From Performance to Power: Leadership Computing Experiences with Argonne National Labs

Lennart Johnsson
University of Houston

UNIVERSITY of HOUSTON

Outline

- Collaborative experiences from the early days of the ACRF – architecture and algorithms
- Collaborative experiences on software
- Recent and current research on energy efficient computer systems design and operation

The Denelcor HEP

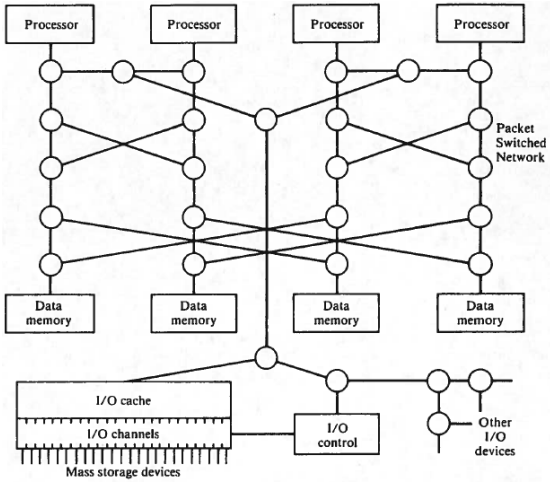

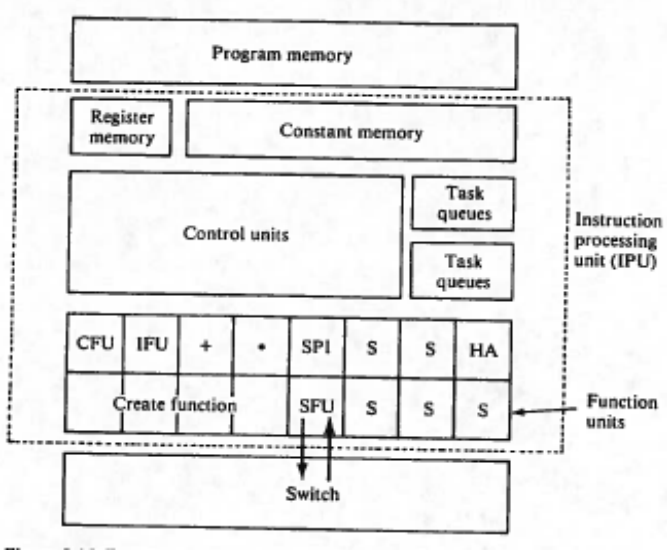

Program Execution Module

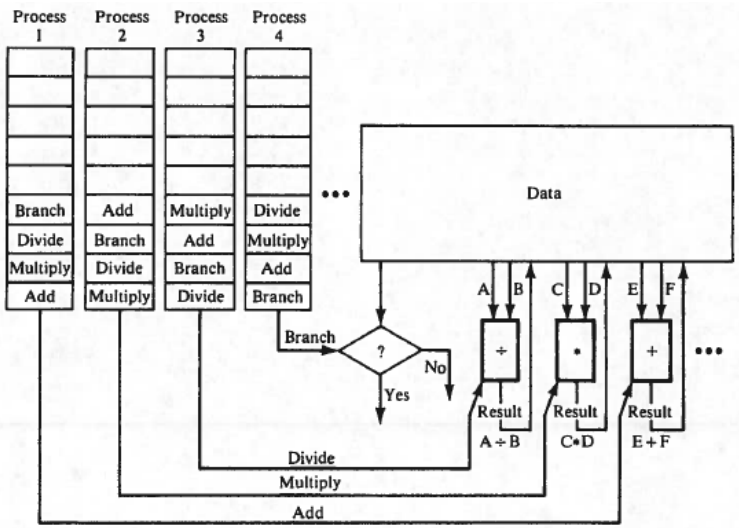

The Denelcor HEP

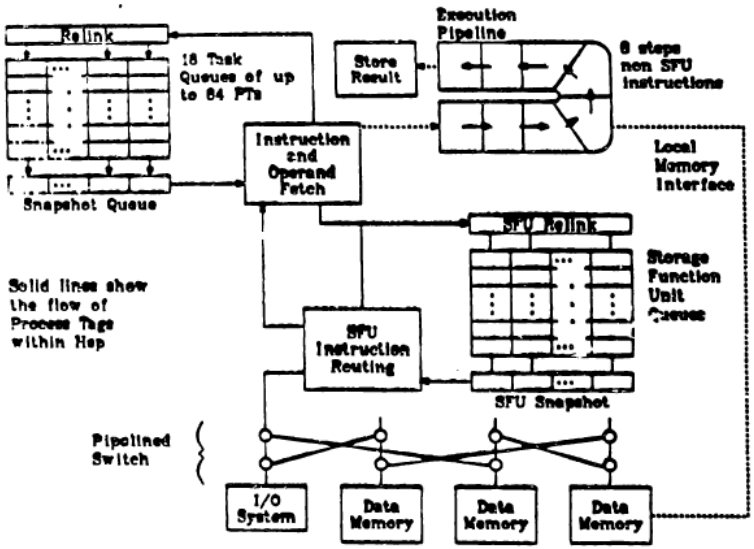

- MIMD
- 64-bit wide instructions
- 2048 64-bit registers
- 4096 Constant Memory locations
- Program Memory 256kB – 8 MB
- Up to 16 10 MIPS processors (100 nsec clock)
- Pipelined functional units, except divide (1700 nsec)
- Up to 128 data memory modules each with 256 kB – 8 MB memory with 80 MB/s bandwidth
