Historical perspective of computing and networks in HEP in Canada

Randall Sobie
IPP/University of Victoria

Rapid and remarkable changes

Rapid and remarkable changes

Grade 10 at UToronto

TRIUMF PDP11 DAQ
Meson Hall (PhD)

Grade 13 – last year of Slide Ruler

MSc thesis

1st-Year (HP35)

A selected history of HEP computing in Canada

Bubble chamber analysis

UA1 event display

OPAL online reconstruction system

Beowulf clusters

The Grid

ATLAS distributed computing and Tier-1 facility

Cloud computing

Research Networks

Computing in Canada (the “Alliance”)

Bubble Chamber Scanning System (J. Prentice et al)

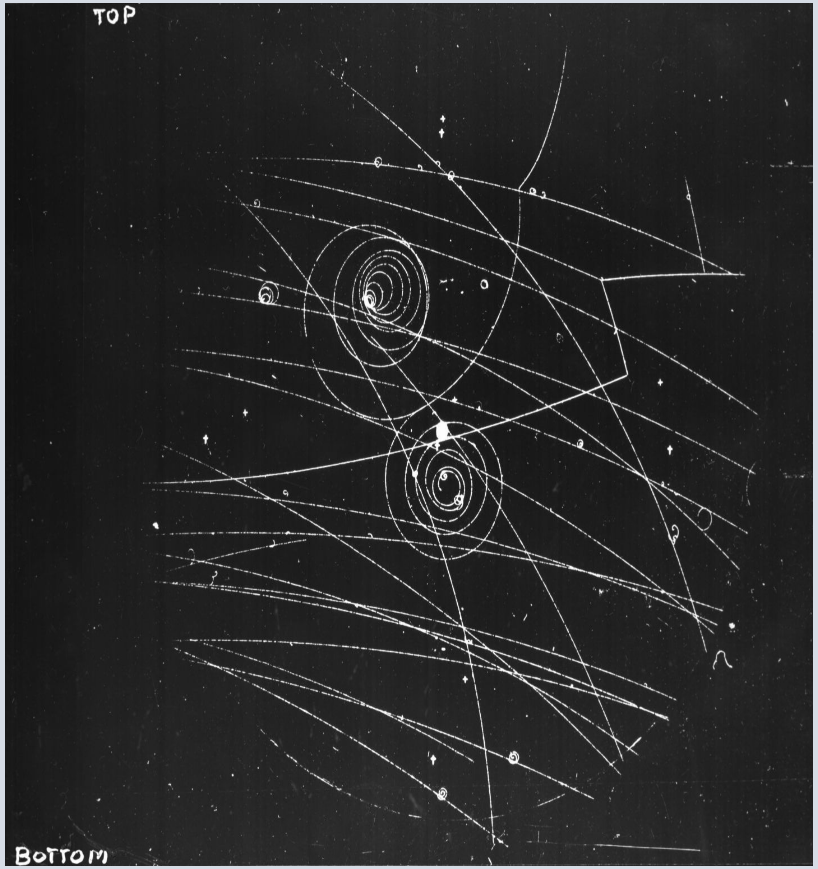

University of Toronto archives:

"…the 1960s up until about 1972, he, along with University of Toronto colleague Dick Steenberg, participated in the Bubble Chamber experiments at the U.S. Argonne National Laboratory, thus becoming Canadian pioneers in the experimental high energy physics.

In 1969, they were responsible for having POLLY, an automatic measuring machine, the first of its kind in Canada, installed at the University of Toronto. "

https://discoverarchives.library.utoronto.ca/downloads/james-d-prentice-fonds.pdf

The images were digitized and are used for a particle physics course: https://www.physics.utoronto.ca/~phy326/hep/

UA1 Experiment (A. Astbury et al)

