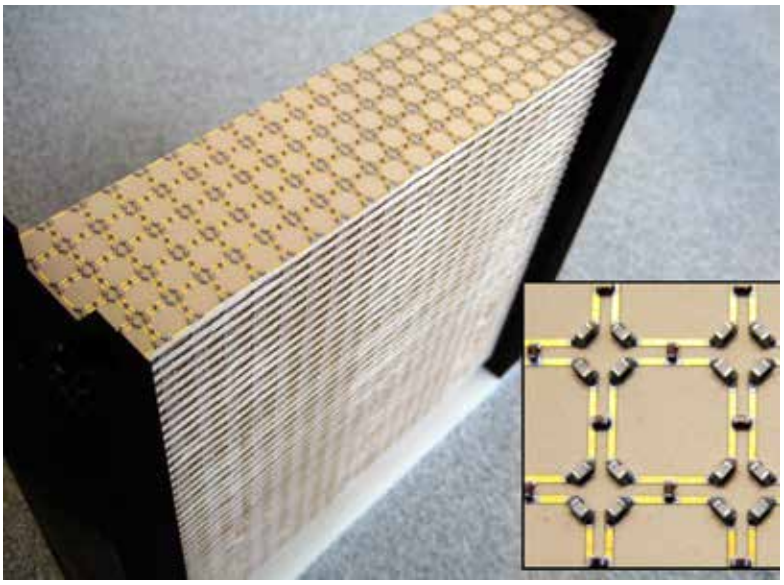

The first three-dimemensional negative-index superlens using transmission-line metamaterials (Lyer and Eleftheriades, 2008).

"At the University of Toronto I have been fortunate enough to have had such excellent graduate students over the years," remarks Eleftheriades. "Together we have made some notable contributions in the area of electromagnetics."

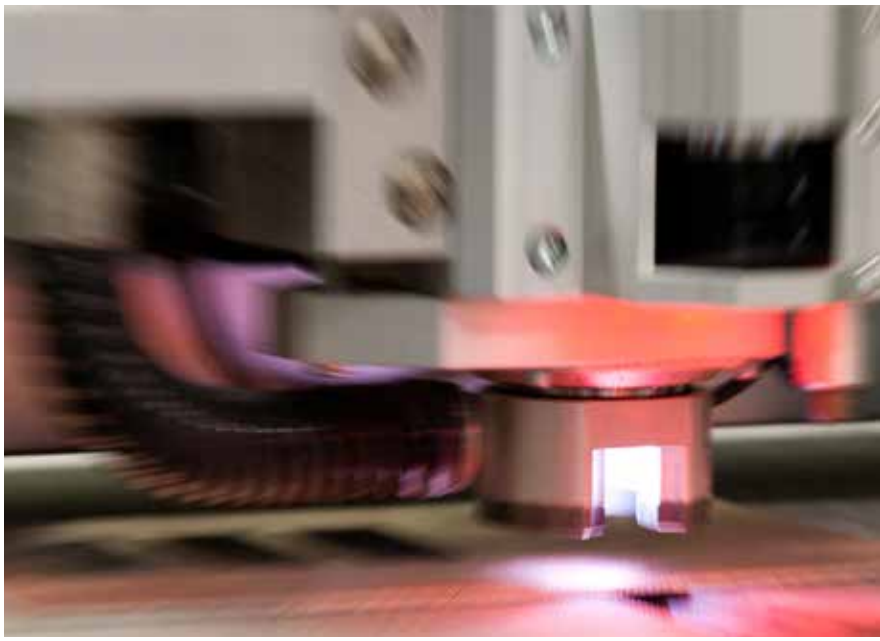

Microfabrication of metamaterials.

Professor Baochun Li

Professor, Bell Canada Chair in Computer Engineering at the University of Toronto, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, University of Toronto.

Revolution Evolution on The Net

HARNESSESING NETWORK DYNAMICS FOR A BETTER TOMORROW

Although there are no textbooks for the area of research that Professor Baochun Li and his graduate students at the University of Toronto are conducting, the topic is of enormous importance globally. In short, Li and his team research Internet dynamics.

"This is not a traditional research topic," says Li. "It's very tough to follow because it moves so quickly. People tend to take the Internet for granted–like it's a given. They don't realize how much research goes into this area, and how much still needs to be done." One of Li's most exciting contributions is bringing network coding to reality. Delivering quality real time multi-media streamed content, over wired and wireless networks.

Working in collaboration with start-up success story, UUSee, headquartered in Beijing China, Li serves as the company's technology guru. "I teach and research full-time, but I also consult with UUSee," says Li. "I do a kind of technology transfer, embedding the technology I develop in the lab here and over there inside the engine at UUSee. In return, I obtain live traces of network coding in operation for research purposes."

UUSee is one of the top three video content providers in China. Li's association with the company provides summer internship opportunities for his students.

"With popular video, every channel has a lot of users simultaneously, like a soccer match or a live concert," explains Li. "The problem we face is how to utilize the upload bandwidth of every user, to allow the servers to sustain a larger number of users as compared to YouTube for example, where everybody streams directly from the server."

Called peer-to-peer streaming, this poses a challenge because there's no certainty as to who's going to come online or what kind of bandwidth each peer can provide. So network coding applied to this scenario is a perfect match because it optimizes the flow of data in a network.

"To get the quality we want, we need more bandwidth on the servers," explains Li. However, "bandwidth is very costly, especially in the Chinese market. The goal is to save bandwidth and maintain the same quality to the end user. Network coding helps substantially."

Another challenge is that users are geographically diverse, with interests that are constantly shifting. This causes changes in demand over short periods of time. "Right now we just create more data centres to accommodate peak demand, but that's not sustainable," says Li. "We're using a lot of energy to operate idle servers waiting for demand to happen."

What Li would like to do is predict user behaviour one or two hours in advance, and then shuffle virtual resources around the world to meet demand. Predictions would be made by gathering information from sources such as social networks where early bird users could be indicators of potential flash crowds that occur when others follow and shift to that topic.

"The problem we have is that flash crowds tip the balance between demand and supply," says Li. "By virtualizing the application servers, we can treat them as a coherent whole and then shift them around to a different cloud or a different data centre within the same cloud. That's what we're working on now. It's very exciting and has huge implications."

Li and his team are focused on delivering quality content to mobile devices such as smart phones. "This kind of platform is something that is really the next generation way of consuming content," says Li.

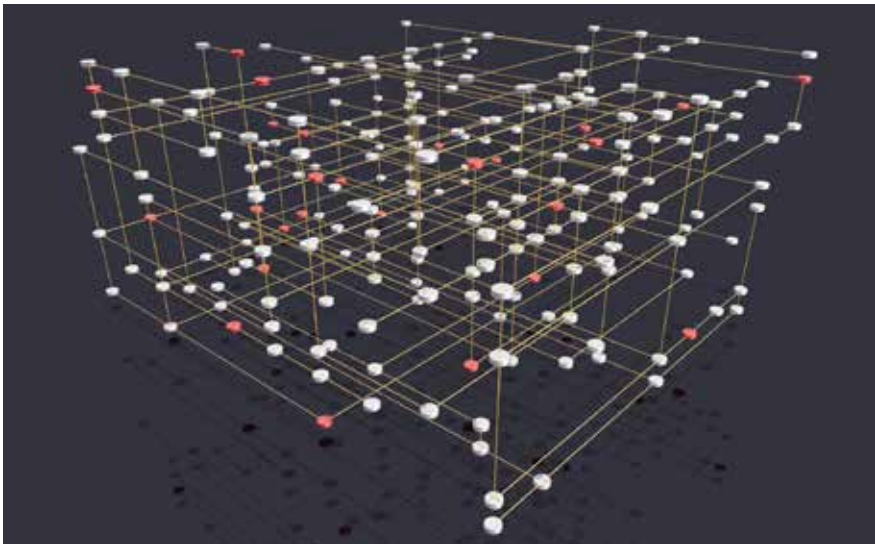

A multi-layered network.