- 32-bit address space

Sparse Matrix Computations (on the HEP)

How to exploit concurrency due to sparsity in factorization (LU, …)

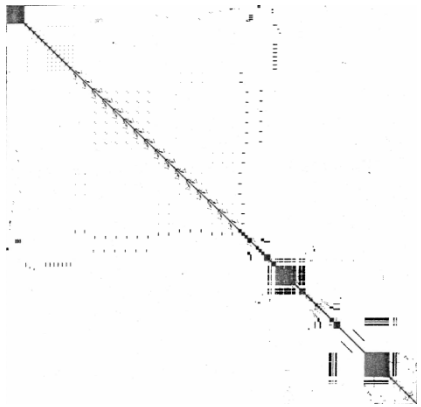

The idea

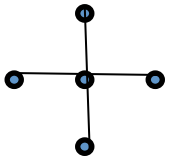

"Bad"

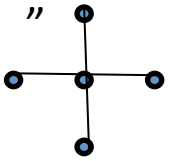

"Good"

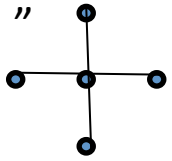

"Step 1"

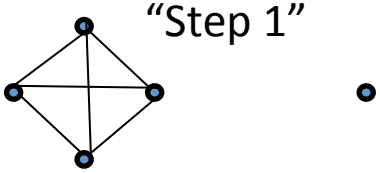

Two forms of concurrency:

- Concurrency in the elimination of a single vertex
- Concurrent elimination of separate vertices

Concurrency in Sparse Matrix Factorization

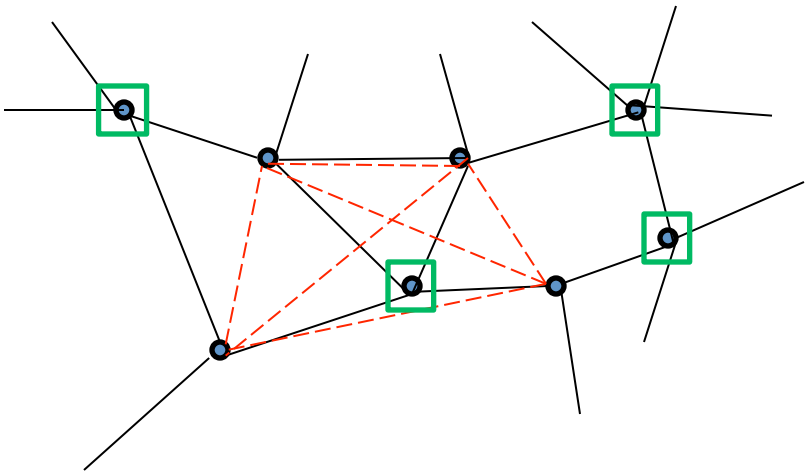

"Submatrix" for i,j,k,l,m

	m	i	j	k	l
m	x	x	x	x	x
i	x	x			
j	x		x	x	
k	x		x	x	
l	x				x

Concurrent Elimination of Nodes

- Row operations can be carried out concurrently
- Column operations can be carried out concurrently
- Row and column operations can be carried out concurrently