Discovery of the W and Z bosons

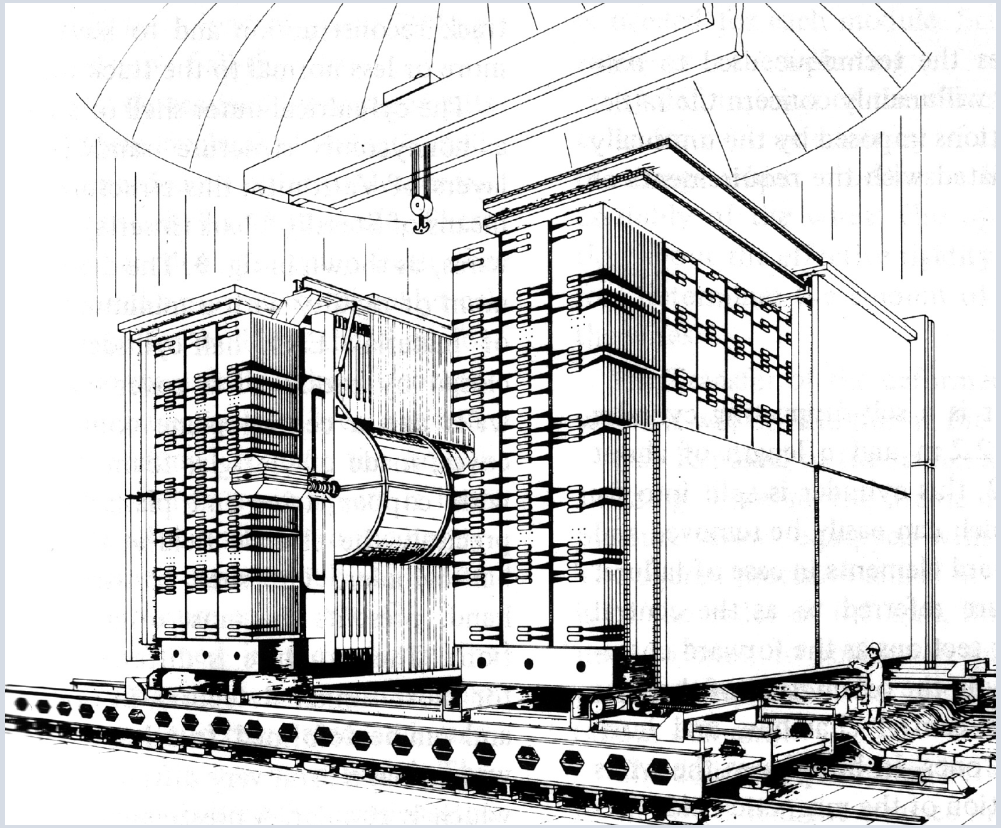

Access to IBM main frame computers
Early use of Monte Carlo in "monojet" search

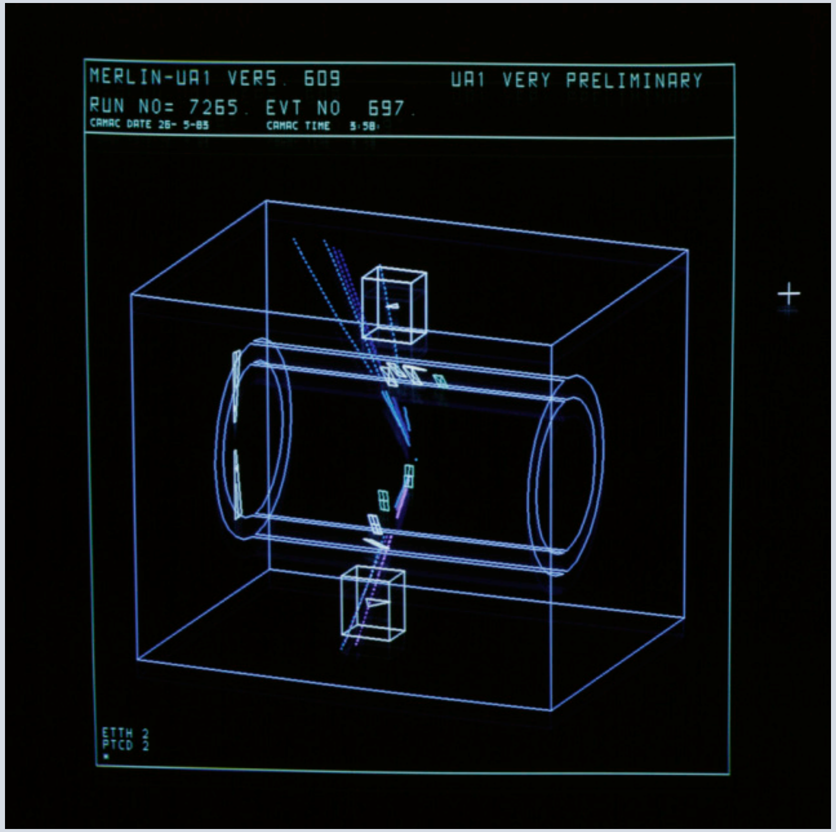

Every event scanned by physicists
3D events images

HBOOK (UA1/UA2 era) and PAW (LEP era)

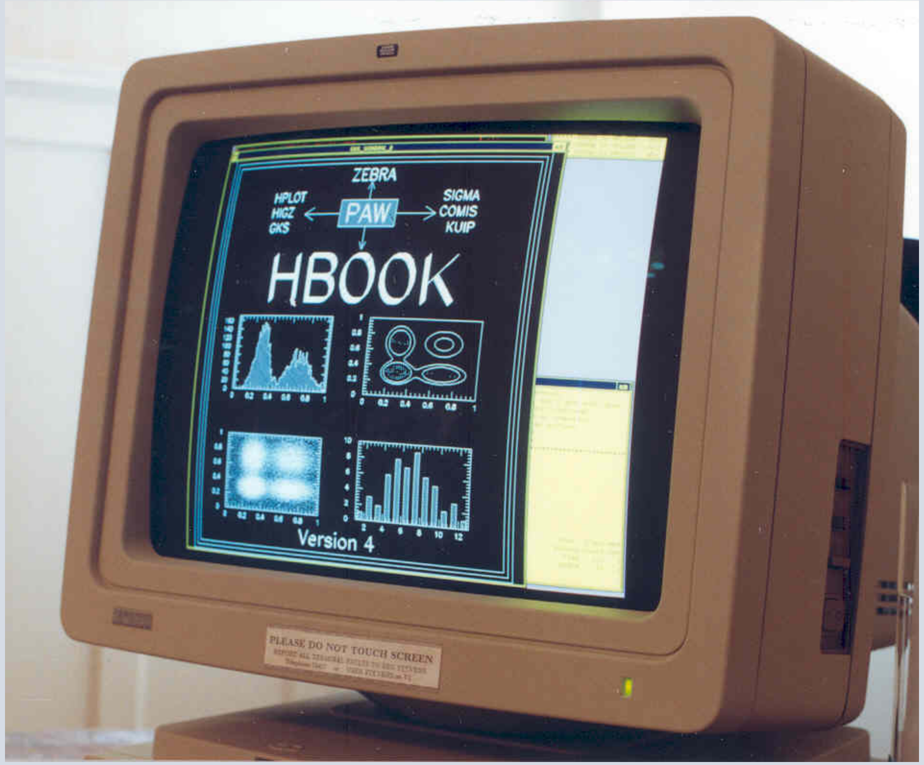

PAW demonstration at Oxford CHEP 1989

Networks enabled new technologies

World Wide Web

The WorldWideWeb (W3) is a wide-area hypermedia information retrieval initiative aiming to give universal access to a large universe of documents.