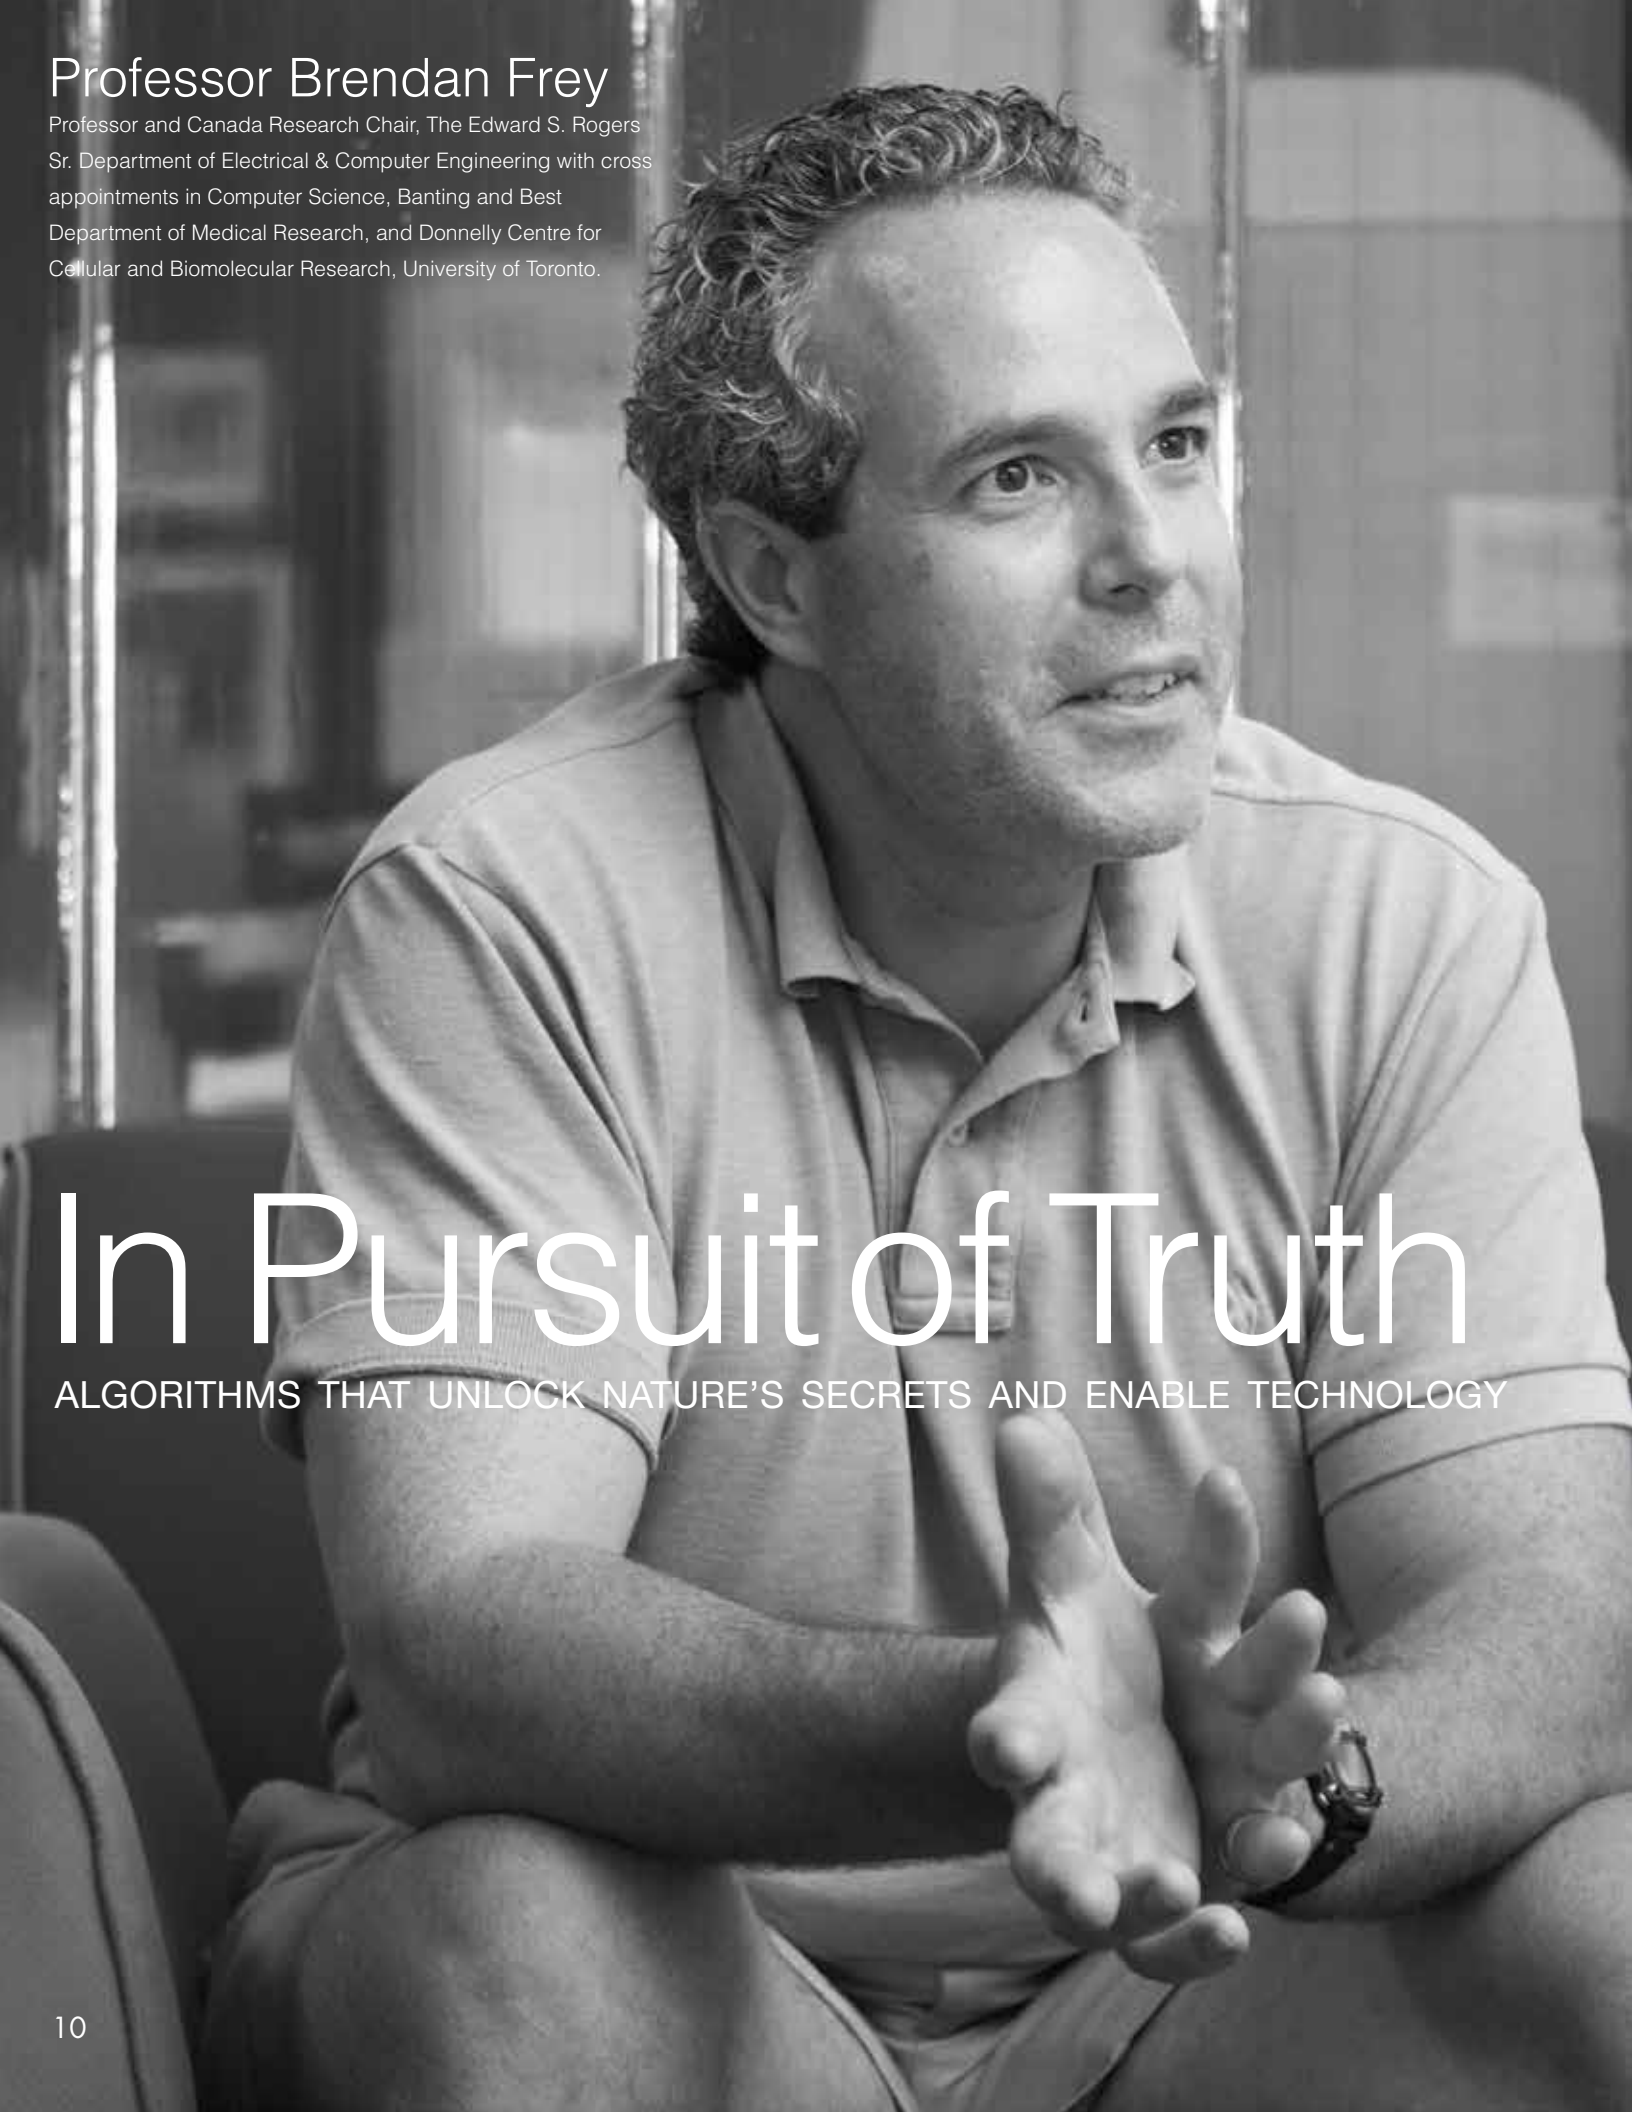

Professor Brendan Frey

Professor and Canada Research Chair, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering with cross appointments in Computer Science, Banting and Best Department of Medical Research, and Donnelly Centre for Cellular and Biomolecular Research, University of Toronto.

In Pursuit of Truth

ALGORITHMS THAT UNLOCK NATURE'S SECRETS AND ENABLE TECHNOLOGY

Utilizing mathematical frameworks and computer algorithms, Professor Brendan Frey is engineering solutions to information and sensory processing problems using high-tech tools created by studying nature.

"Science is about discovering how nature works and modeling it," says Frey. "Engineering is taking that knowledge and developing tools that enable people to do things they couldn't do before. In my group, we reduce what we see in nature to a model or algorithm and then use it to understand nature and develop tools that can benefit humanity."

World-renowned for his contributions to machine learning and information processing techniques that can identify important patterns in large masses of data, among Frey's groundbreaking discoveries are the affinity propagation algorithm (*Science 2007*) and the human genome splicing code (*Nature 2010*).

"We discover algorithms that explain how nature works, or that help us solve problems," says Frey. "These algorithms result in tools that scientists, engineers or doctors can use to solve problems they are interested in. The Affinity Propagation Server for example, has been accessed hundreds of thousands of times by people from across the globe working in areas such as nanotechnology, telecommunications, social and biological networks, genomics, high-energy physics and archeology."

Recently, Frey's work on the splicing code has attracted a lot of attention. Simply put, Frey's group discovered a code that human cells use to rearrange, or splice together, different parts of genes. This enables most human genes to carry out several different tasks, just like a single jacknife can do several different things. Researchers can determine how the process of gene rearranging is controlled, something that was not known before Frey's work. Now, if medical researchers discover a gene that is associated with a disease, they can go to Frey's website and look at that gene and see how it's being rearranged by this process of splicing. Frey's code can be used to understand how that occurs.

Another major area of interest for Frey and his group is that of the human visual system and how it works. Frey is looking at how the human brain adapts itself to make sense of visual input. "We want to know how the human brain interprets the images it receives from the retinas, and how it makes sense of that," says Frey.

When asked about what motivates and inspires him and his students, Frey says that it is a combination of scientific discovery and tool development. "If we were only developing tools, we would be missing out on something wonderful about nature. And, if we were only discovering things about nature without any consideration of potential usefulness, then it wouldn't be nearly as exciting or satisfying."

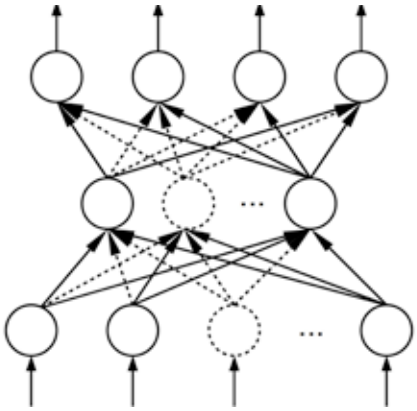

An impressive aspect of many natural and engineered systems is that they manipulate or process information in a way that leads to surprisingly complex and powerful behaviour. Information processing can be viewed as a network of computing units that evaluate and share information (top).

Frey's group studies natural and artificial systems for image analysis and vision (middle) and the processing of genetic information (bottom).

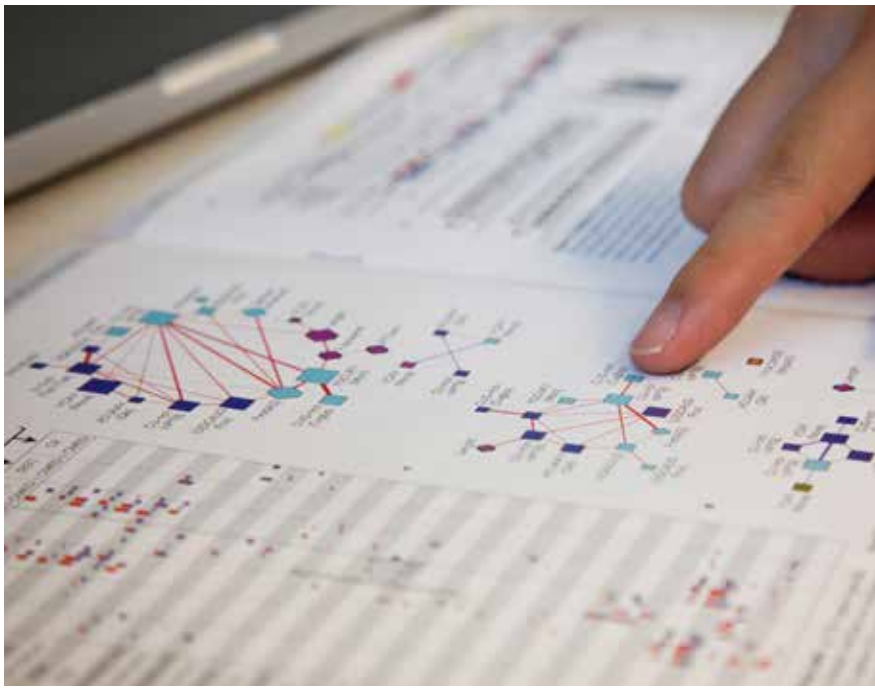

Professor Joyce Poon

Assistant Professor, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, University of Toronto.

Next Generation

OPTOELECTRONIC DEVICES BREAK CONVENTIONAL THINKING

Until Professor Joyce Poon and her team of graduate students at the University of Toronto discovered a design, and a way to drive resonant optical modulators and lasers, no one believed it was possible to simultaneously attain high modulation speeds and achieve exceptional power efficiency in these devices.

"We're really excited about this discovery because we've sorted out an important problem," says Poon. "Now we are trying to demonstrate this work. We've written the theoretical papers about our idea, done the calculations for it, and have completed a proof-of-concept experiment on a large laser. It's time to demonstrate our concept with a microcavity device."