Example: Erisman 1176

Ordering	Min Degree (MD)	Nested Dissection	Generalized MD
Pivots longest path	175	111	82
	1	0.63	0.47
Columns longest path	1	4.63	1.13
Operations longest path	1	0.80	0.75
Speedups compared to sequential Minimum Degree			
Concurrent node elimin.	1.3	2.4	1.8
Concurrent column ops.	20	21	20
Concurrent node+col ops.	40	75	65
Concurrent row+col. ops	88	118	91
Concurrent node+row+col	468	922	912

Node orderings and concurrency in structurally Symmetric Sparse Problems
I.S. Duff, S.L. Johnsson, in Parallel Supercomputing: Methods, Algorithms and Applications, 1989

Example: 30 x 30 Mesh

Ordering	Min Degr. (MD)	Nested Dissection	General. MD	Red-Black +MD	Red-Black + ND
Pivots longest path	124	80	112	119	76
	1	0.65	0.90	0.95	0.61
Columns longest path	1	0.67	1.56	1	0.67
Operations longest path	1	0.75	3.51	1.13	0.72
Speedups compared to sequential Minimum Degree					
Concurrent node elimin.	2.4	4.0	1.7	2.4	3.4
Concurrent column ops.	12	13	23	13	13
Concurr. node+col ops.	46	80	67	50	75
Concurr. row+col. ops	65	79	161	74	71
Concurr. node+row+col	421	799	1070	496	746

Node orderings and concurrency in structurally Symmetric Sparse Problems
I.S. Duff, S.L. Johnsson, in Parallel Supercomputing: Methods, Algorithms and Applications, 1989

The Insight Gained

- For most of the tested matrices from the Boeing-Harwell collection and regular 2-D and 3-D meshes
  - Most of the concurrency to be exploited is in the elimination of a single graph node ("only" concurrency in dense matrix LU factorization)
  - Concurrency in elimination of different graph nodes (the "only" concurrency available for tri-diagonal systems and exploited by e.g. cyclic reduction) of much lower order, but can still contribute to significant speedup for highly parallel execution
- Application based approach to concurrency (as opposed to a DAG based approach)

The Alliant FX Series

- MIMD
- Up to 8 processors (Computational Elements)
  - 1 floating-point adder, pipelined
  - 1 floating-point multiplier, pipelined
  - 1 floating-point divider, pipelined?
  - 8 vector registers of 32 64-bit words each
- 170 nsec clock
- 11.8 MFlops/s peak single-precision (SP) perf. per CE, 5.9 MFlops/s peak double-precision (DP) perf. per CE
- Crossbar connect of CEs to two 64kB caches with 376 MB/s peak bandwidth connected to a 188 MB/s memory bus
- 8 MB memory modules, max 64 MB

The Alliant FX/8

- 8 CEs, 32 MB
- 94 MFlops/s SP peak,
  47 MFlops/s DP peak
- 5.5 kW
- Energy efficiency, nominal
  DP: 0.0085 MF/W