Everything there is online about W3 is linked directly or indirectly to this document, including an executive summary of the project, Mailing lists , Policy , November's W3 news , Frequently Asked Questions .

What's out there?
: Pointers to the world's online information, subjects , W3 servers, etc.

Help
: on the browser you are using

1990

UBC CSC Dept had the first site in Canada

UVic Physics had the next site one week later

Both were one of the first 50 sites in the world

OPAL Experiment (B. Carnegie et al)

INSTITUTE OF PARTICLE PHYSICS

Study of the Z boson at LEP

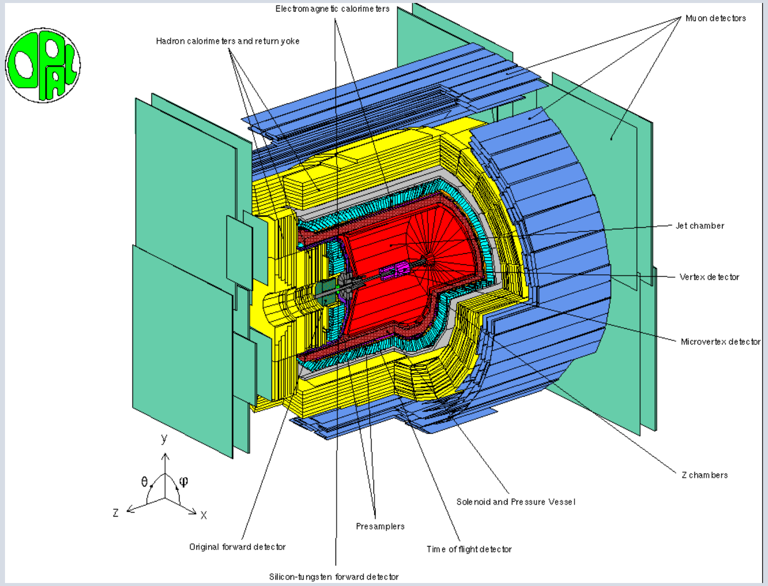

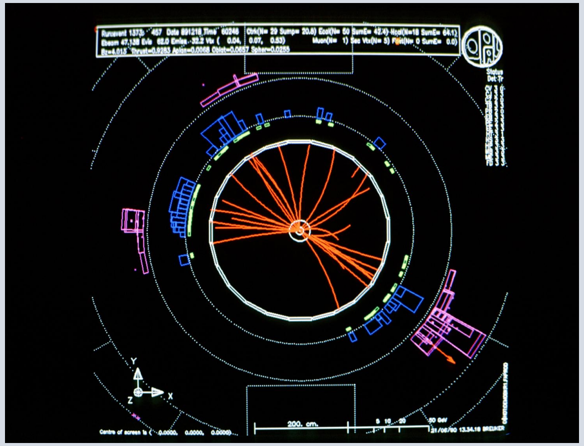

Emergence desktop (workstation) Computing

Wide adoption of PAW and use of MC

Computing resources still at the lab

OPAL reconstructed the data in near real-time

Data written to optical disk and stored in a small robotic library

A Rewritable Optical Disk Library System for Direct Access Secondary Storage

This autochanger system can store up to 20.8 Gbytes of data on-line. Applications include archival storage, automated backup and recovery, and document storage and retrieval.

by Donald J. Stavely, Mark E. Wanger, and Kraig A. Proehl

HEWLETT-PACKARD MANUFACTURES a wide range of computer peripherals. Customers for these peripherals include not only users of HP systems, but also OEM customers and others who use HP peripherals with non-HP host systems.

Supplying peripherals to OEM customers has been a major initiative for Hewlett-Packard and has had a large impact on how we plan and evolve our business strategies. To be successful in the OEM business has required that we develop a broader and more timely understanding of the market than we had in the past. We feel that our experience as a system company gives us valuable insights into how our peripherals work in systems and applications to solve real customer needs.

HP's Greeley Storage Division is responsible for high-end secondary storage devices that are used for backup and archival storage on computers, mainframes, and networks of workstations. Our current product offering is a family of low-cost, autoloading, streaming, 1/2-inch GCR tape drives.[1]

As we looked to the future, we naturally focused our attention on advances in tape technology. Emerging products were using air bearings for media reliability, a thin-film 18-track head for very high transfer rate, and a compact tape cartridge for ease of handling. Initially, this technology seemed a good match to what our current HP and OEM customers needed. Customers were asking for faster backup to reduce planned system downtime—or more accurately, to keep from increasing their downtime as their disk storage requirements grew. They also need ever higher levels of reliability to minimize unplanned system downtime.

Unfortunately, simplistic market research—asking customers what they want—often yields only predictable and simplistic answers. They want what they have now, only faster, cheaper, more reliable, and so on. In other words, customers may be too close to their problems to see them from a new perspective.

We evolved a much more powerful market research process that consists of three steps. The first step is to gain a thorough knowledge of how customers do business. What applications do they run? How much disk space do they have? How do they do backup today? What else do they use tape drives for?

