Personal Overview of Computing

The more things change….

- More complexity
- More security requirements
- Not enough manpower
- Hardware “free” to purchase – costly to own.
- Persistent heat/electrical problems
- Waves of change come rapidly & with more intensity

Disclaimer: Speculation included in presentation…

Computing: Like Growth of Vancouver – only much faster!

Vancouver from Cambie Bridge 1978 & 2003

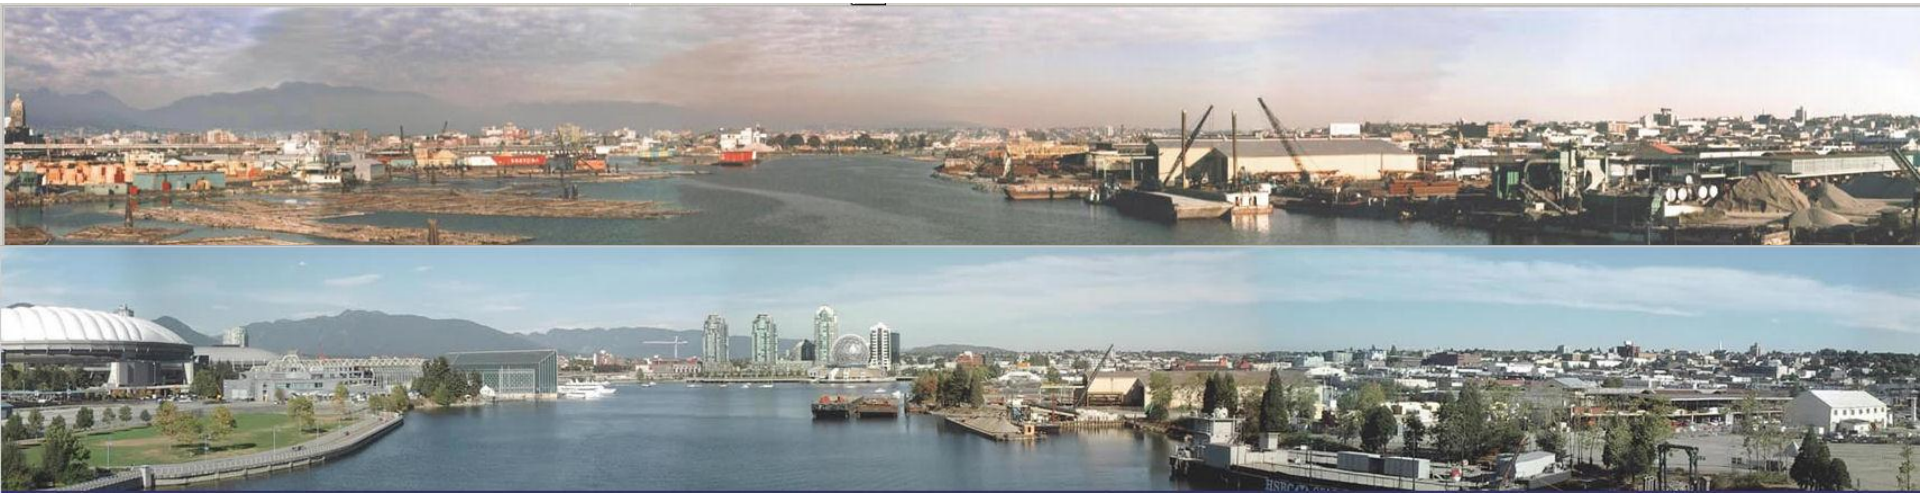

Vancouver from Granville Bridge 1978 & 2003

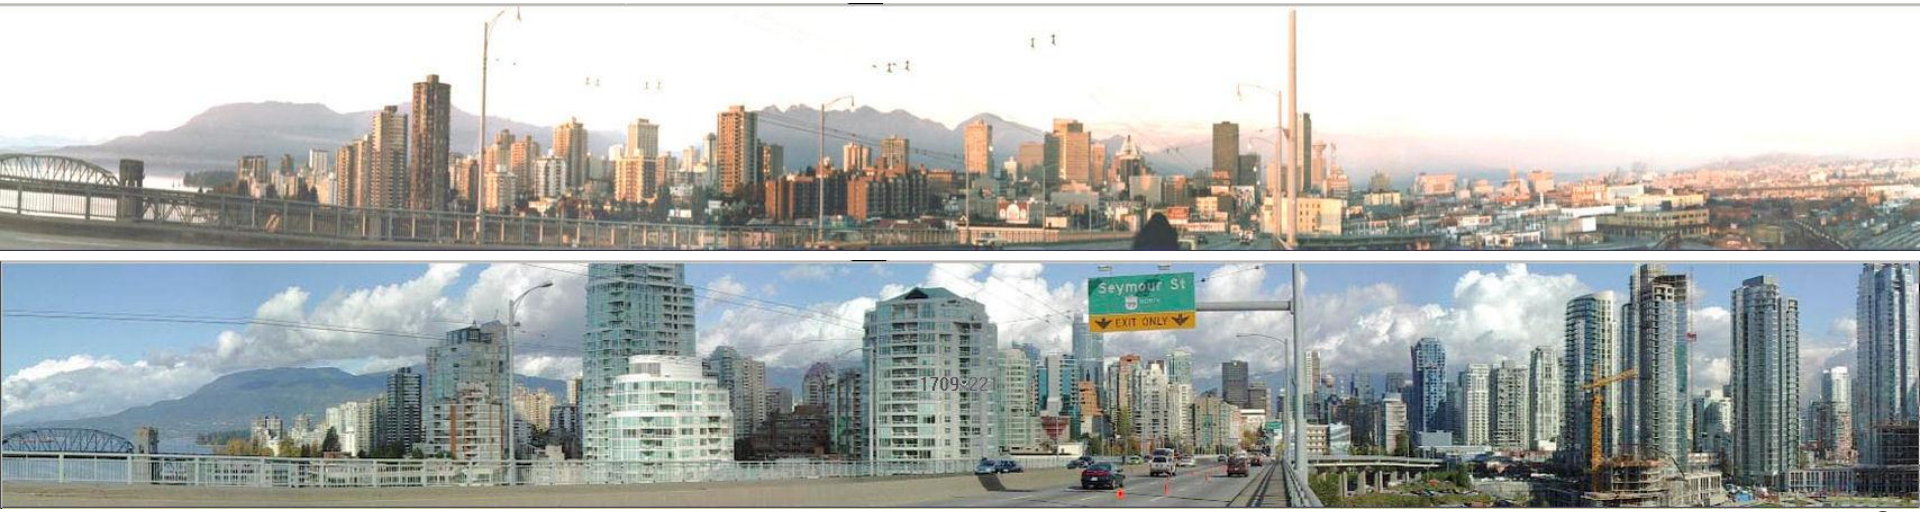

http://vancouver.ca/commsvcs/fade/fadepanorama.htm

PAST, PRESENT… I

THE SEVEN STAGES OF COMPUTATION(a)

1. Manual — up to around 1900
2. Mechanical — circa 1623 to 1945
3. Electromechanical —1902 to 1950
4. Relays — 1939 to 1944
5. Vacuum tubes — 1942 to 1961
6. Transistor — 1956 to 1979
7. Microprocessor — 1971 to present

(a) http://nordhaus.econ.yale.edu/nordhaus_computers_jeh_2007.pdf

PAST, PRESENT… II

Cost to perform a million calculations (constant dollars)

1900: ~ $500

1980: ~ $3x10^{-4}

2010: ~ $3x10^{-11}

2030: $3x10^{-15} (estimated)

→All our calculations done in the past can be done for essentially zero cost 20 years into the future!