Solving Banded Systems on Parallel Processors

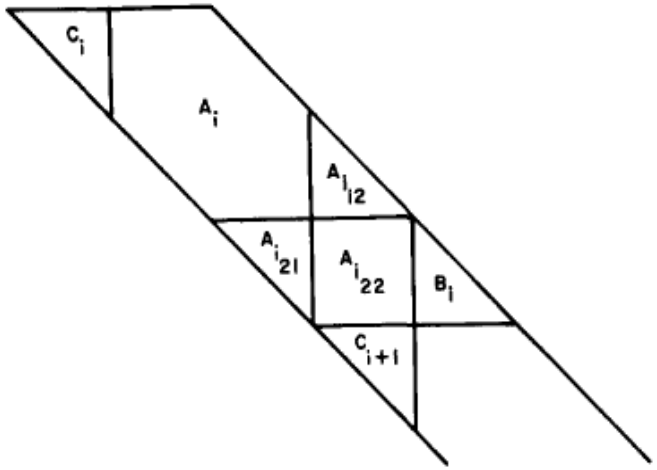

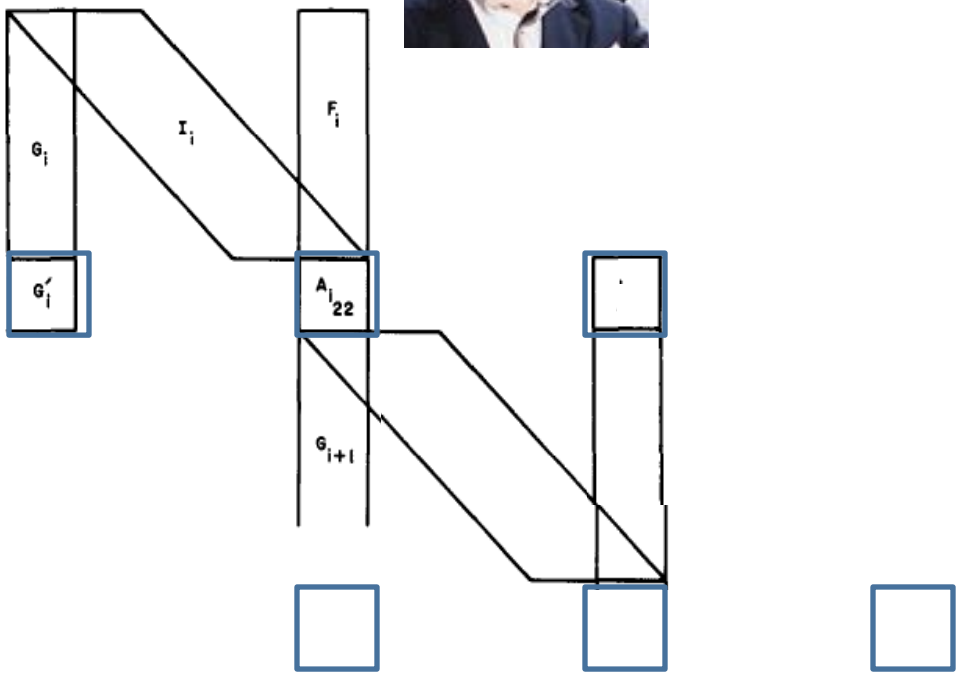

"Solving Banded Systems on Parallel Architectures", J. Dongarra, L. Johnsson, *Journal of Parallel Computing*, Vol. 5, No. 2, pp. 219 – 246, 1987.

Solving Banded Systems on Parallel Processors

Floating-point operations per second (kflops) for factorization on Alliant FX/8, one processor used

m		$P=1$	$P=2$	$P=4$	$P=8$	$P=16$
2	$N=1000$	44	44	40	44	43
	$N=2000$	40	42	42	40	39
	$N=4000$	42	42	42	41	40
	$N=8000$	42	41	41	41	40
8	$N=1000$	303	301	313	305	314
	$N=2000$	308	307	298	301	294
	$N=4000$	306	304	299	298	286
	$N=8000$	304	301	298	293	279
16	$N=1000$	633	626	635	686	794
	$N=2000$	629	641	638	639	643
	$N=4000$	638	640	629	620	593
	$N=8000$	632	631	627	611	575
32	$N=1000$	1164	1202	1335	1725	4039
	$N=2000$	1168	1172	1236	1731	1779
	$N=4000$	1156	1183	1184	1244	1390
	$N=8000$	1123	1110	1106	1074	1008

P=Partition

Peak efficiency: 34%
Parallel efficiency: 2 processors 90%, 4 processors 83%, 8 processors 50%
(scalability data not shown)

"Solving Banded Systems on Parallel Architectures", J. Dongarra, L. Johnsson, *Journal of Parallel Computing*, Vol. 5, No. 2, pp. 219 – 246, 1987.

The Connection Machine CM-2

First Thinking Machines Connection Machine in DOE at ANL (Rick Stevens) Year 1988

