SciNet

PORTAL

vol.1

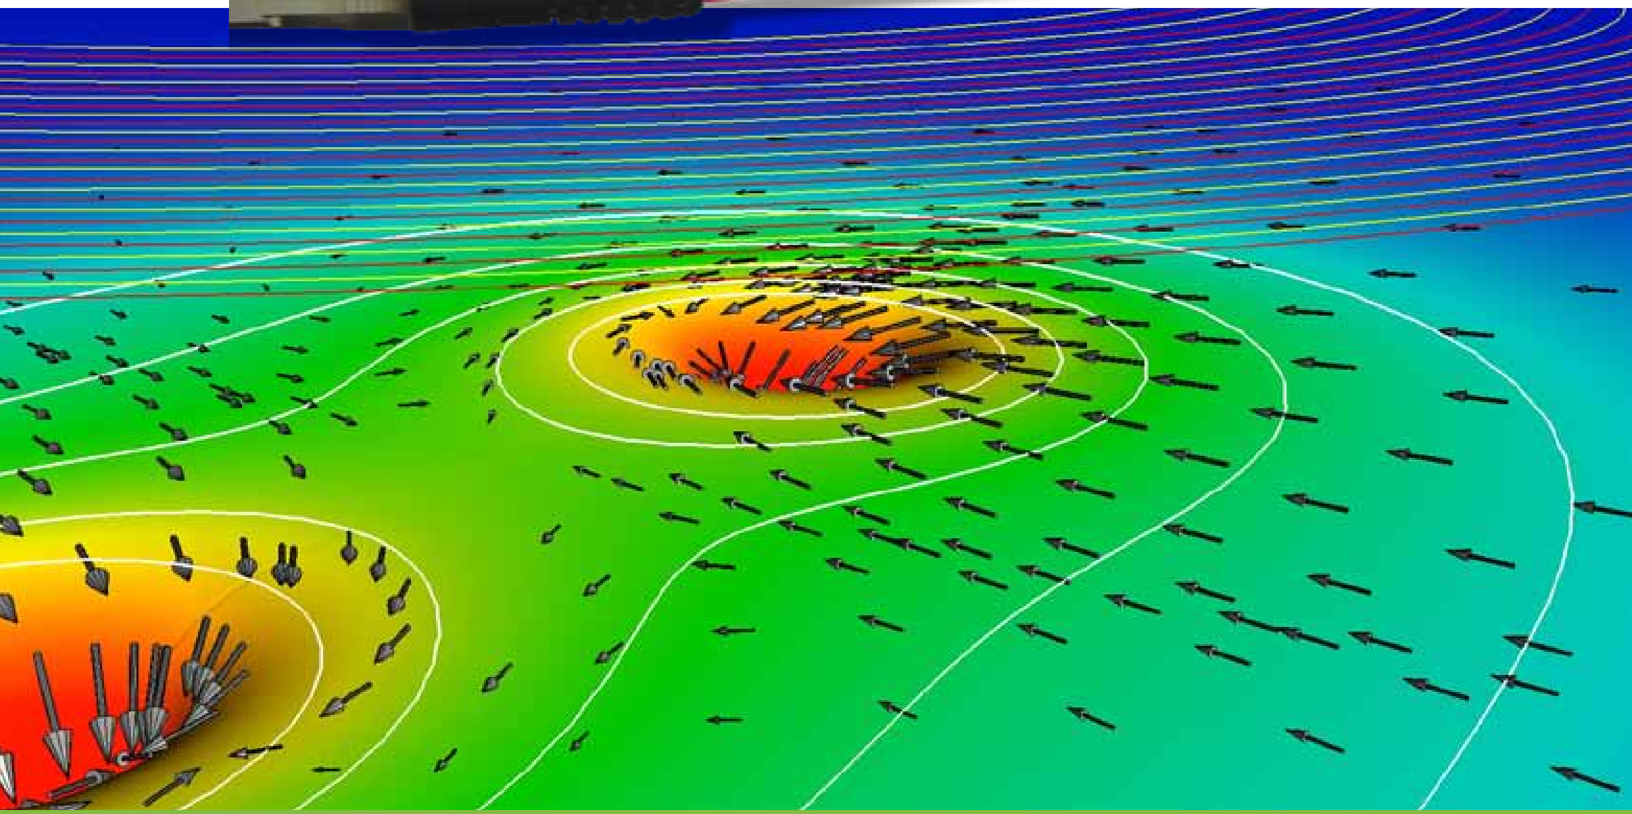

CANADIAN SCIENCE AT SCALE

scinet.utoronto.ca

compute • calcul canada

IN THE PORTAL

Credits

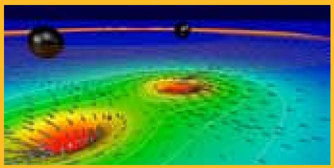

Cover: Harald Pfeiffer, Assistant Professor, CITA; still image from movie of inspiral, merger and ring-down of two non-spinning holes

Right: NASA, JPL-Caltech, WISE Team; WISE: Heart and Soul Nebulas in Infrared

DIRECTOR'S MESSAGE

A Short History of Canadian Supercomputing

The road towards the development of a fully competitive Canadian capability in High Performance Computation has been "long and winding". It began in 1983 with the establishment by NSERC of a Committee on Vector Computer Access which I chaired and which enabled the use by NSERC funded scientists and engineers of the vector facility then operated by Environment Canada at Dorval, Quebec. A next step was the establishment of the Ontario Centre for Large Scale Scientific Computation (OCLSC) at the University of Toronto in 1986. Although this facility was highly successful locally and somewhat successful nationally, the lack of the existence of sufficiently high bandwidth connectivity to the other universities in the province and elsewhere in the country meant that "external" users were inadequately served. This facility was therefore not re-financed at the end of its mandated 5 year period of operation. Recognizing the importance of HPC to the provincial innovation agenda, however, the province established a new organization called High Performance Computing Ontario (HPCO) on July 8 1992, and provided it, in principle, with an allocation of $29.4M over 5 years. It appointed a Board of Directors for the organization, of which I was a member, and charged it to decide upon the equipment to be acquired and to draft a business plan as to how the system would be operated. Unfortunately this Board proved dysfunctional and was unable to decide upon a way forward. The money was left unspent-on the table!

The establishment of the Canada Foundation for Innovation in 1990 re-energized the effort to establish Canadian HPC as it began to provide significant funds with which such equipment could be acquired, a mechanism which proved effective and which led to the acquisition, both by individual universities and Consortia of universities, of a wide range of computing equipment which was dispersed across the Canadian research landscape. At the University of Toronto in both 1995 and 1999 I led applications to CFI that funded three different computing systems. The first award provided funding for systems for the Canadian Institute for Theoretical Astrophysics, the Chemical Physics Group, and for the Atmospheric Physics group. The second CFI grant 4 years later delivered systems for High Energy Experimental Physics, for the Institute for Aerospace Studies and for an upgrade to the system of the Atmospheric Physics Group.