→ In 25 years, home computing will outperform any supercomputer we have today!

Current & Future Trends

- Cloud Computing
- Virtualization
- CPU (multicore)
- Memory
- Storage ← 2007 snapshot CS at TRIUMF →WD 160Gb portable drive
- Network

The Birth of Internet – TCP/IP in 1974

Evolution of the TRIUMF Network

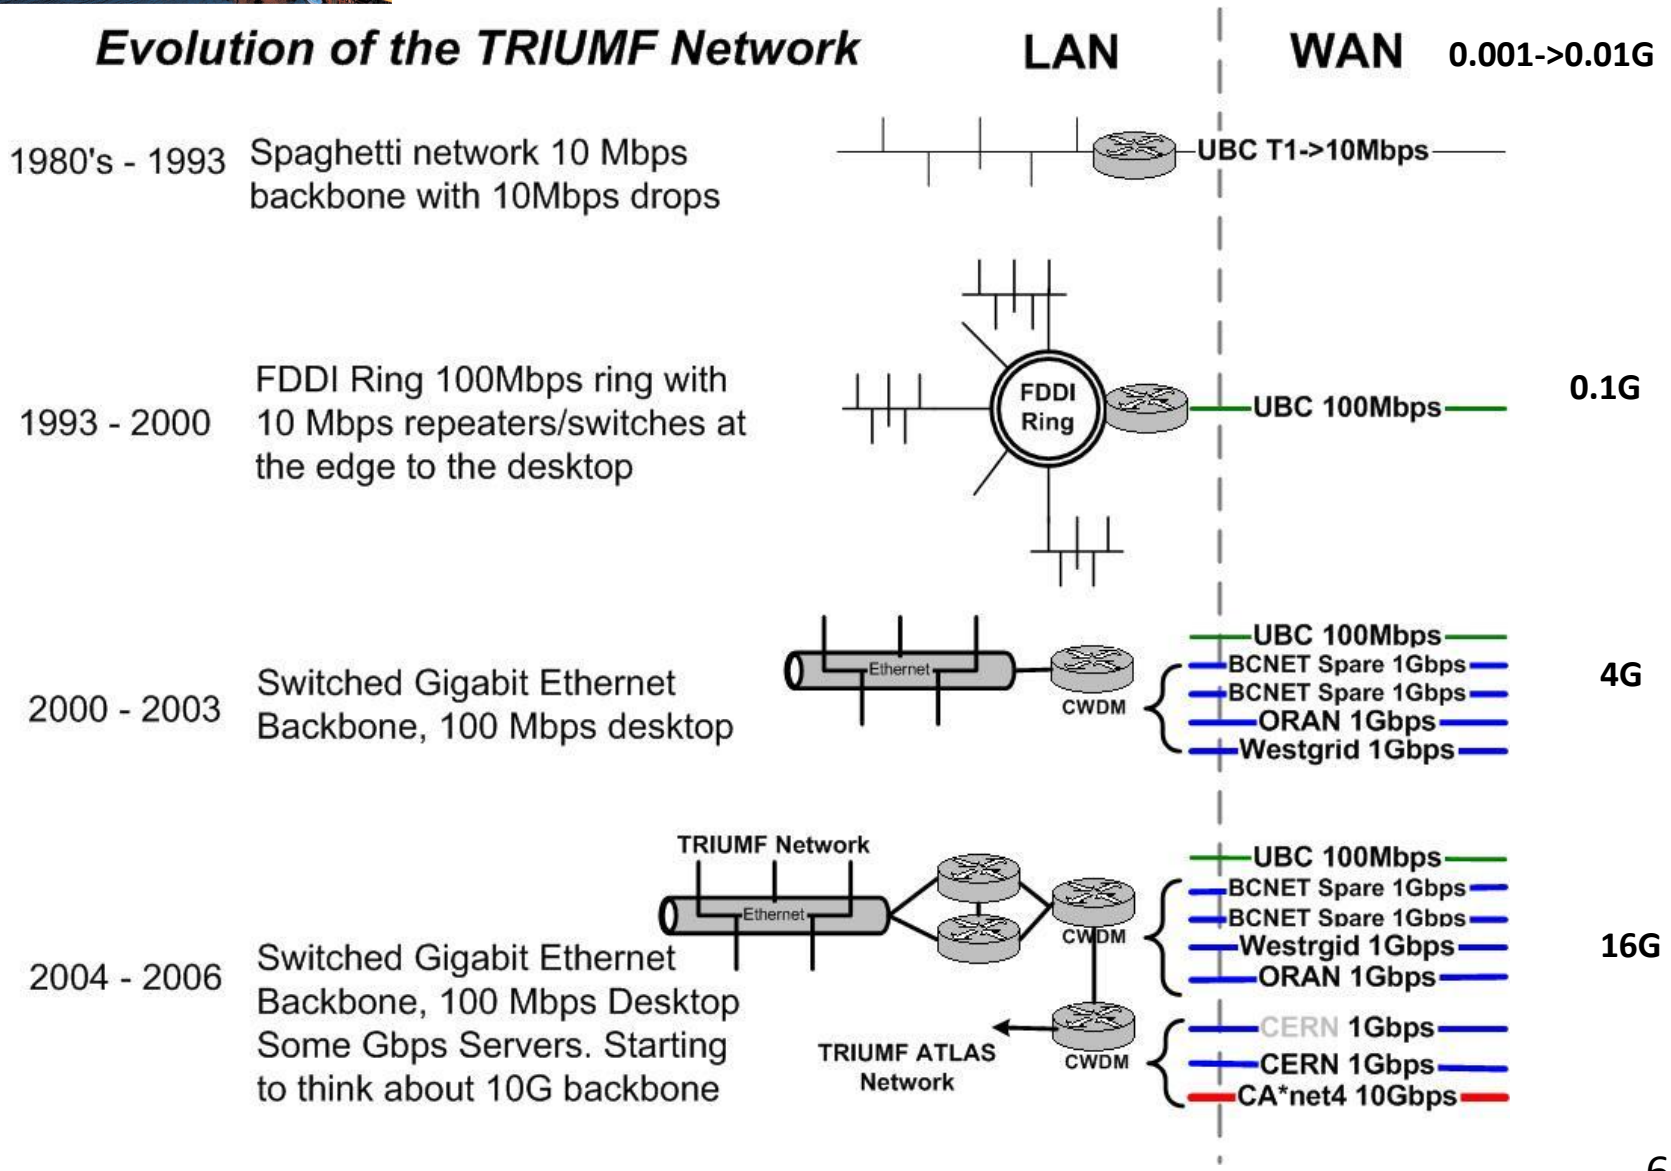

From TRIUMF Site Report by Denice Deatrich

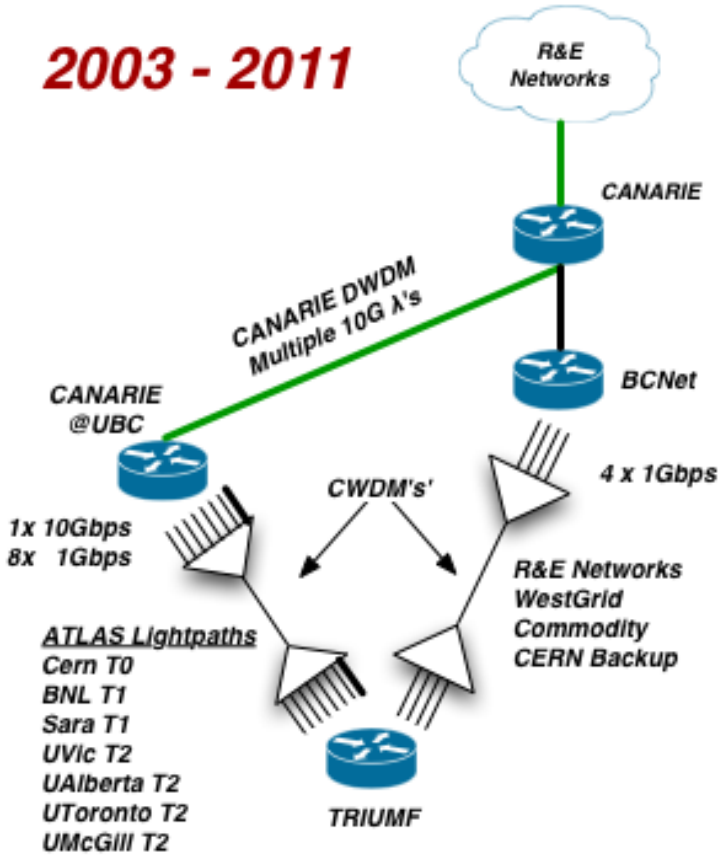

Only 2 pairs of Fiber. All external circuits are multiplexed on to multiple λ's. With exception of one λ per CWDM all other λ limited to 1Gbps

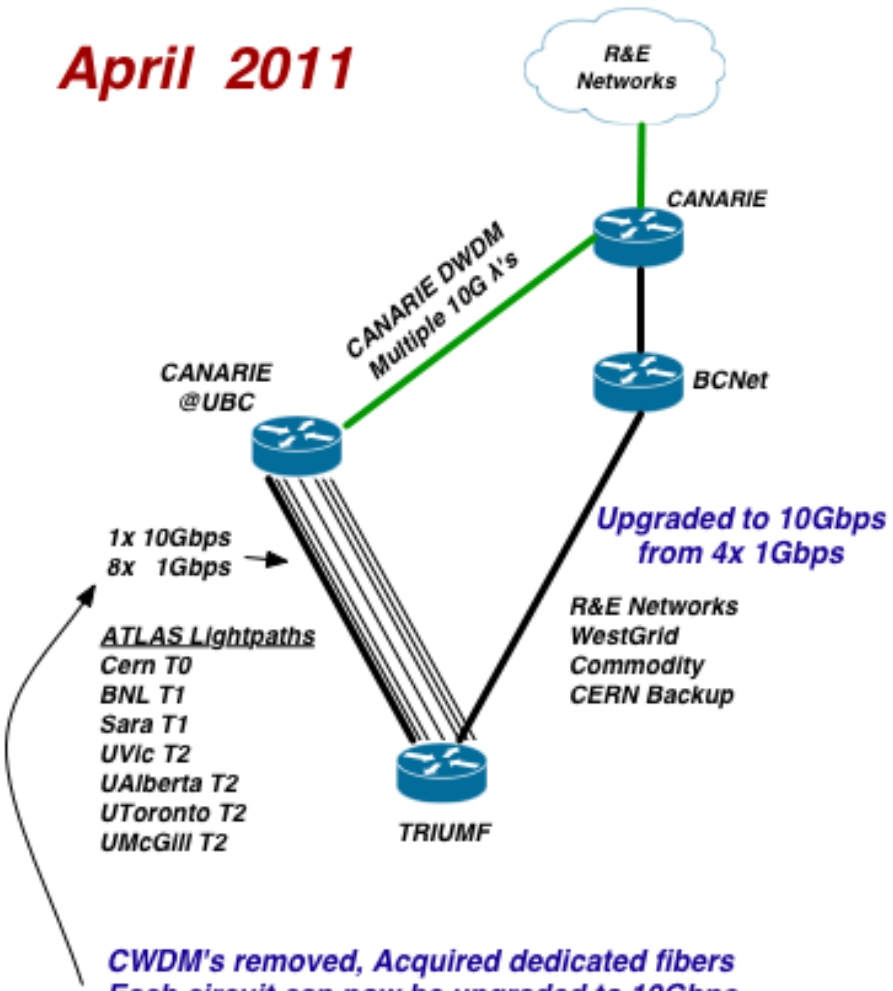

CWDM's removed, Acquired dedicated fibers Each circuit can now be upgraded to 10Gbps

Cloud Computing

- Clouds are like “time-sharing” on computers.
- Inter-cloud impediments (resolved ~5yrs) like pre-internet
- Commercial cloud storage costs < $0.50 / GB / Year
- Proprietary (protocols) clouds could stifle innovation – needs openness, multivendor, & interoperability
- IT departments can “sell” spare capacity to the cloud (eg. SpotCloud by enomaly - http://www.spotcloud.com/)
- “Clouds moving too fast”

livedrive | Cloud Storage for Everyone

Currently $16/month for 2TB
→ $0.10/GB/Year

Virtualization

VMWARE, KVM, HYPER-V, XEN etc…
For example, XEN, is a free software, originally developed by University of Cambridge Computer Laboratory, allows multiple computer operating systems to run concurrently on the same hardware. Saves hardware costs, improves flexibility and efficiency, eases upgrades…

Did you know….

- Disk drive cost to store all the world's music?

  $600

- How many mobile phones in world in 2010?

  5 billion