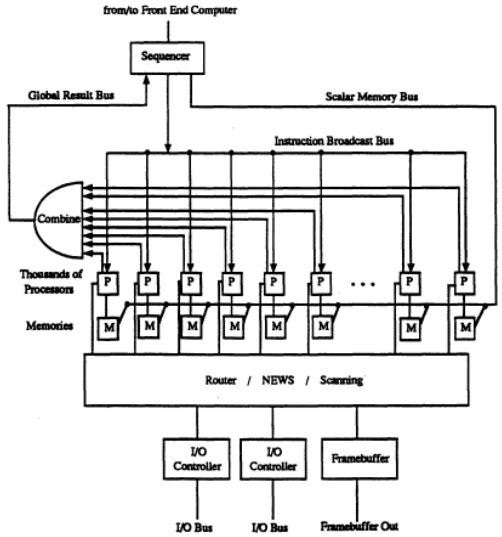

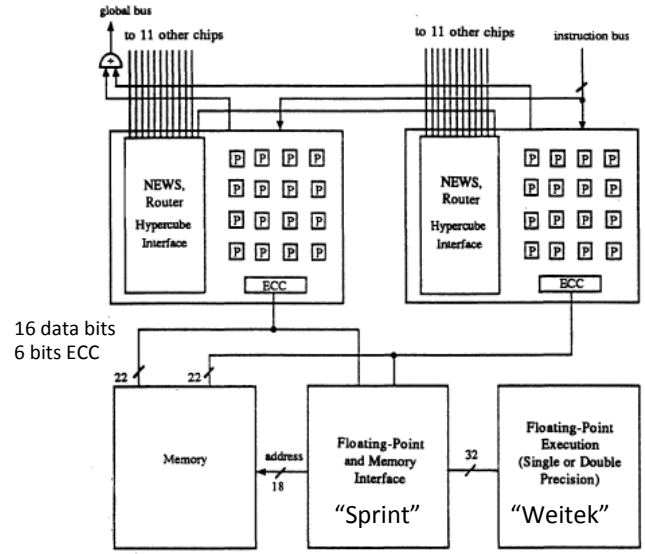

8k – 64k 1-bit processors (16k system at ANL)
16 processors per chip
8kB per 1-bit processor, 256 kbit memory x4 chips;
32kB per processor chip
(later 128 kB with 1 Mbit chips)
12 dimensional hypercube
Instruction set: Paris (PARallel Instruction Set) for bit-serial processing

32 or 64-bit floating-point units
1 unit per 32 1-bit processors
12 dimensional hypercube
Instruction set: CMIS (CM Instruction Set) "Slice-Wise"
6.7 MHz clock, 13.7 GF
28 kW with 512 MB of memory
0.489 MF/W nominally

Connection Machine

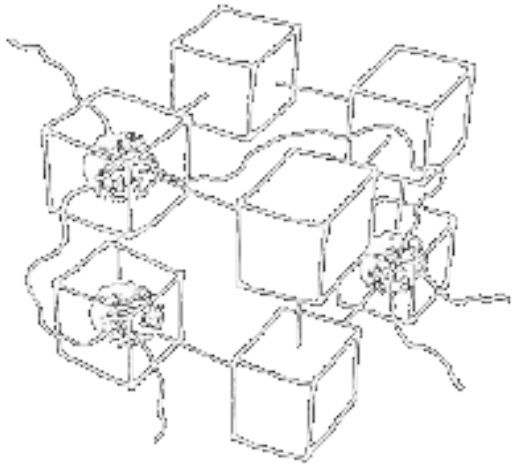

- Many interesting features
- Two programming models
  - Bit-serial
  - 32-bit/64-bit
- Virtual Processors
  - Global memory. Arrays automatically mapped to physical processors through virtualization
- Two addresses
  - "Send" addresses
  - "NEWS" addresses. Embedded emulation of meshes of arbitrary dimensions

Connection Machine

A Virtual Processor

SDI Access application (implemented on the ANL CM-2)

