# The LRP

- content
- recommendations
- next steps
- background

The development of this Long Range Plan was initiated and supported by the Executive and Board of the C3.ca Association Inc.

*"As someone who is not doing grand-challenge computing, I asked "Does this plan offer anything to me?" Other researchers across Canada may wonder the same. I started writing down these points and was quickly convinced."*

*Dr. Andrew Rutenberg*
*Dept. of Physics and Atmospheric Science*
*Dalhousie University*

# The Benefits of the Long Range Plan

## Advantages for the researcher

- Access: clear opportunities for obtaining access to powerful national and regional computing resources
  - streamlined and simplified resource application process
  - ongoing competitive resource allocation opportunities
- Opportunity: support for ambitious, time-critical research projects (mitigates funding-cycle delays and uncertainties)
- Sustainability: access for ambitious projects on a long-term basis (allows planning for research and for personnel development)
- Coordination:
  - ready access to appropriate regional- and national-level infrastructure
  - shared salary-support programs for skilled local technical staff
  - growth of a pool of talented technical support people in Canada
  - development of a national network of technical support available to all users
- Autonomy: local control for powerful desktop and group computing facilities

## Advantages for the federal and provincial granting agencies

- Clarity: clear planning and implementation
- Consistency: clear determination of resource needs and proposal quality
- Efficiency: coordinated use of computational hardware and personnel
- Competitiveness: preservation of an internationally competitive research environment
- Coherence: coordinated dissemination of techniques and results across Canada
- Transparency: appropriate and well-managed injection of resources
- Cost effectiveness: coordinated national plan eliminates the redundancy of independent local efforts
- Financial responsibility: financial and management accountability

## Advantages to industry

- One-stop-shopping:
  - access to significant computational resources for research
  - rapid scientific solutions for problems and for product development
  - a low-cost try-before-you-buy approach to HPC for the SME base
  - a well-maintained researcher knowledge base
  - a technology watch service with timely alerts
  - HQP: access to a guaranteed skilled workforce for the knowledge economy
  - Killer apps: today's leading-edge research builds tomorrow's leading-edge industry

# Thanks to our Sponsors

NSERC CRSNG

CANARIE

Social Sciences and Humanities Research Council of Canada
Conseil de recherches en sciences humaines du Canada
Canada

Canada Foundation for Innovation
Fondation canadienne pour l'innovation

CIHR IRSC

C3.ca

NRC · CNRC

# The Consortia and the CANARIE backbone

# Long Range Plan: Status – June 2005

## C3 Board Final Review Process

- Publication formatting in process, web version available
- Promotion activity in process
- Continuing to meet with major stakeholders
- Full publication in English and French – NRC Press **July**

## The Authors Panel

| | |
|---|---|
| • Kerry Rowe (Chair) | VP Research, Queen's |
| • Jon Borwein | CRC Computer Science, Dalhousie & SFU |
| • Russ Boyd | Chemistry, Dalhousie |
| • Gilbert Brunet | Meteorological Services Canada |
| • Hugh Couchman | SHARCNET and Astrophysics, McMaster |
| • Alan Evans | Brain Imaging, McGill |
| • Martyn Guest | Daresbury Labs, UK |
| • Ian Lancashire | English, Toronto |
| • Jonathan Schaeffer | WestGrid & CRC Comp Science, Alberta |