While Toronto continued to function in this "group centric" way, other universities in the country were organizing themselves into Consortia and began to seek funding from CFI in this collaborative supra-institutional manner. At the same time the C3 organization was established as a lobbying organization to promote the necessity of more focused government investment and produced a very influential Long Range Plan. In 2005 this led to an HPC Town Hall meeting in Ottawa chaired by the CFI President at which the possibility of the community developing an omnibus proposal for national support of HPC was discussed. Given the evident community support for this approach at the meeting, CFI established a "National Platform Fund" and charged a "National Initiatives Committee" consisting of HPC proponents representing existing and newly established consortia to prepare a proposal for funding the national effort in this area. The University of Toronto based SciNet consortium was established in anticipation of the NPF call. The proposal written by the NIC was submitted in the summer of 2006 and the award announced formally in December of 2006. Our current situation derives from the decisions that were made by the NIC in preparing its application to CFI . Because 3.5 of the now existing 7 consortia had only recently been funded, it was decided that this subset of the community would wait for funding from an expected NPF-2 to significantly refresh their equipment. The $120M provided by equally by the provinces and the CFI, was therefore awarded to the newly founded SciNet consortium (University of Toronto and associated teaching hospitals), Westgrid (all research universities from Manitoba westwards), CLUMEQ (McGill University & University of Laval), and RQCHP (University of Sherbrooke & Universite de Montreal). The remaining consortia, Sharcnet and HPCVL in Ontario, and ACENET in the maritimes received no significant funding in NPF round 1.

At the time of writing this history, the new national organization, Compute/Calcul Canada, that has been established to manage the HPC infrastructure acquired through collaborative CFI and provincial funding, is approaching a further important crossroads. We await the decision from CFI as to whether and how it intends to continue to develop the national HPC platform and to maintain the momentum that has been achieved by Compute/Calcul Canada. Following the Town Hall meetings that have been held across the country to seek community input as a lead-in to the second installment of funding for the national platform, and the international review of CCC that has also been conducted, we are preparing to develop a next application for continuing funding of Canadian HPC. Issues currently on the table include the number of data centres that would be optimal to service community needs nationwide, the number of regional consortia needed in the large provinces of Ontario and Quebec, and the question as to how strongly focused our investments should be on the development of Tier 1 capability as the international community moves from the Petacale to the Exascale. It will take real collective effort if we are to be successful in building upon our recent successes.