- Data centres (includes cooling) consume about 1% of global electricity. Home computers/routers etc about 4%(1) This is expected to at least double in 10 years.

(1) http://uclue.com/?xq=724

(2) http://cuppcomputing.com/global-pc-power-consumption/

Leading Edge Computing

FLOPS

YOTTA	10^{24}
ZETA	10^{21}
EXA	10^{18} ← US Pentagon by 2018;
PETA	10^{15}
TERA	10^{12}
GIGA	10^{9}
MEGA	10^{6}

Computing at Home

Folding@Home ~ 4 peta flops

Boinc ~ 5 peta flops

http://boinc.berkeley.edu/

Current fastest supercomputer

10 petaflops (70,000 8 core processors)

Air Force Research Laboratory (AFRL) created a powerful supercomputer (rank 33) by connecting together **1,760 Sony PS3s** which include 168 separate graphical processing units and 84 coordinating servers in a parallel array capable of performing 500 trillion floating-point operations per second. PS3 ($400) compares to alternate systems costing $10,000

I bought a PS3 – not to play games but for all the other things it could do!

Some Trends

- Internet speeds increase 1000-fold (Gig.u or Gig-u.org)

Eg. National LambdaRail

  - 19,000Km fibre-optic network, US University based and owned
  - Uses DWDM (dense-wavelegth-division-multiplexing)
  - Aggregate capacity 1.6Tbit/s
  - Currently 40Gbit/s →160Gbit/s
- Bandwidth of fibre optic cable increasing rapidly – 100 Pbits/sec-km (~ 2TB/s over a single 7000 km fibre). The theoretical limit is apparently 1250 channels (colours) of 10Gbit/s per colour.
- Heat is a pervasive problem in all computing centres
- Operating cost of computing centres 3-5 times capital costs over 20 yrs
- Software bloat - despite far faster machines – performs just as slow.