**Coordinator** - Andrew Pollard, Engineering, Queen's

# The past 60 years

## ENIAC in 1946


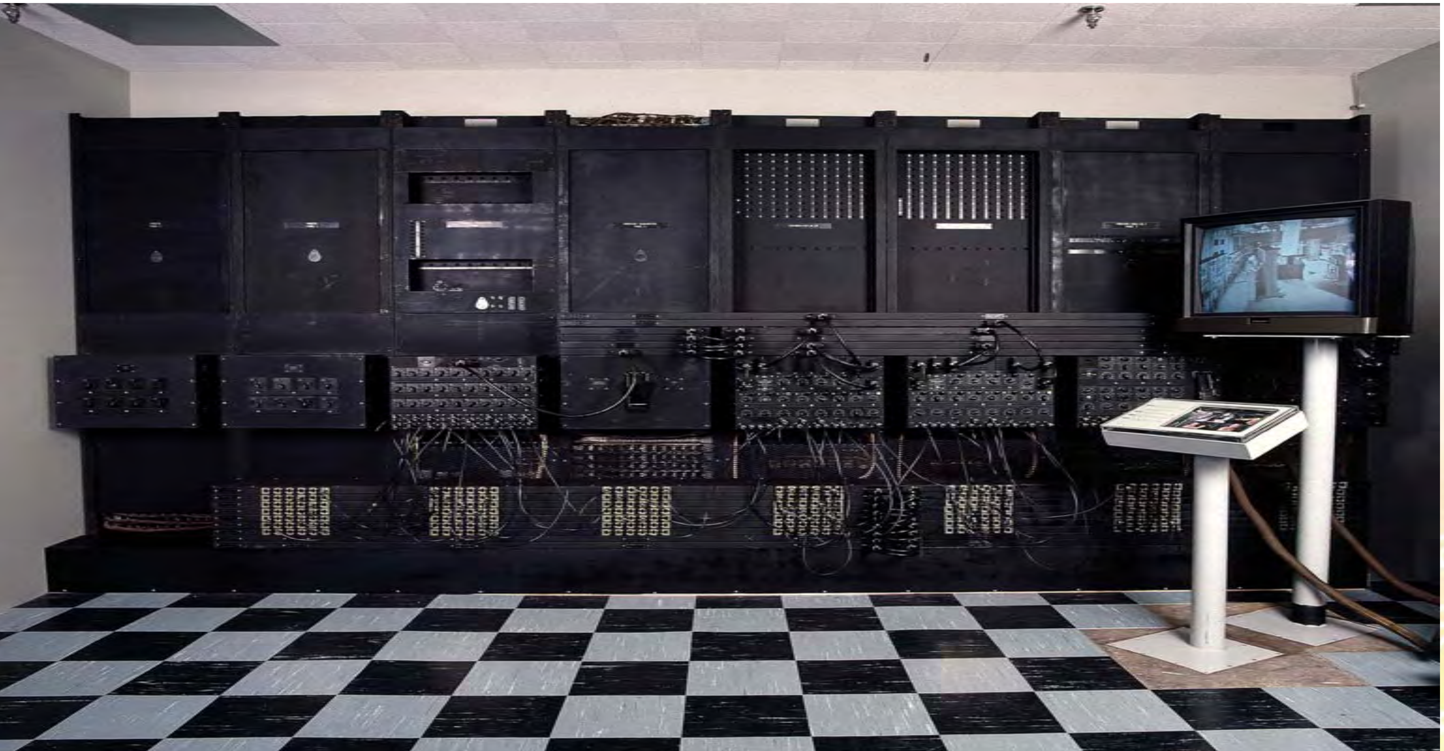

**5Kflop/sec**

Electronic

Numerical

Integrator

And

Calculator

# NERSC's 6000 cpu Seaborg - already 'obsolete' in 2004

A Vancouver Cave

10Tflop/sec

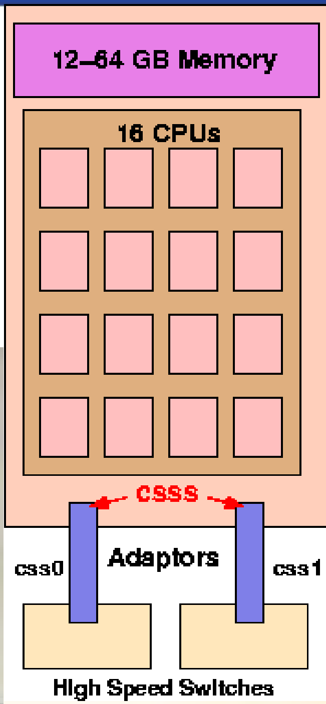


# IBM BlueGene/L system at LLNL

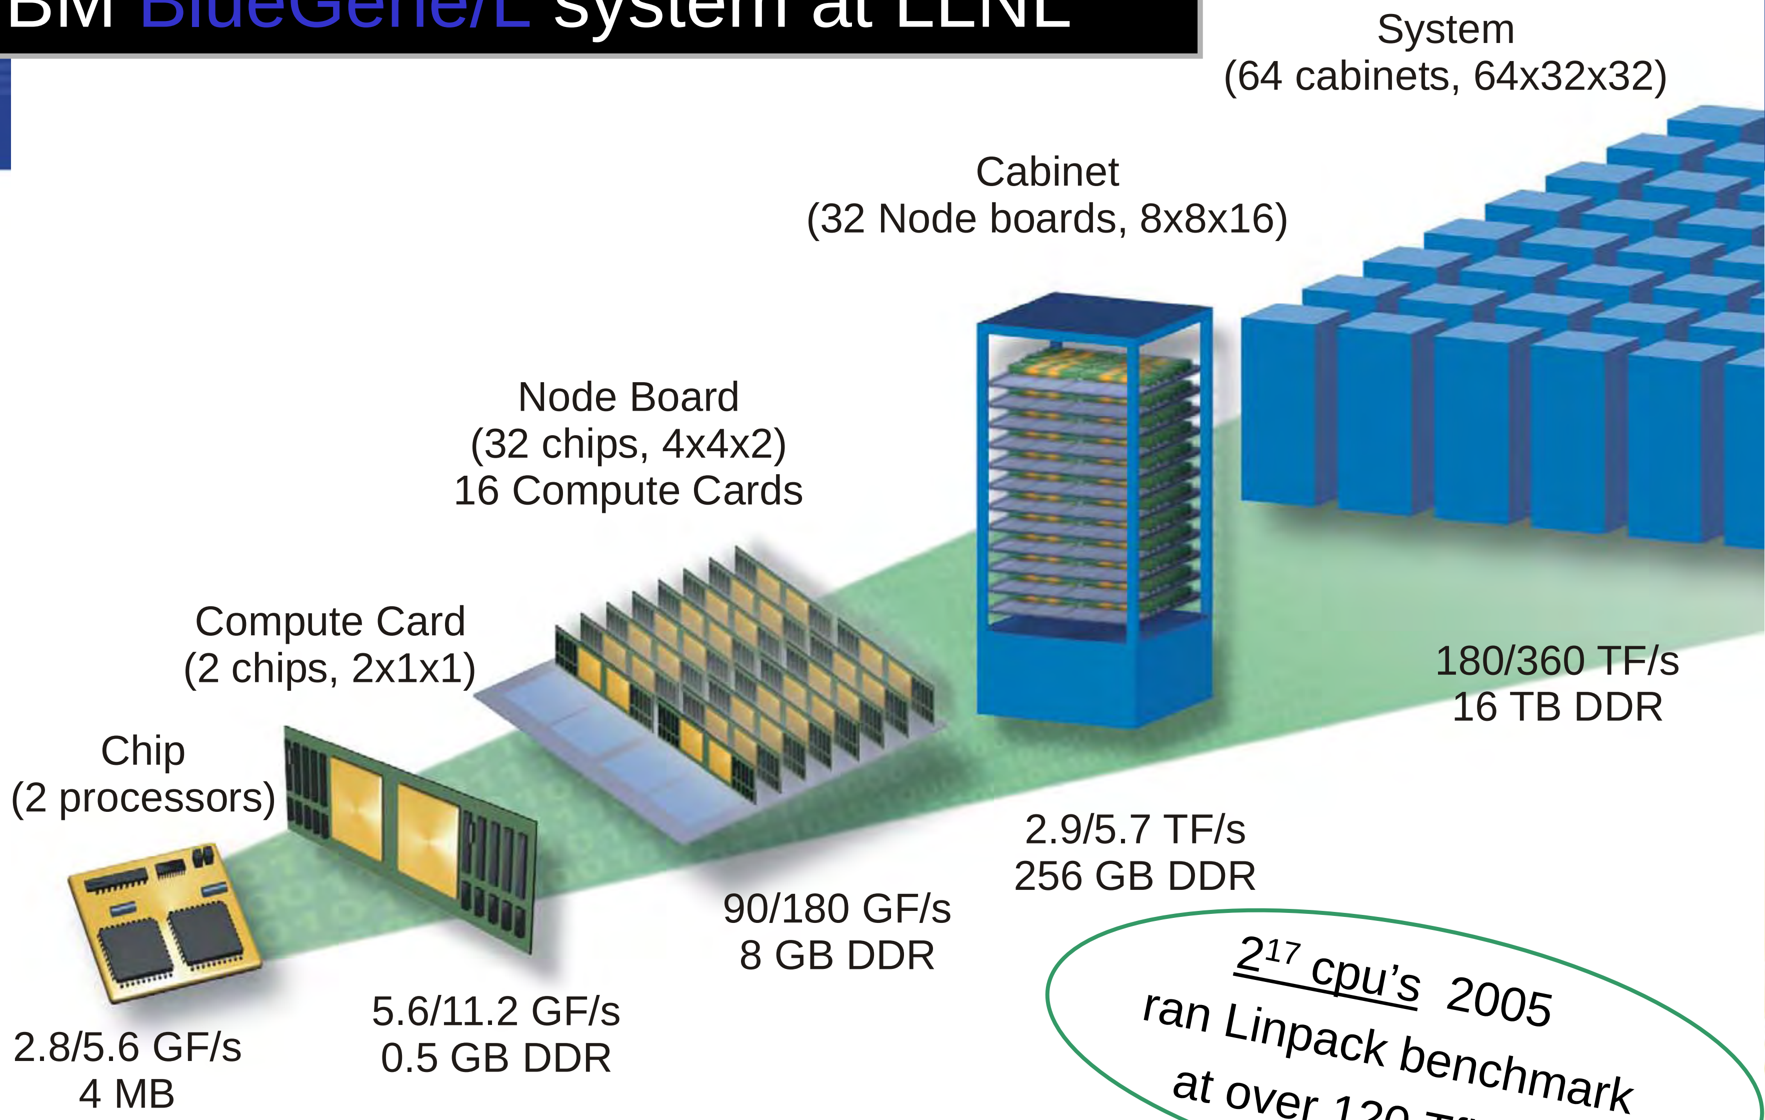


$2^{17}$ cpu's 2005
ran Linpack benchmark
at over 120 Tflop/s

# C3 Consortia members

| | |
|---|---|
| ACENET: | 7 Universities in Atlantic Canada |
| RQCHP: | 2 Universities + 2 Research Centres in Quebec |
| CLUMEQ: | 3 Universities in Quebec |
| HPCVL: | 5 Universities in Eastern Ontario |
| SHARCNET: | 9 Universities and 2 colleges in S. Central Ontario |
| WESTGRID: | 5 (8) Universities + 2 Research Centres in West |

# The LRP consists of …

## A Story

## Executive Summary

## 4 Main Chapters
- Technology
- Operations
- HQP
- Budget

## 25 Case Studies
- Many sidebars

CASE STUDY

*"HPC is an essential research infrastructure for these solutions. Without it, our knowledge society has weak foundations."*

***Dr Ian Lancashire***
***Professor of English***
***University of Toronto***

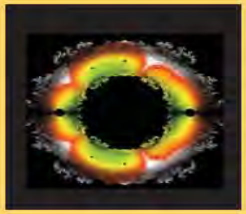

"Numbers are not the only thing that computers are good at processing. Indeed, only a cursory familiarity with fractal geometry is needed to see that computers are good at creating and manipulating visual representations of data. There is a story told of the mathematician Claude Chevalley, who, as a true Bourbaki, was extremely opposed to the use of images in geometric reasoning. He is said to have been giving a very abstract and algebraic lecture when he got stuck. After a moment of pondering, he turned to the blackboard, and, trying to hide what he was doing, drew a little diagram, looked at it for a moment, then quickly erased it, and turned back to the audience and proceeded with the lecture. It is perhaps an apocryphal story, but it illustrates the necessary role of images and diagrams in mathematical reasoning - even for the most diehard anti-imagers. The computer offers those less expert, and less stubborn than Chevalley, access to the kinds of images that could only be imagined in the heads of the most gifted mathematicians, images that can be coloured, moved and otherwise manipulated in all sorts of ways." (Nathalie Sinclair, 2004) [REF 31]

### Real or Fake? Dynamically Validating the Reliability and Authenticity of Digital Records

A paper-only letter today is relatively secure. Its letterhead, its author's signature, and perhaps its watermark breed faith in its reliability as a product of its stated author. Proper pagination and a list of any attachments tell us whether the letter has changed during transmission. Without obviously missing pages or physical damage, we will accept that the document is authentic.

In contrast, few people trust digital documents. Email or web documents often have no digital signature or secure evidence of provenance. A sender's address can be very easily forged, and routinely any attempt to reply is returned for lack of a true address. Digital content can also easily be plagiarized so that the identity of the true creator is hidden. Digital records also are vulnerable to authenticity problems as a result of transmission or use as computer displays may fail to reproduce what was sent. Once digital records are stored, other potential problems can occur. Records can be damaged during routine processing, and no one is the wiser.

So how might one create a level of authenticity and security to instill a level of confidence in digital records creation, management and retrieval? Only by working with grid-based high performance computing technologies can the problems in determining, in

## One Day

High performance computing (HPC) affects the lives of Canadians every day. We can best explain this by telling you a simple story. It's about an ordinary family on an ordinary day: Russ, Susan, and Ceri Sheppard. They live on a farm 15 kilometres outside Wyoming, Ontario. The land first produced oil, and now it yields milk; and that's just fine locally.

Their day - Thursday, May 29, 2003 - begins at 4:30 a.m. when the alarm goes off. A busy day. Susan Zhong-Sheppard will fly to Toronto to see her father, Wu Zhong, at Toronto General Hospital; he's very sick from a stroke. She takes a quick shower and packs a day bag for her 6 a.m. flight from Sarnia's Chris Hadfield Airport. Russ Sheppard will stay home at their dairy farm, but his day always starts early. Their young daughter Ceri can sleep three more hours until school.

Waiting, Russ looks outside and thinks, *It's been a dryish spring. Where's the rain?*

In their farmhouse kitchen on a family-sized table sits a PC with a high-speed Internet line. He logs on and finds the Farmer Daily site. He then chooses the Environment Canada link, clicks on Ontario, and scans down for Sarnia-Lambton. Good, as he hoped:

Periods of morning rain. Risk of a thunderstorm. Wind south 30 km/h gusting to 50 diminishing to 20 in the afternoon. Sunny periods. High 14

One of a million who check that site every week, Russ uses HPC, although he doesn't know he does. Two decades ago, Russ found Environment Canada reliable only two days ahead. Now, the forecasting works pretty well for a week.

Russ thinks, *They've got some smart people at our Meteorological Service.*

But Russ worries. The farming Sheppards go back four generations to the thirties, times of hunger in the evening. A picture in the hall landing reminds him of 'those days'.

He sits a while, staring at the screen. Ten days ago Russ felt his world shake. BSE, bovine spongiform encephalopathy, appeared in one cow in Wanham, Alberta, and now seven countries have banned our beef. Russ gets tired thinking about this because he can do nothing about it.

What would he and Susan do if their cows had to be destroyed? And now there's her father, and he can't burden her with what he thinks.

Russ calls up *www.google.ca* and enters "bse alberta". The news describes the slaughter of a thousand suspect Alberta livestock. But an article from last March details BSE testing in the Animal Health

## Weather Prediction

The accuracy of a five-day forecast in the year 2003 was equivalent to that of a 36-hour forecast in 1963 [REF 1]. The quality of daily forecasts has risen sharply by roughly one day per decade of research and advances in HPC hardware and software. Accurate forecasts transform into billions of dollars saved annually in agriculture and in natural disaster costs. Using a model developed by Dr. Keith Thompson at Dalhousie University, the Meteorological Service of Canada has recently been able to predict coastal flooding in Atlantic Canada early enough for the residents to take preventative action.

