

Disruptive Trends (Accelerators etc.)

Presented at
SOS-11

Steve Poole
Chief Scientist / Director of Special Programs
Computer Science and Mathematics Division

June 14, 2007
Key West, FL.

Short version

Accelerators / Co-Processors have been around for a long time

They will continue to be here in the future

Get over it

We will muddle along and make things work

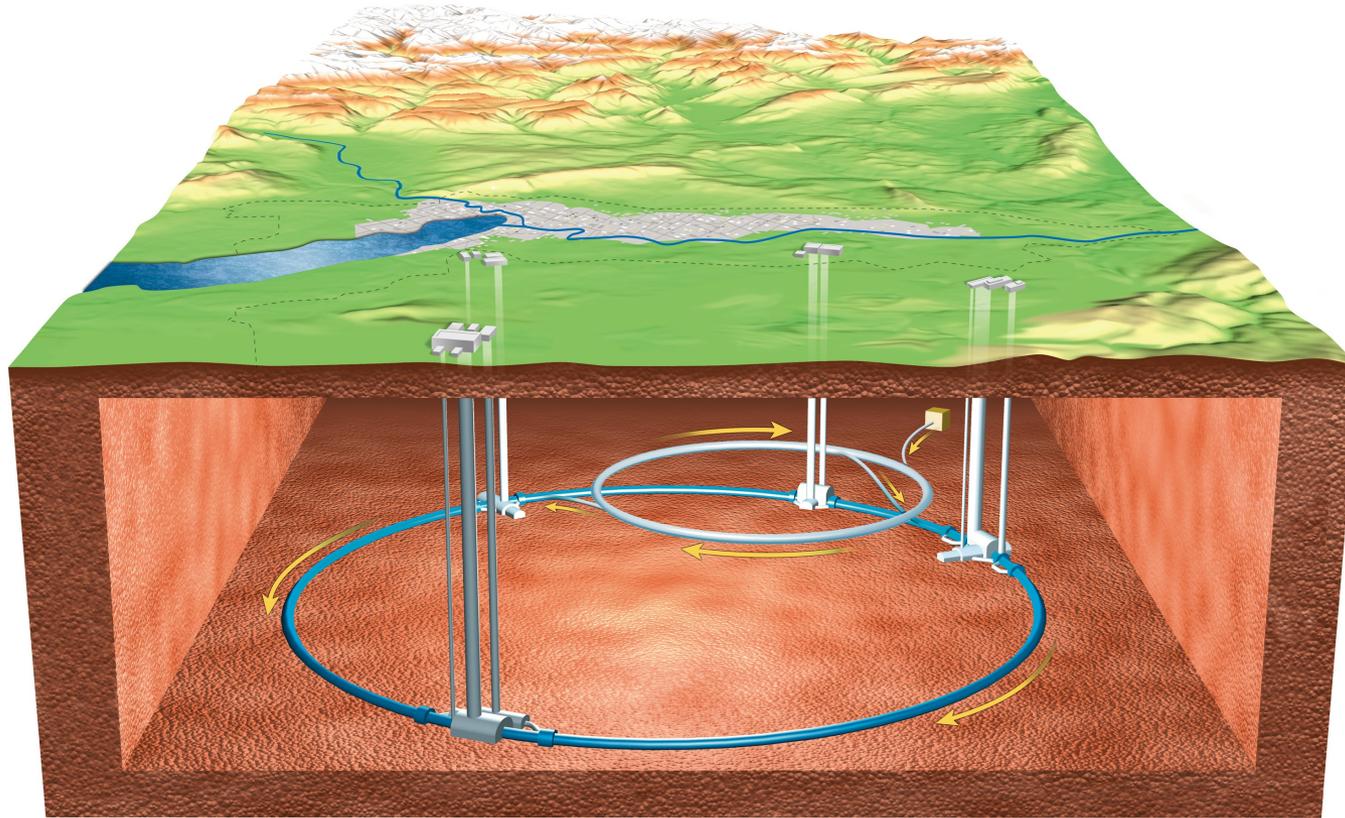
Thanks...