…1991 2011 2021

- Moore's Law – doubling ever 18 months → 10,000 in 20yrs
- My first contact with a computer (1960) was a Bendix G15-D computer installed at University of Manitoba in 1958. Input was by paper tape. It cost about $50,000. (about 2011$400,000)
- Next came the IBM 1620 – FORTRAN!
- By time I got my Phd in 1969 → IBM 360

http://www.columbia.edu/cu/computinghistory/

- Email was the thing that had biggest impact on people in early 90's
- Today the internet/social networking has the biggest impact.
- By 2031 people will say "Why couldn't they do these basic things in 2011?"
- Power/Heat Wall…limits CPU Frequency
  - Frequency ~ V Power ~ V^3 (V^2F)
  - Freq ⇳ factor 2 ←→ Power ⇳ factor 8
- Solution: Optically linked Tiled Lower Frequency Multicores with switch in every core: bandwidth ~ √ Nodes
- IPV4 → IPV6 (everything consuming electricity may eventually be on the grid)
- Storage migrates to online servers → Laptops/Notebooks/Pads – no moving parts

Commodity (&other) trends

	1991	2001	2011	2021 (est.)
Memory ($/GB)	$80,000	$300	$7	$0.07
USB Mem. ($/GB)(3)	-	$1000	$1 ←$64K ?	$0.10
SSD ($/TB)	-	-	$1500-$5000(1)	$40.
Disk ($/TB)	$5,000,000	$6000	$50(5)	$0.50 (no floods)
Tape (TB, MB/s, $/TB)	0.005 1 $8000	0.1 15 $500	1.5 140 $50 (6)	~~20 800 $3~~ ?(7)
CPU ($/GFLOPS)	$100,000	$500	$2.	$0.01
Cores	1	4	64	6400(2)
Transistors/CPU (4)	2 million	60 million	1000 million	20,000 million

(1) For peak 1GB/s and sustained 500MB/s R/W.

(2) Industry -Doubling every 18 months, Academia ~ multiply by 4
There will be need to have/support parallel programming software.
WDM (wave division multiplexing) optical bus will be required.
Another reasonable number: All computing components drop ~ 33%/year
Note that Windows 7 can only handle 256 cores. University of Glasgow has built a 1000 core processor on a chip.

(3) USB1: 1995 USB2: 2003 (480 Mbps) USB3: 2008 (4.8Gbps) transfer 2hr HD movie in ~ 3 min
Capacities: 1GB in 2004 128GB in 2011

(4) Typical or low-cost CPU's

(5) 3TB Seagate Barracuda XT hard drive, 6Gbps SATA, 7200RPM $160.

(6) LTO;
StorageTek: 5TB,240MB/s, $300 scales to 500PB (all native)

(7) Disk vs. Tape: Watts/TB issue ; Fujifilm/IBM Tape prototype →35TB ~ Yr2025

FUTURE of TAPE (Typical)

Eight-Generation Roadmap

Compressed Capacity	Generation 1 200 GB	Generation 2 400 GB	Generation 3 800 GB	Generation 4 1.6 TB	Generation 5 3 TB	Generation 6 8 TB	Generation 7 16 TB	Generation 8 32 TB
Native Capacity	100 GB	200 GB	400 GB	800 GB	1.5 TB	3.2 TB	6.4 TB	12.8 TB
Compressed Data Rate	up to 40 MB/s	up to 80 MB/s	up to 160 MB/s	up to 240 MB/s	up to 280 MB/s	up to 525 MB/s	up to 788 MB/s	up to 1180 MB/s
Native Data Rate	up to 20 MB/s	up to 40 MB/s	up to 80 MB/s	up to 120 MB/s	up to 140 MB/s	up to 210 MB/s	up to 315 MB/s	up to 472 MB/s

Note: Compressed capacities for generations 1-5 assume 2:1 compression. Compressed capacities for generations 6-8 assume 2.5:1 compression (achieved with larger compression history buffer).
Source: The LTO Program. The LTO Ultrium roadmap is subject to change without notice and represents goals and objectives only.
Linear Tape-Open, LTO, the LTO logo, Ultrium, and the Ultrium logo are registered trademarks of HP, IBM and Quantum in the US and other countries.

~ 2 Years / Doubling “Generation” Currently at Generation 5
WORM: Write Once Read Many
Note: 1000PB ~ all the mobile data generate in the US last year(IBM)
Ref: http://en.wikipedia.org/wiki/Tape_drive

Standards

THUNDERBOLT 10Gb/s bi-directional interface is technically PCI Express on a copper cable

http://www.everythingusb.com/comparison.html

Near (~10yrs) Future p1

- **Storage revolution takes place – MASSIVE growth projected**
- Build an input device more efficient than QWERTY keyboard
- Mainframe → “PC” → “PC”+Cloud/Internet (Mainframe)
- Android (open) devices dominate over Apple (closed) devices
- Cognitive (non-von Neumann) computing arrives – search assist
- Break-up of telecom monopolies → dramatically reduce costs
- Research & Education (R&E) networks continue as engines of innovation.

Near (~10yrs) Future p2

- Climate Change disrupts global communications due to massive floods/storms. R&E networks could play a critical role to assist (eg. flood modelling)

- *"Open lightpath exchanges, federated optical networks, enterprise based integrated WiFi broadband wireless networks, brokered commercial cloud services, user controlled or software defined networks, federated identity, collaboration tools, leveraging build out of community based broadband networks (Gig.u) , green IT, etc are other* ***ongoing examples of how R&E networks can continue to engage in disruptive innovation****." – Bill St. Arnaud re: http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/geg-report.pdf*

- 100+Gbps to Scientific Centres, 10 Gbps in labs, 1 Gbps in offices.
- Semiconductor growth slows – only doubling every 3 years by ~ 2013
- Network capacity grows by doubling almost every year.

Far (>20years) Future

- Strong AI → Turing Test + Obvious when you encounter it
- AI make better decisions than humans
- robotics, genetic engineering, and nanotechnology self-replication
- “Will that be Silicon or Carbon Based?”
- Pocket supercomputers and/or cheap cloud computing
- The essential elements of computers have remained unchanged in the last 50 years. Quantum computing MAY change all that. Still there can be “no information without physical representation”
- quantum computation will only work in conjunction with quantum error correction
- factorising large integers by a quantum computer is unlikely in the next few decades
- Moore’s “law” predicted to last for many years
- Future generations of computers (as in past 30 yrs) will run legacy code!

Risk in forecasting…

Computers (1,2)

"Where a calculator on the ENIAC is equipped with 18,000 vacuum tubes and weighs 30 tons, **computers in the future may have only 1,000 vacuum tubes and weigh only 1.5 tons**."
Popular Mechanics, **March 1949**.

"There is **no reason anyone would want a computer** in their home."
Ken Olson, president, chairman and founder of Digital Equipment Corp. (DEC), maker of big business mainframe computers, arguing against the PC in **1977**.

"I have traveled the length and breadth of this country and talked with the best people, and I can assure you that **data processing is a fad** that won't last out the year."
The editor in charge of business books for Prentice Hall, **1957**.

"But what... is it good for?"
IBM executive Robert Lloyd, speaking in 1968 **microprocessor**, the heart of today's computers.