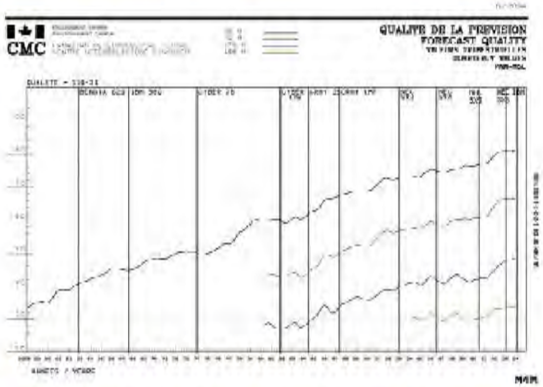


**The figure charts the steadily increasing quality of weather prediction**

**Canada's leading science fiction writer, William Gibson, coined the word "cyberspace" in his novel *"Neuromancer"* which won the Hugo, Nebula, and Philip K. Dick awards in 1984–85. *"Unthinkable complexity. Lines of light ranged in the nonspace of the mind, clusters and constellations of data. Like city light . . ."* [REF 4]**

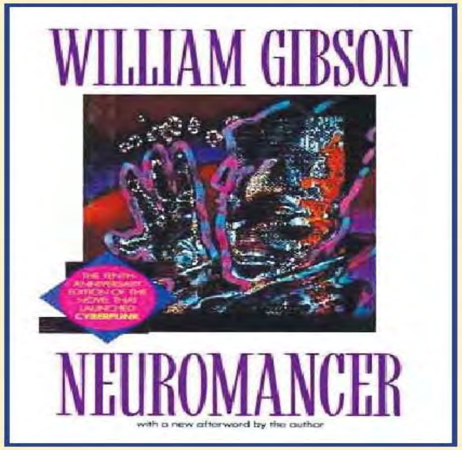


## The Story. Canadian Successes: PIXAR and Ryan

William Reeves, Technical Director of Pixar Animation Studios, studied in Waterloo and Toronto, getting a B.S. in math and an M.A. and a Ph.D. in computer science. In 1980, he joined Lucasfilm and invented an image-synthesis technique called Particle Systems. It generates complex and detailed images through many (normally) simple geometric particles changing stochastically over time. After working on films like *The Empire Strikes Back*, Reeves and two partners created Pixar. Today his technique, taught everywhere in film courses, powers many of Pixar's groundbreaking animated films. As supervising Technical Director on *Toy Story*, the first feature film created entirely by computer animation, Reeves won four technical Academy Awards. Another Canadian, Darwyn Peachey (BSc and MSc from Saskatchewan), has played a key technical role with Pixar, receiving a technical Oscar

Patrick Coleman and Karan Singh, from the Dynamic Graphics Project, the University of Toronto research group of which Reeves is an alumnus, developed nonlinear projection software for Chris Landreth's "Ryan", which received the Academy Award for Best Animated Short film of 2004. The Canada Council for the Arts participated in the development of "Ryan". [REF 5]

**Ryan, a victim of drug addiction and alcoholism, experiences states of mind that affect his perception of the space around him. Computer animation software normally uses the rules of linear perspective, but to show how characters like Ryan see their world, those rules had to be modified. The nonlinear projection system used to create Ryan allows animators to combine multiple points of view in various ways to achieve effects that distort linear perspective.**

**Pixlet (Pitxar + wavelet) is the first studio-grade algorithm for fimmakers. Pixlet provides 20-25:1 compression, allowing a 75 Megabytes/sec series of frames to be delivered in a**

# FROM the LRP

## Grand Computing Challenges

- such problems exist in all areas and their needs are expanding rapidly

| | 5 Tflop (2004) | 20 Tflop (2006) | 80 Tflop (2008) |
|---|---|---|---|
| **Turbulent combustion** | | Modelling of laboratory flames | Practical flames: simulations of environmental and industrial combustion; integrity of structures |
| **Climate modelling** | 100 km atmospheric resolution; 50 km oceanic resolution; multi-century resolution | Include atmospheric chemistry and carbon cycle | Full coupling of atmosphere (50km) and ocean (15km); multi-decade resolution |
| **Astrophysics** | Accurate large-scale cosmic structure | Galaxy formation in cosmic volumes; Supernova simulation; full analysis of cosmic microwave background radiation | |
| **Protein folding** | Folding of 50 amino acids over micro-second intervals | | Folding of proteins over milli-second intervals |
| **Structural biology** | Simulation of sub-cellular systems at the atomic level | | Simulation of full bio-molecular complexes for micro-seconds |
| **Nanoscience: Electronic Structure of materials** | | First principles calculation of electronic structure of a 10,000 atom system (interaction of a small of cluster of nano-particles) | |
| **National-scale economic modelling** | | | Pflop (1000 Tflops) computing: 2009 on the most powerful systems, 2012 on number 10 |
| **Multi-user immersive virtual reality** | | | |

[REF 19]

# From the LRP: HPC at Work

## THE POWER OF HIGH PERFORMANCE COMPUTING - CLIMATE MODELING

**GLOBAL CLIMATE MODELS ON DESKTOP COMPUTER**
Analysis of 500km climate cell takes one hour on one processor, but resolution of analysis is poor

**GLOBAL CLIMATE MODELS ON PARALLEL COMPUTER**
Analysis of same problem still takes one hour on 100 processor system; however, resolution is now 100 times greater (50km climate cells) resulting in a greater accuracy and finer data analysis

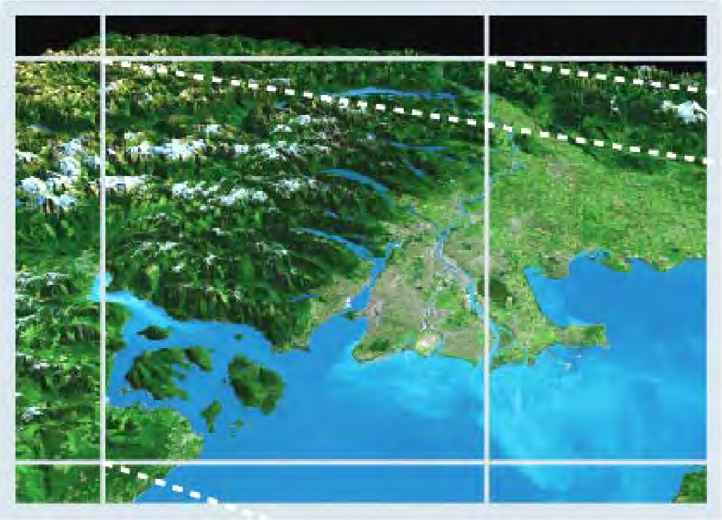

LOWER MAINLAND, BRITISH COLUMBIA

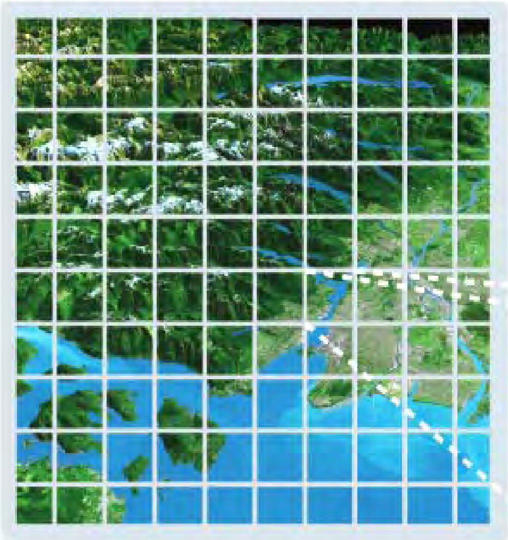

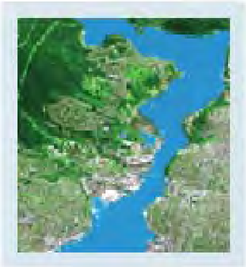

# NEUROLOGY: Mining Memories & Mending Minds

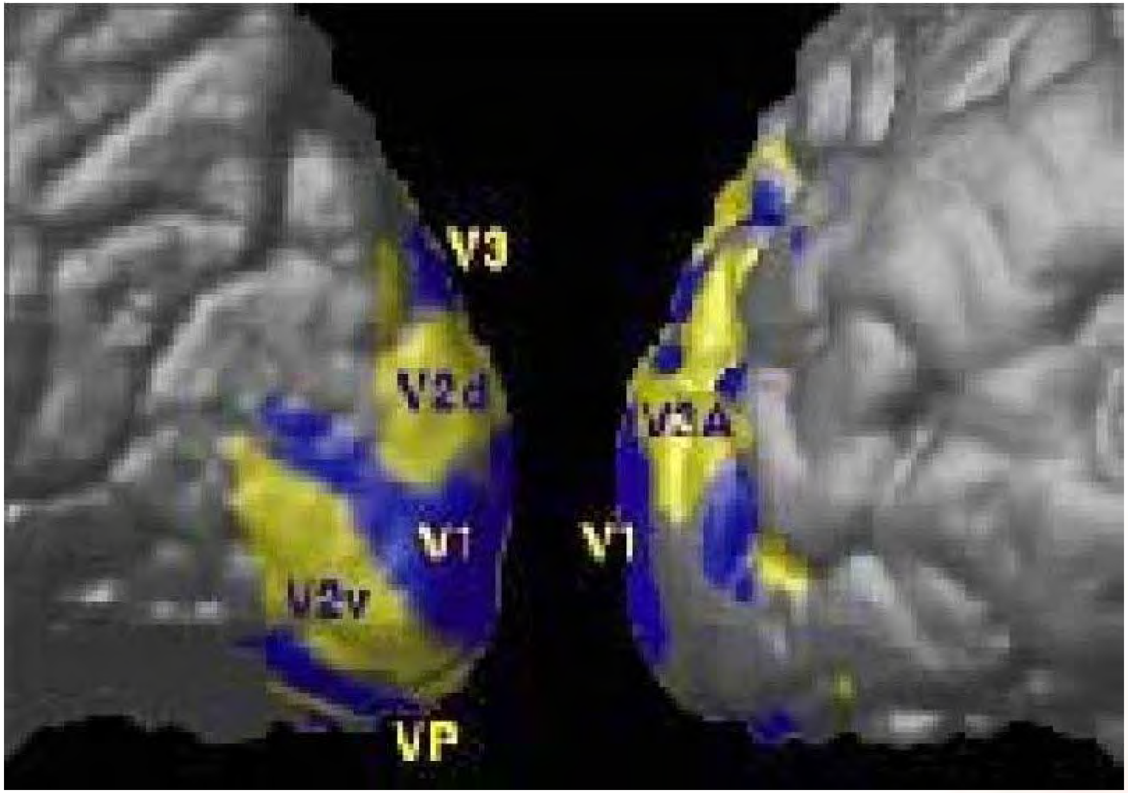


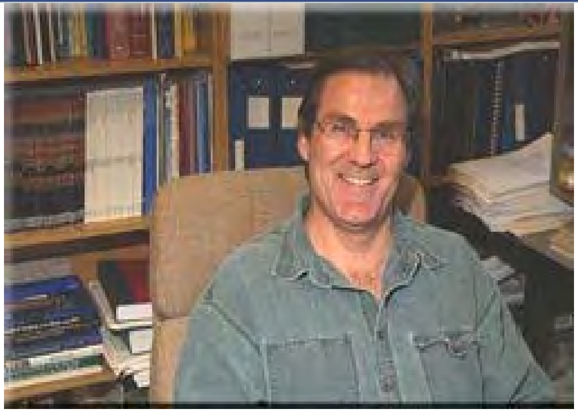

Dr. Alan Evans, McConnell Brain Imaging Centre, McGill University

Field sign map of early visual cortical areas, overlaid on 3-D anatomic reconstructions of medial and lateral views of the right hemisphere of the brain.