Dick Peltier
SciNet and Department of Physics

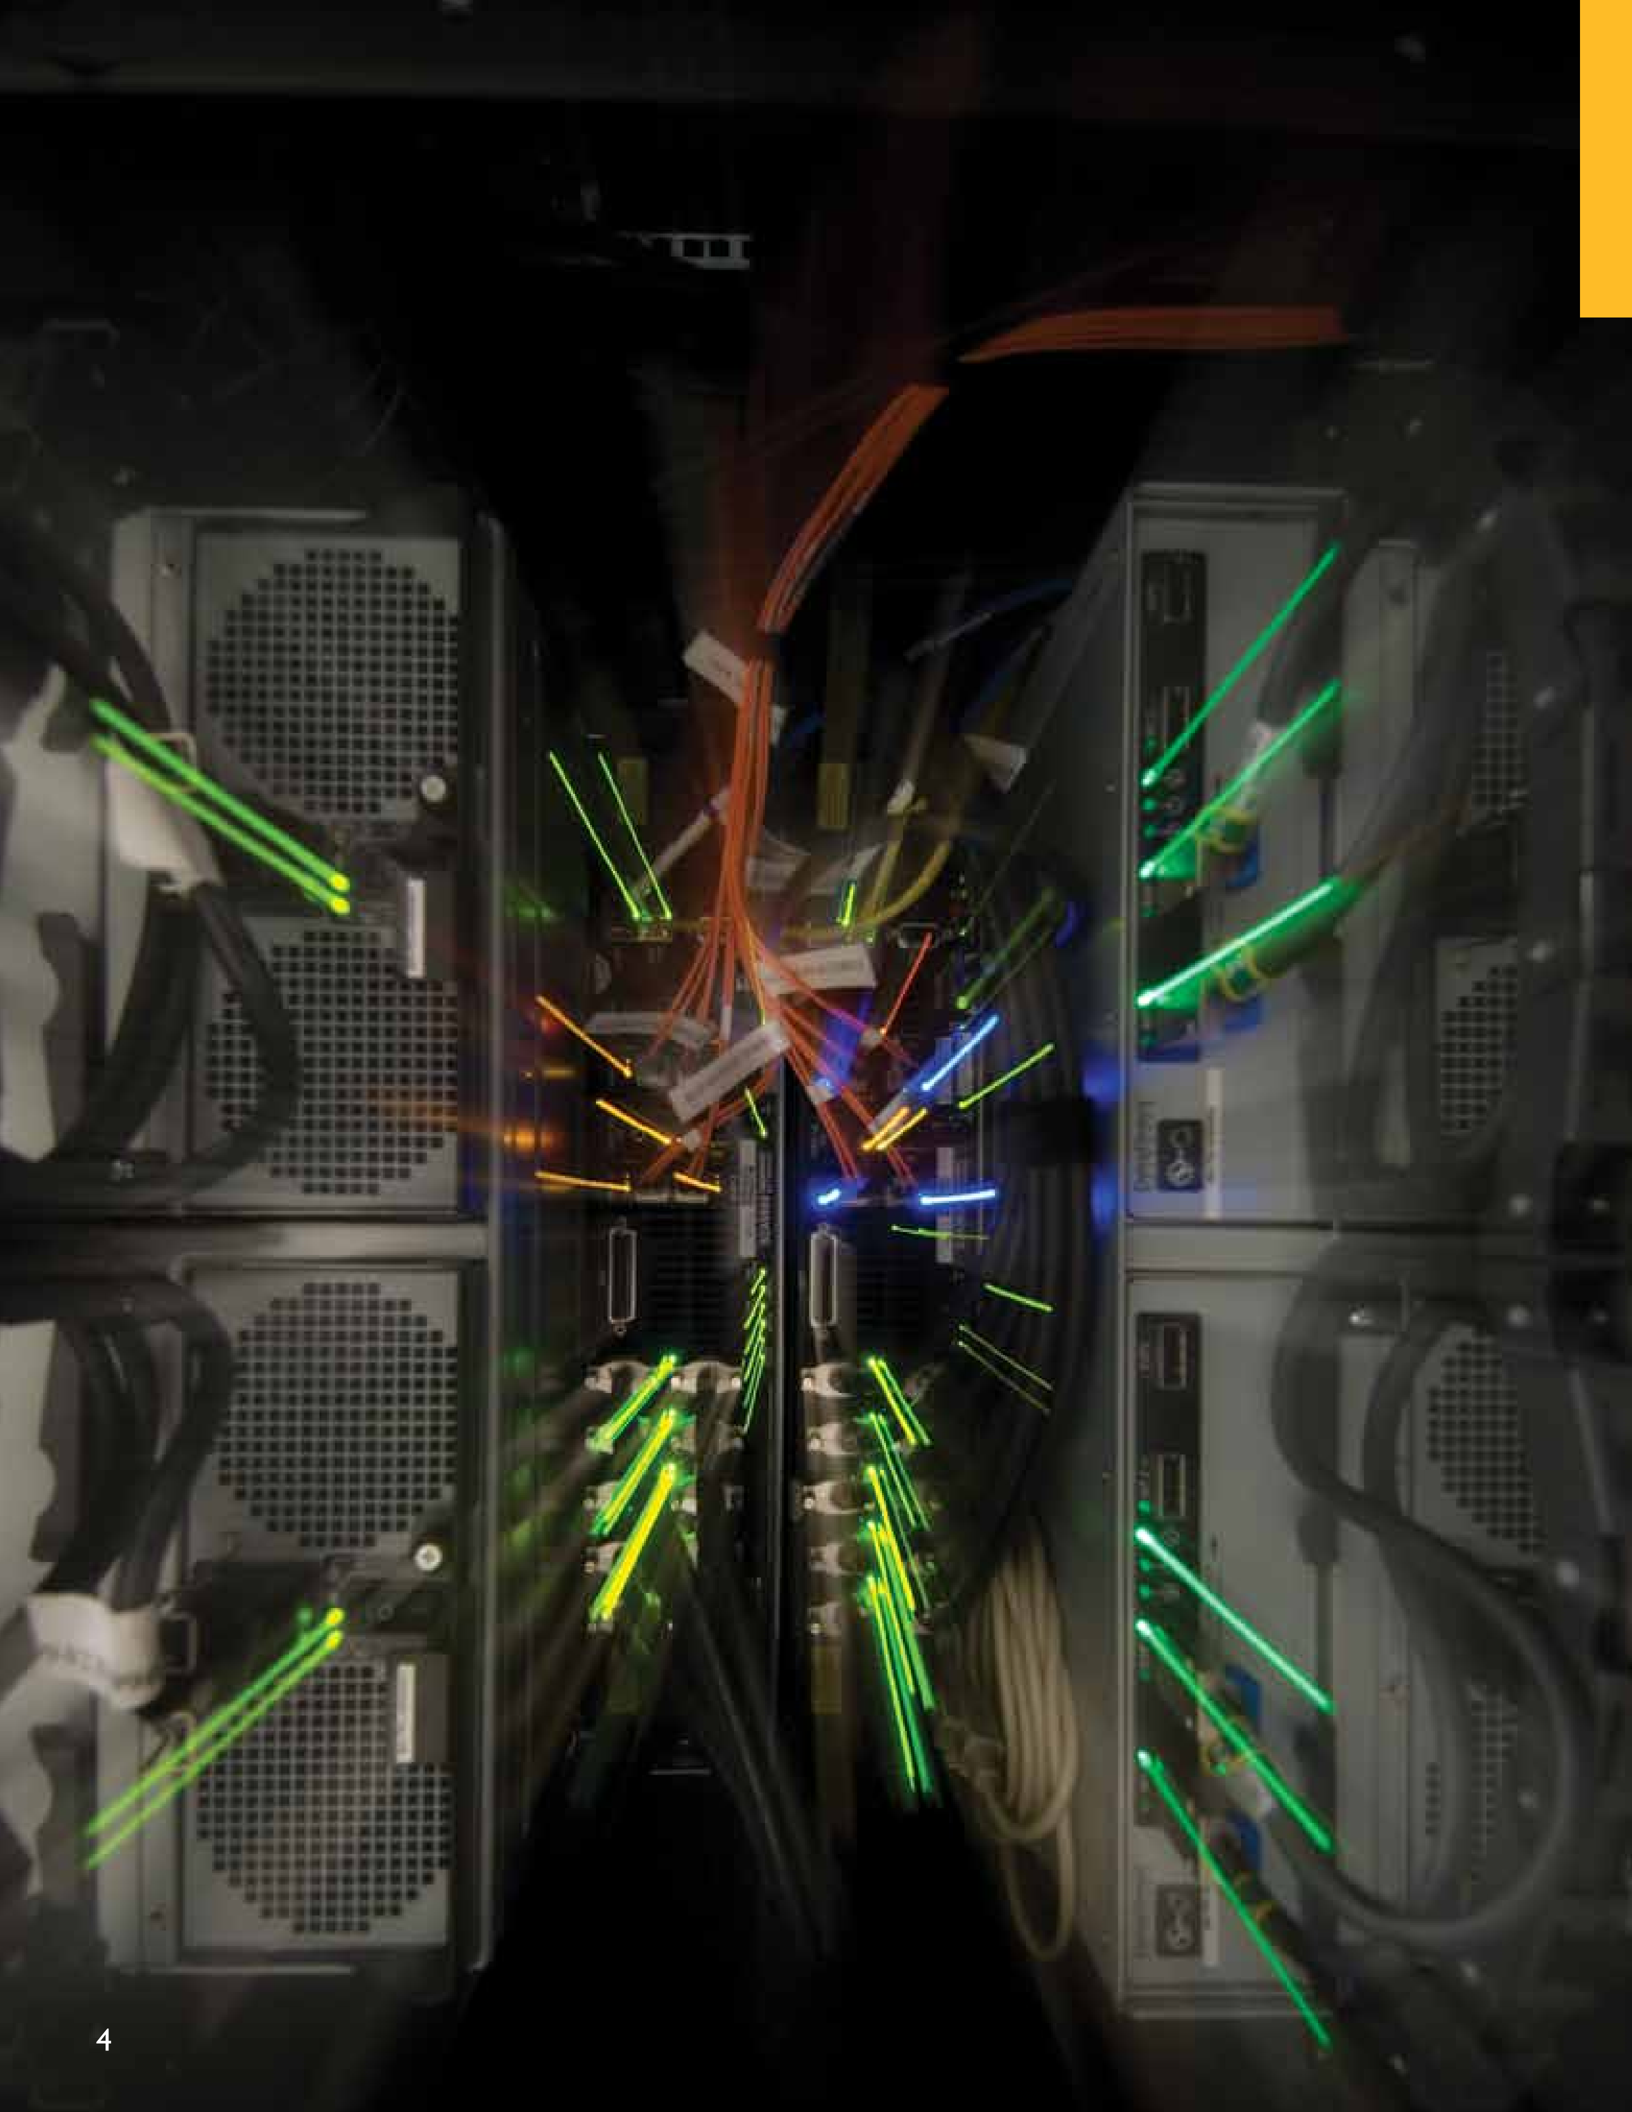

STAFF PROFILES

Systems Administrator

CHING-HSING YU

Ching-Hsing Yu is a Systems Administrator at SciNet, working to maintain optimal system performance. He received his Ph.D. degree in Computational Chemistry from the University of Arkansas in 1999, where he studied conformational properties of small peptide and organic molecules from first principles as well as clay mineral/aqueous solution interfaces using molecular dynamics simulations.

In 2000, he joined Dr. Pomès' research group at the Hospital for Sick Children (Sick Kids) in Toronto, and he was awarded a Restracomp Post-Doctoral Fellowship from the Hospital. As a research fellow, he investigated proton transport mechanisms in membrane proteins, such as D pathway of cytochrome c oxidase and native and chemically modified gramicidin channels where the needed empirical force field parameters were self-developed from ab initio calculations. He has authored 21 peer-reviewed publications, 10 of which as the primary author.