The second step of the process is to try to solve customers' problems in the abstract—matching available technologies to a high-level model of each customer's business. The last step is to present a coherent vision of the future back to our customers. In essence, we are trying to help them look past the limitations of today's solutions and help them architect the solutions of tomorrow. We call this developing an "imaginative understanding of user needs."

Fig. 1. *The HP Series 6300 Model 20GB/A rewritable optical disk library system stores up to 20.8 Gbytes of data on 32 optical disks. An autochanger automatically selects the correct disk and inserts it into one of two internal drives.*

Cluster computing ("Beowulf") for the CDF Experiment

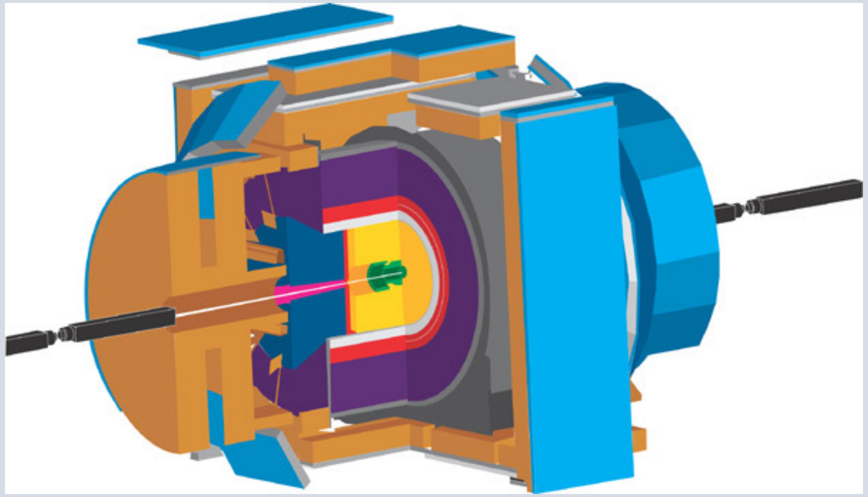

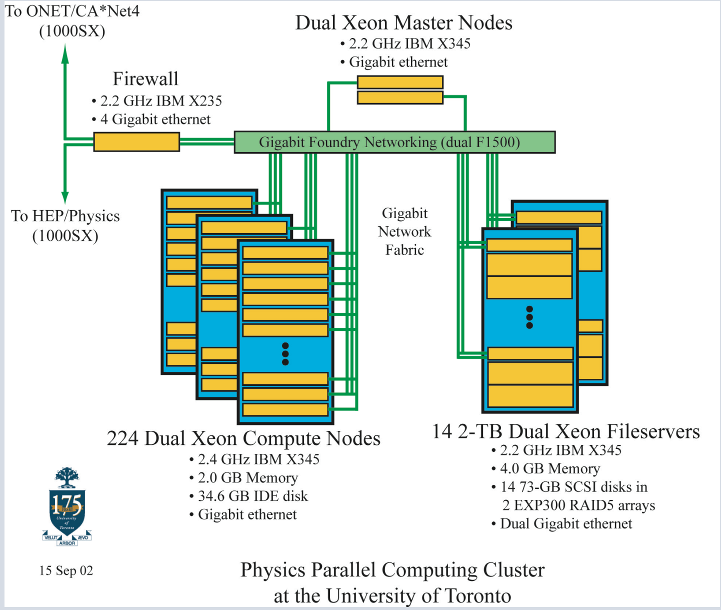

2000 Groups started to build "Beowulf clusters"

".. a high-performance parallel computing cluster from inexpensive personal computer hardware"

Toronto-CDF group build one of the largest clusters in Canada for simulated event (MC) generation

HEP continues to use (commodity) processors as primary resource for computing (e.g. Intel and AMD processors)

Migration from "main frame" computing as one is able to create fast local-area networks to connect servers

Gradual adoption of non-lab based computing resources

Distributed computing – BaBar, CDF, …, ATLAS/LHC era

High speed networks enabled us to use collectively use resources at remote sites away from the host laboratory

Initially, remote sites did Monte Carlo generation or held dedicated samples of select analysis samples

LHC era (early 2000s)
Emergence of technologies to integrate the remote facilities

Computing grid – analogy of the "electrical grid"
Access to computing resources from any source

Still remains the underlying design of the WLCG computing (WLCG World-wide LHC Computing Grid)

The Grid was not adopted by Industry (cloud)
Grid is optimal for continuous computing with similar requirements
Cloud is designed for on-demand computing