# NEUROLOGY: Mining Memories & Mending Minds

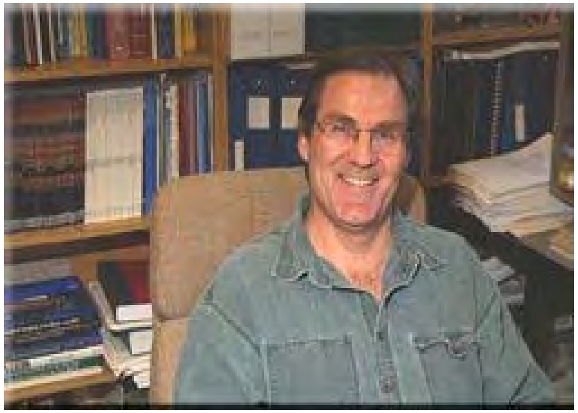

"With competitive HPC resources, Canada will be able to compete, and possibly lead, in applying information technology to brain research. The benefits of such research strength in terms of brain disease treatment are incalculable. With an ageing population, increasingly afflicted by degenerative brain disorders and psychiatric illness, it is essential that we have access to state-of-the-art neuroscience research to develop effective therapeutics."

Dr. Alan Evans,
McConnell Brain Imaging Centre, McGill University

# **ARRYTHMIA**: To the Heart of the Matter

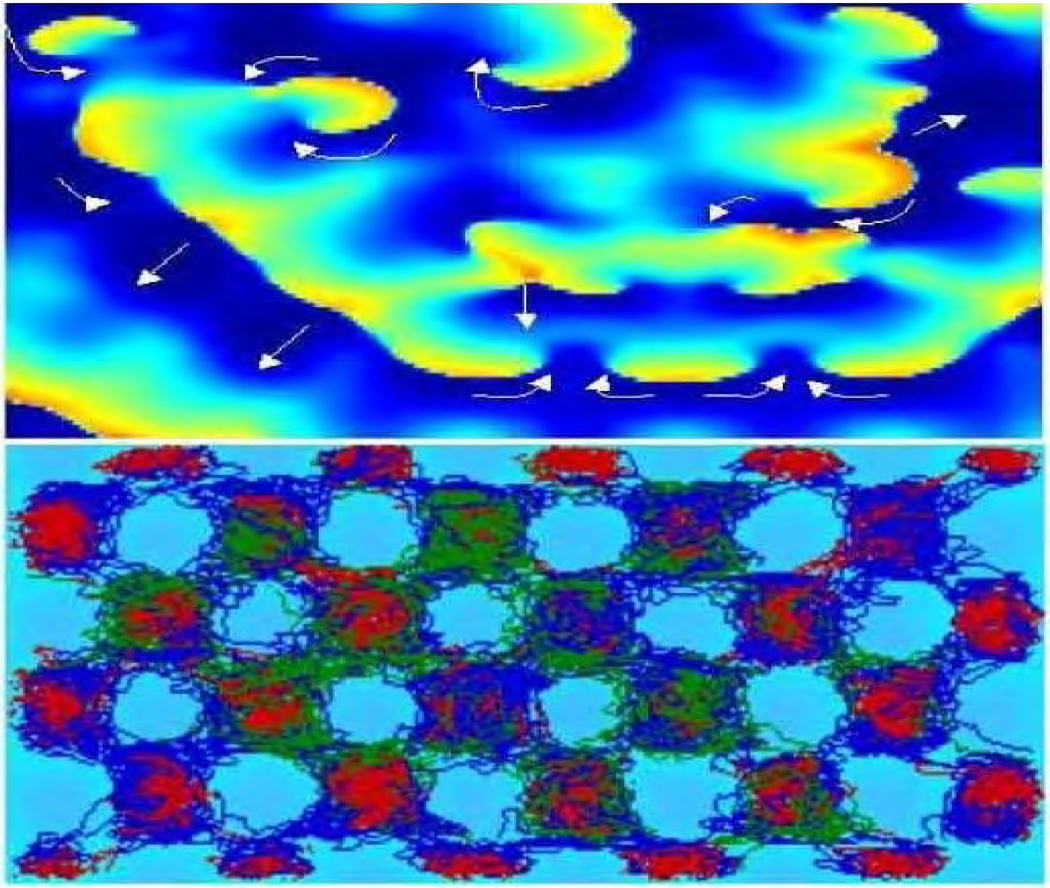

1/1000 of a second snapshot of cardiac electrical activity from a model of 720,000 cells.

Dr. Stanley Nattel
Paul David Chair and Professor of Medicine, University of Montreal and Montreal Heart Institute

# ARRYTHMIA : To the Heart of the Matter

"Present modeling involves up to 720,000 cells, with the activity of each cell determined by 24 equations that have to be solved 80 to 200 times a second, depending on the conditions.  We are working with highly complex models and making them more complex at each stage to reflect the human heart better and better.  It now takes 4 days to model 20 seconds of cardiac activity.  HPC is an absolute must..."

Dr. Stanley Nattel
Paul David Chair and Professor of Medicine, University of Montreal and Montreal Heart Institute

# ASTROPHYSICS: Understanding our Universe

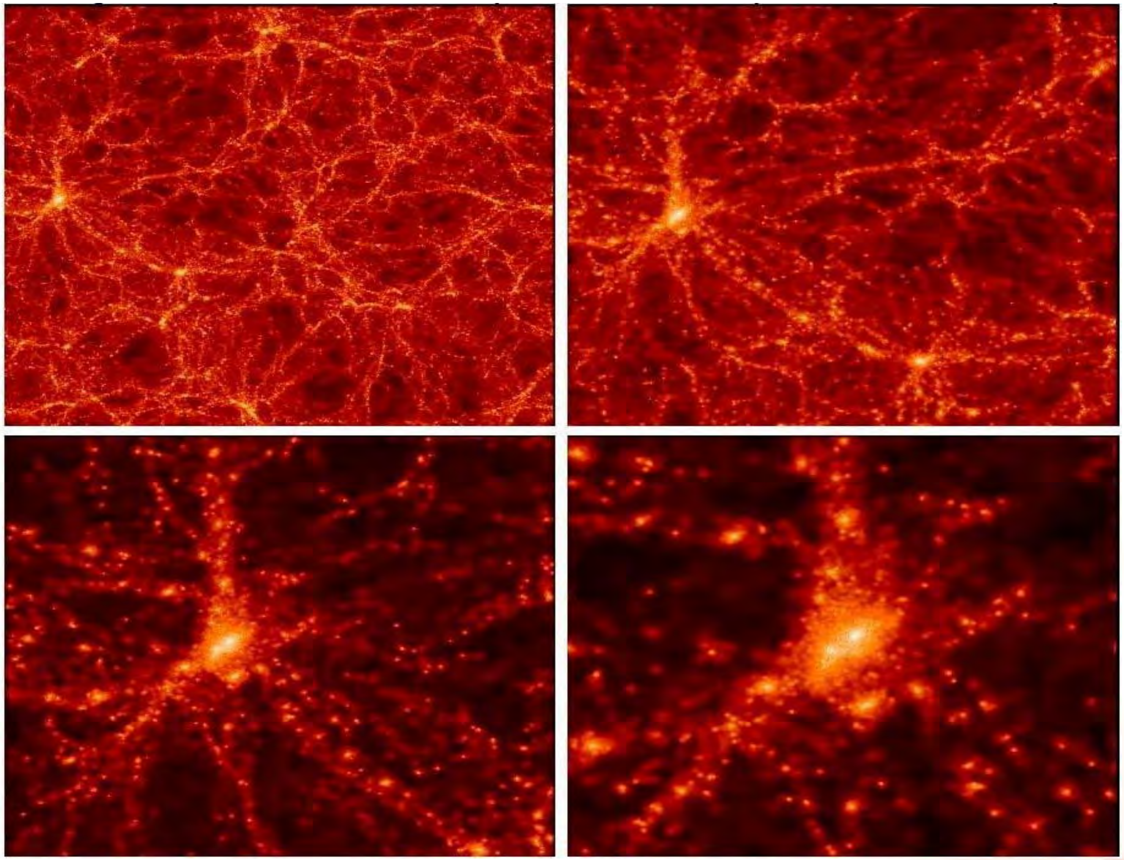

Density plot images from a Virgo simulation showing massive galaxy clusters.

Dr. Hugh Couchman, McMaster University Founder SHARCNET

# ASTROPHYSICS: Understanding our Universe

"Simulating the growth of cosmic structure, from the star clusters gas within galaxies to the large superclusters containing tens of thousands of galaxies and beyond, is a huge computational challenge. Indeed it has been called one of the "Grand Computational Challenges". Canada's continued leadership in this field pivots on access to appropriate large-scale computational resources."

Dr. Hugh Couchman,
McMaster University
Founder SHARCNET

# A Brief History of Growing Success

## 1980 - 1987

- Dorval (Weather Service) peaks at #6 on TOP 500 then quickly drops to #250
- Several short term HPC ventures (OCLSC, University of Calgary, HPCC)
- Pre 1993, NSERC support only to rent time abroad

## 1998 - present

- CFI endows new and growing facilities
- NSERC people (MFA) support stabilizes
- C3.ca founded with NRC help, collaboration includes 20% facility sharing agreement
- Minimal ongoing operating funds available or committed
- Regional Consortia share resources while minimizing overhead
- Strengthened Major Resource Providers emerge

# International Recognition

"Canada has invested wisely in mid-range computing over the last five years and has created the best developed, most broadly accessible, mid-range High Performance Computing facilities in the world."

Professor Martyn Guest, Computational & Engineering Department of UK's Central Laboratory of the Research Councils (CLRC) in Daresbury

# Canada's Standing (a mixed message)

## Human and Operational (Facility) Resources

In common with other major scientific infrastructure, the effective and productive use of the physical infrastructure requires substantial continuing investments in technical management, operational and scientific support, administration, outreach and training: the *human infrastructure resources*. Additionally, there are substantial costs associated with providing power, air conditioning, security, data backup, software acquisition and maintenance, hardware maintenance, uninterrupted electrical power and a host of other support services: the operational or *facility costs*.