They know size

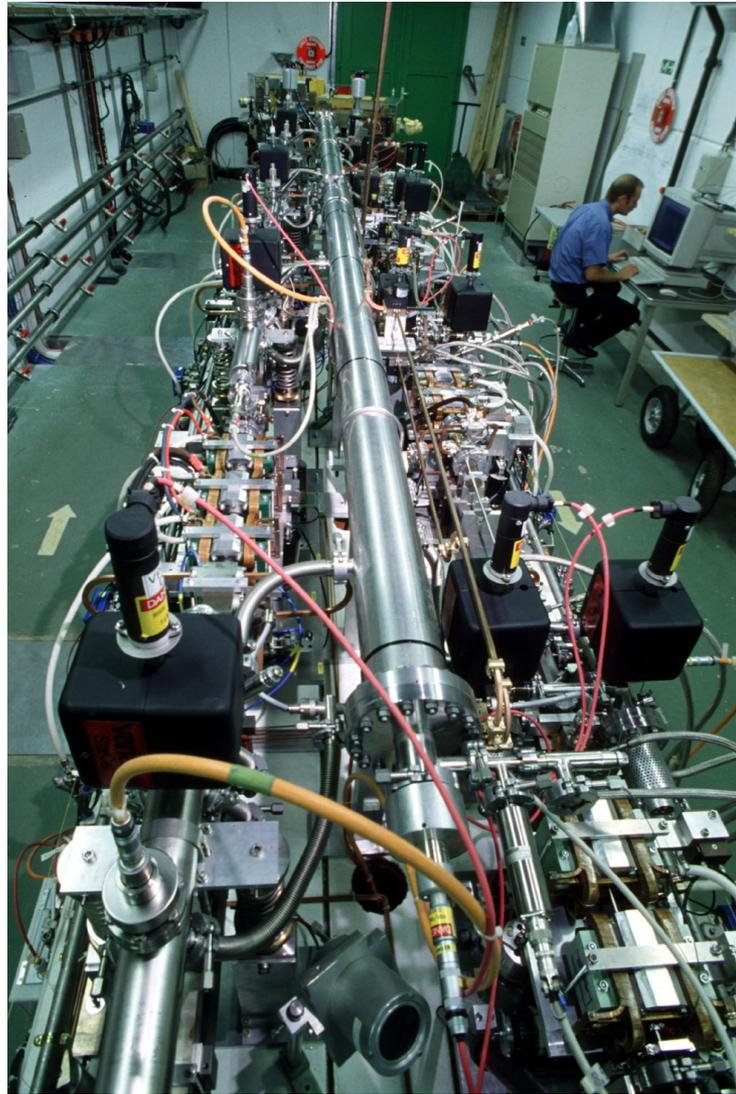


If you Google “Accelerators” this is the type of item you get

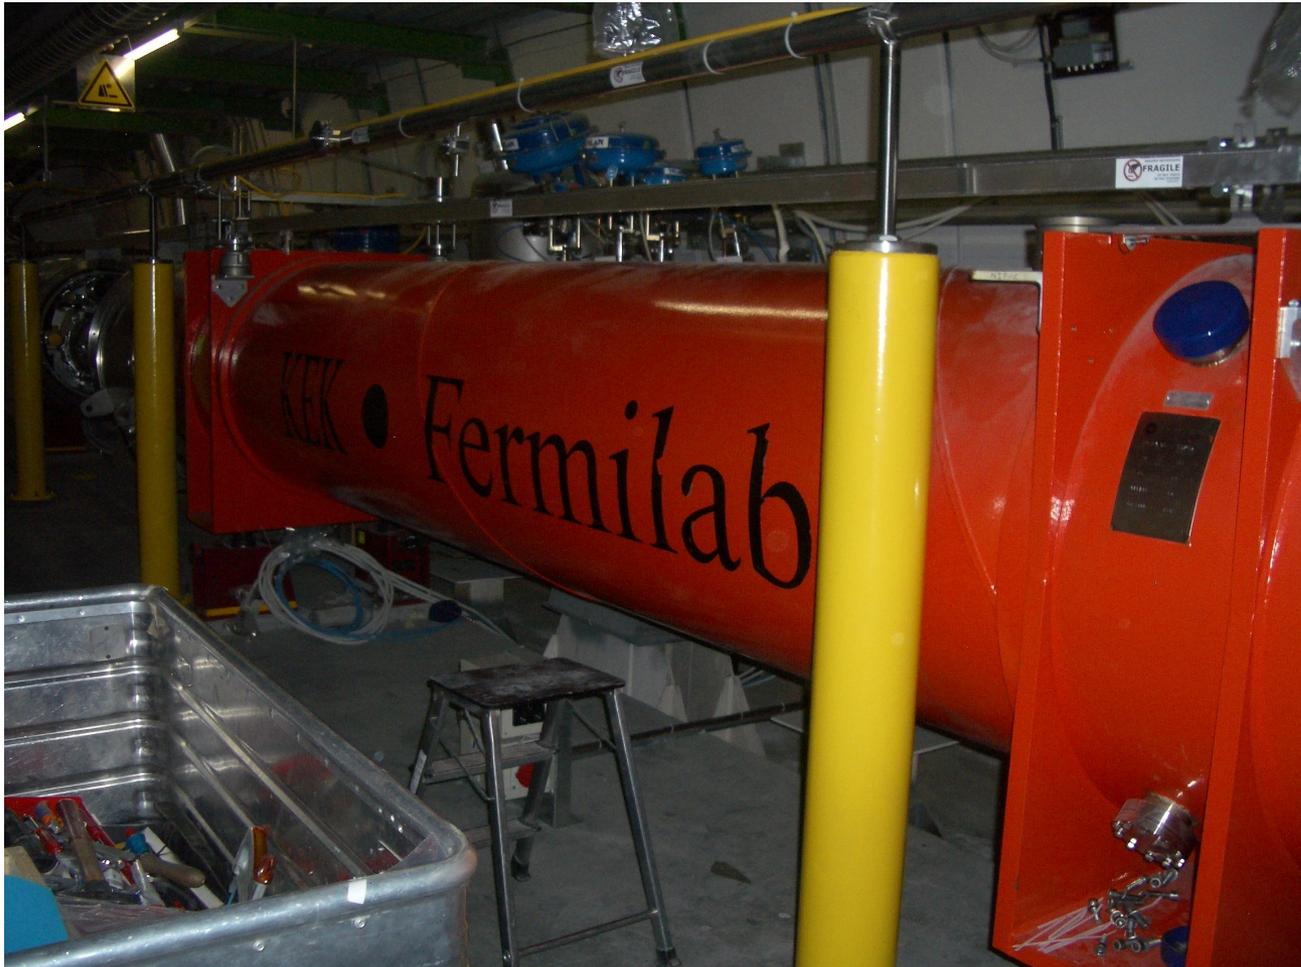
They understand modeling



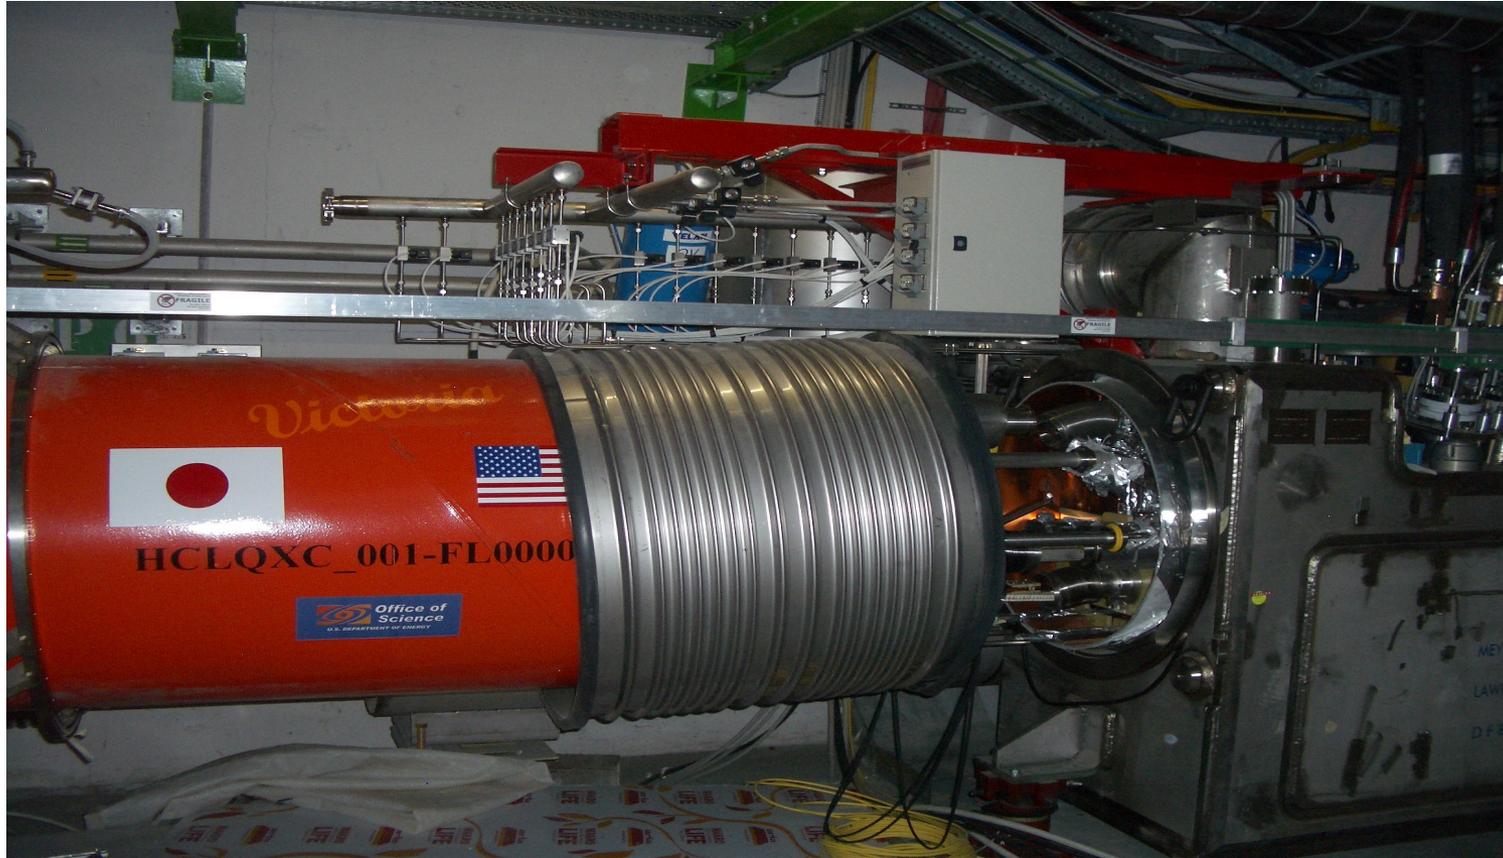
They understand complexity (cooperation)