In 2006, he acquired a position at the Centre for Computational Biology at Sick Kids as a systems administrator and application analyst. In this role, he managed the centre's HPC facility, a cluster of 1000+ CPU cores serving the research needs of over 200 scientists primarily from Sick Kids, as well as provided users and applications support.

SciNet was able to bring Ching-Hsing on to their team in December of 2008.

SCOTT NORTHRUP

Parallel Programming Analyst

As a Parallel Programming Analyst at SciNet, Scott Northrup's primary responsibility centres around helping researchers improve their applications to run and scale efficiently on the SciNet high performance computing facilities. He provides advice and support to users on topics such as code optimization, application porting, and parallel implementation. He also works closely with the systems administrators, helping benchmark and test new hardware and software configurations, in an effort to continually improve system performance and reliability. He is very interested in evaluating emerging architectures, and programming models designed to allow HPC at peta and exascales.

Scott has undergraduate degrees in applied mathematics and mechanical engineering from the University of Western Ontario and conducted his masters and Ph.D. research at the University of Toronto Institute for Aerospace Studies (UTIAS). His research has focused on the development of highly scalable efficient parallel numerical algorithms for computational fluid dynamics, specifically applied to combustion. While at UTIAS he gained extensive HPC and systems administration experience designing, deploying, and maintaining multiple generations of hardware including distributed memory clusters.

Scott also has considerable engineering and software development experience outside of academia, working for companies such as Pratt & Whitney Canada as well as for Defence Research and Development Canada (DRDC).

GPC USAGE

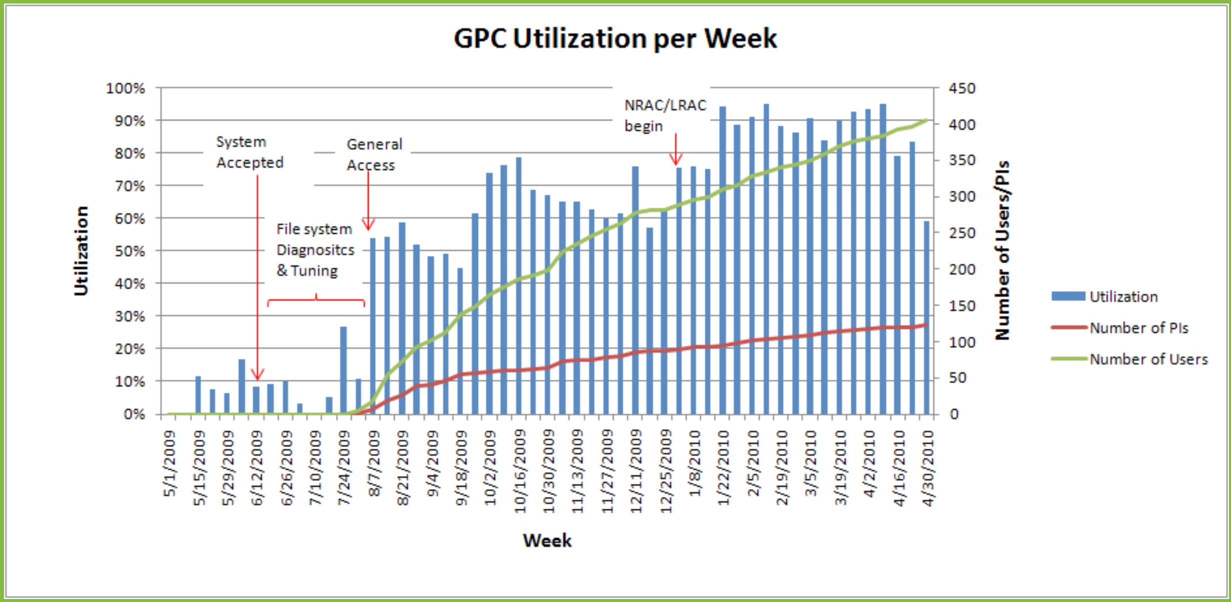

TCS USAGE

SciNet News & Events

September 3 & 10
Intro to SciNet Classes: a class of approximately 90 minutes where you will learn how to use the systems. Experienced users may still pick up some valuable pointers during these sessions.

Second Wednesday of each Month
SNUG + TechTalk
TechTalks are 1/2 hour mini tutorials held as part of the SNUG meetings.
Sept. 8th, The GPFS File Systems
Oct. 13th, Version control - svn, git, mercurial
Nov. 10th, Debuggers & parallel debugging- gdb, ddd, padb
Dec. 8th, Performance and profiling - gprof, scalasca, peekperf

September 22
Full day Intro to Parallel Programming Course

October 6
Full day Parallel I/O Course

May 2011
Exact dates TBD, 5-day Parallel Programming Course: OpenMP, MPI and more will be covered. The aim is for you to leave the course able to start parallelizing your code.

HPC FACILITIES

SciNet has two supercomputers, housed in an energy efficient datacentre.

The General Purpose Cluster (GPC) consists of 45 racks of IBM iDataPlex nodes using Intel's new Nehalem architecture CPUs and runs a linux operating system. This system has been Canada's fastest supercomputer since installation in June 2009. Key specs:

- 306 TFlops peak speed
- #16 in world on June 2009 top500 list
- 30,240 cores
- 60 TB of RAM

The Tightly-coupled Capability System (TCS) is based on IBMs Power 6 CPUs which run at an industry-leading 4.7 GHz. It was Canada's fastest supercomputer in Nov 2008 and is now ranked third in the country. Key specs:

- 60 TFflops peak speed
- #53 in world on Nov 2008 top500 list
- 3,328 cores
- 13 TB of RAM

DataCentre

The SciNet systems are housed in a 12,000 square foot custom-built, state-of-the-art, green datacentre which minimizes energy use through a number of innovative measures including the use of water-cooling, an evaporative cooling tower (eliminating the need for mechanical cooling for one-third of the year) and management software which automatically powers off nodes when not in use. The carefully integrated design of the datacentre and the computer systems saves energy equivalent to that used by 500 households.

Acknowledgements must be made to CFI, NSERC, Ontario, ORF-RE and UofT deans for their funding and support. Without these agencies the existence of Canadian super computing would not be possible at such a scale.

SciNet is one of seven HPC Consortia in Canada. These consortia are part of a national HPC organization called Compute/Calcul Canada. Working in collaboration, Compute/Calcul Canada and the university-based regional HPC consortia provide for overall architecture and planning, software integration, operations and management, and coordination of user support for the national HPC platform.

GALAXY CLUSTERS

From Comets to Cosmology: How the solar system helps explain a mysterious radio glowing of distant galaxies

~ Christoph Pfrommer, Sr. Research Associate, CITA
~ L. Jonathan Dursi, Sr. Research Associate, CITA; Parallel Programming Analyst, SciNet

In 1986 the European Giotto satellite, passing Halley's commet, observed a drape of magnetic field around our most famous visitor. As the ancient comet sweeps through the solar system, the magnetized solar wind that fills space around our Sun accumulates in a sheath around it; the magnetic field lines are stretched around the body like rows of elastic bands. This effect has been seen around other bodies in the solar system -- Mars, Venus, and the Earth; these magnetic drapes might even help bodies without magnetic fields of their own to be partially shielded by damaging radiation from space.