Poon's research in the area of optical science and photonic technologies centers on novel optoelectronic devices for high-performance communication and computing systems. The work that she and her team are currently conducting will lead to advancements in fibre optic communications, computation, and optical sensing and monitoring. In their research, they explore physical phenomena and then design and control them to invent new devices.

"The reason why our modulator and laser research is important is because computers are running out of juice," says Poon. "Increasingly more data is being processed by large-scale computing systems, such as super computers and data centers, that are taking up too much power. It will soon become unrealistic to run these systems. Optics has a role to play here because the energy for communication in these systems can go down dramatically when the information is transmitted with light."

The process involves sending information using miniscule semiconductor devices that convert between electrical and optical signals. At high data rates, optical communication can require significantly less energy than communication with voltage signals transmitted across copper wires.

Poon's near-term goal is to have a chip-scale microcavity version of their modulator design that will enable them to demonstrate all of the properties and advantages they have predicted. To fabricate the devices we will use the nanofabrication facilities on campus and work with collaborators to accomplish our goal," say Poon.

Poon explains that the diversity of the topics inherent in her group's projects makes the research they do at the University of Toronto so interesting and fun. "What we do is between the worlds of science and engineering. I think this is very powerful because the ideas and devices we come up with can change the entire structure of electrical and computer engineering. The devices we are developing here in our lab, are the basic building blocks for a myriad of systems and applications to come. The potential to make a tremendous impact in the field and in society is very exciting for the students and extremely rewarding for me."

"Increasingly more data is being transmitted around the world by these super computers that are taking up so much power. It's becoming too expensive to run these systems," says Poon. "Optics research is so exciting because the energy consumption goes down exponentially when information is transmitted with light."

Exchange student, Hasitha Jayatilleka, prepares a solution for chemical etching.

Professor Sorin Voinigescu

Professor, and Stanley Ho Professor of Microelectronics,
The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, University of Toronto.

Exceeding

Professor Sorin Voinigescu's significant contribution to the industry and groundbreaking development of the first single-chip 10 Gbps transceiver led him to the University of Toronto. Today he and his team of graduate students are on the brink of launching the next generation transceiver, the 100 Gbps Ethernet. Even more impressive however, is that they are world leaders in the millimetre wave radio, radar and imaging field through the work they do at Peraso, a venture-backed start-up company created by Voinigescu, as well as in joint research projects with Robert Bosch GmbH and STMicroelectronics.

"When we set out in 2002 to work on the 60 GHz radio application, there was one other group at Berkley that was ahead of us," says Voinigescu. "We quickly caught up with them and for the past three years have reported the highest frequency radio, radar and imaging circuits in silicon technology of up to 170 GHz. So, we're actually three times higher than our original frequency goal. I was surprised by that, and also that we were able to transmit and receive from a single, extremely small chip silicon transceiver with on-die antenna. What was once unthinkable and impossible now seems so natural."

"That 60 GHz radio application is now the target of the product we're focused on at Peraso," says Voinigescu. "What we're developing is something that is extremely low power, very compact and very cheap, yet extremely powerful in terms of data rate transmitted, 5 Gbps, to your cell phone and from your cell phone to a kiosk or to a monitor on the wall."

Peraso is a success story modeled after a University of California program in which the professor helps graduate students start their own company. Voinigescu, who has extensive start-up and industry experience prior to teaching, wanted to create a self-sufficient company in Toronto that would enable graduate students to continue their research in advanced topics. "We were losing too much talent south of the border," says Voinigescu. "I wanted a place here in Canada where we could prove that our research is applicable."

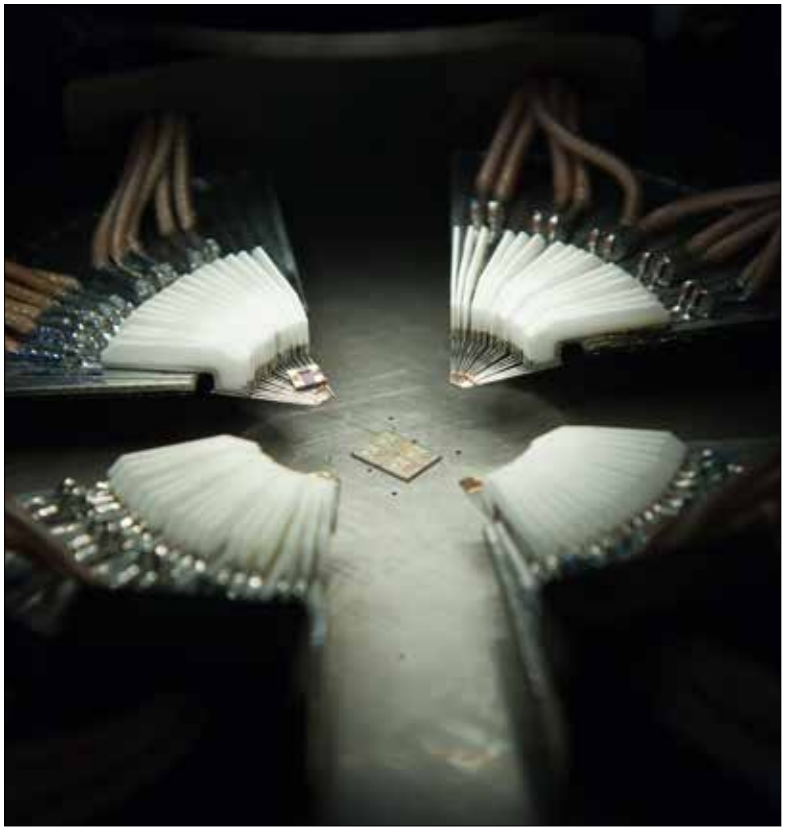

A close-up image of a 60-GHz single-chip radio die mounted on probe station.

Expectations

WORLD FIRSTS IN SILICON WIRELESS TRANSCEIVERS AND APPLICATIONS

Voinigescu and his team at U of T are currently working on higher frequency applications above 100GHz for industrial, radar, and imaging applications as well as environmental monitoring.

In the healthcare sector, Voinigescu is setting up a collaboration with McGill University to look at remote measurements of thoracic capacity by surrounding the patient with an array of millimetre wave sensors.

"Because we are able to put the antenna on such a small die means that it is extremely lightweight, and consumes very little power, for very little cost," says Voinigescu. "They're wearable and can be applied everywhere."

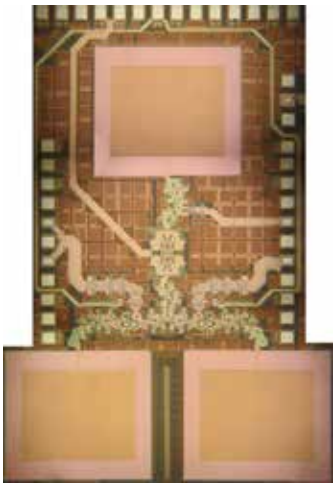

Voinigescu and his team have been able to transmit radar signals over five metres from a single chip silicon transceiver whereby the antenna is so miniscule that it actually fits on a tiny silicon die measuring only two by two millimetres in size.

Professor Alberto Leon-Garcia

Professor and Canada Research Chair, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, University of Toronto.

Next Generation

SMARTER

ECONOMICAL

GREEN

Internationally recognized for his work that has led to significant improvements in the design of algorithms and hardware for network switches, including a revolutionary all-optical switch, Professor Alberto Leon-Garcia's current focus is on designing the future network architecture.

"What I'm looking at is an evolution where it's really hard to distinguish between whether the networks are computers or the computers are networks," says Leon-Garcia. "I see a ubiquitous, intelligent, electronic skin that will spread across the earth to allow communication, and enable all kinds of applications."

This involves what Leon-Garcia calls smart infrastructure, whereby everything from power grids to intelligent highways, with integrated transit systems and networked vehicles, are constantly connected and cost-efficiently managed by automation. Incorporated into the mix are next generation data centres that perform intelligent algorithms able to answer much more complex queries than by today's standards.

Leon-Garcia's Virtual Application Networking Infrastructure (VANI) is a Canadian test bed system with unique elements. VANI provides the infrastructure for the networking that combines cloud computing and applications, and involves a wide use of sensors. The sensors are designed to monitor systems such as traffic flow or power consumption, and management systems then apply controls to influence demand.

"VANI works, and is ready to be accessed by the outside," says Leon-Garcia. "We've moved from a position of asking for funds to build prototypes to one of deployment of a national test bed. VANI also serves as a sandbox for experimentation that promotes entrepreneurship among our students and keeps us competitive."

Another focus for Leon-Garcia and his team is in the area he pioneered a decade ago–optical networking. "The largest data centres have a quarter of a million computers in a warehouse," explains Leon-Garcia. "You have to interconnect

Internet

these at a very high speed in order to do an online search for example. What's happening as we approach data centres with millions of computers is that we're getting into problems just managing the number of cables between these computers. Physically there are just too many. And they use way too much power. If we want to continue growing we have to switch to optics."

Optics, or light, has the ability to transfer information at much higher speed than through wire. One of the challenges Leon-Garcia has presented to his graduate students is to design highly compacted networks using optics. "That means the number of cables get reduced, the sizes of the computer boxes get reduced, and the power required gets reduced dramatically," says Leon-Garcia. "What this means is that we will be able to continue to grow and build large data centres while keeping the envelope, in terms of how much power and space we use, compact–maintaining the greenness."

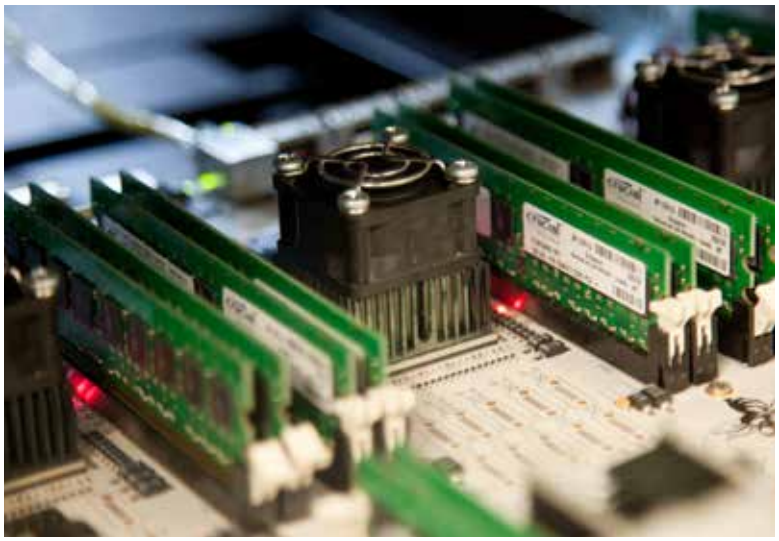

Programmable hardware in VANI.

An entirely new area of research, Leon-Garcia and his team of graduate students at the University of Toronto are investigating the interplay between traditional networking and data centres to enable applications, while concerned with cost efficiencies and renewable sources of energy without carbon emissions.

Testing VANI in the Network Architecture Lab.

RESEARCH DIRECTORY

Each research category has a corresponding colour. Complete a quick-search by colour to locate the lead researcher and then refer to the alphabetical listing that follows.

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

Computers
Computer Communications
Computer Software/Hardware

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

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

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

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

Quick-Search by Colour Coded Listing

Communications

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

Computers

- Computer Communications
- Computer Software/Hardware

Energy

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

Quick-Search by Colour

Engineering/Sciences

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

Life Sciences

- Biomedical Engineering
- Human Health
- Life Sciences
- Life Sciences, including Biotechnology

Information Technology

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

Quick-Search by Colour

Information Technology

Information and Communications Services
Information, Computer and Communication Technologies

Information Systems and Technology
Information Technology

Alphabetical Directory by Lead Researcher

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

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 amount 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 to 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 (hence, 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 are 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 by utilizing Extremely Large Datasets (ELD). ELDs allow for greater accuracy in extracting information from images and videos, but in return require 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). (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 classification the vast amount of untagged images and videos. In turn, the computational load can be addressed through FPGA-based hardware acceleration which would enable the classification of an image or video to be performed in real-time.

Abdelrahman, Tarek — 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.

Abdelrahman, Tarek — Compiler Support for GPGPU Programming

www.eecg.utoronto.ca/~tsa

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.