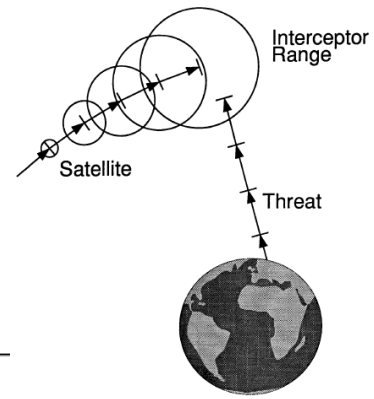

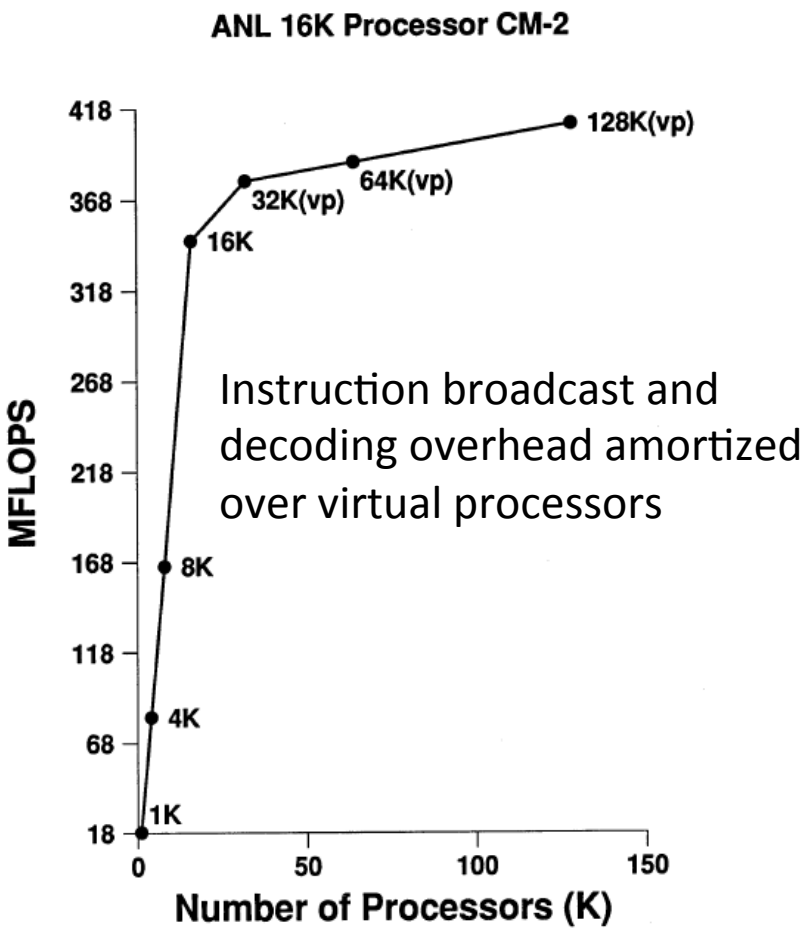

Instruction broadcast and decoding overhead amortized over virtual processors

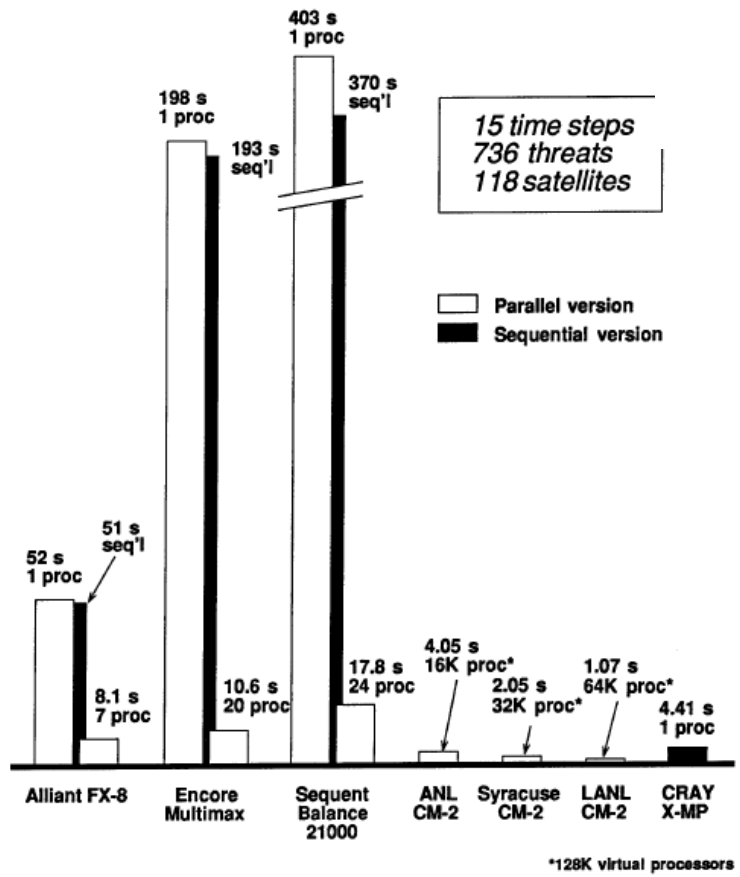

http://www.dtic.mil/cgi-bin/GetTRDoc?Location=U2&doc=GetTRDoc.pdf&AD=ADA344828

QCD on the Connection Machine

"In 1990-1992, the Connection Machine was the most powerful commercial QCD machine available" (@ 5.3 Gflops/s on a 64k CM-2)

http://www.netlib.org/utk/lsi/pcwLSI/text/node40.html#SECTION00637000000000000000