### Current Support

Prior to the 2002 CFI competition, there was no CFI funding for HPC operations — researchers had to cobble together operational funds from a variety of sources (some provinces, universities and industry partners provided funds). The 2001 and 2003 CFI competitions supplemented successful grants by providing an additional 30% of the CFI award to be used towards operating costs. The provinces did not follow suit, which meant that the operating funds applied to only 40% of the total project budget — or about 10% of the total budget.

Experience with the infrastructure that was put into place as part of the 2002 CFI awards indicates that the CFI funds are sufficient to provide a minimal operational level of human infrastructure support, but no more. This assumes that the facility overhead costs (such as electricity and air conditioning) are paid for by the universities. This is becoming an increasingly problematic situation, as this type of facility does not attract federal indirect costs of research support. There are, therefore, no funds left over to provide proactive support services — analysts to help develop, enhance, port, and optimize HPC applications—and this is a serious deficiency. In the future, this matter will be compounded if CFI is not able to provide operational funding at the current 30% of their portion of the capital infrastructure cost.

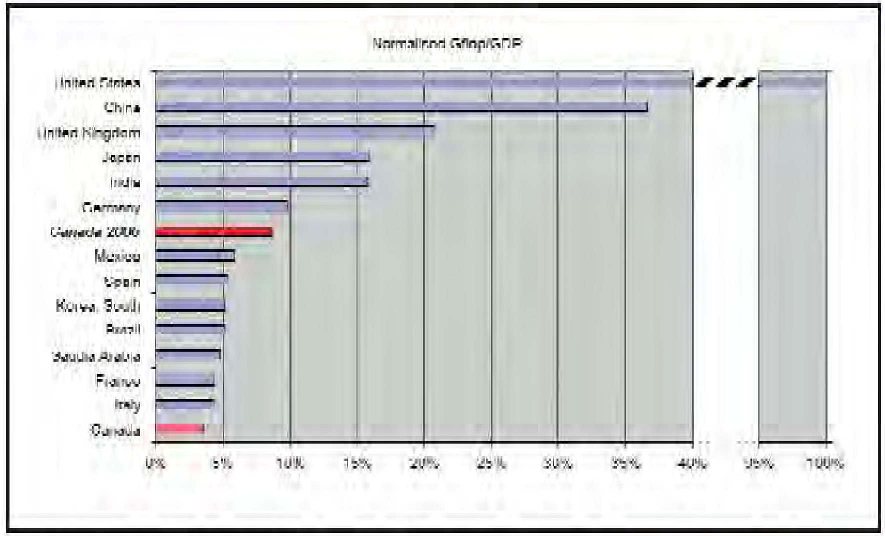


Figure 5.1: Normalized graph showing GFlops per GDP (per capita) and where we would be positioned in 2006 if the recommendations in this LRP are adopted [51]

*The Department of Energy High-End Computing Revitalization Act of 2004 authorizes DOE to establish supercomputer user facilities that will provide U.S. researchers access to some of the world's most advanced computers on a competitive peer-reviewed basis.*

# LRP **Recommendation 1**: Infrastructure

- **Priority 1** - *Sustained* funding for advanced 'mid-range' HPC
  - Building on existing, successful, consortia/MRP model for facilities and people
  - Meet growing needs of most HPC users

- **Priority 2** - Establish high-end facility
  - *Sustained* TOP 20 world class facility
    - ✓ Three being planned in EU !
  - A pan-Canadian model
  - Serves relatively few critical applications

# HPC in the US

## Supercomputing Bill Heads To President For Signature

- November 18th Breaking News: The House of Representatives gave final approval to a bill that will strengthen U.S. competitiveness by revitalizing domestic computing capabilities and supporting the development of the world's fastest supercomputers …. at least one R&D center devoted to the development of software for supercomputing applications

- Science Committee Chairman Sherwood Boehlert (R-NY) said, "Supercomputing capability is increasingly becoming a vital component of the efforts of industry and academia to remain global leaders."

# Things we can not do here include:

## DOE Computational Science: Accelerator Physics

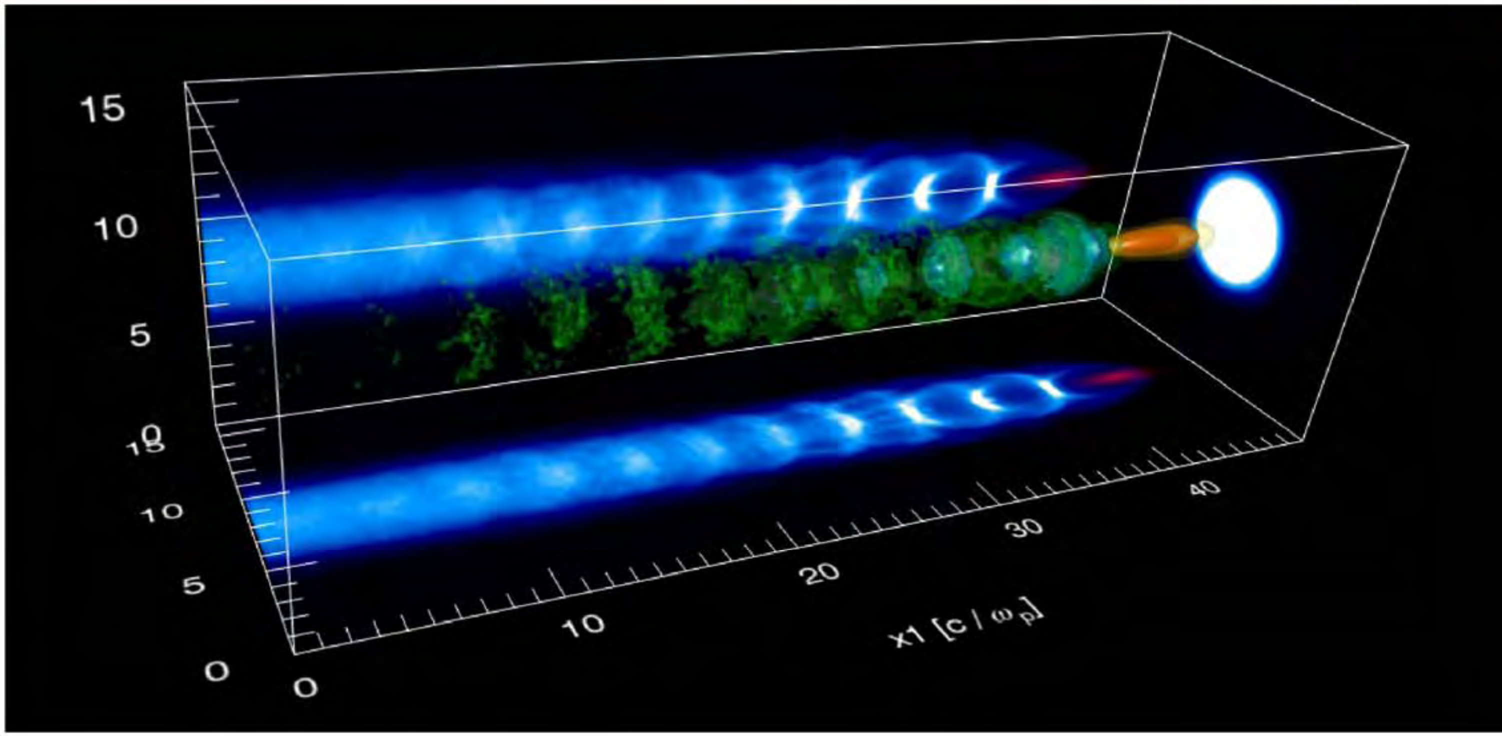


Recent breakthrough Wakefield Accelerator experiments have produced 100 MeV beams with significantly improved beam quality.

These have been made possible by 3D simulations (such as ORIRIS shown here). Results show electron density (in green) from intense laser beam in Lithium gas cell. 512x128x128 grid, 3100 time steps. [Cover of Nature]

A virtual heart model containing millions of "virtual cells". Researchers can investigate the virtual heart through the use of a 3D cave shown here. This allows for heart investigations to go from viewing the surface cells of the heart to viewing all cellular and tissue levels when monitoring heart activity.

## Heart Images And Sars Model



*"The rapid dissemination of the SARS genomic sequence through the Internet did allow groups all around the world to participate in the computational analysis of the genome and the 3D modeling of the protein sequences. Ultimately, resulting in an unprecedented number of scientific publications and vaccine initiatives just over one year after the discovery of this emerging disease".*

***Dr Steve Jones***
***Head of Bioinformatics***
***Genome Sciences Centre***
***Director of Bioinformatics at Genome BC***
***University of British Columbia***

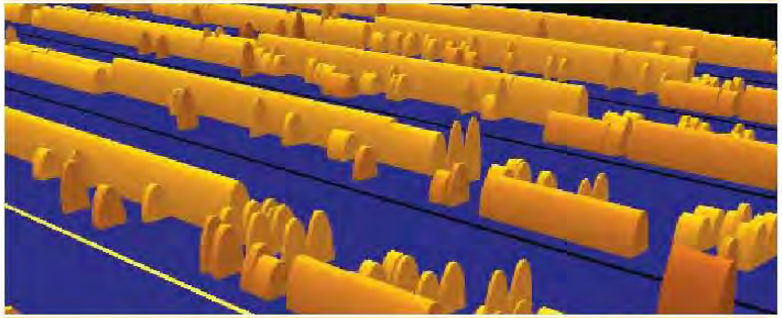

This diagram displays the genetic makeup of the human chromosome illustrating genes and other critical information. This graphs allows scientists to review genetic information from human chromosomes in hopes to better understand and develop vaccines for infectious diseases, like SARS, which has drastically affected the human population in the last two years.

# … HPC in the US The view from Congress

- "High performance computers are central to maintaining U.S. leadership in many scientific fields," said Chairman Biggert. "With House passage of this bill, American researchers are one step closer to gaining the tools they need to remain the world leader in the development and use of supercomputers. Our nation's scientific enterprise, and our economy, will be stronger for it."
- Congressman Lincoln Davis (D-TN), lead Democratic sponsor of H.R. 4516, said, "With today's vote Congress has made a decision to build the fastest and most efficient computers in the world. This sends a strong message of America's commitment towards leading the world in the research and development of supercomputing technology…."