They understand implementation



They understand failure/bumps



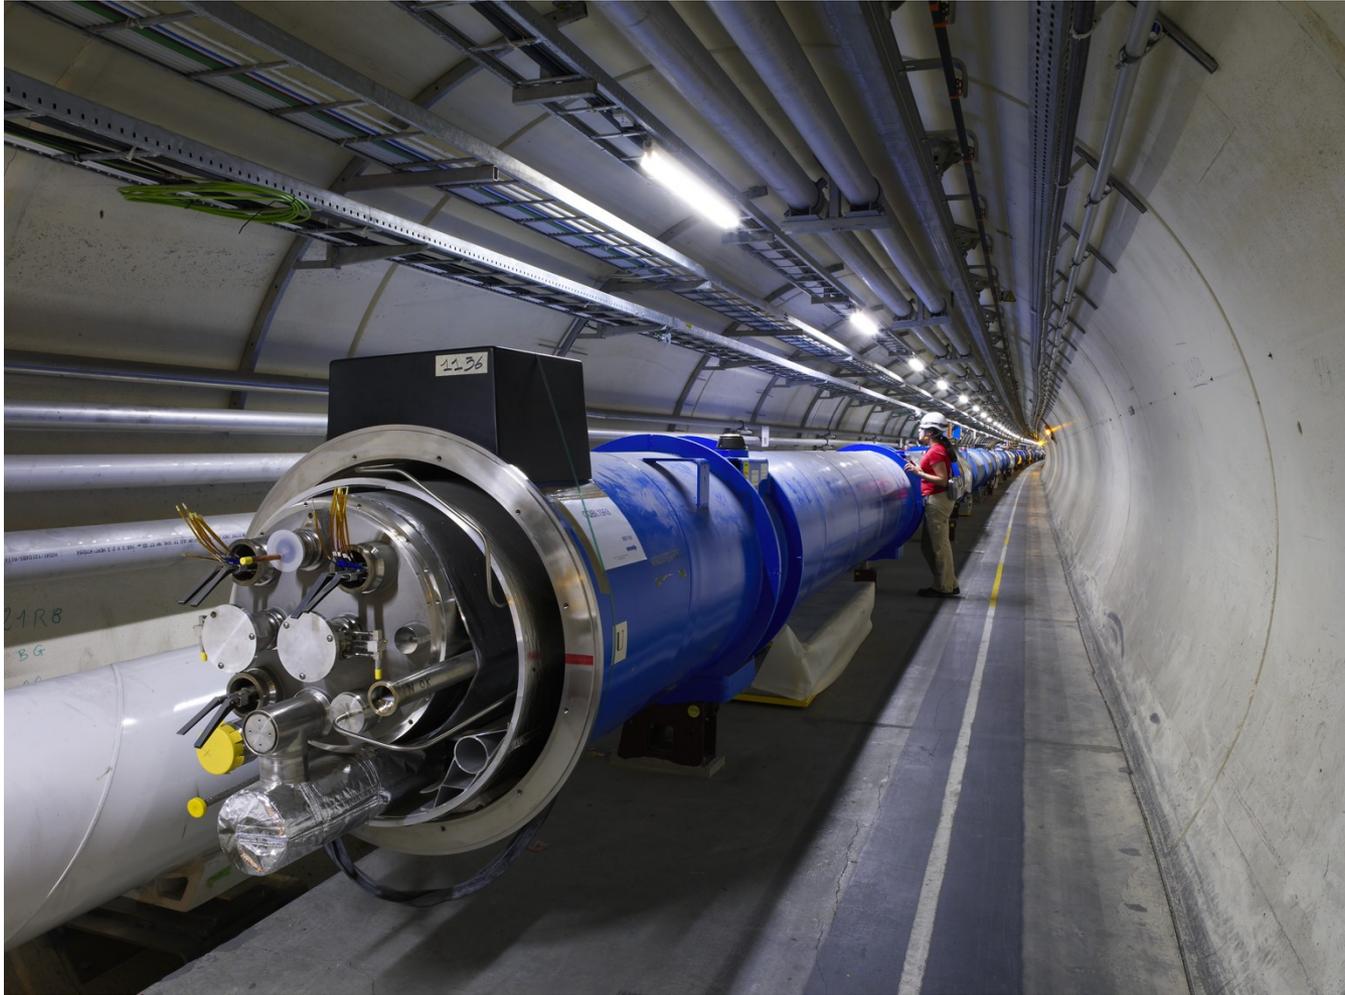
They understand long range



In multiple directions/dimensions



They understand RAS



Accelerators/M* Cores are coming !! (Oops, too late)



OR, Are they still here ?

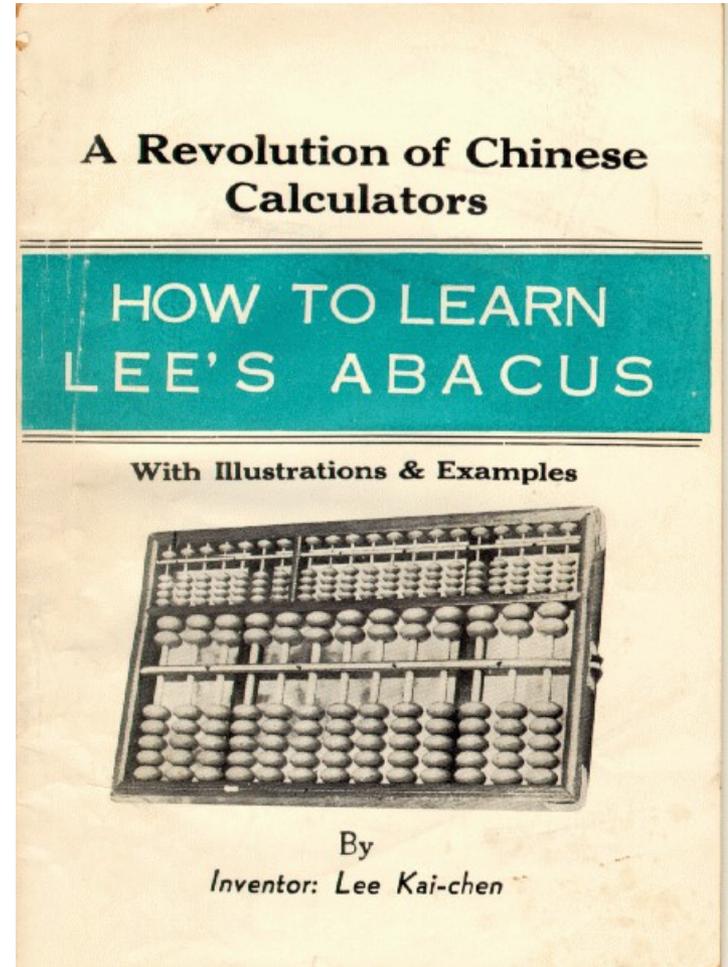
Early Accelerator



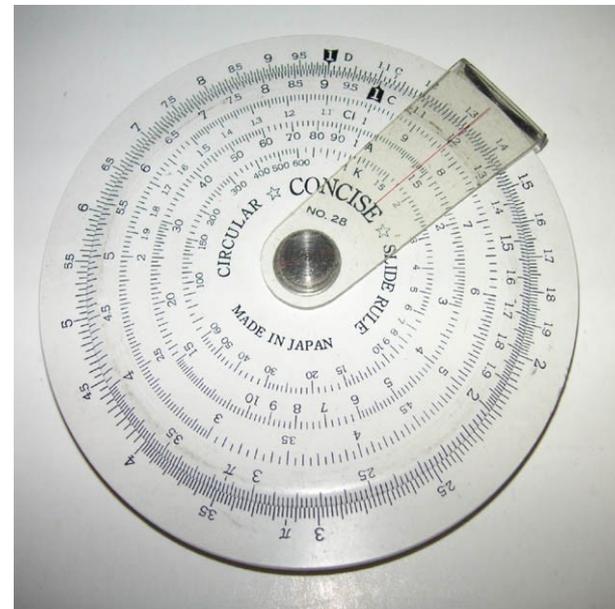
Industrial Design Accelerator & Programming Manual



Skip ahead a few years



Analog Accelerators ??