While magnetic fields can't be directly seen -- and there is no way to sprinkle iron filings over a comet! -- their effect on charged particles like electrons can. On encountering magnetic fields, free electrons begin to spin around the field lines and emit `synchrotron radiation' -- radio waves, at frequencies similar to those used for the Wi-fi connection on your laptop.

Christoph Pfrommer is a researcher at the Canadian Institute for Theoretical Astrophysics, at the University of Toronto. His expertise is in galaxy clusters, closely-knit groupings of galaxies, from fifty to several thousands, orbiting a common centre of mass at thousands of kilometers per second. Galaxy clusters are the latest and biggest structures that formed during the evolution of the Universe; so these objects encode a wealth of information critical for nderstanding our cosmic history. Our Milky Way and its neighboring Andromeda galaxy, along with some 30 smaller ones, form what is known as the Local Group of galaxies. This entire group is hurtling toward the center of the nearby Virgo cluster at one million miles per hour.

One of the remarkable facts about galaxy clusters is that most of their ordinary matter is not the stars in the galaxies themselves; it is in the form of a hot, wispy, magnetized gas that fills the cluster. It has temperature of tens of millions of degrees Celsius, so that it emits X-rays which cools this gas.

Dr. Pfrommer was introduced to a galaxy cluster mystery at a conference in Montreal; radio observationsof spiral galaxies in the Virgo cluster showed very strange synchrotron emission at their edges, taking the form of ridges on one side of the galaxy. He explains, ``as these spiral galaxies move through the cluster, they must be splayed by the magnetic field threading the cluster gas, like so many flies on a windshield of a car - and the ridges must be from that magnetic drape, lit up by electrons coming from the galaxies' own stars!"

But how to show this? Jonathan Dursi, a technical analyst at SciNet, knew the answer. ``Using GPC, the largest computer in Canada, it suddenly became possible to do theroretical `experiments' -- to see what would happen if we ran a model galaxy through material like the intercluster medium, and to see if it would generate radiation in the same patterns seen by the observers."

Such a set of `experiments' could only be done at a machine like SciNet's. "These simulations required exquisite resolution", said Dr. Dursi. "The drape formed is very thin, and gets thinner as the galaxy moves faster. The more resolution you need, the bigger the memory requirements become; and the compute requirements grow even faster."

Without Scinet, it's hard to see how this could have been done in Canada; we generated over 32 TB of data, and required almost a half-million CPU hours of computation. Even once the simulations were done, just to generate some of our figures from that data required using 256 processors at once."

Pfrommer and Dursi's simulations not only showed that these `synchrotron ridges' were easily explained by magnetic draping of the cluster gas; there was a more exciting, and completely unexpected result. "Once you see how the magnetic fields are draped over the galaxy, and if you know how the galaxy is moving, you know where the original magnetic field lines were pointing" explained Pfrommer. "And that turns out to be very powerful. From that, we could infer from the ridges not only that draping was occurring, but how the magnetic field in the Virgo cluster is pointed at the location of those galaxies."

"I was blown away when Christoph showed me the plot", said Dursi. "Everyone expects that the magnetic field in a galaxy cluster should be a tangled mess. There was a theory earlier in the decade that suggested the field lines should be pointed more or less outwards, but no one really expected that to be completely true. But there it is; we can see the field lines at these points, and they are shooting straight outwards, Virgo some kind of a cosmic magnetic hedgehog."

Pfrommer and Dursi's results were published in July's Nature Physics, one of the most prestigous physics journals in the world.

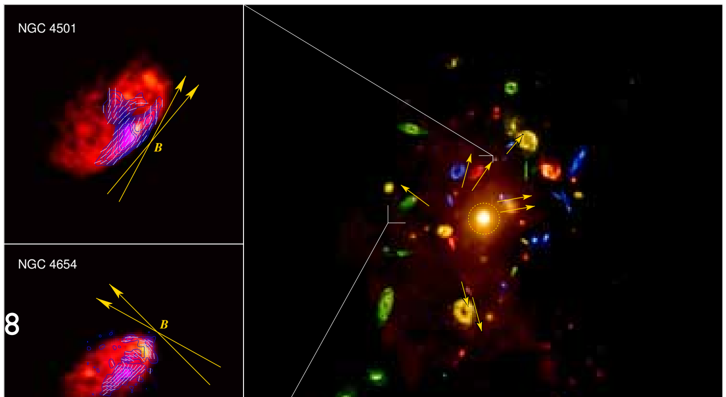

< Magnetic Field in the Virgo Cluster

Self-assembly, defect motion, and rheology of mesophases

~ Jorge Viñals, Professor of Physics , McGill; CRC in Nonequilibrium Materials; Director, CLUMEQ

Macromolecular self assembly of block copolymer mesophases is being actively investigated for the development of a wide variety of nanostructured materials, or as templates, also at the nanoscale. When materials or systems spontaneously self assemble, they adopt configurations or states that have specific properties, spatial configurations, or structures without any external intervention. Such a capability is especially important in tailoring materials with specific properties at the nanoscale.

Spontaneous self assembly, for example, provides an alternative route to conventional lithography, potentially allowing features in electronic devices in the sub-15 nm range, while retaining the reliability and mass production ease that are required in the semiconductor industry. High density storage devices that are built on block-copolymer templates are also being investigated by creating small magnetic wires held by a block copolymer matrix. The density of wires is high enough that if each post could be addressed individually, a maximum storage density of 1 TB per square centimeter of material could be achieved. A third different realm of applications is based on the periodic variation of the index of refraction in a macroscopically order copolymer, leading to the production of photonic band gap materials by using polymer chains of fairly standard length. Devices based on this principle are begin developed with the aim of achieving "all optics electronics". Using light without intervening electronics improves speed and reduces power consumption.

While self assembly is quite efficient at creating locally ordered ("assembled") configurations at the nanometer scale, achieving long range order over macroscopic distances is very difficult. Macroscopically, samples are still disordered because they contain a large number of defects. Such disordered samples are not suitable for the applications outlined above. Controlling the long range microstructure of bulk samples or thin films is therefore the central challenge that needs to be overcome before widespread use of self assembly can be contemplated.

Our computational studies of self assembly focuses on both the specific symmetries into which the materials self-assemble and how the process of self assembly is modified by flows and external stresses, both present during processing of the materials. The goal is to elucidate the lows that govern the formation of order over large distances, how to control or eliminate defects, and in general, how to speed up the formation of uniform samples of controlled symmetry and properties. The figure shows a transient configuration of a so called lamellar phase of a block copolymer during the ordering process. The final equilibrium structure should be a uniform, layered, configuration, with perfectly flat and parallel layers. The intensity of the color shown is proportional to the local polymer composition. In this transient configuration in which the material is ordering spontaneously, locally ordered lamellar regions appear and grow, but with a range of orientations. Hence different partially oriented regions meet at defected boundaries. Theoretical and computational studies are helping us elucidate the types of defects that are possible, how they move in time to reduce the energy of the system, and what external interventions are possible to speed up the appearance of a fully order, uniform, phase. Since the periodicity of these materials is of the order of tens of nanometers, even thin films would contain billions of layers interacting with each other.