# Things we can not do here include:

## Self-Assembled Wires 2nm Wide
**[P. Kuekes, S. Williams, HP Labs]**

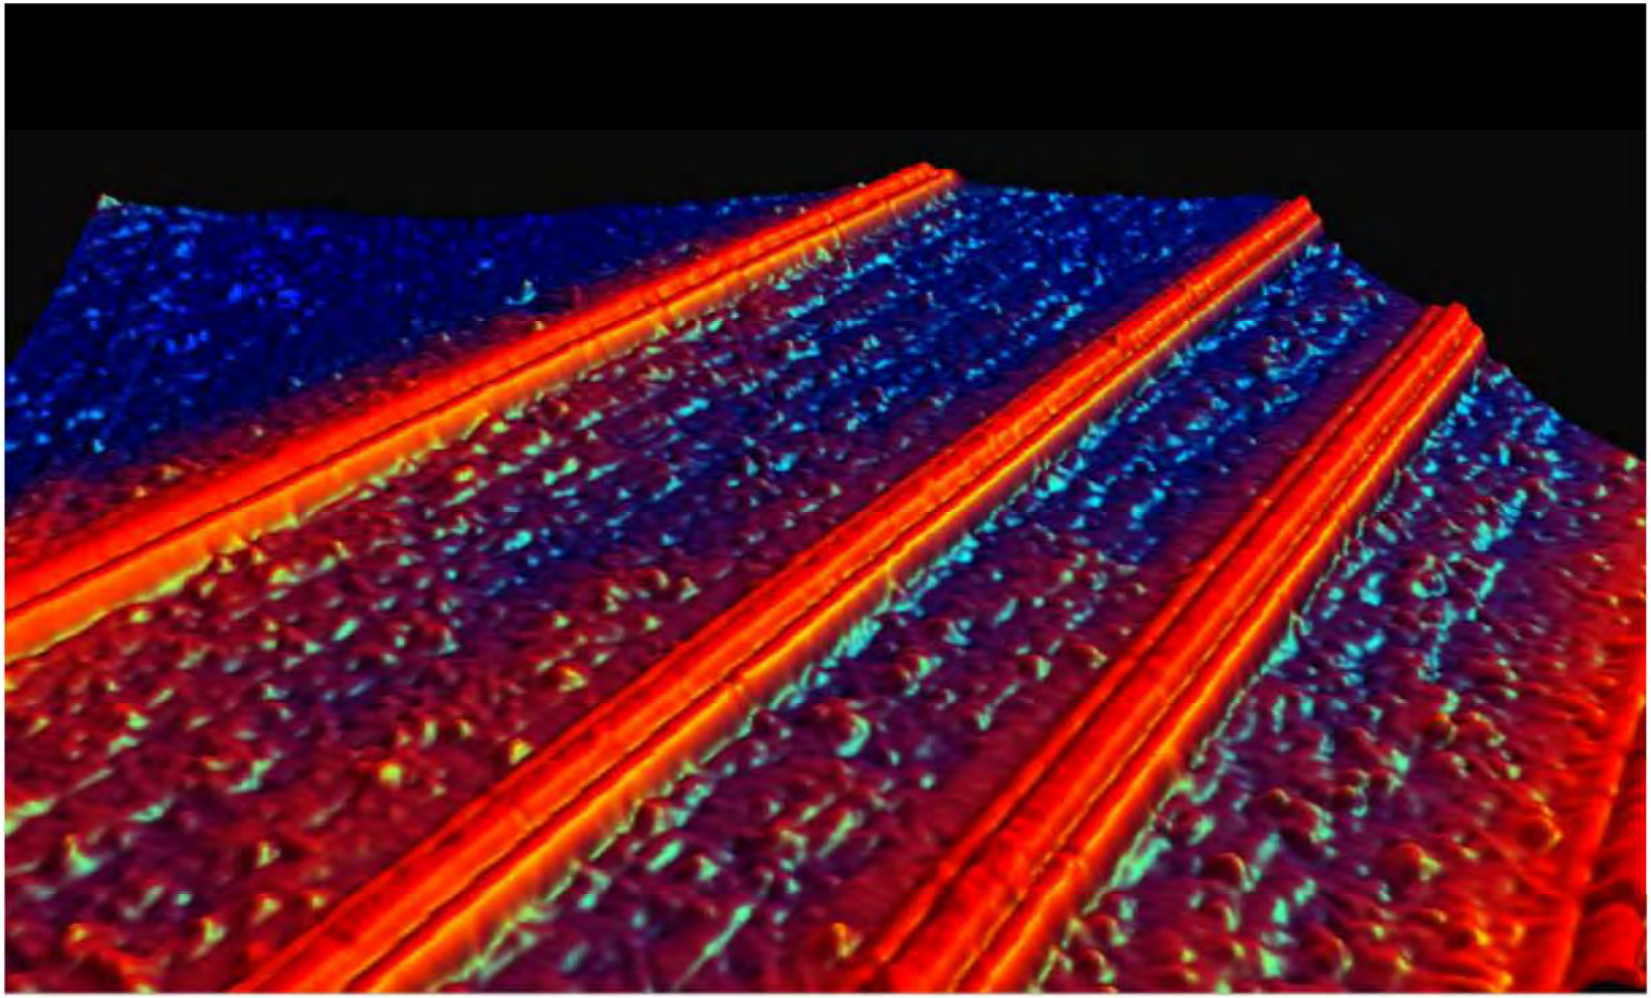

# HPC in the EU: Tuesday, Nov 16 @ 06:57 PST

"The Wissenschaftsrat, the German Science Council, an agency of the federal government, has recommended that three supercomputers of the highest performance be created in Europe

- One of them to be located in Germany. On the current list of the Top 500, there are 128 supercomputers in European countries. [7 in Canada]
- To allow Europe to compete scientifically in future with the US and Japan and in light of recent developments also with China, it was essential that the supply of computing resources be expanded continuously in both quantitative and qualitative terms.

# HPC in the EU: a Consortia approach

"The Science Council recommends that three computers of the highest performance class be created in future by combining national resources at the European level

- On account of the five-to-six-year life time of a supercomputer the creation of three European peak-performing machines of this sort made sense, the Council declared.
- In the event that attempts to situate one such European supercomputer in Germany proved to be successful, the federal government and the German federal states could be expected to face costs of up to 30 to 50 million euros, the Council observed."

# TOP 500 HPC RANKINGS by COUNTRY (June 2005)

www.top500.org

"**The TOP500 project** was started in 1993 to provide a reliable basis for tracking and detecting trends in high-performance computing. Twice a year, a list of the sites operating the 500 most powerful computer systems is assembled and released. The best performance on the Linpack benchmark is used as performance measure for ranking the computer systems. The list contains a variety of information including the system specifications and its major application areas."

List : June 2005

Group by : Countries

Generate Table

| | Count | Share % | Rmax Sum (GF) | Rpeak Sum (GF) | Procs Sum |
|---|---|---|---|---|---|
| Australia | 5 | 1 | 19502 | 28380 | 4492 |
| Belarus | 1 | 0.2 | 2032 | 2534 | 576 |
| Brazil | 4 | 0.8 | 12409 | 23669 | 3844 |
| Canada | 7 | 1.4 | 22281 | 43958 | 7236 |
| China | 19 | 3.8 | 53479 | 90340 | 15906 |
| Egypt | 1 | 0.2 | 1232 | 2419 | 336 |
| Finland | 1 | 0.2 | 1170 | 2253 | 512 |
| France | 11 | 2.2 | 19785 | 36013 | 8172 |
| Germany | 40 | 8 | 80482 | 117311 | 22056 |
| Hong Kong | 1 | 0.2 | 1186 | 2005 | 358 |
| India | 8 | 1.6 | 13995 | 24726 | 4212 |
| Israel | 8 | 1.6 | 14721 | 24067 | 3902 |
| Italy | 11 | 2.2 | 20064 | 32244 | 7012 |
| Japan | 23 | 4.6 | 128626 | 180676 | 33602 |
| Korea, South | 14 | 2.8 | 28867 | 44545 | 7428 |
| Mexico | 5 | 1 | 8619 | 15509 | 2726 |
| Netherlands | 4 | 0.8 | 33320 | 43177 | 13760 |
| New Zealand | 3 | 0.6 | 7719 | 19834 | 3256 |
| Portugal | 1 | 0.2 | 1210 | 1536 | 256 |
| Russia | 3 | 0.6 | 7852 | 11795 | 1788 |
| Saudia Arabia | 5 | 1 | 9300 | 19286 | 3168 |
| Singapore | 1 | 0.2 | 1191 | 1690 | 384 |
| South Africa | 2 | 0.4 | 2848 | 4224 | 704 |
| Spain | 4 | 0.8 | 32432 | 48104 | 5428 |
| Sweden | 3 | 0.6 | 7239 | 8896 | 1526 |
| Switzerland | 3 | 0.6 | 25008 | 37058 | 10792 |
| Taiwan | 2 | 0.4 | 4089 | 6784 | 1184 |
| United Kingdom | 32 | 6.4 | 85559 | 173196 | 32978 |
| United States | 277 | 55.4 | 1039518 | 1584367 | 382486 |
| Venezuela | 1 | 0.2 | 1198 | 1536 | 256 |
| All | 500 | 100 % | 1686933 | 2632133 | 580336 |

# Three Rings: Canadian HPC Needs

## Canadian High Performance Computing Needs

The array of canadian research projects each have unique high performance computing requirements.