Floating-Point Systems AP-120B (6 MHz, 1975)

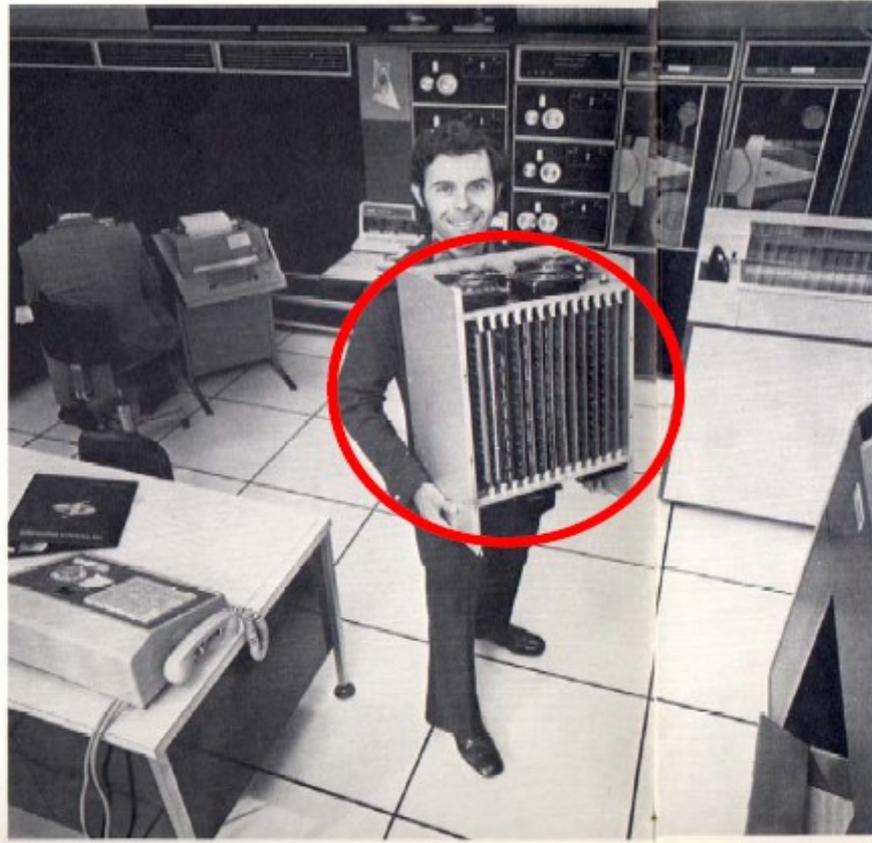
POWER YOU CAN HANDLE

Introducing the AP-120B, the unique peripheral floating-point array transform processor.

You'll look at the specifications and see the computational speed and accuracy of a mega-dollar mainframe.

Then you'll look at the cost, size, weight and ease of programmability, and realize what a breakthrough the AP-120B really is.

The unique culmination of years of floating-point and array processing expertise, the AP-120B offers a completely new source of convenient *dedicated* computational power.



Source: Dan Reed, David Culler

More Accelerator Areas

Intel 80387 Floating-point Coprocessor
(16+ MHz, 1987)