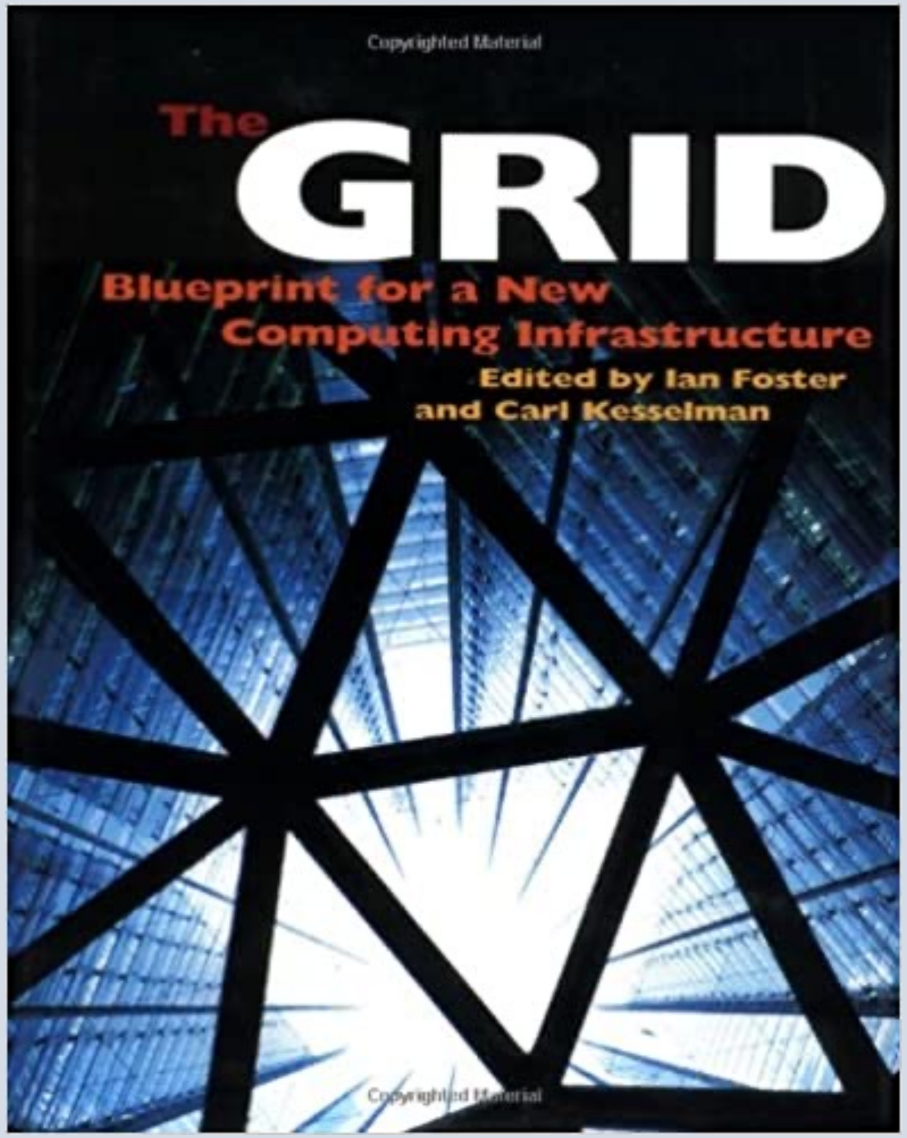

Open-source software

SLAC (BaBar) and CERN (LHC) initially adopted Objectivity for storing events (ODBMS – commercial database product)
Viewed as a way to exploit commercial development and reduce manpower effort/costs

See CMS Computing Technical Proposal https://indico.cern.ch/event/408139/contributions/979837/attachments/815772/1117804/chepdb.pdf
BaBar and Objectivity: Operational Aspects of Dealing with the Large BaBar Data Set (CHEP2004)
https://www.slac.stanford.edu/pubs/slacpubs/9750/slac-pub-9970.pdf

Data retrieval was found to be very slow as the number of users accessing the DB increased

Shifted back to R.Brun etal software package ROOT in mid-2000s

HEP stores its data in ROOT-format (robust, reliable, community supported, ..)

Recognition that HEP is best served by open-source software
(operating systems, databases, analysis tools, ..)

INSTITUTE OF PARTICLE PHYSICS

Computing in High Energy Physics Victoria 2007

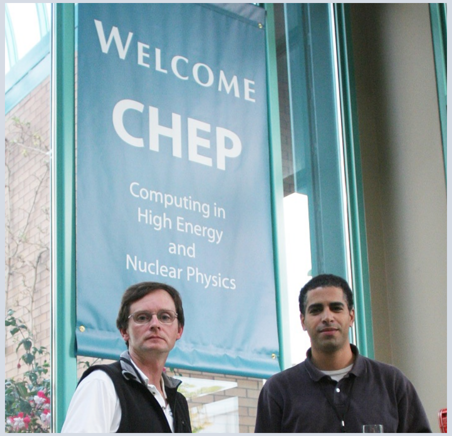

LHC COMPUTING FOCUS

LHC computing stability emphasized at CHEP '07

As preparations for the LHC proceed, this year's main conference on computing in high-energy physics focused on getting ready for the fast-approaching onslaught of data.

A CHEP '07 plenary session. This year, the conference was held in Canada for the first time. (Courtesy Albert Pace.)

2-day WLCG Workshop (200 attendees)
5-day CHEP Conference (470 attendees)
(fun with wireless networks and food)

Just prior to the first LHC beams 2008
WLCG (Les Robertson)
CMS (TJ Virdee)
Summary (M.Kasseman)