**Ring 1**
Desktop Computers
(1-64 processors)

**Ring 2**
Small Cluster System
(64-300 processors)

**Ring 3**
Supercomputers /
Terascale System

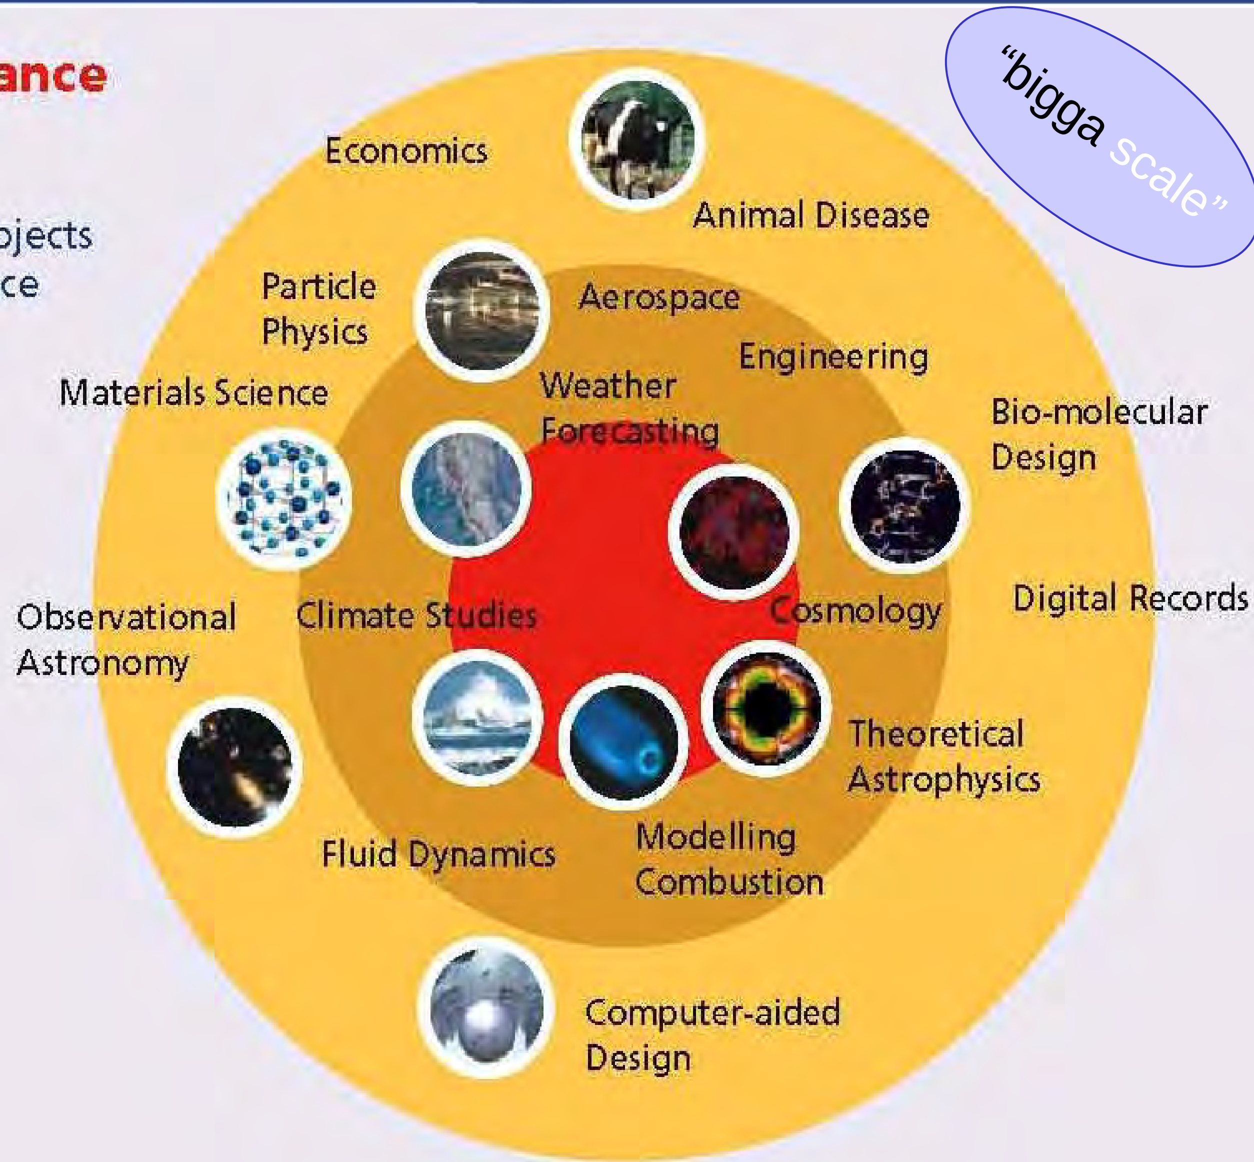

# LRP **Recommendation 2**: Outreach

- Establish a national initiative, **ImPACT Canada**, to
  - Ensure HPC awareness across all Canadians
  - Keep Canada Competitive
  - Assist in leadership and coordination between MRP's
  - Maintain the national HPC overview
  - Encourage growth and training of HQP
  - Establish mechanisms and opportunity for industry access and use of facilities and HQP

# LRP Recommendation 3: Government Funding

| | Current | 2006 | 2009 | 2012 |
|---|---|---|---|---|
| **HPC Capital Infrastructure** | | | | |
| Consortia/MRP | 40 | 44 | 49 | 54 |
| High-end Facility | 0 | 10 | 12 | 14 |
| **HPC Operations** | | | | |
| Human Infrastructure | 1 | 13 | 15 | 17 |
| Facilities | 2 | 8 | 9 | 10 |
| ImPACT Canada | 0 | 1 | 2 | 2 |
| **Totals** | **43** | **76** | **87** | **97** |

# LRP Recommendation 4: HPC Supporting Research

- Computationally-based research and people to support this research need to be formally and actively recognized

  within the research plans of universities

  - within support programmes for universities and research being provided by Government
  - by the public

CASE STUDY

*"In the long term, nanotech is all about computing. It's the only tool we have to bridge the chasm between theory and experiment for hard problems that resist analytical treatment. Which, of course, means almost all problems . . ."*

***Dr. Mark Freeman***
***Canada Research Chair***
***in Condensed Matter Physics***
***Professor of Physics, University of Alberta***

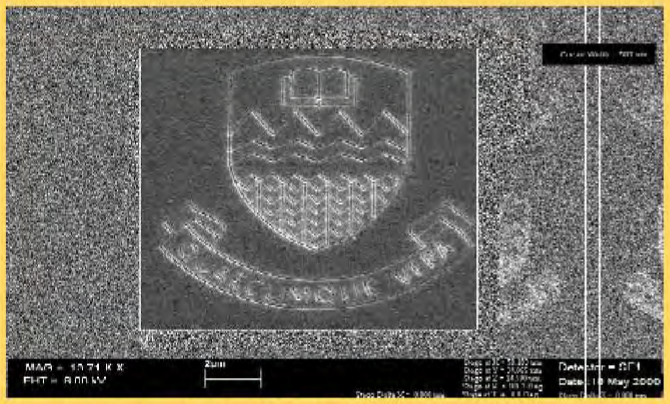

Using nanoscale instruments, researchers were able to reconstruct the University of Alberta logo in nanoscale form. The logo now fits on the head of a pin.


Dalhousie Distributed Research Institute and Virtual Environment

# East meets West:
# Collaboration goes National

**Welcome to D-DRIVE whose mandate is** to study and develop resources specific to <u>distributed research</u> in the sciences with first client groups being the following communities

- High Performance Computing
- Mathematical and Computational Science Research
- Mathematics and Science Outreach:
  - ✓ Educational
  - ✓ Research

IRMACS (SFU, above)
D-DRIVE (Dal, below)

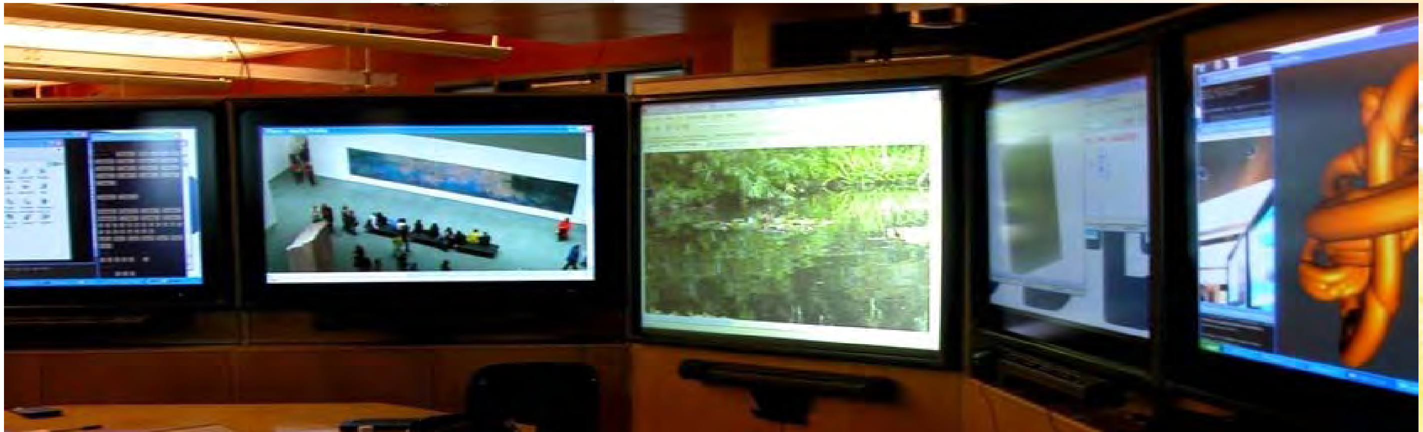

**The 2,500 square metre IRMACS research centre**

# Centre seen as 'serious nirvana'

**Trans-Canada Graduate Seminar**
11.30 PST - 3.30 AST on Thursdays

**April 07, 2005 , vol. 32, no. 7**

By Carol Thorbes

Move over creators of Max Head-room, Matrix and Metropolis. What researchers can accomplish at Simon Fraser University's IRMACS centre rivals the high tech feats of the most memorable futuristic films.

The $14 million centre's acronym stands for interdisciplinary research in the mathematical and computational sciences. The centre's expansive view of the Lower Mainland from atop Burnaby Mountain echoes its limitless potential as a facility dedicated to fostering interdisciplinary research among scientists whose primary lab tool is the computer.

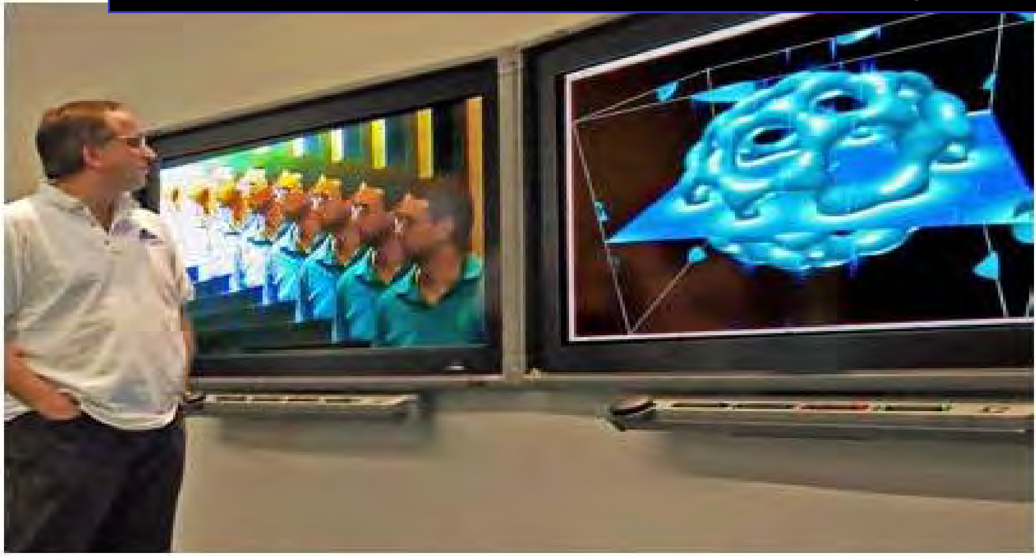

*SFU mathematician and IRMACS executive director Peter Borwein (left) communicates with IRMACS collaboration and visualization coordinator Brian Corrie. To the right of them another plasma display portrays a 3D image of a molecular structure.*

A newly constructed 2,500 square metre space atop the applied sciences building, the centre has eight labs, five meeting rooms and a presentation theatre, seating up to 100 people. They are equipped with easily upgradeable computational, multimedia, internet and remote conferencing (including satellite) technology. High performance distributed computing and clustering technology, designed at SFU, and access to WestGrid, an ultra high-speed, interprovincial network, with shared computing and multimedia resources, make IRMACS unique in Western Canada.