GPU Outside



Crypto Accelerator

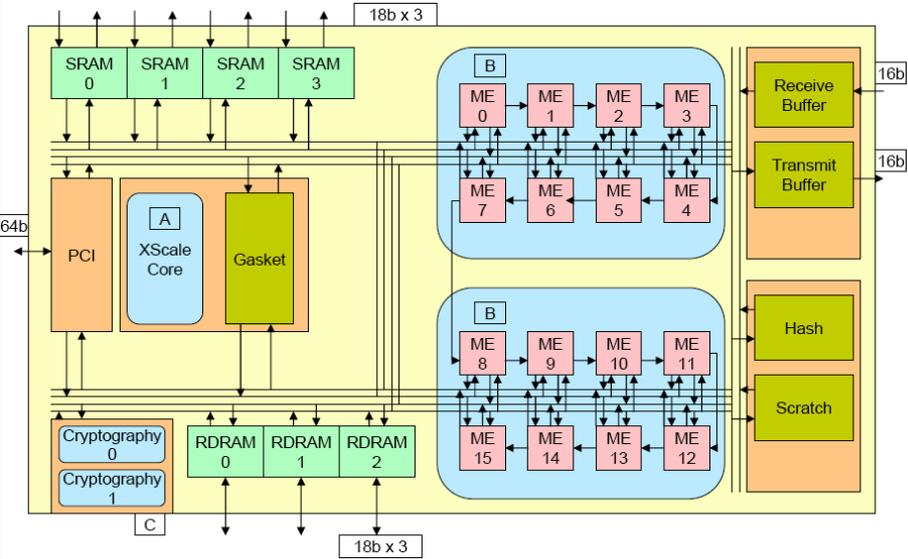


Cavium Networks, OCTEON-CPB, MIPS



Terari XML Accelerator

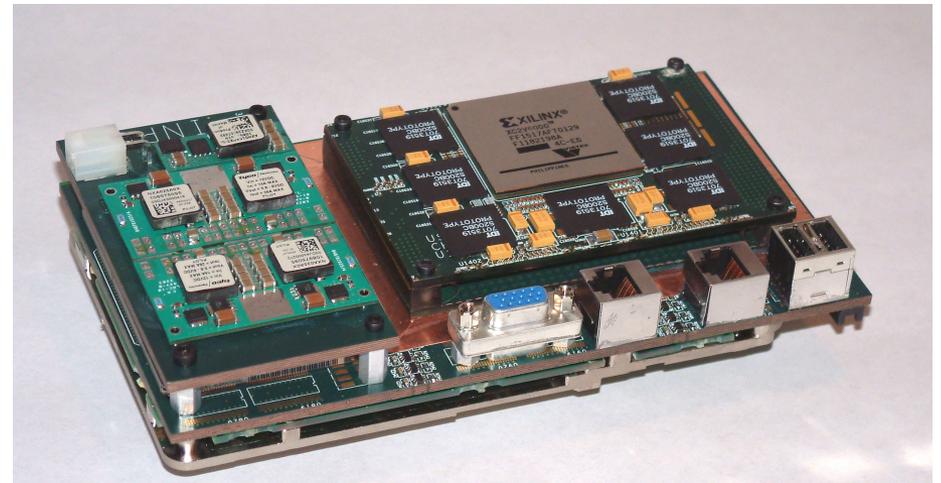
NPU / IXP-2850



Courtesy Intel / University of Illinois

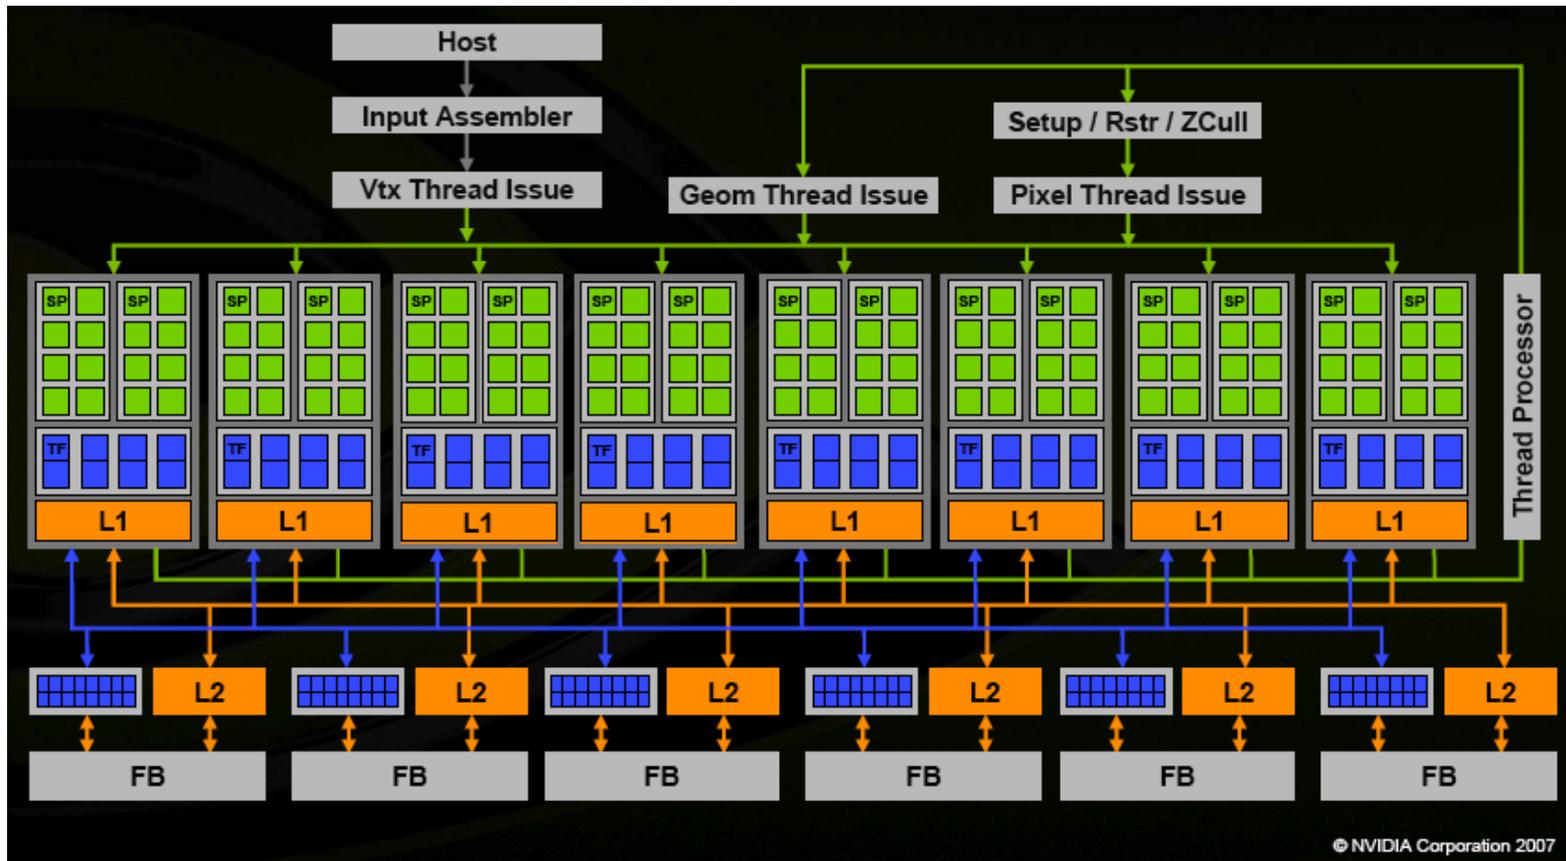


ClearSpeed



FPGA / DSP / Torenza

nVidia 8800



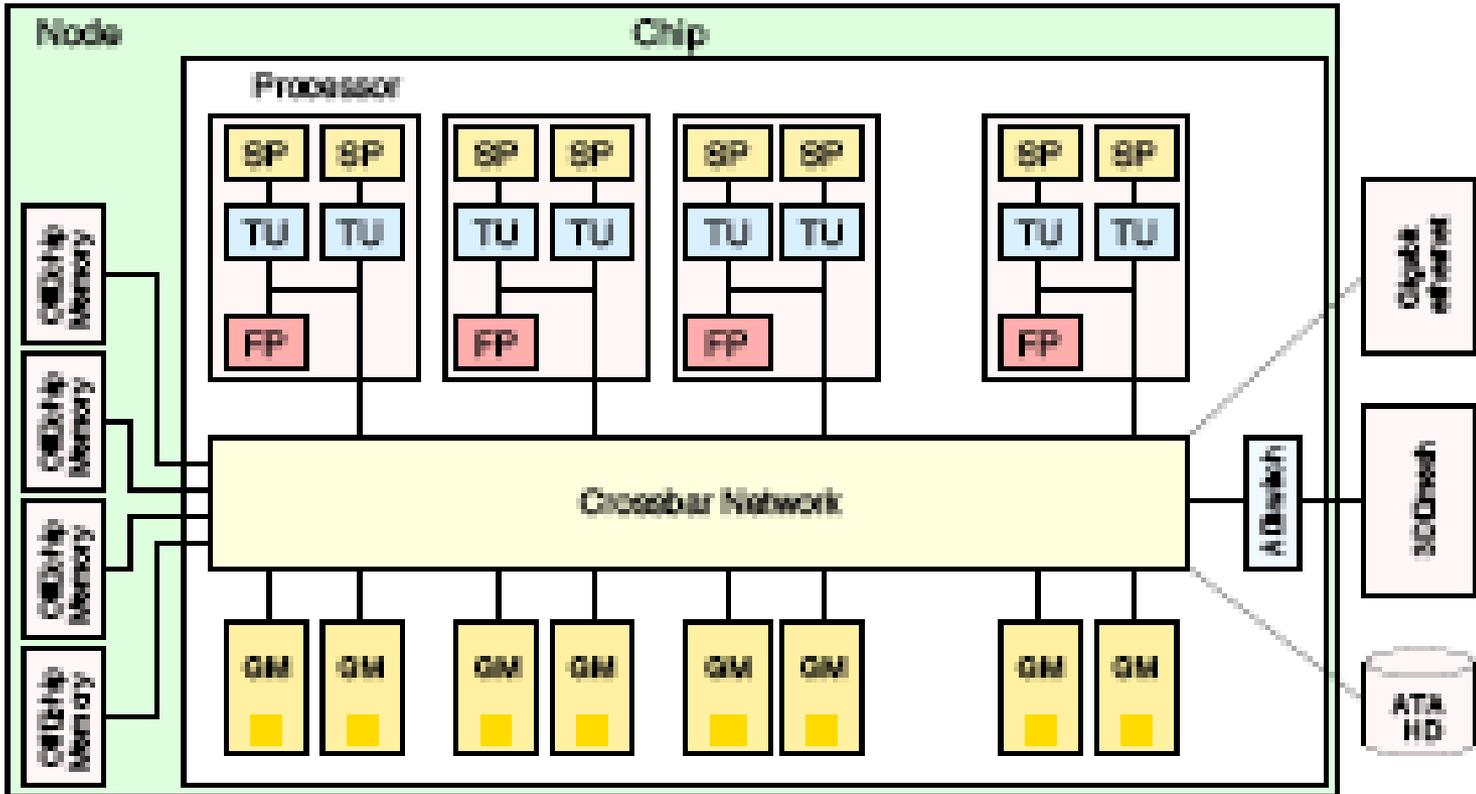
128 P, 103 GB/s Memory BW

ATI 2900, 320P, 512 bit memory, 128b FP

Cell

Intel (L*)

Cyclops Processor Design



IBM to Build World's First Cell Broadband Engine™ Based Supercomputer