"Spam will be a thing of the past in two years' time"
Bill Gates, **2004**, speaking at World Economic Forum in Davos

(1) http://www.2spare.com/item_50221.aspx

(2) http://www.makeuseof.com/tag/8-spectacularly-wrong-predictions-computers-internet/
http://news.bbc.co.uk/2/hi/business/3426367.stm

HEPiX Future

- History is prologue – great value in co-ordination/communication/sharing via HEPiX
- Relevancy of community history
- Value of sharing experience → promoting open competition/innovation
- Evolution vs. Revolution
- Software dev. (needs many people) vs Hardware dev. (needs fewer people)
- High-end users push leading edge of hardware/networks
- By attendance figures HEPiX is very healthy!
- HEPiX meetings are still needed every 6 months due to rapid change
- Some Pictures at https://picasaweb.google.com/105666128349182844611/HEPiXVancouver2011

THANK YOU!

MERCI!

http://qrcode.kaywa.com/

Shrek's Law(1)

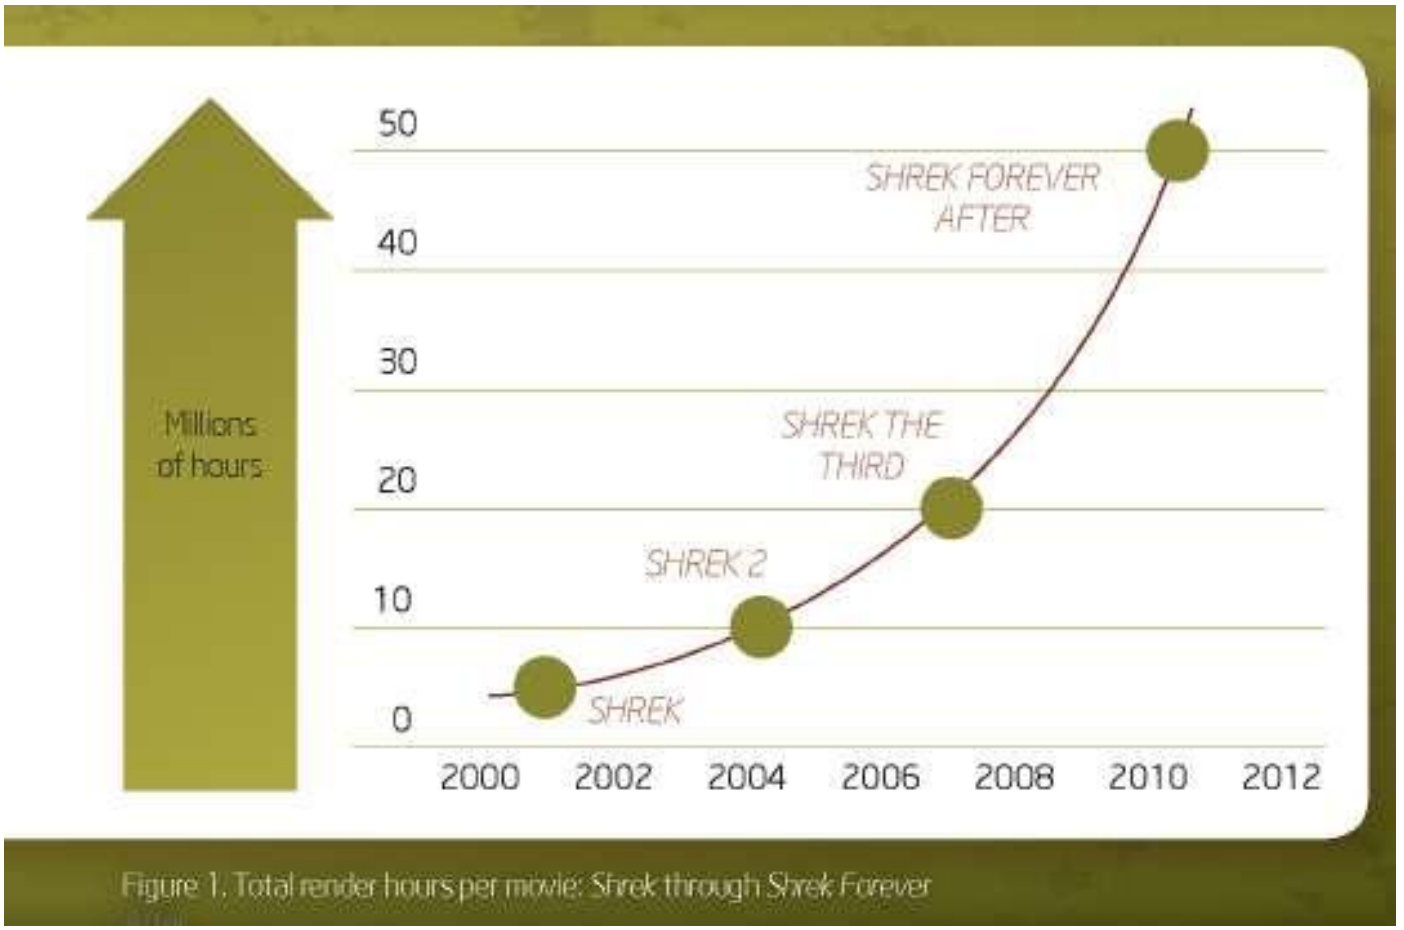

Figure 1. Total render hours per movie: Shrek through *Shrek Forever*

Although CPU-HRS increased computing costs remained stable over the 10 years

(1) http://software.intel.com/sites/billboard/article/moving-animation-high-performance-computing-cloud