Abdelrahman, Tarek
www.eecg.utoronto.ca/~tsa

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 that fact that traces are short straight-line segments of code, which makes them more amenable for 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 cooperation 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 wavelength 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 engineering 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 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, 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 trouble shooting 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 on the efficiency of this industry, by 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 on-line performance modeling and anomaly detection algorithms and tools that form the basis for self-configuration, self-tuning, and self-healing servers. We use these techniques in our data center laboratory towards automatically providing quality-of-service for a range of dynamic content services such as, e-commerce, on-line bidding and massively multi-player games.

Amza, Cristiana
www.eecg.toronto.edu/~amza

System Support for Parallel and Distributed Software Transactional Memory

Due to 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. libTM 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.

Anderson, Jason
www.eecg.toronto.edu/~janders

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

Field programmable gate arrays (FPGAs) are programmable semiconductor chips that are part of that 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, breaking it into smaller problems, whose solutions can be rapidly looked-up in a library, and stitched together to form a good solution for the complex problem.

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

Beam Formation in Anisotropic Plasma and Anisotropic Metamaterials

Study of electromagnetic beam phenomena in highly anisotropic media, by means of experiment and simulation. The two classes of materials of interest are magnetized plasmas in space and in the laboratory together with planar two-dimensional anisotropic periodic structures known as metamaterials. For the first class of materials (anisotropic plasmas) the proposed work relates to luminous discharges seen in the 1995 OEDIPUS-C sounding rocket experiment as well as in more recent laboratory experiments. Their formation and shape will be the subject of this research in order to better understand the processes that led to their occurrence. This includes re-evaluation of data from the OEDIPUS-C experiment followed by computational simulations that were not possible in the past, and eventually laboratory experiments. For the second class of materials (metamaterials) the work will focus on two-dimensional wire and patch grids without ground planes loaded orthogonally with capacitors and inductors. If fed by a point source, so-called resonance cones will emerge from the source as lines along which electric and magnetic fields are extremely high. It has been predicted in previous

work that these structures have waveguiding properties and provide special opportunities for making microwave components. Another aspect of the proposed metamaterial research is the scattering of electromagnetic waves incident on these and similar structures. If a conducting object is located underneath the grid, resonance cones can be excited indirectly thus leading to a means of finding hidden or buried objects.

Bardakjian, Berj — Bioengineering of the Brain

heart.ibme.utorontxo.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 (1) to 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.

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.

Broucke, Mireille — Patterned Linear Systems

www.control.utoronto.ca/~broucke

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 system, 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 Center, 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, as well as being of practical use when implemented in an industrial quality CAD tool. My current research effort is in the area of CAD flows for FPGA devices, and specifically in the development of new methods of efficiently compiling high-level language code (such as C code) into circuits that can be implemented in FPGA devices.

Chan Carusone, Anthony — Integrated Circuits and Systems

isl.utoronto.ca

Research in my Integrated Systems Laboratory strives to extend the rate, reach, and reliability of our communication links–whether they are between a satellite and an earthbound receiver, or between two integrated circuits just a few millimetres apart. Specifically, we develop integrated electronic systems that process information with high energy efficiency and/or speed. Projects combine signal processing concepts with integrated circuit design.

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, 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.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

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: e.g. in chemical engineering, electrical power systems, aerospace systems, transportation systems, building temperature control systems, large flexible space structures, pulp and paper control systems, as well as in other areas such as in management science and in biological systems.

Dawson, Francis — Improving Energy Efficiency of Energy Conversion Processes

www.ele.utoronto.ca/~dawson

The general research interests are in the area of modeling 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 best system architecture and control philosophy that leads 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 modeling of thermoelectric and piezoelectric devices.

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. Example devices include small antennas and multi-functional RF/microwave components, sub-diffraction imaging lenses and probes, invisibility cloaks and related 'transformation optics' lenses, plasmonic optical circuits, and nano antennas. Research includes both experimental work as well as fundamental theory.

Enright Jerger, Natalie — Interconnection Networks for Heterogenous Multi-Core Systems

www.eecg.toronto.edu/~enright

In recent years, embedded and mobile devices have proliferated 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 of increased functionality. Modern chips may contain several processing cores, 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 general purpose devices; we are developing novel architectures to meet the power and real-time performance constraints of these systems.

Enright Jerger, Natalie — Semantically Rich Networks for Many-Core Architectures

www.eecg.toronto.edu/~enright

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. With a semantically-rich interconnection network, we propose to provide the programmer and compiler with information about the communication capabilities on-chip to ease programmability. Today's on-chip interconnection networks are largely oblivious to the needs of the components they connect and serve the sole purpose of shuffling bits around the die. This work proposes to embed additional functionality, such as novel message passing support and support for reduction operations, both in the interconnection network as well as further integrate communication into other layers of the system.

Francis, Bruce — Distributed Robotics: Application

sites.google.com/site/brucefranciscontact

During 2007-2009 we conducted a major application of our theory in collaboration with Professor Tim Barfoot of the University of Toronto Aerospace Institute and Jared Giesbrecht of Defense 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. The resulting design and field trials were successful enough to appear in the premier journal on field robotics.

Francis, Bruce — Distributed Robotics: Theory

sites.google.com/site/brucefranciscontact

The methodology of control theory is to begin with a practical problem; to abstract the central issues and formulate an idealized, hypothetical problem; to develop, if necessary, new mathematical methods for its solution; and to work out a rigorous solution. Then one has a framework on which to do real applications.

My students, co-supervisors, and I began by formulating the hypothetical problem of rendezvous for point robots. This is entirely analogous to birds flying in a flock; the heading angles are all equal. The rendezvous problem is a challenge because there are no leaders, the robots do not have a common map and preferably should have identical stored programs, and there can be no human intervention.

We began with the simple strategy of cyclic pursuit. The theory of robot formations relies heavily on graph theory. We derived a new control-theoretic result: a necessary and sufficient condition on the graph for rendezvous to be achievable. We studied the cyclic pursuit strategy applied to unicycles. We found that stable formations could be achieved, where the unicycles are moving in a circle in the same direction. We confirmed this experimentally. We moved to the problem of formation control, for example getting four rovers to form a square, derived control laws, proved stability of the formation, and verified by experiments on real rovers. In our major theoretical work we studied the rendezvous problem when the robots can see only a fixed distance away. We gave the first mathematical proof of an algorithm for this case.

Frey, Brendan — Algorithms for Inference and Machine Learning

www.psi.toronto.edu

We develop 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. Professor 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 we developed 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*).

Frey, Brendan — Data Analysis and the Affinity Propagation Algorithm

www.psi.toronto.edu

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. We 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 we developed was accessed over 100,000 times by over 3000 users (unique IP addresses), 28,600 of which were from Canada. Google returns over 10,000 hits for the search term 'affinity propagation'. This method has been applied to solve problems in biology, genetics, genomics, medicine, physics, chemistry, telecommunications, electronics, archeology, economics and social networks.

Frey, Brendan — Deciphering the Human Genetic Code

genes.toronto.edu

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 cells comprising all human tissues. In the words of the famous genomics researcher Eric Lander, "Genome: Bought the book, hard to read." Recently, we 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, I describe research our research. We 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 was reported on in the *Globe and Mail*, the *Toronto Star*, CBC Radio, BBC Radio, and a variety of other national and international news.

Frey, Brendan — Image and Video Analysis

www.psi.toronto.edu

We explore computational techniques for analyzing images and videos so as to automatically extract a representation of what's going on in the image or video. The methods developed are motivated by neuroscience and make use of large scale graphical models and inference algorithms, including those based on optimization, Markov chain Monte Carlo, and variational techniques. A movie showing the input and output of one of the algorithms developed in our laboratory is available here: http://www.psi.toronto.edu/images.old/figures/cutouts_vid.gif.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Genov, Roman — Portable, Wearable, and Implantable Sensory Biomedical Electronics

www.eecg.utoronto.ca/~roman

I am 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 — Data Recovery

isl.utoronto.ca

This research aims to design computer systems that enable efficient data recovery (i.e., maintain data integrity and never lose data) in the presence of attacks and catastrophic failures. Our approach consists of systematically monitoring all components of the system, and capturing critical data that will allow analysis and recovery of compromised systems.

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 know as Lattice Reduction that 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 particularly attractive for low-power implementations.

Future work is focussed 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 is also characterized by enhanced privacy and confidentiality. The proposed security architecture, namely, '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, 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 which are among the most fundamental requirements of modern and future health/home care services.

Hatzinakos, Dimitrios — 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 essentially 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, which 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, as well as 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, accommodating various user needs successfully in a wide range of Scenarios. Depending on the application, a pre-selected 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 spectralefficiency, which will enable high-data-rate wireless multimedia communication to be an affordable and practical reality.

Hatzinakos, Dimitrios — 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, due to 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 in exact localization of wave boundaries. To avoid fiducial points detection, the signal is processed holistically, using second order statistics. Our autocorrelation based method is a very simple and effective approach that does not require any waveform detection. It depends on estimating and classifying the significant coefficients of the Discrete Cosine Transform (AC/DCT) or the Linear Discriminant Analysis (AC/LDA) of the autocorrelation of heartbeat signals. The AC/LDA algorithm has been incorporated to a prototype system developed at the BioSec.Lab, the HeartID. HeartID is a Matlab based software with various functionalities, such as user enrollment, database handles, security level adjustment and identification / verification modes of operation.

Hatzinakos, Dimitrios — 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, as well as the general Canadian society, 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, 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. U of T 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 — A Universal Trace Sensing Platform for Nanoparticles and Biological Samples in Solution and Gas Forms

photonics.light.utoronto.ca/helmy/research

This project is concerned with developing a universal platform for detecting and analysing substances in solution and gaseous forms in dilute concentrations using enhanced Raman spectroscopy. This capability can greatly enhance fields such as detection of toxic content, sensing biological samples in dilute concentrations as well as nanoparticles detection in the environment amongst others.

Raman spectroscopy is much more powerful than more commonly used optical detection techniques such as absorption and Photoluminescence. Raman provides unparalleled specificity to the detection as it conveys wealth of information about the chemical composition of the specimen under test. Raman spectroscopy provides a weak signal however. Numerous techniques have been developed to enhance the Raman signal emitted from a molecule particle or molecule. Most of these techniques either require complicated setup that does not lend itself to mobile sensing or they involve contaminating the specimen under test by adding contaminants such as metal nanostructures.

The technique provides orders of magnitude enhancement in the Raman signal while achieving this goal in a compact rugged platform that lends itself to mobile and bed-side diagnostics.The enhancement is achieved using a hollow core photonic bandgap fibre (HC-PCF). These fibres work on enhancing the interaction length between the substance under test and the laser pump which induces the Raman signal.

Helmy, Amr — High Power Diode Lasers

photonics.light.utoronto.ca/helmy/research

Bragg reflection waveguide lasers (BRLs) are essentially one-dimensional (1D) photonic bandgap (PBG) p-i-n structures, 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 at Bragg reflection waveguide (BRW) mode, not the conventional total interface reflection (TIR) mode. BRLs have been predicted to enable the realization of high power lasers and amplifiers due to single mode waveguides with larger mode volumes, high gain coefficient and strong mode discrimination. Moreover,

this class novel PBG lasers have also shown the potential in applications related to¬ nonlinear frequency conversion, monolithically integrated optoelectronic integrated circuits (OEIC). Our group has demonstrated the first edge-emitting BRW laser-diode with low threshold current and high characteristic temperature. We are interested in developing high performance single-mode BRLs, second-harmonic laser, monolithically electrically injected optical parametric oscillators (OPO), and compact spontaneous parametric down conversion (SPDC) single-photon sources by integrating BRLs with nonlinear BRWs.

Helmy, Amr
photonics.light.utoronto.ca/helmy/research

Photonic Intergated Circuits Using Second Order Nonlinearities

Nonlinear frequency conversion based on second-order nonlinear interactions in semiconductor optical waveguides is an attractive approach for creating compact, widely tunable, monolithically integrated, and efficient coherent light sources. In comparison, the wavelength coverage of available laser sources is restricted by inherent electronic transitions, leaving gaps in the spectrum, where no suitable material exists. The GaAs/AlGaAs material system is of particular interest for its large second-order nonlinearity, broad transparency window (0.9–17 mm), and well-established fabrication technology. Recently the group was able to achieve phase matching of second order nonlinearities in a promising novel platform using Bragg reflection waveguides (BRWs). BRWs offer exact PM between the lowest order modes, hence enabling maximum power utilization among the interacting frequencies. The platform demonstrated second-harmonic generation difference frequency generation and sum frequency generation with very attractive properties. In addition the platform also demonstrated laser action when quantum wells are added to the waveguides in a p-i-n structures.