Supercomputers are the only tools at the moment to open up a window into the details of the collective motion of these nanophases. In fact, this research project has made use of more than 7 million hours of CPU time on the GPC cluster in just 5 months - on a typical home PC those same calculations would have taken more than 400 years!

v below and top right: images that show a transient configuration of a so called lamellar phase of a block coploymer during the ordering process

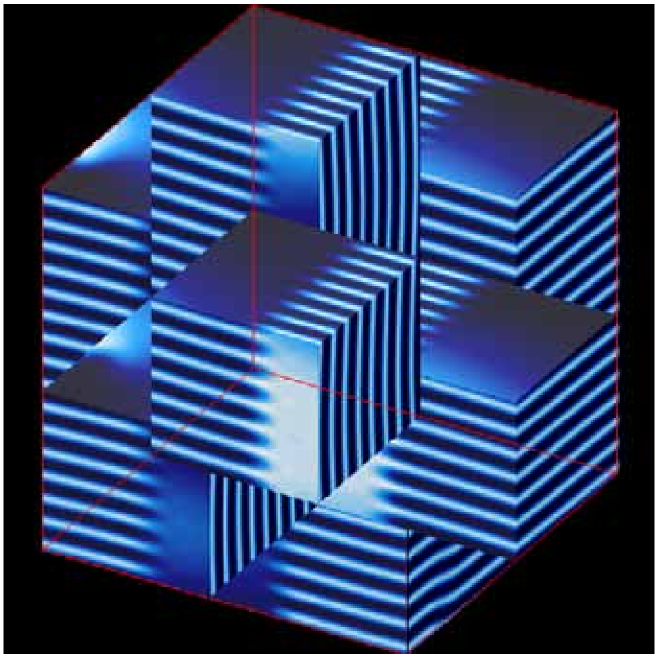

BIO-JET FUELS

Clearing the Air: How HPC is Helping Improve our Air Quality

~ Meghdad Saffaripour, Seth Dworkin, Murray Thomson,;
MIE, University of Toronto

Air travel is a globally common mode of transport that humans have come to depend upon; however, with the rise of fossil fuel costs and growing environmental awareness, the viability of current jet fuel production methods is being called into question. Atmospheric emissions associated with air travel are of growing concern and are receiving a lot of attention due to their adverse environmental impact. Combustion generated carbon dioxide (CO_2) and particulate matter, commonly known as soot or black carbon, are two of the most prominent contributors to global climate change. The question is, what can be done? Researchers at the University of Toronto think at least some of the answers may lie in alternative fuel sources..

With SciNet coming online, a new realm of studies can be performed that are of interest to the commercial aviation industry and of strategic importance to Canada and the environment. While the power generation and ground transportation sectors explore alternative energy sources (hybrid-electric, fuel cells, wind, and solar), commercial airliners will rely on the burning of liquid fuels for the foreseeable future. Renewable bio-jet fuels (derived from plant life) are an attractive alternative to fossil fuels since their associated plant growth removes carbon dioxide from the atmosphere so that they emit significantly less net CO_2.

Three researchers in the Mechanical and Industrial Engineering department at the University of Toronto, Meghdad Saffaripour, Seth Dworkin, and Murray Thomson, are using SciNet to study what would happen if conventional jet-fuels were to be replaced with bio-jet fuels. By using the supercomputing facilities to perform massively parallel simulations of jet-fuel combustion, the emissions characteristics of various fuel blends can be compared. The team, working with fuel and aircraft companies, are interested in knowing the resulting black carbon production if a bio-jet fuel – with subtle but important chemical differences from fossil fuel – were to be implemented for commercial flight.

The research group has spent nearly a decade carefully developing algorithms that are specifically well-suited to parallelization. This has been achieved by coupling a highly-detailed computational fluid dynamics (CFD) module to an implicit chemical reactor simulator wherein the flame is subdivided into thousands of tiny cells (the reactors) in which combustion occurs and particulate matter may be formed. The physically accurate coupling of combustion and fluid dynamics is achieved via an iterative process that globally refines the fluid properties, and permits network communication to share the results of each chemical reactor simulation. The resulting semi-implicit algorithm is scalable and can effectively handle the stiffness of the highly-nonlinear set of partial differential equations that govern fluid dynamics and combustion. In the last 25 years, combustion simulation has been confined to the realm of simple fuels, those with one to three carbon atoms per fuel molecule, including methane, ethylene, and propane. This is partly due to the difficulty in developing the libraries of chemical reactions that are involved in combustion, and to the relationship between fuel and computational cost; the more complex the fuel, the larger the chemistry library, and the more computationally intensive the calculation becomes. Working with combustion chemistry experts at the CNRS in Orléans, France, the group was able to generate a chemistry library that tracks 304 distinct chemical species – twice the size of that used in comparable calculations to date.

With the chemical reaction library in place, the timing of SciNet's arrival proved crucial to advancing the study. In the fall of 2009, the group embarked upon a series of computations, the first of which ran nearly continuously for three months on 192 processes of the Tightly Coupled System. The resulting 43 CPU-year simulation, unique in size and accuracy in the field of combustion to date, predicted the entire flame structure of jet-fuel including temperature profiles, chemical species distributions, and emissions production. The figure depicts the temperature profile of the jet-fuel flame (left side), and the regions and quantities of black carbon production (right side). By lessons learned from the original calculation, and subsequent optimization techniques that have been developed, computation time has been cut in half, down to four to six weeks per computation. The group is currently performing complimentary computations, investigating the effects that bio-jet fuel components might have on the overall efficiency, viability, and emissions of commercial airline travel.

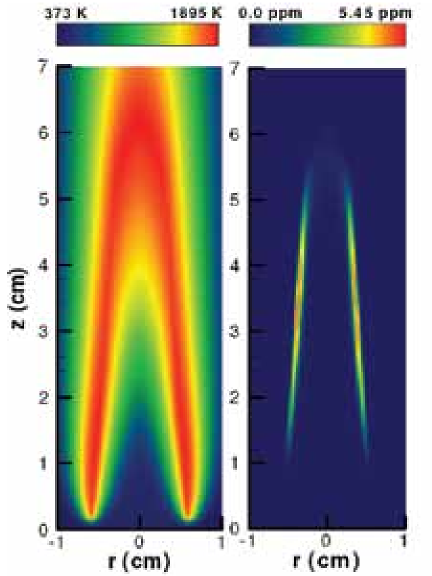

< Computed profiles of temperature and black carbon (soot) formation for a jet-fuel flame.

BIOMOLECULES

Molecular Miracles

~ Régis Pomès, Senior Research Chair, Sick Kids: Associate Professor, Biochemistry, University of Toronto; CRC, Proteomics, Bioinformatics & Functional Genomics

What if surgeons had access to a 'smart' new material that had all the elasticity and strength of human tissue? Such an alternative would result in more resilient prosthetic devices such as artificial skin for burn patients or vascular grafts for heart patients – and provide a life-changing miracle for their recipients.