WLCG - World-wide LHC computing grid

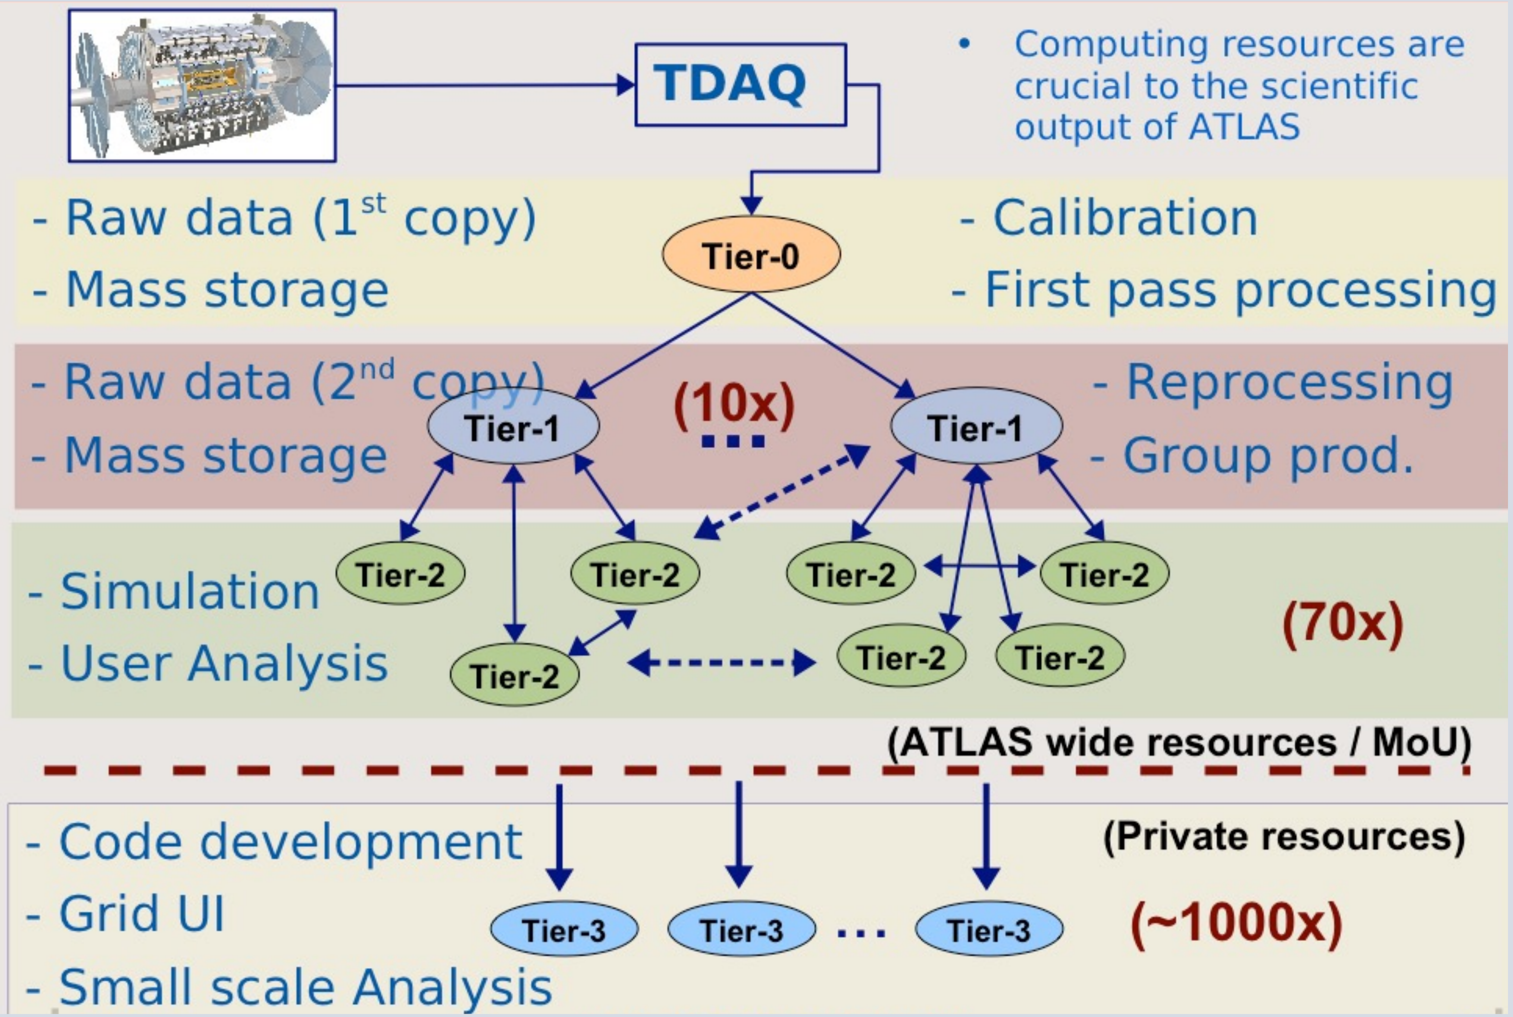

First global computing infrastructure for research

Canadian role:

ATLAS Tier-1 at TRIUMF

Tier-2 computing using the Compute Canada resources across the country

ATLAS Tier-1 Computing Centre

Prototyping began to help secure funding in 2002

MoU signed with CERN in 2006

ATLAS-Canada consortium led by SFU

- $25.6M in capital funds (CFI and BCKDF)
- $7.1M for operation

Full-scale 24/7 operations with dedicated network to CERN

Initially located at TRIUMF, moved to SFU in 2018

Capacity In 2022: 18 PB disk, 31 PB tape, and 7824 cores

Tier-1 cluster in 2006

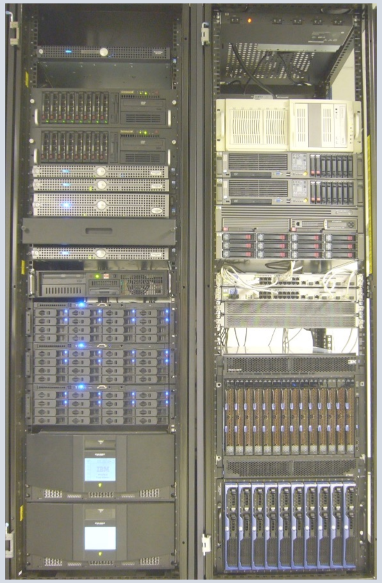

Emerson/Liebert XD cooling

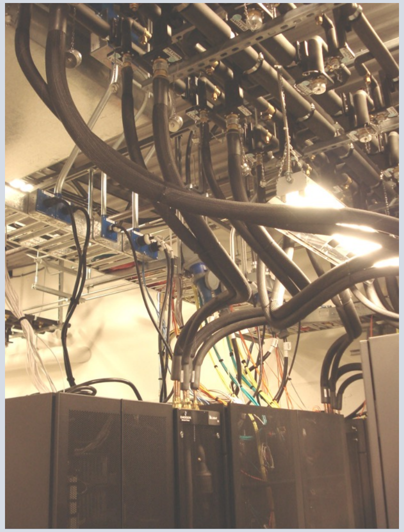

Special cooling infrastructure

refrigerant-based, not chilled water

Cloud computing

Adopted by Industry

- Many providers around the world(Amazon EC2, Microsoft Azure, Google GCE)
- On-demand services; ideal for users with sporadic (time-dependent) workloads

Research computing

- “the Grid” still suits the HEP environment (constant workloads)
- ”In-house” clouds provide 5-10% of the HEP computing resources
  - Canadian astronomy is cloud-based for image analysis
- Commercial clouds less used than 5-10 years ago; cost is still high
  - Canadian medical research community are large commercial cloud users
  - ISED wants the Alliance to increase commercial cloud use

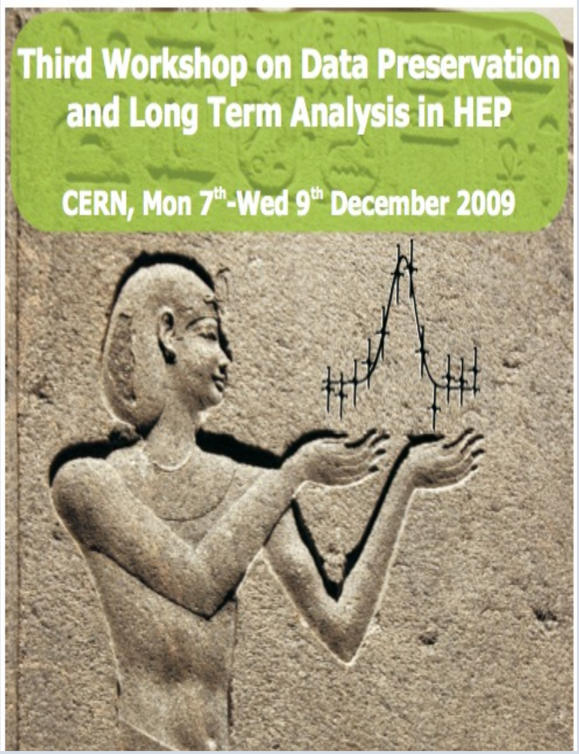

Canadian cloud computing

- Distributed cloud for Belle II, ATLAS, BaBar and Dune (6000 cores on 3 continents)
  - Mainly in-house clouds but commercial clouds with in-kind awards
  - Ideal for the BaBar data preservation system (GridKa and Victoria)
- Belle II Raw Data Centre (“Tier-1”) will be cloud-based

Networks have enabled our distributed computing systems

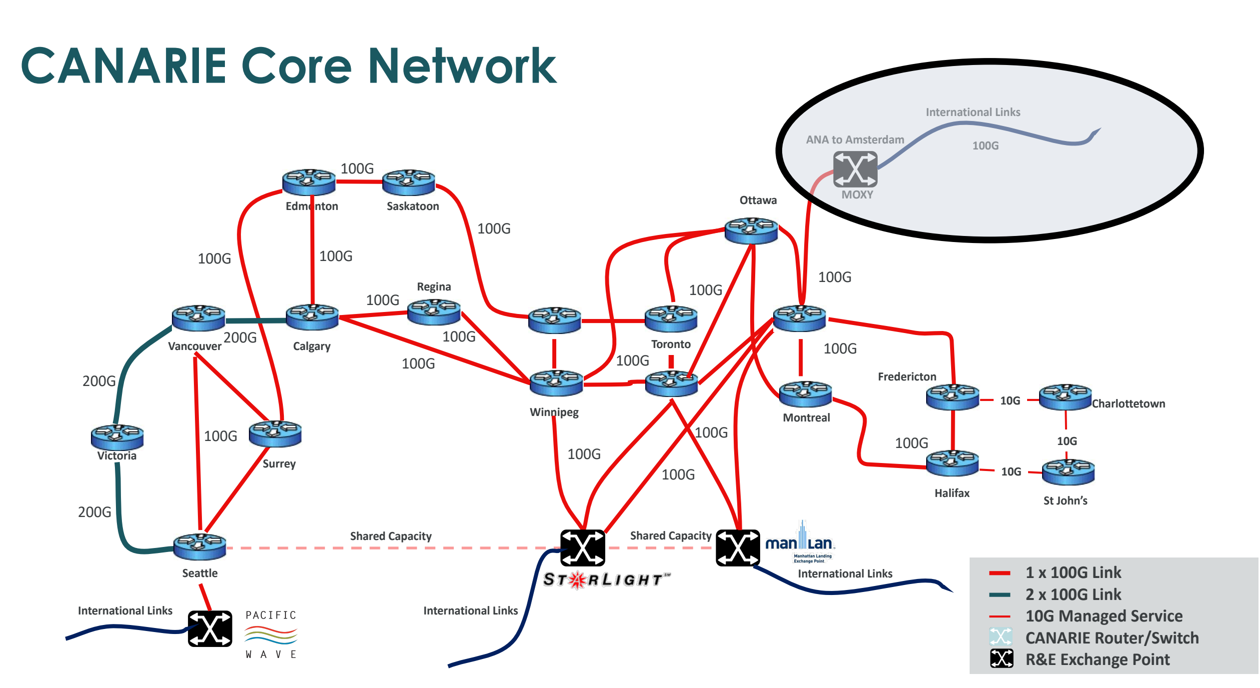

CANARIE Established 1993

Canada's national research network organization

Provides the national network backbone

200+ Gbps across Canada today

International connections (multi-100Gbps)