This work demonstrates the possibility for the first time to realize monolithic integrated photonic circuits utilizing active, passive and nonlinear devices based on ultrafast second order nonlinearities. These circuits can provide groundbreaking advances for applications in novel monolithic tuneable laser sources, all optical signal processing circuits and novel devices to modulate the quantum mechanical properties of photons.

Herman, Peter
photonics.light.utoronto.ca/laserphotonics

3-D Laser Fabrication: Enabling Nano-Optics for the Nano-Sciences

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, mechanica– 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. Nanooptics 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 3-D nano-optical systems. Near-field and phase-shifting techniques will be exploited in multi-level diffractive optical elements to design 'intensity defects' within 3-D periodic interfering laser patterns. Photosensitive optical materials exposed to these modified laser patterns will see nanooptic devices precisely embedded at the critical points of a 3-D periodic lattice to enable the nanofabrication of compact 3-D photonic crystal circuits, 3-D 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 3-D nanomolding– of significance to Canada's optics, biophotonics, and nanotechnology industry.

Herman, Peter
photonics.light.utoronto.ca/laserphotonics/

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 nano-scale manufacturing methods for a new generation of smart medical devices, bio-sensors, and Telecomm 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 3-D manufacturing of Telecom optical circuits, optical fiber assemblies, smart medical catheters, optical sensors, and other high value photonic systems for our Canadian partners.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Hum, Sean — Reconfigurable Antennas for MIMO and Compact Handsets

www.waves.utoronto.ca/prof/svhum/research.html

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.

Hum, Sean — Reconfigurable Aperture Antennas

www.waves.utoronto.ca/prof/svhum/research.html

This project focused on the development of aperture antennas producing highly selective antenna beams that can be electronically manipulated in a cost-effective manner by tuning elements composing the aperture. Aperture types include reconfigurable reflectors, lenses, and conformal surfaces. This technology has applications in satellite communications, RADAR, and other long-range communication links.

Hum, Sean — Ultra-Wideband Antenna Arrays

www.waves.utoronto.ca/prof/svhum/research.html

This project aims to develop reconfigurable beam-forming networks for ultra-wideband communication systems operating in the 3-10 GHz range. The basic idea is to use electronically agile networks that process signals in both space and time, in conjunction with an array of antennas, to realize high-gain antenna patterns for receivers which maintain their characteristics over an ultra-wide frequency band. Ultra-wideband antenna arrays can be used in communications, radar, microwave imaging, wireless localization, and many other applications.

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

www.ele.utoronto.ca/~iravani

This work includes research and development of analytical and time-domain simulation tools, 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, (2) microgrids with high-depth of penetration of distributed generation and storage units.

Jacobsen, Hans-Arno — Middleware Systems

msrg.org

This research aims to ease the development of scalable, reliable, and secure distributed enterprise applications. In pursuit of these objectives, he engages in basic research on event processing, publish/subscribe, service-orientation, aspect-orientation, and green middleware. In research and development engagements with various companies, he pursues projects on business process management, service delivery models, service and infra-structure management, and e-energy. Selected research projects include the PADRES Events & Services Bus (padres.msrg.org) for effective business process management, the eQoSystem project (eQoSystem.msrg.org) for declarative monitoring and control of SLAs (service level agreements) in enterprise applications and business artifacts, and the AspeCtC (ACC) project (aspeCtC.net) for increasing modularity in systems software and embedded systems.

Ongoing research seeks to add and improve enterprise-grade qualities of the middleware. The PADRES system is a distributed content-based pub-lish/subscribe middleware with features built with enterprise applications in mind. These features include: (1) Intelligent and scalable rule-based routing protocol and matching algorithm ; (2) Powerful correlation of future and historic event; (3) Failure detection, recovery and dynamic load balancing (4) System administration and monitoring. As well, the PADRES project studies application concerns above the infrastructure layer, such as; (1) Distributed transformation, deployment and execution; (2) Distributed monitoring and control; (3) Goal-oriented resource discovery and scheduling ; (4) Secure, decentralized choreography and orchestration. A publish/subscribe middleware provides many benefits to enterprise applications.

Janischewskyj, Wasyl — Lightning Studies at the CN Tower

wasyl.janischewskyj@utoronto.ca

The CN Tower, the tallest free-standing structure instrumented for observation of LIGHTNING, has been used by the CN Tower Lightning Studies Group (LSG), based in the ECE Department, for the last more than 40 years to record, significant to engineering, parameters associated with lightning strikes. Staff and students of the University of Toronto as well as those of Ryerson University and McMaster University, and occasionally Visiting Scientists and Visiting Professors from various parts of the world are members of the Group. Further to collecting data through measurement, graduate courses are offered and research is conducted in the area of Lightning Protection. Findings of the Group have been adopted in Canadian Lightning Standards and in documents of International Working Groups dealing with questions of Lightning in such technical societies as IEEE, CIGRE and IEC. In August of 2001 an International Workshop was held at University of Toronto at which the International Project on Electromagnetic Radiation from Lightning to Tall Structures (IPLT) was launched. There are now 12 countries participating in the Project. Our Group represents Canada and Wasyl Janischewskyj has been Chairman of the IPLT Steering Committee from 2001 to 2009.

Activities of LSG are continuing. Recent research findings include development of the Five-Section Model for representation of the CN Tower in computational processes, proposal of a procedure to account for variable speed of return-stroke propagation within the lightning channel and introduction of Cartesian coordinates in computation of the electromagnetic field radiated by a return lightning stroke to a tall structure. This latter development provides a convenient tool for computation of the electromagnetic radiation by a tortuous path of the lightning channel.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

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 nano scale electronic circuits to sense physical quantities. Some recent examples of commercial applications for MEMS are pressure sensors used as microphones in devices such as cell phones and hearing aids, inertial sensors used in airbag deployment as well as positional control in handheld games and cell-phones, 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.

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 range of 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;
(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.

Kherani, Nazir — MicroPower Sources and Sensors

www.ecf.utoronto.ca/~kherani

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.

One approach that we are currently pursuing is the development of a continuous vibrational piezoelectric energy harvester using a tritium-occluded-in-silicon beta-emitting source. Integration of the beta source enables a constant source of vibration thus overcoming the limitations of ambient vibrational sources. The vibrational energy harvester is based on aluminum nitride piezoelectric material.

Current sensor development uses the aluminum nitride platform, investigating surface acoustic wave detectors with the use of various sensitizers.

Kherani, Nazir — Optical Coatings and Photonic Materials

www.ecf.utoronto.ca/~kherani

On the basis of diamond-like carbon (DLC) films we are (1) exploring innovative thin film integrations for 'energy conservation in buildings' applications, and (2) developing novel DLC integrations for photonic applications.

Diamond-like carbon film, a silicon-compatible material, has 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 DLC as a viable optical-photonic platform material through the research of advanced thin film synthesis techniques, development of DLC-based optical coatings technology for energy applications, and investigating passive and active photonic devices.

Kherani, Nazir — Photonic Crystal – Photovoltaics

www.ecf.utoronto.ca/~kherani

The aim of this project is to investigate photonic crystal - photovoltaic integrations with the aim of creating high-efficiency, economic, third-generation solar cells. The novelty of the research lies in novel integrations of nano-materials 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 ~1mm 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. Much larger path length enhancement factors, on the order of 103 to 104, however 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 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.

Khisti, Ashish — Signals, Multimedia and Security Laboratory

www.comm.utoronto.ca/~akhisti

Our research focuses on two areas that are crucial for the success of future wireless systems: Security and Multimedia. In traditional architectures, wireless links are the weakest links in security. Our research challenges this notion and strategically exploits the unique characteristics of wireless links at the physical layer to develop new security mechanisms. Our group is also investigating new wireless networks optimized for multimedia rich services that will play a key role in the future of human communications.

Kschischang, Frank — Error Control in Network Coding

www.comm.utoronto.ca/frank

Network coding, proposed by a small group of researchers ten years ago, has the potential to greatly improve the efficiency of information transmission in packet networks. The key idea of network coding is to generalize the operation of intermediate nodes in the network, changing their operation from routing to coding. Traditional network routers treat packets as fragile and distinct pieces of a message, to be switched along appropriate network pathways, and reassembled into a message at the receiver. Network coders, on the other hand, treat packets as robust and indistinct lumps of "evidence" which can can be mixed together (and not just routed) at intermediate nodes. Receivers gather evidence, and can infer which message was sent when sufficiently many clues have been received. The great advantage of allowing evidence to be combined at intermediate nodes is that this evidence can then better squeeze through network bottlenecks, allowing for greater transmission efficiencies than can be achieved through routing alone. In this work, we have developed several novel methods for error control coding in the context of network coding. In one interesting approach, information is encoded using a "codebook of vector spaces". Together with Ralf Koetter and Danilo Silva we have developed a family of Reed-Solomon-like codes, provided decoding algorithms, developed connections to rank-metric codes, and developed schemes to provide security against wiretapping. Our initial paper on this subject, published in 2008, received the 2010 IEEE Communications Society and Information Theory Society Joint Paper Award. Ongoing work seeks to develop these ideas at the "physical layer", with the potential for improved and robust transmission protocols in wireless mesh and relay networks.

Kschischang, Frank — Advanced Coding and Signal Processing in Fiber-Optic Communication Systems

www.comm.utoronto.ca/frank

Fiber-optic transmission systems are evolving at a rapid pace towards achieving greater spectral efficiencies. Coherent detection is supplanting noncoherent detection, 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. A first aim of this project is to investigate practical signal processing algorithms that mitigate channel impairments that dominate in optically- routed systems that operate at high spectral efficiencies in a WDM environment. Such systems are likely to employ inline chromatic dispersion mapping, but such mapping alone is unlikely to be able to fully compensate for fiber nonlinearity at high spectral efficiencies. Electronic pre- and post-compensation is thus likely to be needed. In particular, electronic back-propagation is a promising approach that may lead to an equivalent channel for which new coding and modulation techniques may be applied to practically achieve the desired spectral efficiencies.

A second aim of this work is to investigate the synergy that results from a judicious selection of coding and modulation techniques in conjunction with the receiver-side signal processing. It is well known in wireless communication that joint channel processing and decoding is beneficial (both from a performance and from an overall complexity point of view) compared with a system in which these receiver functions are separated. A central issue here is code design: how should the error-control system be designed to meet communication objectives while simultaneously providing assistance to signal processing functions at the receiver? A potential outcome of this work is a new approach to joint coding, modulation, and signal-processing that leads to the practical achievement of greater spectral efficiencies in optical transmission systems.

Kwong, Raymond — Dependability and Security in Control and Multimedia Systems

www.control.utoronto.ca/people/profs/kwong

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

Prof. 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.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Leon-Garcia, Alberto — Autonomic Service Architecture
alberto.leongarcia@utoronto.ca
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.

Leon-Garcia, Alberto — Design of Converged Communications and Computing Infrastructure
www.vani.utoronto.ca
We are developing future network architectures that are application oriented, that have the flexibility to simultaneously support multiple protocol stacks, and that operate over a converged computing and communications infrastructure. In order to enable at-scale experimentation with novel protocols and applications, we have developed application oriented routers that provide virtualized computing and communications resources, and that when interconnected provide a network test bed. We are also developing a service-oriented control and management system that allows the creation of multiple, simultaneous converged networks for testing and developing new protocols and applications.

Leon-Garcia, Alberto — Optical Networks for Ultrascale Datacenters
alberto.leongarcia@utoronto.ca
We are designing optical networks that can provide connectivity for future datacenters 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 Bio-Sensors and Biomedical Imaging Systems
biophotonics.utoronto.ca
Our research interests include developing biomedical imaging systems and optical bio-sensors based on semiconductor devices and nano-structures, and their application to bio-medical diagnostics, in vivo imaging, and study of bio-molecular interactions. The goal of our work is to integrate sensor components into miniature functional bio-sensors and apply them to novel biology and bio-medical applications. As such, our research is interdisciplinary and include semiconductor device physics, optics, micro- and nano-fabrication, chemistry and applications in biomedical diagnostics, cancer studies and neurobiology.