Trend in Chip Size

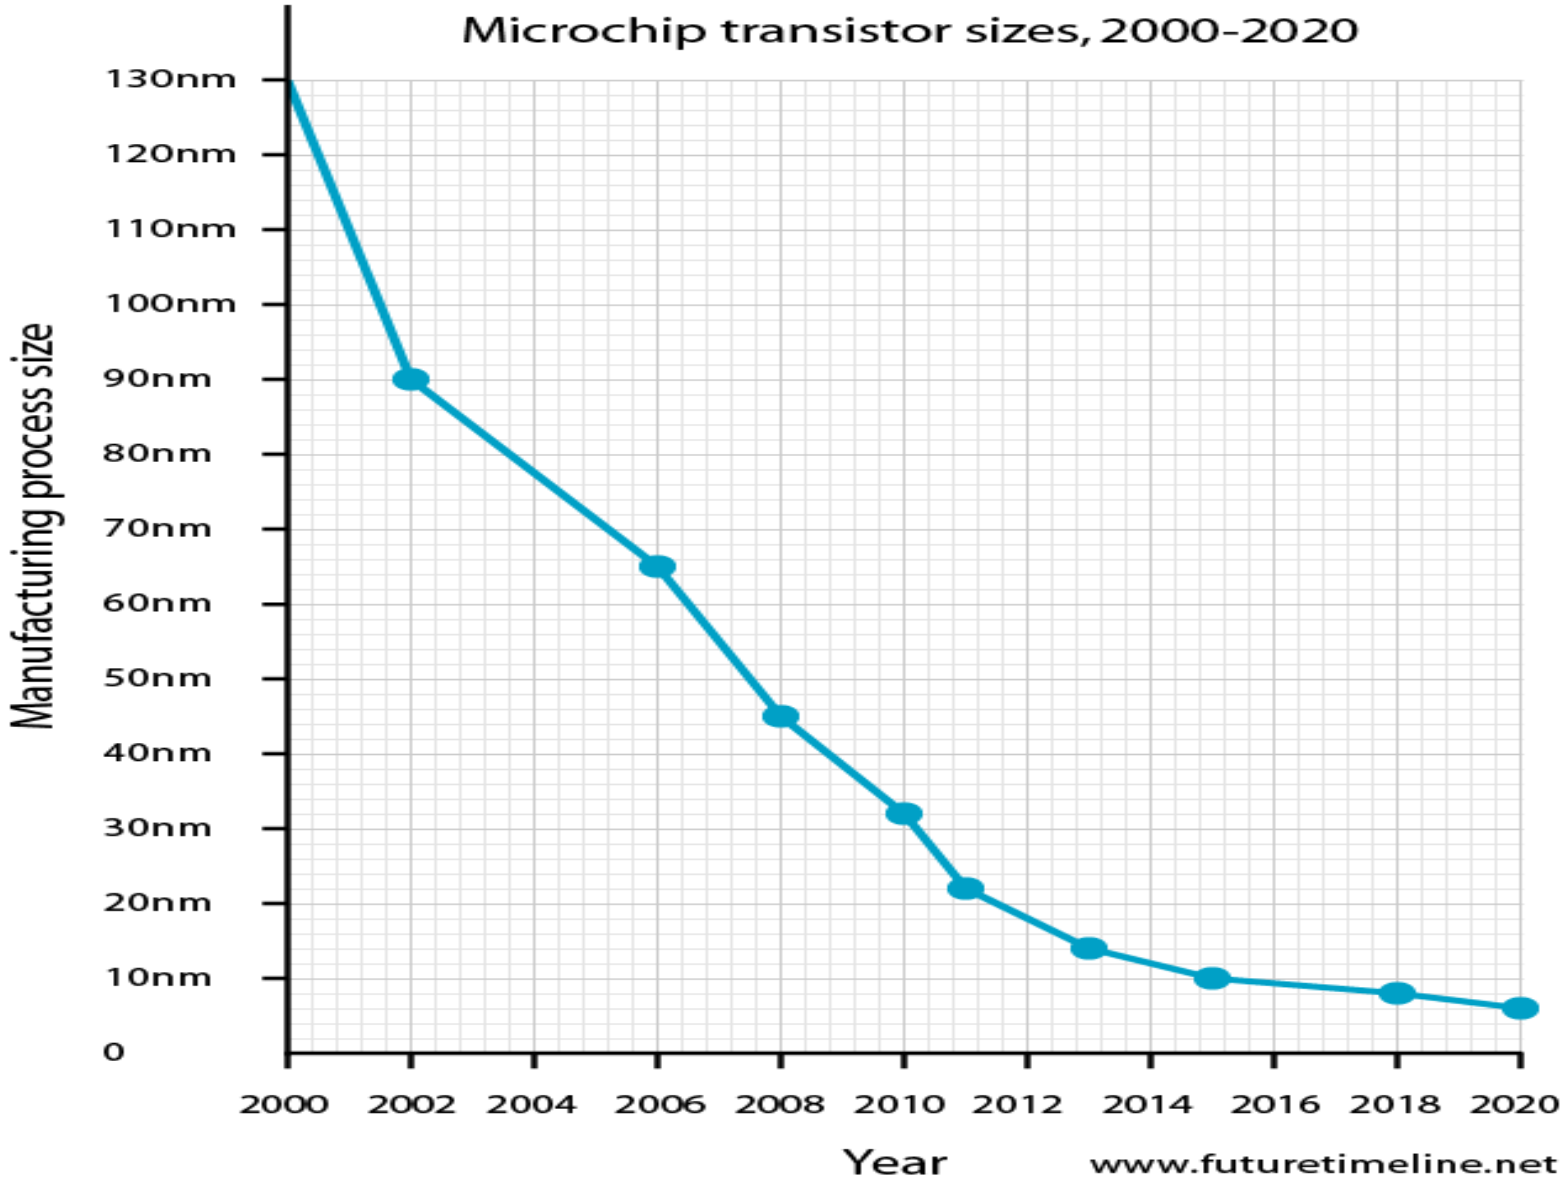

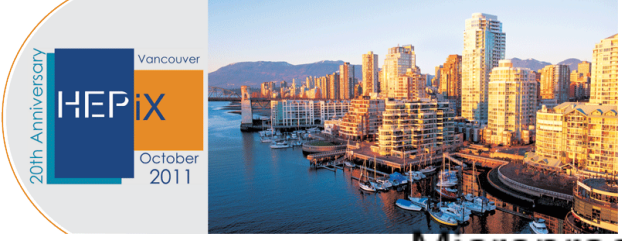

Microprocessor Transistor Counts 1971-2011 & Moore's Law

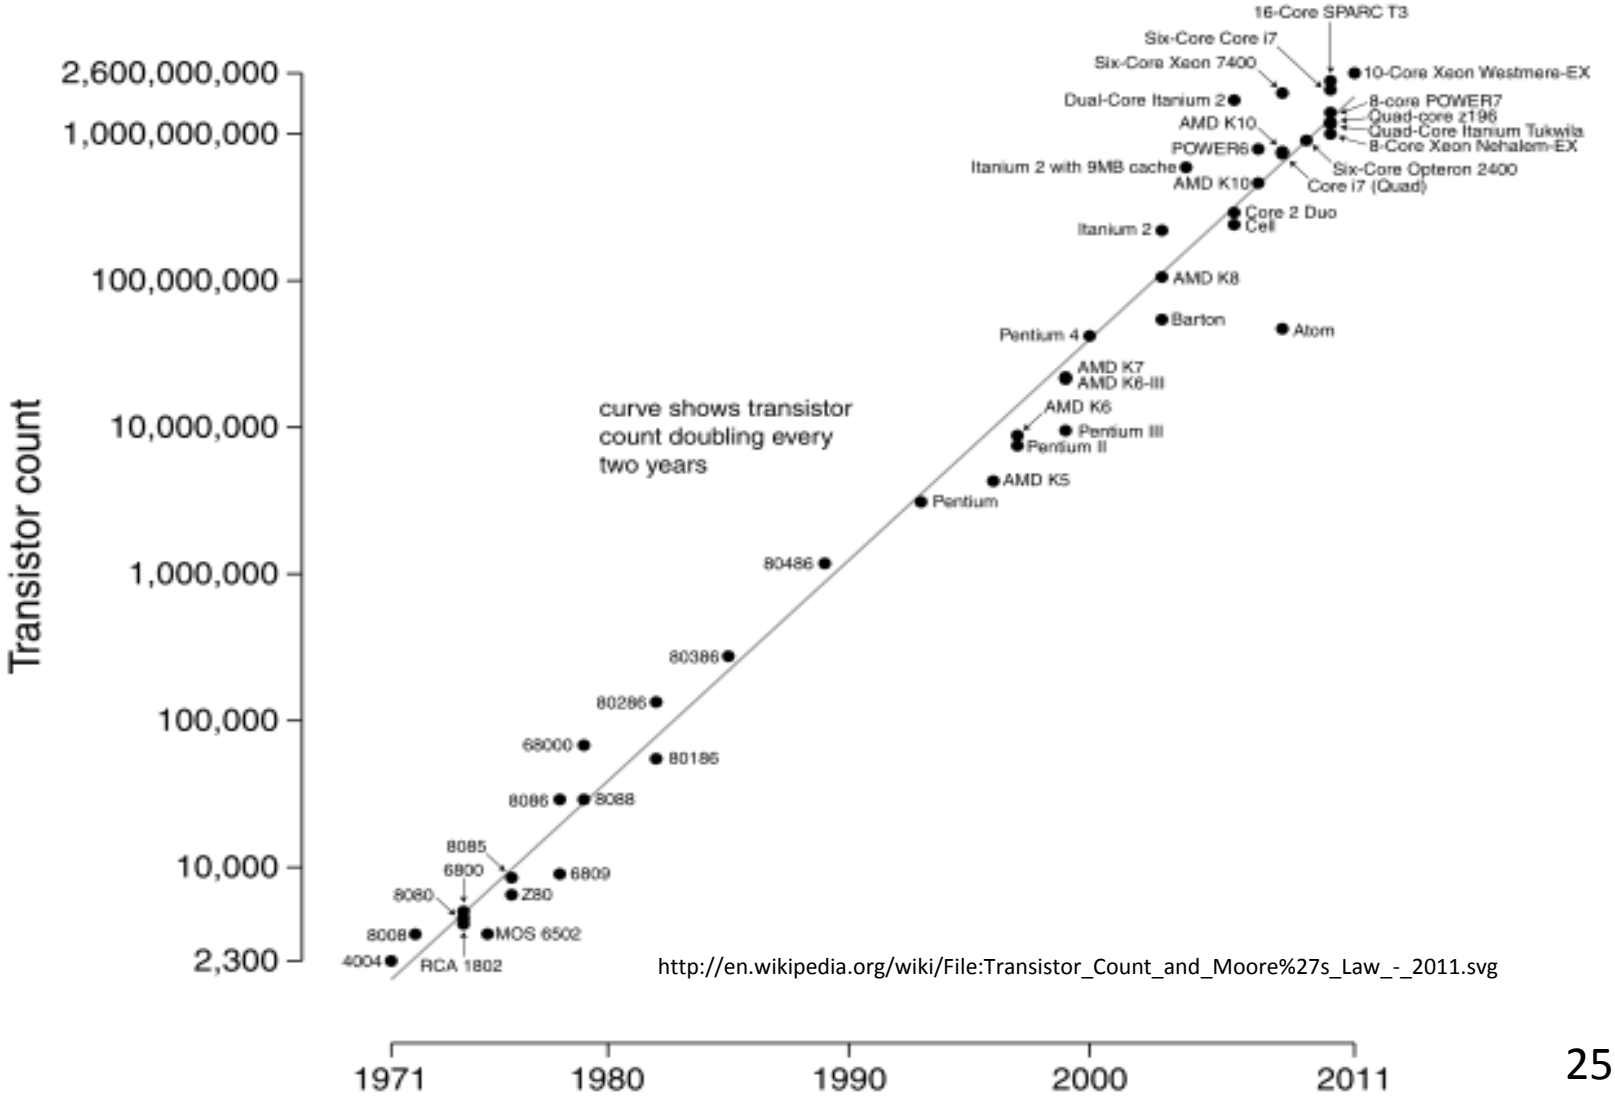

http://en.wikipedia.org/wiki/File:Transistor_Count_and_Moore%27s_Law_-_2011.svg

Caveat on Parallel Computing