It all starts with understanding molecules. Dr. Régis Pomès, Senior Scientist, Molecular Structure and Function Programme, Hospital for Sick Children describes his research as basic science -- asking how things work at a fundamental level. "We are like detectives," he says. "Our work is about unknowns, not certainties. In that process, you find new and different paths, and different questions to ask."

Dr. Pomès' approach is grounded in statistical mechanics, which is used to examine the structure, function, and dynamics of biomolecules; the folding, solvation, aggregation, and binding equilibria of proteins; and the transport of ions across biological membranes.

For example, elastin is the molecular basis of elasticity in biological tissues such as skin, arteries and lungs. Understanding the molecular forces underlying elastin's mechanical properties will ultimately advance the rational design of artificial skin and vascular grafts. What gives tissue its propensity for extension and recoil? What is found in the microscopic composition of these molecules, how do they arrange themselves – and how do they react under mechanical stress and strain?

Dr. Pomès is getting answers to these questions faster than ever before, due to the new HPC facilities available. "We use SciNet to generate computer simulations so we can study biological systems at the molecular level," Dr. Pomès says. "it has completely changed our perspective on our own work. We are no longer making just incremental progress; instead, we are considering new approaches we weren't even dreaming of before."

"Thanks to the powerful computational infrastructure, we will continue to expand the scope of what we're doing," Dr. Pomès says. Scientific breakthroughs are not far behind.

Dr. Pomès is working with a research effort at U of T that has discovered a small molecule which appears to eliminate Alzheimer's symptoms in mice. A promising therapeutic approach is to develop small-molecule inhibitors as a potential therapy for this disease. By manipulating the genetic makeup of mice to induce Alzheimer's-like symptoms, then feeding the mice this small molecule, a marked improvement is seen in the animals' ability to remember and learn. The formerly erratic behaviour of the mice becomes indistinguishable from a control group. Dr. Pomès is examining the molecular basis for this effect.

He cites the example of a colleague at the Toronto Hospital for Sick Children who, for the past decade, has been stymied by a set of incomplete data. "Simulations were needed to understand it, but we didn't have number crunching power to do the calculations," Dr. Pomès recalls. "Last fall, we asked for dedicated resources with SciNet to do that study. Now, the study is almost complete! In fact, we have about 100 times as much data as we would have had without high performance computing."

"Before, we knew the questions, but didn't have the capabilities to begin answering them. Now, we have the capabilities, and the possibilities are endless," Dr. Pomès enthuses. "What can we do now that we hadn't thought of doing? It is tremendously exciting! Having access to SciNet not only speeds up research, but opens up unexplored realms of investigation."

www.pomeslab.com

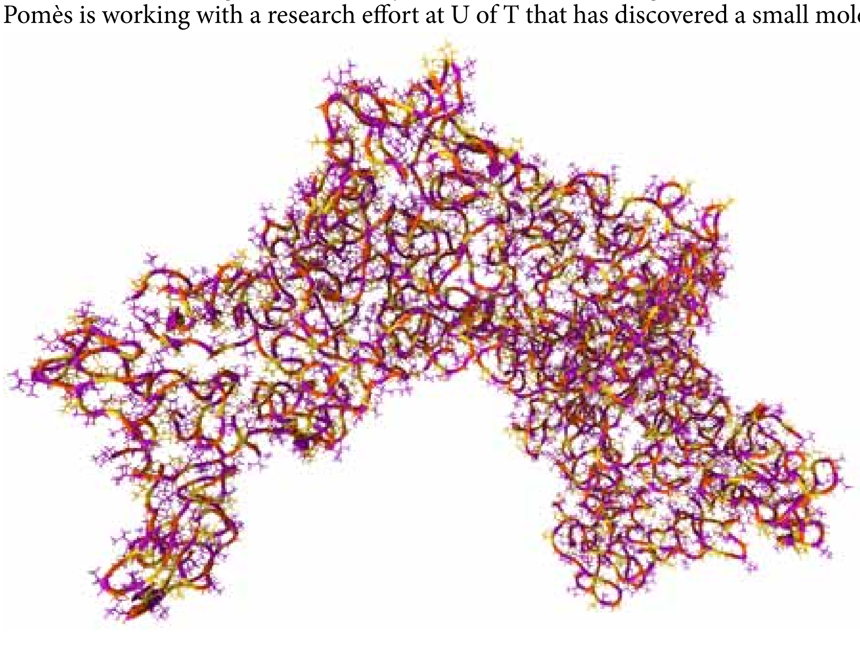

<

A glimpse into the molecular structure of elastin. Thin lines connect individual atoms." (C) Sarah Rauscher and Régis Pomès