- ✓The building is a also a 190cpu G5 Grid
- ✓At the official April opening, JMB gave one of the four presentations from D-DRIVE

irmacs

interdisciplinary research in the mathematical and computational sciences

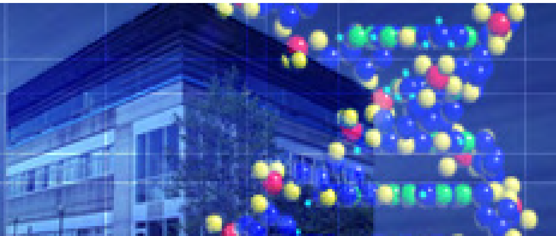

# LRP **Recommendation 5**: Research Asset Management

- Requests for computing resources in support of research should be evaluated against existing and planned HPC facilities and infrastructure

- Balance needs, demands and costs

  - ✓ We are pragmatic dreamers

  - ✓ Low-end funding from grant councils remains essential

- Consistent with new **CFI National Platforms Fund** envelope

# CFI Response to LRP: February-March 2005

"The National Platforms Fund provides generic research infrastructure, resources, services, and facilities that serve the needs of many research subjects and disciplines, and that require periodic reinvestments because of the nature of the technologies. The Fund is established to deal first with High Performance Computing, but may be applicable in other cases."

# LRP **Recommendation 6**: Supporting Infrastructure

- Support for advanced networks and integrated network applications such as "Grid" must continue through CANARIE Inc. and Grid Canada, NSERC, SSHRC etc
  - ✓ 'trusted' secure networks
  - ✓ independent of fate of US funded initiatives
  - ✓ social and scientific impacts

*Intelligent infrastructure*
requires an
*Intelligent computing infrastructure*

CASE STUDY

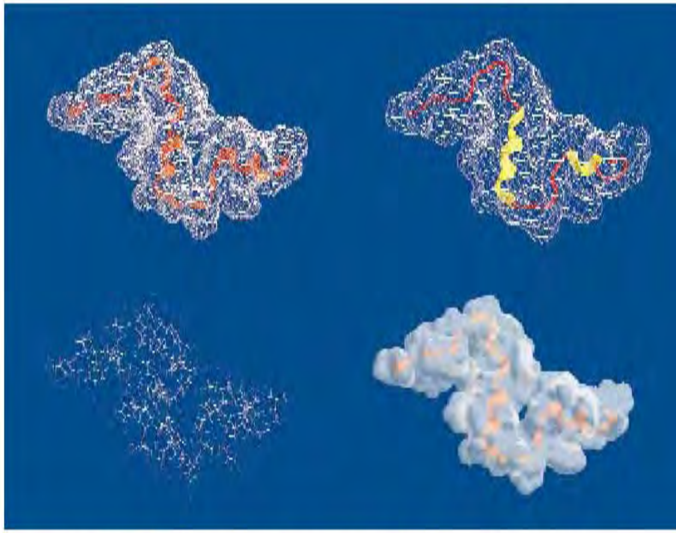

One of the leading approaches to drug design for Alzheimer's disease is to engineer molecules that can bind to the beta-amyloid peptide (shown here). Since this peptide has never been effectively crystallized, such structural studies are totally dependent upon molecular modelling and simulation studies within a high performance computing environment. These computer-aided studies are crucial in enabling the design of drug molecules that interact selectively with discrete disease-causing molecules.

# Additional Benefits

- ✓ HPC is the engine of discovery that enhances multi-discipline research
- ✓ HPC and LRP ensures a stable collaborative research environment
- ✓ LRP will enhance research and allow scientific agencies to accomplish more with their precious research dollars (complementarity)
- ✓ HQP in all disciplines will be HPC literate

Strong message from funders

*"We must all hang together, or assuredly we shall all hang separately."*

(signing Declaration of Independence)

# Critical Extra Need: Framework for Large Science

**HPC as National "Engine for Discovery"**

- To assure that other major science projects have the HPC and IT resources they need to assimilate and to propagate information coming into and out of their project and community
- The Light Source, TRIUMF, Earth Observation Satellite….

> "Overlooking IT has a long and sad history in international science. For example, NASA's Earth Observing System has been crippled by an inability to retrofit data handling issues that were neglected before launch and it is not clear that the Canada Light Source is much better prepared."

# Frequently Asked Questions (FAQs)

The issues raised by the following questions are dealt with in detail through the body of this report. They are listed here to illustrate the essential rationale behind the report. References to sections dealing with these issues are provided (e.g., E2 – Section 2 of Executive Summary; C2 – Chapter 2).

### *If HPC is so beneficial, why can't industry support the infrastructure ?*

Industries develop mature applications derived from basic, curiosity-driven academic research. Few corporations can afford to invest in this critical early discovery phase. Also, many HPC applications are in the public domain: infrastructure (transport), big science (meteorology, astronomy, physics), government jurisdictions (fisheries, health, education). Increased support for HPC in the public domain would be primarily a government responsibility.

(Relevant Sections: Neurochem Case Study, "One Day" Story)

### *Why can't the researchers apply for HPC funding through existing mechanisms ?*

Some funding mechanisms exist for mid-range HPC infrastructure (e.g. CFI) but there is no long-term coordination for a national HPC strategy that oversees equipment, personnel and infrastructure, and provides the linkages between government, academe and industry. There is no mechanism for securing the large investment needed for a top-end HPC facility in Canada. Finally, the rapid timescale required for refresh of the technology is not well suited to current funding mechanisms.

(Relevant Sections: E2 ; E3 ; C3-Funding for HPC Support Staff ; C4 All ; C5 Introduction)

### *Why can't we just link up lots of increasingly-powerful, cheap desktop computers to handle the big problems ?*

This is possible if the problem can be broken down into many self-contained components, each of which can be handled by one desktop computer. However, there are many important research applications which simply cannot be scaled in this manner, since they have CPU, memory, data, or communications needs that far exceed that of a desktop computer. Such problems can only be solved by HPC equipment.

(Relevant Sections: E1 ; C2 - Capability and Parallel Applications sidebars ; C2 - Central Rings section)

### *Isn't Canada too small a population to compete in this field ?*

A number of smaller countries have invested in HPC strategies, recognizing this is a key engine of future economic development. Moreover, Canada has unique characteristics that make it ideally suited for HPC leadership. Canada is not so small that it lacks the resources (human, financial, knowledge) to make a significant impact on the world stage. Neither is it so large that national efforts are so daunting that they tend to break up into incompatible, regional initiatives. This principle has already been demonstrated with CANARIE's CaNet*4 network. Canada's vast extent but limited number of major cities makes it ideal for a coordinated computational grid which combines the high-speed data communication of CANARIE with the HPC resources described in this document. Much like the railways of the last century, such a national grid would bind together Canadian science and society in a manner that few other countries could match.

Moreover, external metrics demonstrate that Canada's research community ranks second among G8 countries (behind the UK) in research productivity. Canada is well-placed to be a leader not a follower.

(Relevant Sections: C1-Canarie Sidebar ; C2 Networks and Grids section and Sidebars ; C3 Scientific Productivity Sidebar)

### *Why invest in expensive HPC ? Why not simply wait for other countries to do the basic research while we concentrate on development and production phases ?*

The principal benefits of information technology accrue to those who develop it first. If Canada waits for other countries to make the breakthroughs, then we will inevitably fall behind in the race to bring new knowledge to the marketplace. This has never been more true than in this field. Imagine if Silicon Valley had grown up in a Canadian city…

(Relevant Sections: C1 Introduction)

### *Why can't we just buy time on computers in other countries?*

This is ultimately more expensive and counter-productive to the development of Canada as a leader in high-end research and development. As discussed in Chapter 2, the true cost of running one application on a tera-scale facility may run from $100K's to $M's. This is all money lost to Canada. An offshore HPC strategy, if scaled up to a national level, would cost more than the in-house strategy proposed here. The use of the Japanese Earth Simulator to model a hurricane weather system is an illustration of the potential benefits and pitfalls. The true cost of this work was $5-10M; the Canadian team was able to provide technical resources instead of payment but still suffered from low priority access and onerous international bureaucracy.

After 9 months of use the work is still only about 40% complete due to low priority. Buying CPU time, if it could be secured, would be a signal that Canada is not a significant player in the G8. We would lose many young computational scientists to other countries since they could not compete from Canada due to lower priority access than their peers in the country providing the facility.

(Relevant Sections : C2 The Central Rings section and footnote ; C2 Earth Simulator Sidebar, C4)

### *What about provincial support ?*

It would be desirable for provincial governments to participate in this national strategy. This would allow them to augment the national HPC infrastructure in ways that address specific provincial HPC priorities. Addressing these individual perspectives was beyond the scope of this plan and would require discussions between the two levels of government.

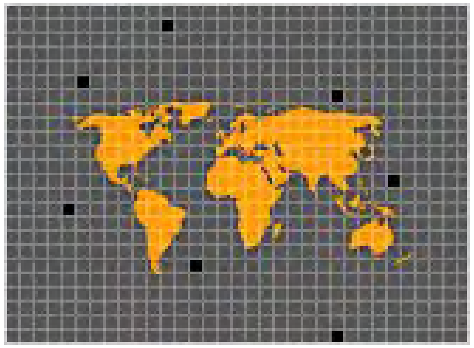


HPC • CHP
ENGINES OF DISCOVERY: THE 21ST CENTURY REVOLUTION
THE LONG RANGE PLAN FOR HIGH PERFORMANCE COMPUTING IN CANADA
March 2005
Thank you
c3.ca