http://en.wikipedia.org/wiki/Amdahl%27s_law

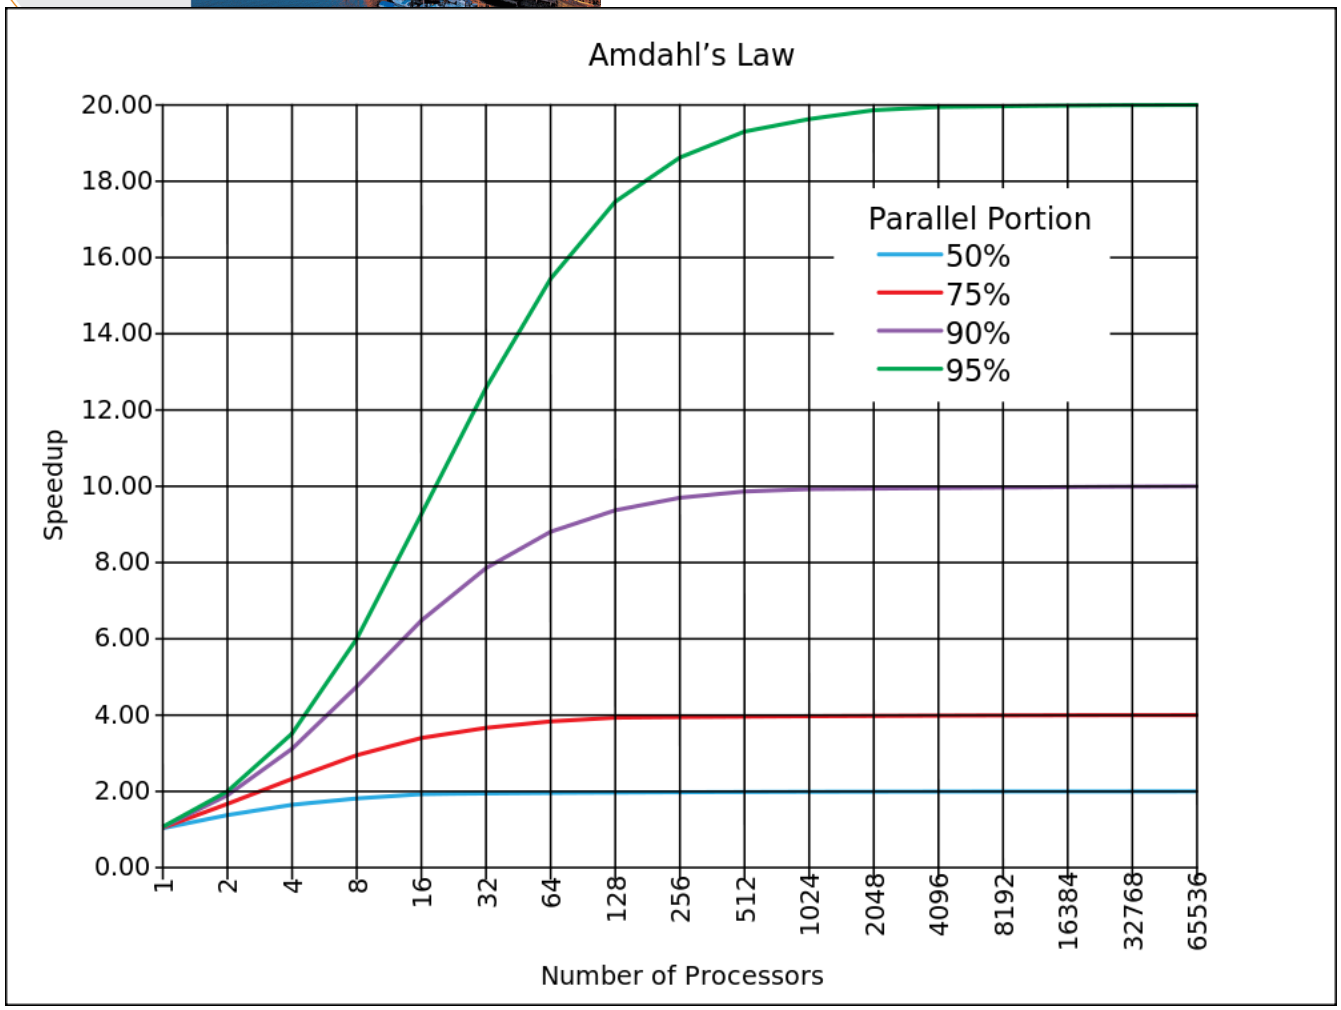

The Good News:
In many applications, particularly with very large data sets the amount of sequential code is close to 0%, (parallel portion is close to 100%) as essentially every data element can be processed independently.