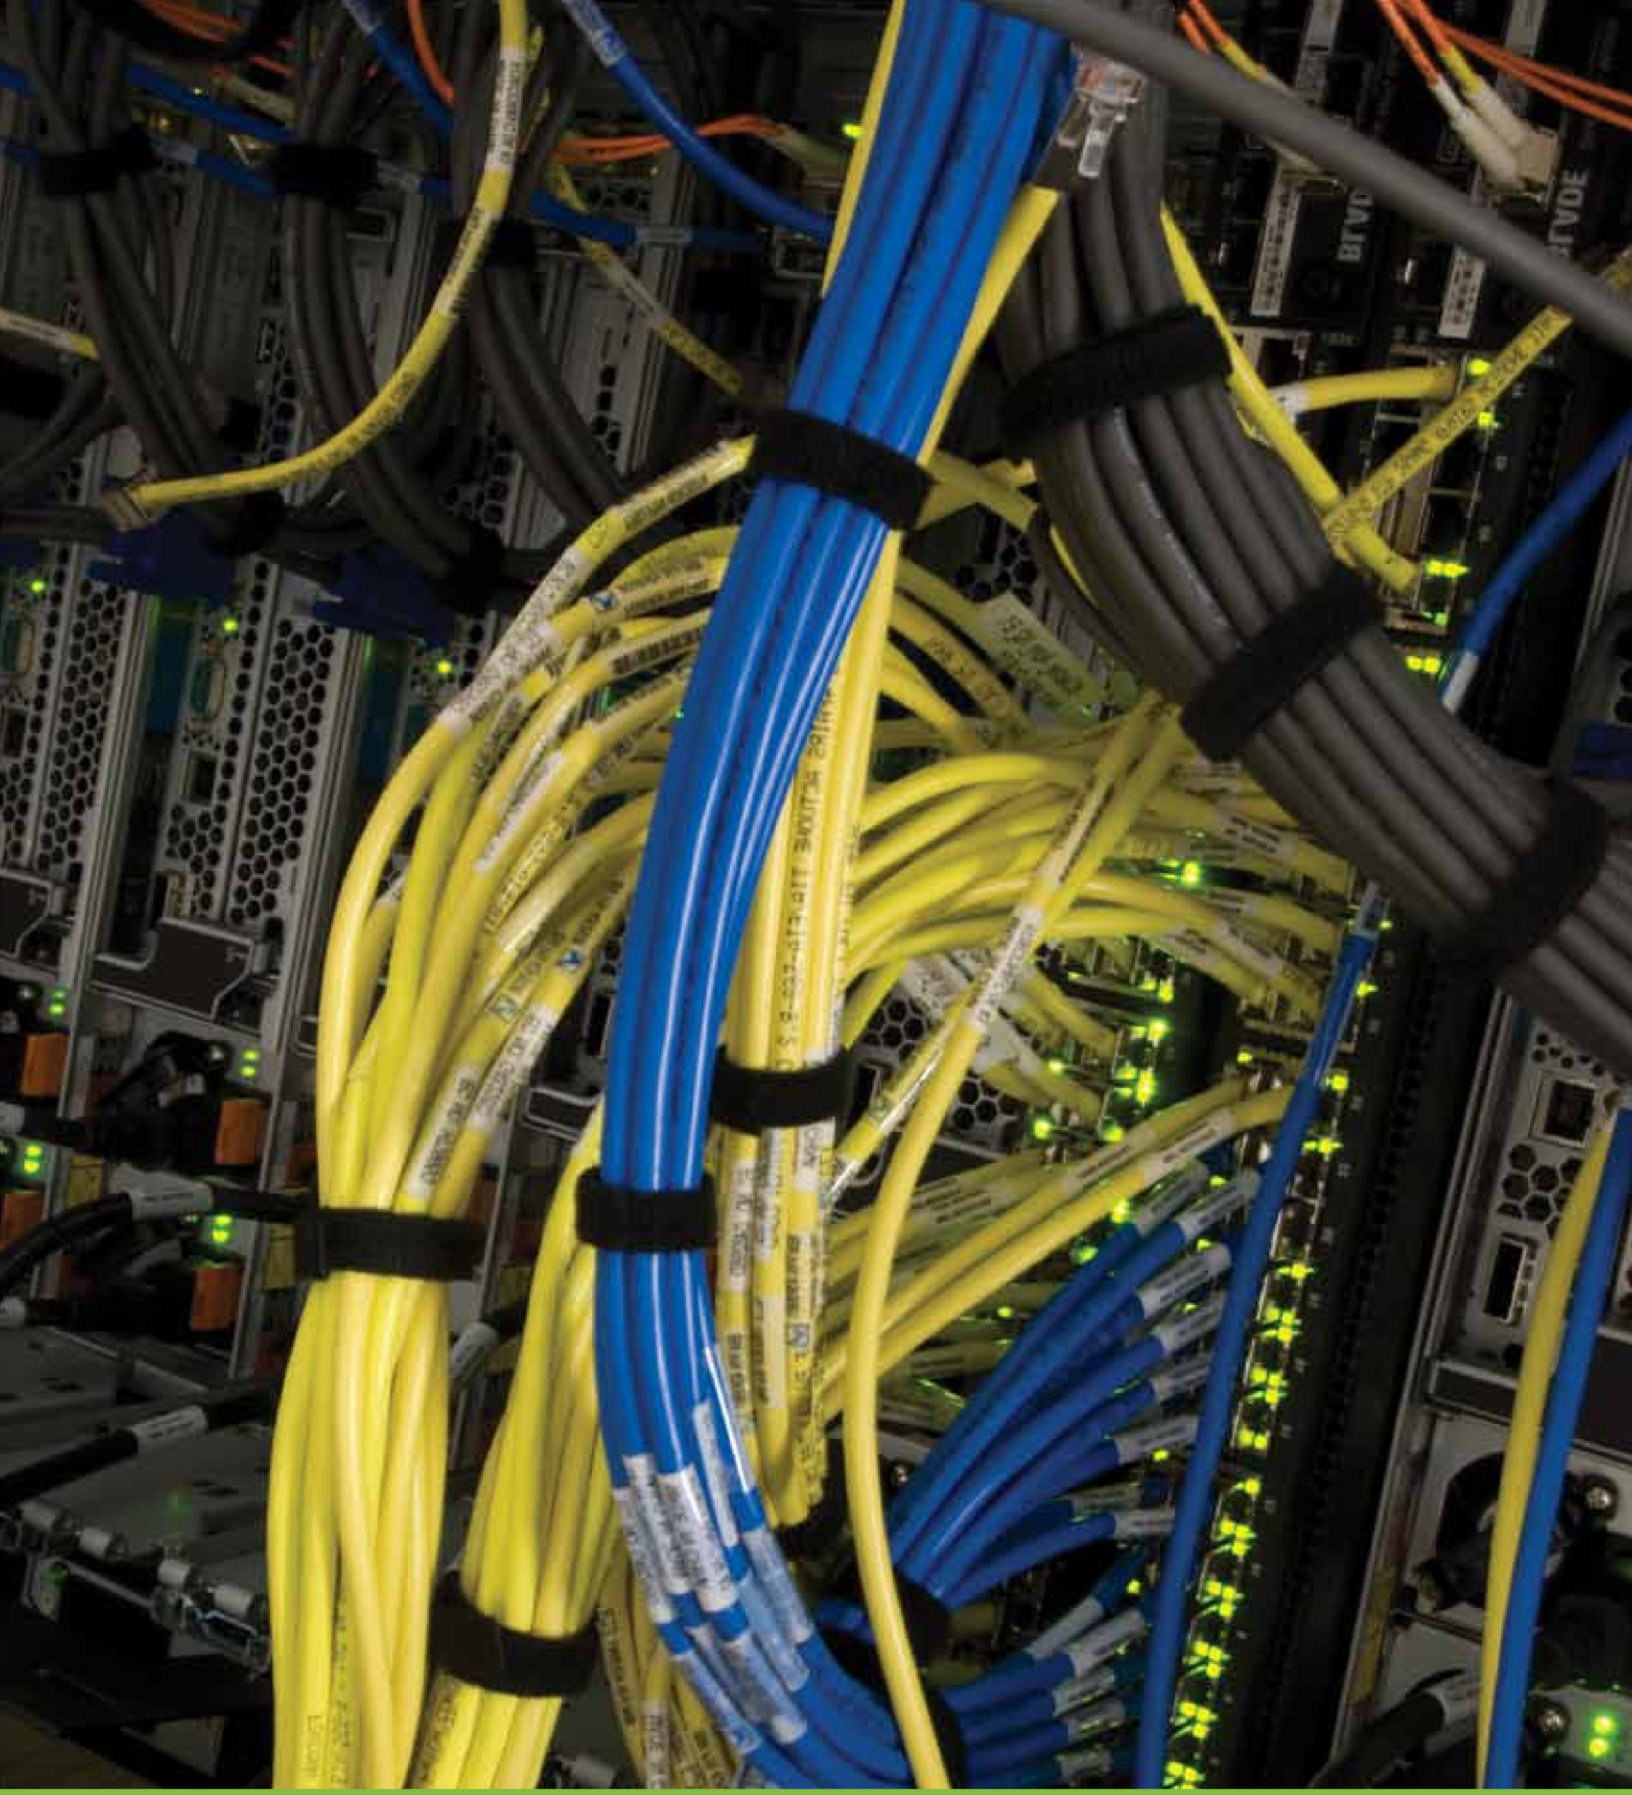

256 McCaul Street, Toronto ON M5T 1W5
E. info@scinet.utoronto.ca
P. 416-978-2922   F.416-978-2995