Li, Baochun — Agile Cloud Computing via Multi-Level Performance Inference
iqua.ece.toronto.edu
Cloud computing systems need to be judiciously managed to scale up to multiple geographically diverse data centres, each with tens of thousands of servers, networking switches, and storage systems. At this scale, energy costs have already become more than a third of the overall costs of maintenance and upkeep. As clients are served with a large number of service-oriented applications — such as on-demand media streaming — running on such cloud computing systems, it is crucially important to operate these systems at the maximum energy efficiency, to achieve the vision of green cloud computing. In this project, we will establish and design a fundamental inference framework that forms a basis towards these goals. Leveraging time-sequential graphical inference models, this framework can (1) forecast the demand in each service instance, (2) predict the workload at each individual server node, and (3) characterize the dependencies between the quality of service when delivered to clients and the incurred bandwidth and energy costs. Underlying all these inference tasks is a simple intuition — a node's future behavior is predictable through observations of its current (or historical) states. Our inference framework only needs each server node and client to collect direct and limited observations from its neighborhood, and can be re-trained online.

Li, Baochun — Magellan: Charting Large-Scale Commercial Peer-To-Peer Streaming Systems
iqua.ece.toronto.edu
In the Magellan project, we seek to extensively measure and analyze a large-scale peer-to-peer (P2P) streaming system, in collaboration with UUSee Inc., a leading P2P streaming solution provider in China. Magellan is based on large volumes of traces (more than a terabyte over a period of one year) collected from the UUSee streaming system, involving millions of users, with a snapshot of the entire system every ten minutes. The methodology and scale of this work is unprecedented in P2P streaming research. Our research in Magellan investigates P2P streaming topologies, characterizes P2P streaming flows, provisions server capacities, and predicts the scale and performance in the near-term future.

Li, Baochun — Nuclei: Many-Core Network Coding on the GPU
iqua.ece.toronto.edu
In the Nuclei project, we have successfully implemented an accelerated multi-threaded implementation of network coding, to take advantage of both multiple CPU cores with aggressive multi-threading, and SSE2 and AltiVec SIMD vector instructions on x86 and PowerPC processors. We have shown a 20x performance improvement over the baseline implementation. Further, by taking full advantage of modern Graphical Processing Units (GPUs) with hundreds of computational cores, we have managed to push the performance envelope of network coding by showing a further 10x performance improvement over the best possible performance without GPUs.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Liang, Ben
www.comm.utoronto.ca/~liang

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 aggravate 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 cooperative 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 compute infrastructure. It provides users a pay-as-you go model, and allows users 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 give 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 smart phones have more computing power than yesterday's work stations. 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 security 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 shown to be a disruptive technology, that 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.

Lim, Teng Joon
www.comm.utoronto.ca/~limtj

CUDA for Wireless Communication Simulations

Wireless communications systems are growing ever more sophisticated in order to provide increasingly demanding users with the service quality they expect. As a result, the signal processing and other functions in a transmitter and receiver are very complex, e.g. it is now standard for systems to adapt their modulation/coding formats to the channel, employ error control codes that are theoretically close to optimal, and so on. In order to design such systems, it is necessary to simulate them in software before prototyping them in hardware. However, due to the system complexity, end-to-end performance simulations using the popular software package Matlab is very time-consuming, often requiring a few days of continuous computations for even quite routine simulations. Sometimes it would be simply infeasible to perform these simulations for realistic parameter values and so researchers would resort to simulations of very small systems (two users only for instance). In this project, we have been using Nvidia's CUDA platform to leverage the parallel processing feature of graphical processing units (GPUs) to perform these simulations. The goal is to provide a toolbox of often-used functions that other researchers can apply in their simulations. With speed-up factors up to several hundred over C programs run on conventional serial processors, this approach is very attractive and is fast catching on in many computationally intensive research fields.

Lim, Teng Joon
www.comm.utoronto.ca/~limtj

Downlink Transmission with Noisy Channel Information

A base station transmitting to multiple users can perform much better when it knows the channel to each user, e.g. it can process the signal it transmits so that each user receives its own signal without interference. However, having accurate channel knowledge is practically impossible, as the channel must be estimated with finite accuracy, and those estimates used after a finite delay during which time the channel may have changed. The characterization of the performance of transmitter precoding in the presence of inaccurate (or noisy) channel information, and the design of precoders that explicitly account for these inaccuracies, are the objectives of this project. Its outcomes will be design guidelines that tell system designers when precoding using noisy channel information is beneficial, and a superior design that always performs better than conventional methods because it models the noise in the channel information explicitly to a reasonably accurate approximation.

Lim, Teng Joon
www.comm.utoronto.ca/~limtj

Opportunistic Communications through Spectrum Sensing

The finite band of electromagnetic spectrum that is suitable for wireless communications is growing more crowded every day, with much of it already dedicated to licensed users such as radio and TV stations, cellular service providers, and emergency services. However the nature of many of these communication systems make it impossible for them to use their assigned bands efficiently, e.g. data communications occurs in bursts with silent

periods in between, and TV channels are not used in some parts of the country due to low demand. In this project, we explore the possibility of opportunistic communications through intelligent sensing of the spectrum (e.g. sensors acting cooperatively). In particular, we design schemes for sensing spectrum availability in the shortest possible time, and we calculate the throughput and probability of interfering with the licensed user of various spectrum sensing methods. The outcome of the project will be an understanding of the feasibility of opportunistic communications for various application scenarios, and enhanced spectrum sensing methods to maximize system performance.

Lo, Hoi-Kwong — Quantum Cryptography: from Theory to Practice

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

We seek to build high-speed (> 1 Gbit per sec) unbreakable secure communication systems based on quantum mechanics. “The human desire to keep secrets is almost as old as writing itself.” With the advent in electronic businesses and electronic commerce, the importance of secure communications via encryption 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 security of practical quantum key distribution systems. We do so through system building and studying hacking strategies and counter-measures.

Maggiore, Manfredi — Advanced Motion Control in Robotic Systems

www.control.utoronto.ca/~maggiore

Motion control problems are ubiquitous in industrial manufacturing and advanced robotics. Examples include the design of autopilots for unmanned aerial or land vehicles, the control of haptic interfaces, and the control of medical exoskeletons for patients with severe leg injuries. What these problems have in common is the fact that their solution does not involve the stabilization of an equilibrium or the tracking of reference signals, as is the case in classical manipulation. Rather, these problems require the enforcement of certain constraints between the states of the system. Classical control theory is not equipped with methods to solve this new class of problems, and in fact existing algorithms for motion control are the result of ad hoc considerations for specific systems. Without general methods to synthesize motion controllers, modern robotic systems will not attain their full potential. The objective of this research is the development of systematic methodologies and algorithms to solve motion control problems for a large variety of electromechanical systems. One of the potential applications is the control of walking motion for biped robots.

Maggiore, Manfredi — Development and Control of an Autonomous Co-axial Helicopter

www.control.utoronto.ca/~maggiore

In the aerospace industry there is a growing interest in unmanned aerial vehicles (UAV) for a wide variety of civil and military applications, as well as for scientific activities. In particular, unmanned helicopters are attracting significant attention in this industry because their ability to hover makes them ideal for monitoring and surveillance applications. Helicopters, however, are difficult to control because, unlike airplanes, they are naturally unstable. Such instability is reflected in the fact that if the pilot does not perform continual corrections, the helicopter will tend to roll over and lose control. In collaboration with Quanser Inc., we are developing a small co-axial helicopter which is highly maneuverable and completely autonomous. Leveraging recent theoretical advances in the field of nonlinear control, we will also design an autopilot to make the aircraft perform complex maneuvers without any human intervention.

Mann, Steve — Augmented Reality Will Never Work, and That’s Why We Need Mediated Reality

www.eyetap.org/research/medr.html

Augmented Reality, whether by handheld iPhone applications that we’ve developed in our lab and elsewhere, or by eyePhone (electric eyeglasses), has been shown to be problematic by causing 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. This research emphasizes fundamentals of physics, computer science, and engineering. This research is also coupled closely with the undergraduate and graduate course, ECE516 http://wearcam.org/ece516/.

Mann, Steve — Brain-Computer-Interaction (BCI) and EEG-based Cyborg Technologies

InteraXon.ca

Brain-Computer-Interaction (BCI) systems developed as part of the wearable computing and cyborg technologies efforts 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 http://www.eyetap.org/publications.

Mann, Steve — Comparametric Equations and High Dynamic Range (HDR) Imaging

wearcam.org/comparam.htm

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’s applications in extending the dynamic range of imaging devices such as electric eyeglasses and portable cameras and cellphones, etc.. See ‘Comparametric Equations with Practical Applications in Quantigraphic Image Processing’, IEEE Transactions on Image Processing, Vol 9, No.8, August, 2000, which you can download fromhttp://www.eyetap.org/publications.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Mann, Steve
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 display for computer-mediated reality that achives augmented reality but also goes beyond it, to not only augment, but to also modify, and thus help people see better, find their way better, etc.. Not only the wearable face-recognizer that puts virtual nametags on people, etc., but also the mediated vision that helps people see better and improves their personal safety. See http://www.eyetap.org and http://glogger.mob.

Mann, Steve
glogger.mobi

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 fisheye lens and various sensors (http://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. Others such as DARPA, HP Labs, and Nokia have also been building on this lifeglogging work which is also known by many other names such as lifelogging, lifeblogging, CARPE, or lifestreaming. We now have a community of more than 80,000 'cyborgs' on http://glogger.mobi and research continues into the mobile multimedia iPhone apps, as well as versions built inside the eye socket of the blind.

Mann, Steve
FUNtain.ca

Musical Instruments and other Human-Machined 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 hydraulophone; take a look at this vid and then imagine the possibilities when we miniaturize it and use air instead of water: http://vimeo.com/14018088.

Mann, Steve
wearcam.org/absement/examples.htm

Physics-Based Modeling 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 http://wearcam.org/absement/.

Mojahedi, Mo
www.waves.utoronto.ca/prof/mojahedi/mo.html

Engineering the Electric and Magnetic Dispersive Responses of Artificial Media

Much 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, materials and systems' behaviours 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 delays, indices, or velocities to characterize the same dispersive effects. Despite these different nomenclatures, we may note that fundamental and important relations exist among the various delays, indices, and velocities. Dispersion Engineering paradigm formulates our attempts to control and manipulate the aforementioned various delays, indices, or velocities–the dispersive effects–by synthesizing artificial materials and designing novel systems, which 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 use 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.

Mojahedi, Mo
waves.utoronto.ca/prof/mojahedi/mo.html

Nano-Plasmonic and Nano-Photonic Devices

Rapidly performance of computers is expected to reach its fundamental limits in terms of speed, bandwidth, power consumption, and electromagnetic interference. Partially, problem lies in degrading performance of electrical interconnects. Unlike transistors, for which their functionality increases by miniaturization, electrical interconnects functionality degrades substantially with miniaturization. It has been suggested to replace the electrical interconnects with optical interconnects where the latter do not suffer from signal latency, limited bandwidth, or high power consumptions as 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, is 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, similar to 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

Bandwith Efficient DRAM Controllers in Non-Coherent Systems

Embedded and mobile hand held devices have been proliferating enabling applications that were not possible 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, and digital photography. As their capabilities are increasing, novel applications will be possible such as health monitoring. 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 and at what energy cost the memory can provide. 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.

Moshovos, Andreas
www.eecg.toronto.edu

Exploiting Multi-Megabyte On-Chip Memory Hierarchies

Several technology and application trends favor 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 multi-program and parallel/multithreaded workloads. They can also be used as the building blocks for shared multiprocessors (SMPs). 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, multi-program 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 multi-megabyte caches.

The key questions addressed by this research are: (1) How do we manage these multi-gigabyte caches; are the techniques that are currently used still adequate, or is there room or need for rethinking these decisions? (2) Can we exploit this tremendous wealth of on-chip storage to further optimize performance above and beyond of 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 for achieving 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 behavior that can be used to anticipate and optimize for future program demands.