	Calc. time	Comm. time	Total time	GFlops
***Lisp**	8.7	4.5	13.2	0.9
Inner loop in CMIS	3.3	3.9	7.2	1.6
Multi-wire CMIS	1.7	0.5	2.2	5.3

Fermion update time (sec) on a 64k CM-2.

Machine/Compiler	GFlops/s
Cray XMP/CF 77 (1 proc)	0.13
Cray YMP/CF 77 (1 proc.)	0.21
CM-2/ CM Fortran	1.87
CM-2/CM *Lisp	1.95
CM-2/CMIS (WTL3132)	8.50
CM-2/CMIS (WTL3164)	10.20

Performance of mult. of a 3x3 by 3x1 matrix

6.7 MHz CM-2 @ LANL 13.7 GF peak
Matrix-multiplication peak efficiency 74.5%, overall efficiency 38.8%
0.365 MF/W based on installation power spec.

Code Generator for generation of straight-line matrix multiplication kernel code

QCD on the Connection Machine: Beyond *Lisp, R. G. Brickner, C. F. Baillie, S.L Johnsson, Computer Physics Communication, Vol. 65, Issues 1–3, 2 April 1991, Pages 39–51, http://www.sciencedirect.com/science/article/pii/001046559190152B

Top500 System Performance Evolution

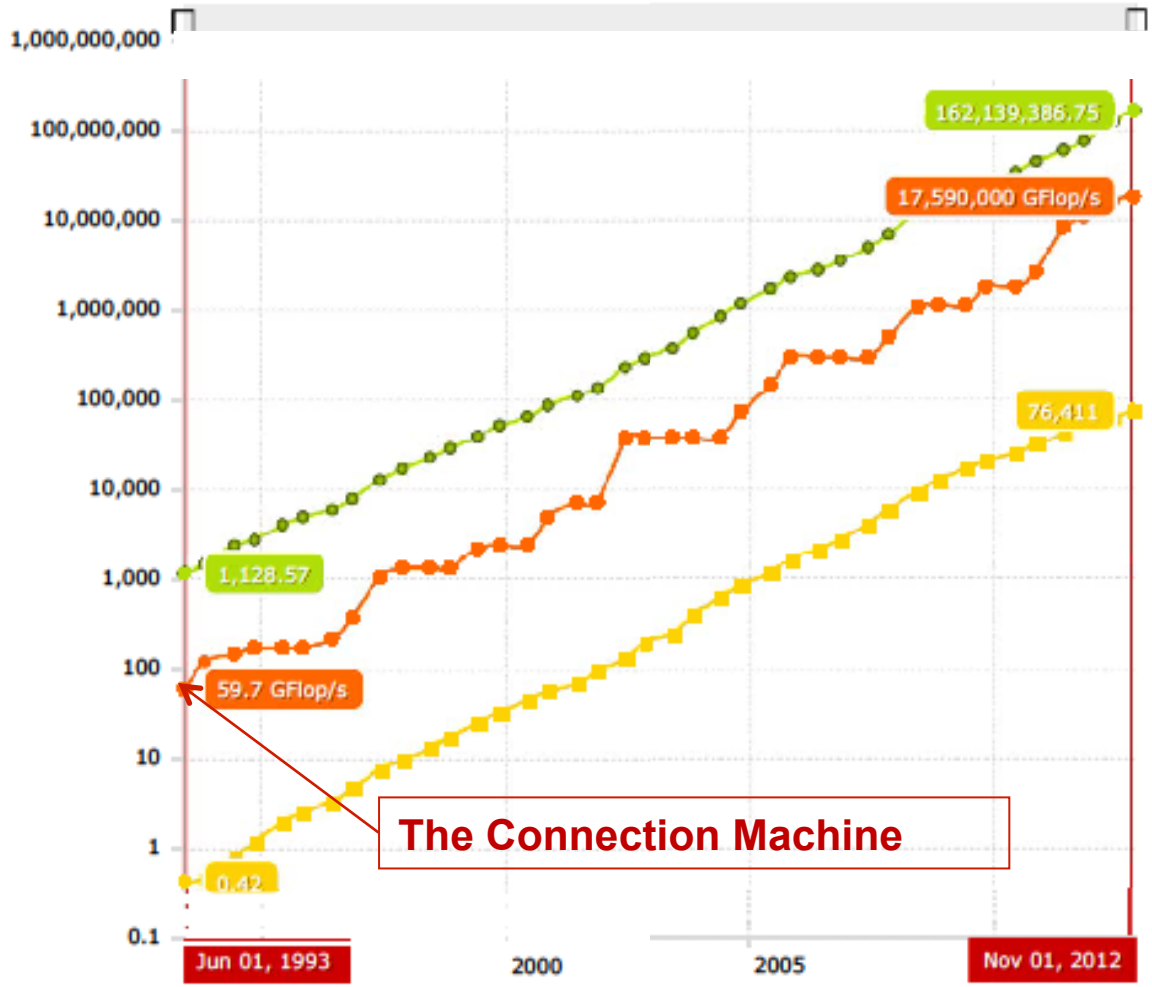

www.top500.org

Performance doubling period on average:

No 1 – 13.64 months

No 500 – 12.90 months

Five of the Top 10 systems CM-5's on the 1st Top500 list