Revolutionary Hybrid Supercomputer at Los Alamos National Laboratory Will Harness Cell BE Chips and AMD Opteron™ Processor Technology

x86 Linux
Master Cluster
AMD Opteron™
processor
(x3755 4U, 4 Socket)



Cell BE
Accelerator Cluster



HTx Infiniband Cluster Interconnect



Los Alamos
NATIONAL LABORATORY



Cluster Design Points

- ⇒ 1.7 PetaFLOPs peak (Double Precision)
- ⇒ 1.0 PetaFLOPs sustained
- ⇒ >16,000+ AMD Opteron™ processor cores
- ⇒ >16,000 Cell BE procs. (8,000 Blades)
- ⇒ ~360 racks
- ⇒ ~12,000 sq. ft (~3 basketball courts)
- ⇒ ~6 MegaWatts



Enhanced Cell BE engine

- ⇒ "Supercomputer & Network on a Chip"
- ⇒ 1 PPE + 8 SPE Cores
- ⇒ 102.4 GFLOPs/chip
- ⇒ 25.6 GB/s off-chip memory bandwidth/chip
- ⇒ Element Interface Bus (EIB) @ 300+GB/s
- ⇒ 1 Cell BE Blade = 2 Cell BE Procs = 16 SPEs

Cell Technology Highlights

New Cell BE Technology

- 65 nm BE chip
- Double Precision Floating Point, 100 GFs / chip peak
- DDR memory interface (replaces XDR)
- Existing I/O interface (uses Axon southbridge)
- Cell - 2 ??

New Cell BE Blade – Cell Blade 2

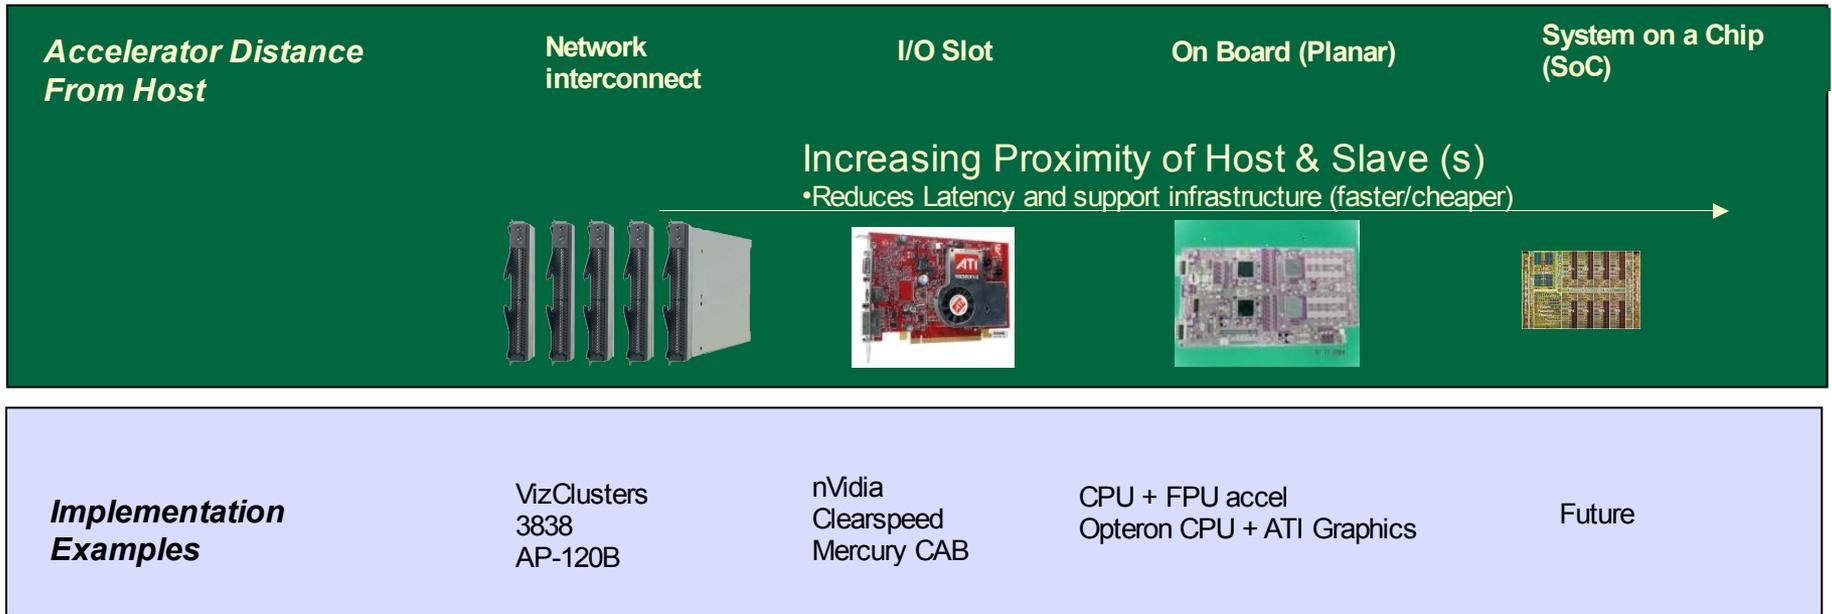
- 2 Enhanced Cell BE, 2 Axon, 4GB memory per Cell chip