Moshovos, Andreas
www.eecg.toronto.edu/~moshovos

Power-Aware Cache-Based Structure Design

Computing devices comprise processing elements that process digital information and memory elements for storing digital information. Due to 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 devices. 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% of a modern processor) and due to semiconductor technology trends, their power dissipation is expected to increase in relative terms. Accordingly, there is a need for developing techniques to reduce their power 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 accesses to memory from devices without intervention from the processor (this is necessary to achieve adequate performance). Accordingly, there is a need to understand and develop caching mechanisms for supporting this aspect of virtualization. In addition this project investigates the caching and communication architecture for fused CPU and graphics processor systems.

Nachman, Adrian
www.eecg.utoronto.ca/~sorinv/mm_wave_lab.html

Millimetre-Wave Imaging System

This is a joint project with Prof. S. 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.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Nachman, Adrian — 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 Prof. M. 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 — Power Grid Verification

www.eecg.utoronto.ca/~najm

With increased power dissipation and reduced supply voltage, large microprocessors today 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 these extremes of 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 do not lead to functional failures. We are aware of at least two industrial instances, a DSP core and a large microprocessors, 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 is very time-consuming and expensive, so that it is often incompletely done, or not done at all. We are developing efficient techniques for verifying that the voltage on the power grid is safe and within user specifications and, in case the grid is found to be unsafe, for redesigning and optimizing the grid to achieve safety.

Najm, Farid — Variations-Aware Timing Verification

www.eecg.utoronto.ca/~najm

In advanced IC design, a myriad of sources of manufacturing process variations have become significant enough to affect chip performance. Likewise, variations of other (non-process) variables, such as supply voltage and temperature, have to be included, leading to the acronym PVT which represents process, voltage, and temperature variations. Of paramount importance is the impact of PVT variations on circuit timing, and the resulting yield loss when circuit timing exceeds certain limits for some parts. While ASIC manufacturers need to meet the specified performance at all corner settings of the PVT parameters, and microprocessor companies can afford to do speed binning around the nominal PVT design point, both camps are concerned about the variability of circuit timing, and the resulting timing yield loss. We are working on verification of circuit timing in the presence of timing variability.

Ng, Wai Tung — Smart Power Integration and Semiconductor Devices

www.vrg.utoronto.ca/~ngwt

Our focus is on the integration of power devices, smart power integrated circuits, and power management systems. We have worked extensively on the development of CMOS compatible HV fabrication processes for automotive and consumer applications in the 40-100V range. Our VLSI power management work includes the demonstration of Dynamic Voltage Scaling systems using various integrated DC-DC converters with soft-switching topology, predictive dead-time control and with dynamically adjustable power transistor size for power conversion efficiency optimization.

Pavel, Lacra — 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, re-configured and re-adjusted with minimal human intervention. The system theoretic 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.

Plataniotis, Konstantinos N. (Kostas) — Privacy Enhancing Face Recognition

www.dsp.utoronto.ca

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 privacy risks, strengthen regulatory oversight, and promote public confidence.

Plataniotis, Konstantinos N. (Kostas) — Privacy Protected Video Surveillance

www.dsp.utoronto.ca

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 with the provision of a secret key, thus negating the need to store the unprotected original video footage in case an incident investigation occurs.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Plataniotis, Konstantinos N. (Kostas) — Signal and Image Processing for Stereoscopic Cameras, Biometric Sensors and Laser Radar Applications
www.dsp.utoronto.ca

We seek to provide support for development efforts in the areas of: (1) Signal-image processing for 3-D imaging applications in the frequency regimes of Visible and Infrared stereoscopic camera systems and Biometrics Sensors, (2) Image and Data fusion for the above multi-sensor systems, (3) Implementation 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 — Microcavity Optical Modulators and Lasers
www.ecf.utoronto.ca/~poon

We are researching microcavity modulators and lasers that can be modulated at rates greatly exceeding the cavity photon decay rate. The devices can be simultaneously ultra- low-power, high-speed, and compact. They may find applications in chip-scale optical interconnects. We are implementing the devices in combinations of silicon, silica, silicon nitride, and InP-InGaAsP.

Poon, Joyce — Nanocavity Optical Modulators and Lasers
www.ecf.utoronto.ca/~poon

We are investigating the physics of operation and implementations of ultra-small optical modulators and laser sources that use metallic structures. Our goal is to realize modulators and lasers that are at least several orders of magnitude smaller in volume compared to dielectric waveguide devices. The small sizes can enable a tight integration with electron devices and high integration density.

Poon, Joyce — Novel Optical Waveguides
www.ecf.utoronto.ca/~poon

We are developing structures and methods of fabrication for optical waveguides. We are designing waveguides in combinations of silicon, silica, and silicon nitride with reduced temperature sensitivities for chip-scale interconnects. We are producing ultra-low-loss, tapered optical fibers with sub-wavelength waist diameters using acid flows driven by surface tension. The fiber tapers have applications in device coupling, sensing, and nonlinear optics.

Prodic, Aleksandar — Power Management and Integrated Switch-Mode Power Supplies
www.ele.utoronto.ca/~prodic

Traditionally, low-to-medium switch-mode power supplies (SMPS), used in cell phones, computer systems, communications, vehicles, medical devices, and other application consuming power from a fraction of watt to several kilowatts have been controlled by analog means. Mostly, due to the 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 SMPS, and is currently focusing on fully utilizing the digital control advantages as well as on the development of novel converter topologies.

Examples include digital controllers that offer novel features such as plug-and-play operation, system identification, ultra-fast response to load transient minimizing overall system size, and on-line efficiency optimization.

Qian, Li — Fiber-Optic Sensing
www.ecf.utoronto.ca/~qianli

We utilize photonics technology to create instrumentation for fiber-optic sensing and metrology. Our frequency-shifted interferometry technique has been demonstrated to have a variety of applications, such as dispersion measurement, fiber length measurements, multi-point optic sensing for stress and/or temperature sensing used for civil structures, multi-point chemical gas sensing for environmental monitoring as well as industrial monitoring in hazardous environments, and liquid level sensing in cryogenic environment required in space applications.

Qian, Li — Nonlinear Optical Devices– Ultrafast Switching and Frequency Conversion
www.ecf.utoronto.ca/~qianli

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 fiber-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.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Qian, Li — Quantum Communication

www.ecf.utoronto.ca/~qianli

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 implementable with existing fiber-optic technologies. These include a specialized optical homodyne detector for quantum key distribution, a fiber-based entangled photo pair source, and fiber-based quantum key distribution systems.

Rose, Jonathan — Acceleration of Protein Co-Evolution Detection for Determining Interactions in Software and on FPGAs

www.eecg.toronto.edu/~jayar

The MatrixMatchMaker algorithm was recently introduced to detect coevolution between proteins using pairs of phylogenetic distance matrices, and has numerous advantages over existing methods of coevolution detection. Co-evolution detection is a method for determining which proteins interact. We have developed a revised algorithm, which recasts the coevolution problem as multiple maximum clique subproblems on a graph of protein pairs, and achieves a 400x speedup with comparable accuracy to the original algorithm. We are proceeding to accelerate this algorithm further, by building special purpose computation hardware on FPGAs.

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 $1000. As well, 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 as the key technology for medium-volume systems.Our research explores FPGA architectures, focusing on heterogeneous architectureswhich 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.

Rose, Jonathan — 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 software to become a complete flow from the Verilog Hardware Description Language through to placement and routing, and timing analysis. 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.

Rose, Jonathan — Computer Vision Acceleration with FPGAs on Mobile Devices

www.eecg.toronto.edu/~jayar

Modern Mobile devices are revolutionizing how we interact in our environments. A key technology that would enhance this interaction would be to have a smartphone that could see in the way that humans do, as is the goal of the Computer Vision research field. Much of that research does not seek methods that operate usefully in real-time on powerful computers, much less the computers on mobile devices. In this project we seek to explore the use of FPGAs as computational accelerators for vision problems - such as object recognition or object tracking - on mobile devices.

Rose, Jonathan — Overlay Architectures for FPGAs

www.eecg.toronto.edu/~yiannac/VESPA

There are two crucial issues facing the use of FPGAs today: first, the amount of time it takes to compile a design into the FPGA, and second, the difficulty of learning the programming model implied in typical digital systems design. This broad research project is engaged in several methods of dealing with these two problems. First, we are looking at many methods of building soft processors (processors built on an FPGA fabric) that are much easier to program, as it is a programmable overlay on top of the programmable fabric. Second, we are beginning to explore methods of synthesizing pre-synthesized, placed and routed modules that reduce the need for time-consuming optimization.

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

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.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Sargent, Edward — Low-Cost High-Efficiency Photovoltaics

light.utoronto.ca

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 — Computational Electromagnetics

www.waves.utoronto.ca/prof/sarris

For many years, research on computational electromagnetics has focused on improving the modeling of fundamental building blocks for complex, real-world systems. However, translating our detailed knowledge of component parts to their impact on the behaviour of the systems they belong to, defines a new frontier for computational physics. This frontier of multi-scale modeling involves the integration of potentially heterogeneous models over the wide range of scales present in most physical problems. Our research is inspired by these challenges to address fundamental questions, devise novel techniques and investigate their application in critical areas of electromagnetics, optics and communications. Hence, we conduct basic research on novel multi-scale/multi-physics computational methods for electromagnetics from microwaves to the nanoscale. Applications of interest include wireless channel modeling, wave-propagation in meta-materials, plasmonics, carbon nanotube/graphene based structures, electromagnetic compatibility/interference (EMI/EMC) problems and modeling under uncertainty.

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 so as 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 if fixed. The other one is a thin layer, also called free layer, whose magnetization can be switched between a direction that is either parallel or anti-parallel to that of the fixed layer. This corresponds to storing a digital 1 or 0 in the cell. Reading the stored bit is achieved by examining the resistance of the cell. This resistance is low for the parallel state and high for the anti-parallel state.

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

Sheikholeslami, Ali — High-Speed Wireline Signaling

www.eecg.utoronto.ca/~ali

This research targets circuit design for high-speed chip-to-chip signaling, backplane signaling, and optical communication. It includes circuit design for transmitter blocks (such as muxes, pre-emphasis, and drivers), and receiver blocks (such as equalizers, and clock and data recovery).

Smith, Peter W. E. — Ultrafast Photonics

www.ecf.utoronto.ca/~upl

My teaching and research interests have centered 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 deployment in a random manner by users and the network infrastructure is built and grows organically. This approach to wireless network deployment will greatly reduce the cost of base stations/access points and result in networks with much greater capacity which is required for the emerging broadband wireless services. This vision imcompases the current industry developments referred to as femtocells. The work is also related to what is referred to as cognitive radio.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Steffan, Gregory — Making Programming Multicores Easier

www.eecg.toronto.edu/~steffan

The microprocessor technology road-map predicts a future with tens to hundreds of processors per chip and beyond, but with limited clock frequency improvements and potentially simpler individual processors. We are developing compiler and analysis tools for making these multicore processors easier to program, in particular by exploiting "optimistic parallelism" via support for "transactional memory" (TM). TM provides an easier, optimistic alternative to locks for critical sections–allowing programmers to avoid deadlock and fine-tuning when synchronizing code, and also allowing critical sections to execute in parallel whenever they operate on independent data. We have designed and evaluated new software and hardware systems for TM support, targeting both conventional multicore systems as well as FPGA–based multicore systems.

Steffan, Gregory — Overlay Architectures for FPGAs

www.eecg.toronto.edu/~steffan

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 their economics are often vastly superior to the use of 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. We have developed several scalable, customizable, and easy-to-program overlays for different computing domains, including customizable soft processors, soft multiprocessors for packet processing, soft vector processors, and current work on a GPU-like engine that will support programming with OpenCL.

Stumm, Michael — System Software Performance Optimizations

www.eecg.toronto.edu/~stumm

Our primary objective is to make improvements to operating systems so as to 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 exception-less system calls to Linus, 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 for 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 seconds) and secondly, testing of the UKF estimator with acmeasurements.

Tate, Joseph (Zeb) — Phasor Measurement Unit Data Characterization and Compression

www.ele.utoronto.ca/~zeb

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 to be made 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.