Canada-Amsterdam link (MOXY)

PACWAVE (Seattle), Starlight (Chicago), MANLAN (NYC)

CANARIE contributes to the cost of the international links

(HEPNET used to pay $100K/year for a shared link with our US colleagues)

Canadian physicists achieve 100 gigabit/second transatlantic transmission, enabled by CANARIE and its global partners

HEPNET/Caltech had exclusive use of the Canadian 100G link between Ottawa and CERN

CANARIE and particle physics

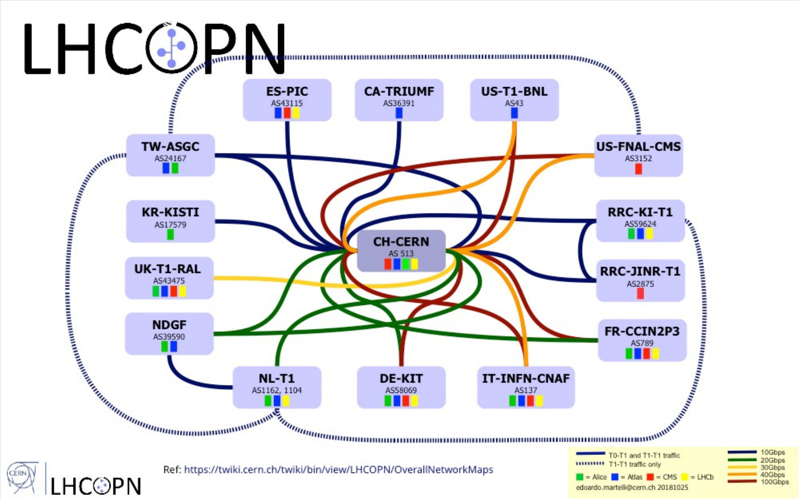

CANARIE provides the dedicated link to the ATLAS Tier-1
LHCOPN Network

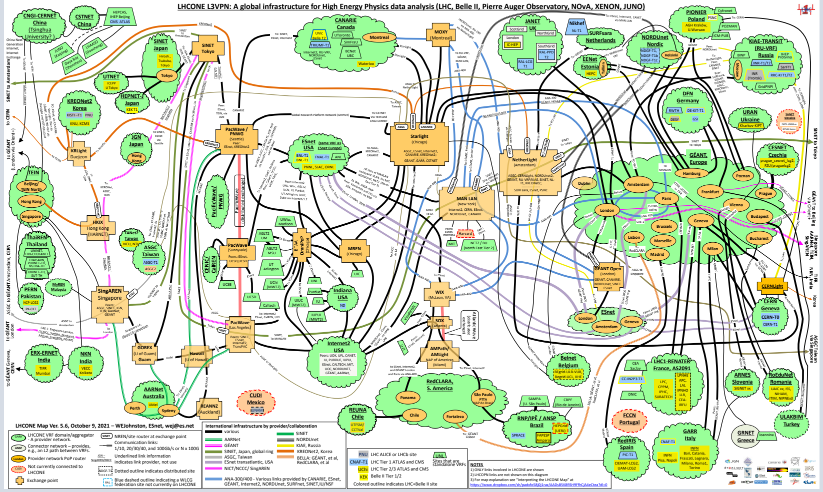

CANARIE provides the particle physics routed network
LHCONE Network
Used by LHC and non-LHC experiments

Research computing in Canada

INNOVATION
Canada Foundation for Innovation
Fondation canadienne pour l’innovation

CFI has provided most of the funds for research computing in Canada 2000-2022

Individual researchers/organization
Continues funding “specialized” computing systems. e.g. ATLAS-Tier1 Genome Canada, ..

Regional consortia created around 2005

WestGrid dissolved 2022

Compute Canada established 2012

Digital Research Alliance of Canada

Digi Home search Alliance of Canada | Alliance de recherche numérique du Canada

Established 2020 to unify computing, research data management and software for Canadian researchers
(expanded role compared to Compute Canada)

Funded by ISED (directly) in partnership with provincial governments, consortia and institutions
(complex in-kind matching funding requirements)

Replaced Compute Canada in April 2022
(facilities and personnel)

Preparing a budget and contribution agreement for FY 2023-2024
(using existing funding awarded in 2019 budget)

CANARIE remains responsible for research networks but given a mandate to develop cybersecurity strategies

Outlook for HEP Computing

Amazing developments over the past decades, what might happen in the next 10-20 years?

Distributed systems

Continue evolution to decoupled systems (compute-storage-personnel)

Co-locating power intensive systems are green power generating stations

Adoption of new and emerging technologies

GPUs, ARM (low cost, low power) processors

Expanded us of opportunistic resources (cloud and HPCs)

Better software and new techniques

Multithreaded applications that fully exploit the 256+ cores on a single server

Using of virtualization and container technologies, and increasing use of AI and ML

Faster local networks

Getting the data from the storage into the servers (solid state technologies)

Personnel

Will we find a career path for “computer-physicists”?