Host/Accelerator programming model

- 1 accelerator chip per host core
- 1:1 ratio of host memory vs accelerator chip memory
- Programming model developed with LANL

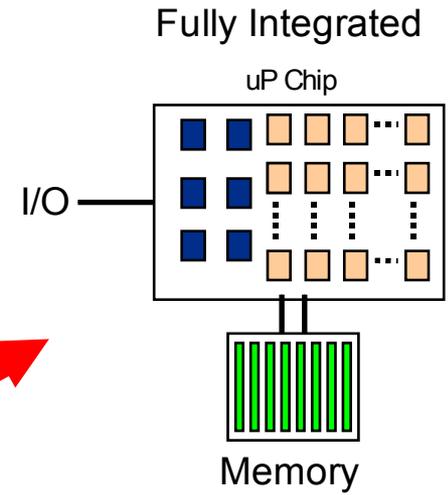
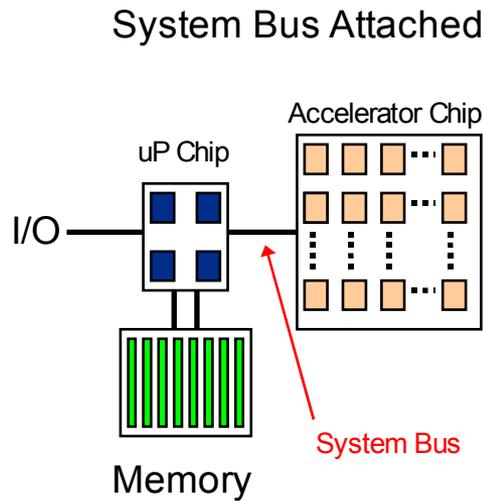
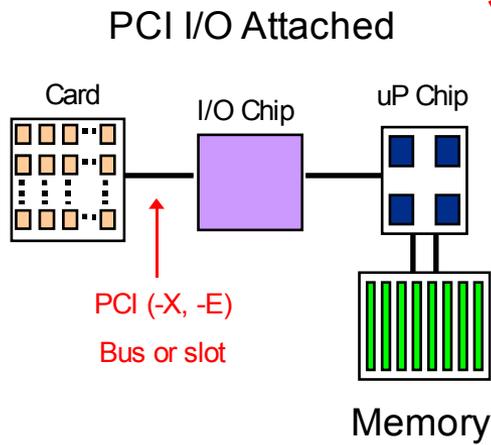
Accelerators move closer to hosts over time

Accelerator programming model choice dictated by the distance and corresponding latency and bandwidth performance between host and slave



This is a continuous life cycle for accelerators

Accelerator Evolution



Why the Interest in Hybrid Computing

Observation that traditional clusters are straining to reach PF scale:

- Processor core performance slowing

- Practical limits on network size and cost

- Programming challenges with 100Ks of nodes

- Technology discontinuity driving price/performance

 - Accelerators offer promise of 10x reduction in \$/MF

Did we mention ExaFLOPs ?

- Future Systems all look like they have accelerators

 - IBM , AMD , Intel , Cray, SUN...

Why Accelerators?

Specialized engines can perform selected tasks more efficiently (faster, cheaper, cooler, etc.) than general-purpose cores.

Development of hardware (e.g. PCI-express) and software (e.g. Linux) standards provide convenient attachment points.

Parallelism (e.g. SIMD) exploits increasing number of gates available in each chip generation.

Accelerators have been around for a long time!

FPS AP120B, IBM 3838, CSPI, TI-AATP, Intel 80387/80487, Weitek, Atari ANTIC, S3 911, ILLIAC-IV, TI-ASC, BSP, INPAC (IBM), SGI (TPU, EPU) CM*, PIM , QCD , HEP (3D) (and the list goes on)
(Google “Dead Computer Society”)

Some RR History

After DH (DoE no fund) , CHEAP, Not YACC , DRC , ASC-PI (Family of proposals)

Just a small matter of SW and \$\$

Libraries

Potential new LANGUAGES.

A robust / good compiler(s) would be nice (Difficult)

They have helped so far, but !!

HPCS Languages *might* help (Jury is still out)

Challenges are a good thing

Look to HPEC

They live for / with accelerators

Libraries

Data Flow ?

Heterogeneous

Sensor Model ?

Conclusions

Things will continue to be exciting

Only a few miracles need to happen

DNA computing and Quantum computing will take a while, but that is OK. They will be accelerators. -)

Watch for large Cell phone clusters.

It is the Software!!!

Backup