The Bad News: The speedup of a program using multiple processors in parallel computing is limited by the sequential fraction of the program. For example, if 95% of the program can be parallelized, the theoretical maximum speedup using parallel computing would be 20× as shown in the diagram, no matter how many processors are used.

SIMD vs. MIMD

SIMD: Single Instruction, Multiple Data
computers with multiple processing elements that perform the same operation on multiple data simultaneously. Ubiquitous in video games. They cannot switch to MIMD mode.

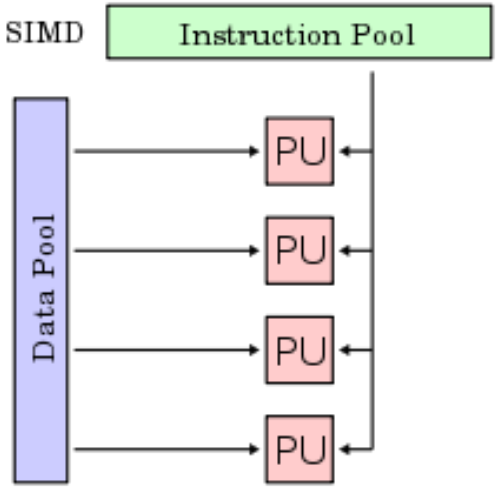

MIMD: Multiple Instructions, Multiple Data
Computers have multiple processors that function asynchronously and independently. They can switch to SIMD mode.

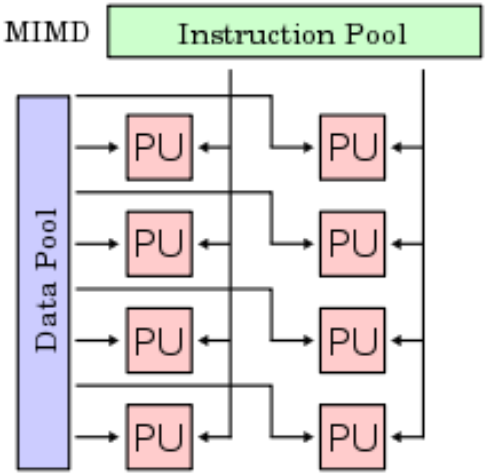

ATLAS CANADA TRIUMF-CERN LIGHTPATH DATA TRANSFER TRIAL FOR IGRID2002

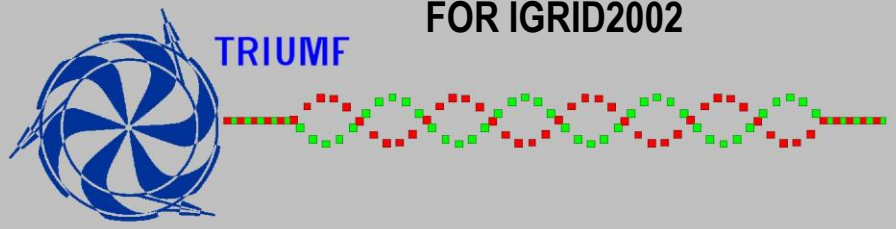

Two 1Gigabit optical fibre circuits (colours)

What was accomplished?

- Established relationship with "grid" of people for future networking projects
- Demonstrated a manually provisioned 12,000Km lightpath
- Transferred 1TB of ATLAS Monte-Carlo data to CERN (equiv. to 1500 CD's)
- Established record rates (1 CD in 8 seconds or 1 DVD in <60 seconds)
- Demonstrated innovative use of existing technology
- Largely used low-cost commodity software & hardware.

Participants
- TRIUMF
- University of Alberta
- Carleton
- CERN
- Canarie
- BCNET
- SURFnet

&

Acknowledgements
- Netera
- Atlas Canada
- WestGrid
- HEPnet Canada
- Indiana University
- Caltech
- Extreme Networks
- Intel Corporation

ATLAS CANADA TRIUMF-CERN LIGHTPATH TEST FOR IGRID2002

TRIUMF SERVER

THUNDER - 1.3TB FILE SERVER

INTEL

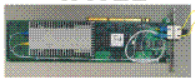

10 GigE

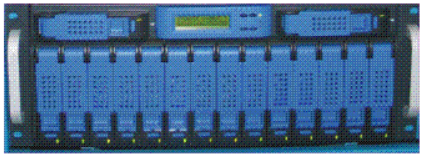

2.5 TBYTE Storage
16 * 160GB IDE
SCSI Attached

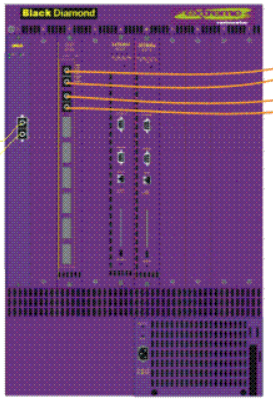

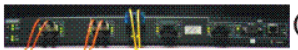

CWDM

TRIUMF to BCNET ~ 20Km

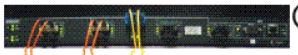

CWDM

BCNET

CISCO
ONS 15454

CA*net4

SURFNET

AMSTERDAM

CHICAGO

TRIUMF to CERN ~ 12000Km

CERN

CISCO
ONS 15454

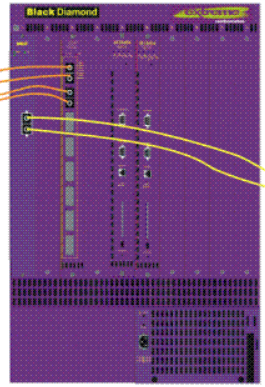

CERN SERVER

Intel

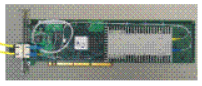

10 GigE

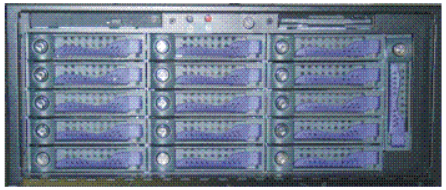

CERN - 1.3TB FILE SERVER

Steve McDonald, TRIUMF
Corrie Kost, TRIUMF
Wade Hong, Carleton U.
Bryon L. Caron U of Alberta

References

- http://legacyweb.triumf.ca/hepix96/
- https://www.hepix.org/
- http://wwwhepix.web.cern.ch/wwwhepix/meetings/cern92.html
- http://w3.hepix.org/afs/hepix.org/project/hepixmeetings/
- http://w3.hepix.org/afs/hepix.org/project/strack/
- http://w3.hepix.org/afs/hepix.org/project/ptrack/

HEPiX(/HEPNT) in Vancouver

- Apr 1996 http://legacyweb.triumf.ca/hepix96/hepix.html
- Oct 2003 http://legacyweb.triumf.ca/hepix2003/
- Oct 2011 http://www.triumf.info/hosted/HEPIX2011/index.html

Hepix/Pictures Sources

- http://legacyweb.triumf.ca/hepix2003/Pictures/
- http://legacyweb.triumf.ca/hepix2004/Oct/Pictures/
- http://www.rhic.bnl.gov/hepix/talks/041022am/bird.ppt
- https://www.racf.bnl.gov/Facility/hepix/pictures/ ← oct 2004

History of Agendas

- All https://www.hepix.org/mtg/meetings_html
- Fall 2004 BNL https://www.racf.bnl.gov/Facility/hepix/agenda.shtml

WINDOWS 8: http://www.techradar.com/news/software/operating-systems/hands-on-windows-8-review-1025259