Tate, Joseph (Zeb) — Power System Simulation using Programmable Graphics Processing Units

www.ele.utoronto.ca/~zeb

Modern power grids are comprised of millions of individual nodes, which are interconnected in a variety of ways (e.g., transmission lines, transformers, switches, power electronic converters, etc.). At each node, there may be anywhere from zero to hundreds of sources and sinks of electric power, each with its own complex model (e.g., consider a modern household which may include slowly varying electrical loads such as clothes washers 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 behavior 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 & operators to consider a wider variety of events and/or more detailed modeling of power system components, and this has led to a renewed interest in algorithm design & implementation for power system simulations. This project focuses on the development of power system simulation software which targets a particular parallel computer architecture-programmable graphics processing units (GPUs). These processors pose several unique challenges due to 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 solve 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 the utilization of 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 which is based on GPU-accelerated computing.

Tate, Joseph (Zeb)
www.ele.utoronto.ca/~zeb

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, 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., 4 hours into the future) so that operators can have a better understanding of potential problems on the network. Because the potential impacts on the grid are heavily dependent on both the levels of wind generation & 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.

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 light-weight 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 based battery technology until very recently; despite vastly superior energy density compared to Ni-MH batteries used in the first generation 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 3 times higher then 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 been recently developed as a complimentary energy storage technology to 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 then Li-Ion batteries. On the other hand, the 6 Wh/kg specific energy of these Ucaps is at least 10x worse then Li-Ion batteries, leading to the concept of using a hybrid storage system comprised of a smaller Li-Ion battery and an 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 are to develop new models, control schemes and power electronic converters to extract the maximum performance from modern EV energy storage systems.

Trescases, Olivier
www.ele.utoronto.ca/~ot

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

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Trescases, Olivier — Power Converters for High-Efficiency LED Lighting

www.ele.utoronto.ca/~ot

Electric lighting accounts for approximately 11 % of the world 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 (SMPS) to optimize the lighting efficiency and chromatic properties in HB LED systems using closed loop thermal and optical feedback. The targetted controller can efficiently regulate the electrical-to-optical energy conversion process under a wide range of operating conditions.

Trescases, Olivier — Power Electronics for Photovoltaic Applications

www.ele.utoronto.ca/~ot

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 price 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 due to government incentives and multi-disciplinary technological advances. The exponential growth of PV technology presents a 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.

Truong, Kevin — Computational Tools for Protein Sequences, Structures and Networks

apel.ibbme.utoronto.ca

Cells are composed of protein signaling 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 signaling networks that will allow us to one day manipulate cell biology with the same precision as electrical circuits and computer networks. To accomplish this goal, my proposal will focus on developing computational tools for studying protein sequences, structures and signaling 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. Next, to study the protein signaling kinetics within cells, fluorescent protein biosensors are powerful tools but the design of these biosensors is often trial and error. Using a computational tool to model the conformational space of protein biosensors, I improved the design however the tool was not quantitative. To address that problem, I will include molecular factors that select preferred biosensor conformations. Lastly, 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 signaling networks. Together this work will yield insights into protein sequences and their networks that will ultimately aid in developing therapies for human illnesses.

Truong, Kevin — Live Cell Imaging and Control of Caspase Kinetics Using Engineered Proteins

apel.ibbme.utoronto.ca

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, it accentuates 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 addressing a series of three aims: first, targeting caspase biosensors to subcellular organelles (Aim 1); second, imaging caspase cascades in living cells (Aim 2); finally, engineering proteins to control caspase activation based on XIAP (X-chromosome-linked inhibitor of apoptosis protein) and a Ca2+ binding protein called calmodulin (CaM) (Aim 3). This work will pioneer designs for engineered proteins that will provide new tools for fundamental studies in cell biology.

Valaee, Shahrokh — Localization of Wireless Terminals in Indoor Environment

www.comm.utoronto.ca/~valaee

Location based services (LBS) are emerging as new applications on mobile phones. In LBS, the main challenge is to locate the user, especially in indoors and covered areas where the GPS service is not available or has unacceptable error. In this research we estimate the location of a mobile phone using the strength of signals arriving from WiFi access points. We have designed and developed the system on three WiFi-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 Canadian National Institute for Blind (CNIB). The system has a tracking and navigation system that uses voice instruction to guide individuals with visual impairments to find their way in indoor environments.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Valaee, Shahrokh — Wireless Communications in Vehicular Environment

www.comm.utoronto.ca/~valaee

In 2000, an estimated 6,394,0000 motor vehicle crashes were reported to the police in the US. Based on a series of in-depth investigation of police reports and on-scene investigations, human factors were found to be the definite cause for 70.7% of the crashes. Most of those accidents could have been prevented if there were proper devices installed in cars that allowed vehicle-to-vehicle (V2V) and vehicle-to-roadside (V2R) communication. 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, Test 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 in part attributed to advances in the Electronic Design Automation industry community that develops Computer-Aided Design (CAD) tools to aid the engineer designing such complex high-performance devices. The research of our group centers around 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 it has been awarded in premiere conferences for its impact to the community.

Venetsanopoulos, Anastasios — Research in Telecommunications, Signal and Image Processing, Multimedia and Biometrics

anv@dsp.utoronto.ca

Four research topics: Multimedia (image compression, image and video retrieval); Digital Signal/Image Processing (multichannel image processing, nonlinear, adaptive and M-D filtering, knowledge based processing and recognition, 3-D imaging, biomedical applications); Telecommunications Biometric Research. http://en.wikipedia.org/wiki/Anastasios_ Venetsanopoulos

---

Voinigescu, Sorin — Digitally-Enhanced Analog Equalization Techniques for 50-110 Gbps Wireline Applications

www.eecg.toronto.edu/~sorinv

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 to realize arbitrary waveforms for transmitter equalization directly at frequencies in the 50-110 GHz range. In receivers,oversampling with 100+ GHz clockwill be employed to recover the clock and data information and to recondition the signal before transmission.

Voinigescu, Sorin — High Efficiency mm-Wave Transmitter Array

www.eecg.toronto.edu/~sorinv

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 45nm SOI technologies are being investigated for the practical implementation.

Voinigescu, Sorin — Silicon SoCs in the 100-500 GHz Range

www.eecg.toronto.edu/~sorinv

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 un-chartered 100-500 GHz band.

Envisioned applications are in industrial sensors, automotive radar, active and passive imaging, remote sensing, environmental monitoring, radio astronomy and very high data rate 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 three integrated systems that have already been demonstrated in the lab are: 1) a 120 GHz high resolution position sensor with above IC antenna fabricated in a SiGe BiCMOS process and which works over 2-5m, a) 160 GHz radio transceiver array with on-die antennas operating as a 4Gb/sradio and as a Doppler radar, and 3) a 165 GHz passive imaging receiver with sub 0.3 K temperature resolution and consuming 82 mW.

Wong, Willy — Sensory Neuroengineering

individual.utoronto.ca/willy

My interests are in the area 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.

LEAD RESEARCHER	RESEARCH TITLE	COLOUR KEY

Wonham, Murray — Supervisory Control of Discrete–Event Systems

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 modeling 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 pushbutton systems, and industrial diagnostic systems.

Yu, Wei — Cooperative Wireless Cellular Networks

www.comm.utoronto.ca/~weiyu

My research focuses on the use of cooperative communication techniques for capacity improvement in wireless cellular networks. We design future wireless networks in which base-stations can cooperate with each other and where relays are deployed within each cell to enhance the performance of mobile terminals. From a theoretical perspective, we investigate the information theoretical capacity of relay networks and networks with transmitter or receiver cooperation. From a practical perspective, we design optimal scheduling, beamforming, power control, and multiuser detection strategies for multicell networks. Cooperative techniques promise to enhance cellular coverage and to mitigate interference, thereby improving the overall performance of the cellular network.

Zhu, Jianwen — Compiling Software to Silicon

www.eecg.toronto.edu/~jzhu

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

Research students in the lab with Professor Joyce Poon (centre), Assistant Professor, The Edward S. Rogers Sr. Department of Electrical & Computer Engineering, University of Toronto.

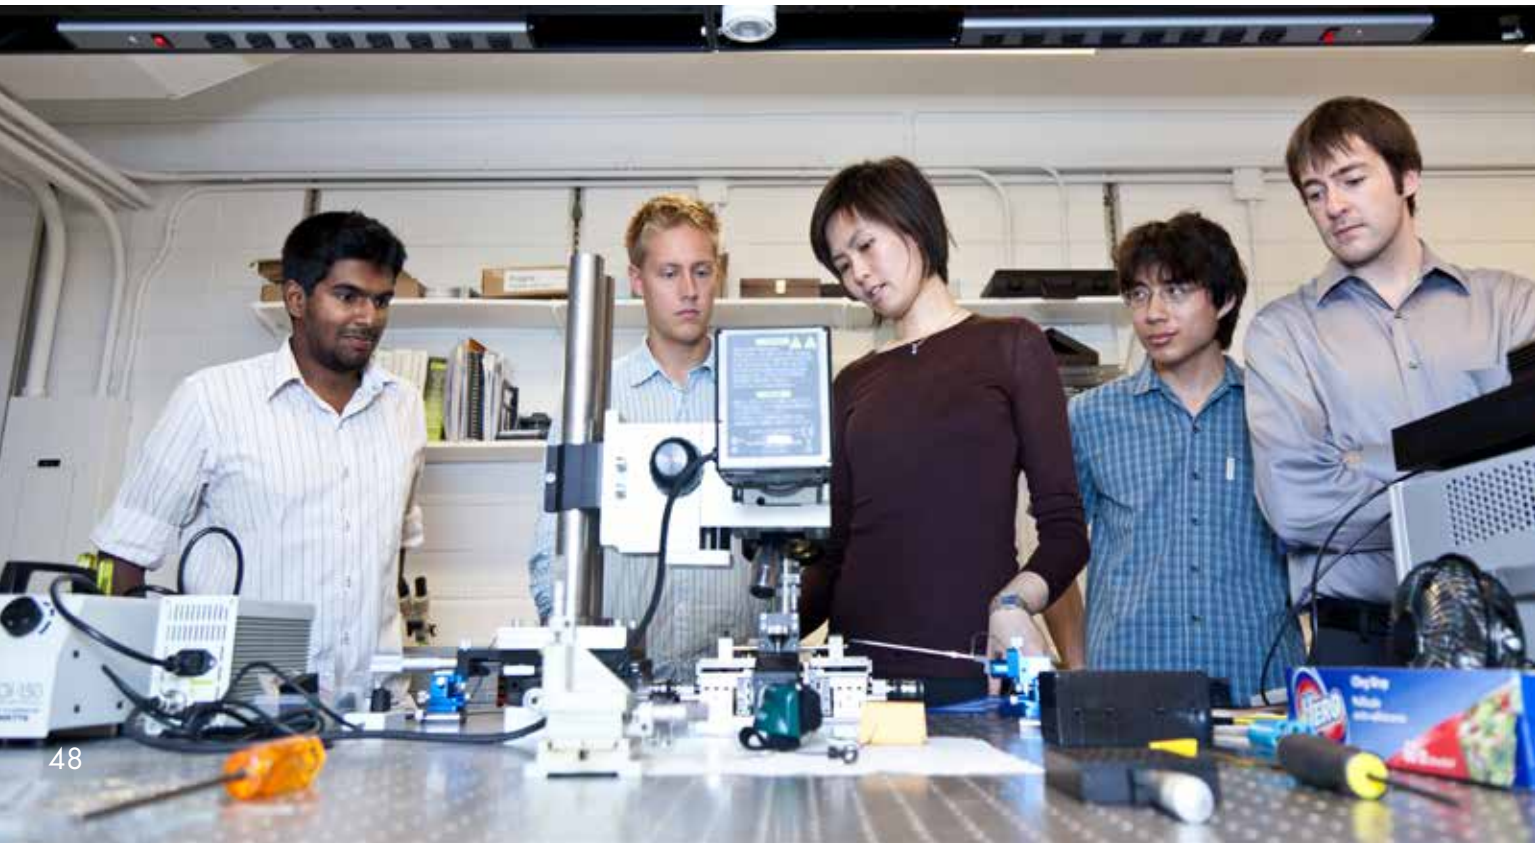

The Edward S. Rogers Sr. Department
of Electrical & Computer Engineering
University of Toronto
10 King's College Road
Toronto, Ontario, M5S 3G4
Canada

Telephone: 416-978-1801

www.ece.utoronto.ca

General Inquiries: susan.grant@utoronto.ca

Design: Iconography Art+Design
laurielabelle@mac.com

Photography: Jason KB Photo

University of Toronto campus cover photo:
Mark Balson,
Engineering Strategic Communications

Quantum Corral image was originally created
by IBM Corporation.

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