Connection Machine Scientific Software Library

- Multiple Instance Functionality
- Two levels of concurrency
- "Overloaded" library calls – multiple algorithms dynamically chosen
- Run-time algorithm choice as a function of argument sizes and shapes, machine size and data allocation
- Created through the use of code generation (assembly) from higher level description
- For matrix operations with non-uniform operations for data elements load-balance by execution order, or data allocation, to avoid data reallocation for load-balance

HPL Scalability on CM-5

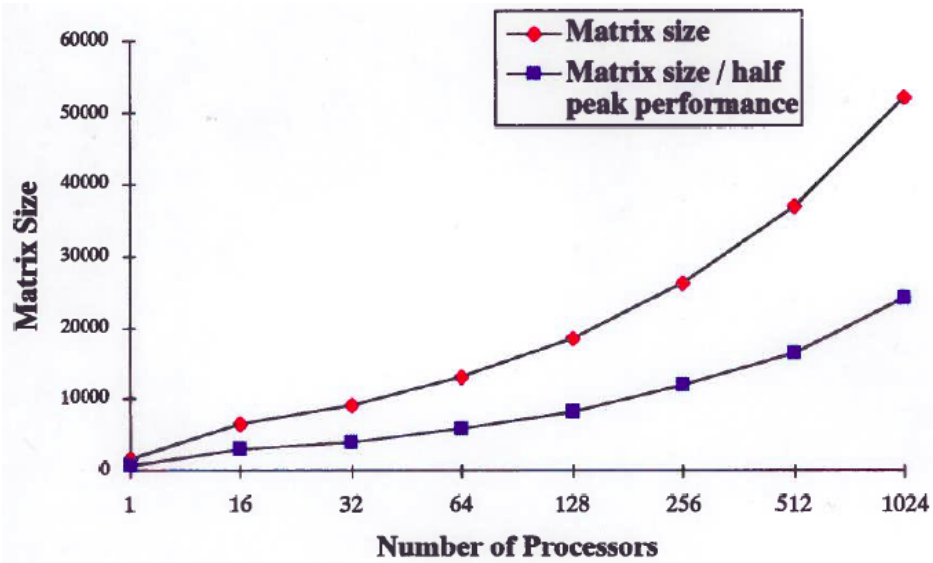

CM-5 nodes	Matrix size	MF/s	MF/s/node	$N_{1/2}$
1	1632	68	16.99	672
16	6528	978	15.28	3008
32	9216	1937	15.13	4096
64	13056	3809	14.88	6016
128	18432	7649	14.94	8192
256	26112	15058	14.71	12032
512	36864	30402	14.85	16384
1024	52224	59667	14.57	24064

No. 1 system on 1st Top500 list, 1993

Peak efficiency: 45.5%
Parallel efficiency: 85.8%
Energy efficiency: 0.25 MF/W estimate

Out-of-Core LU Factor and Solve

CM-5, 64 nodes, 8 GF/s peak, complex data, scalable disk array

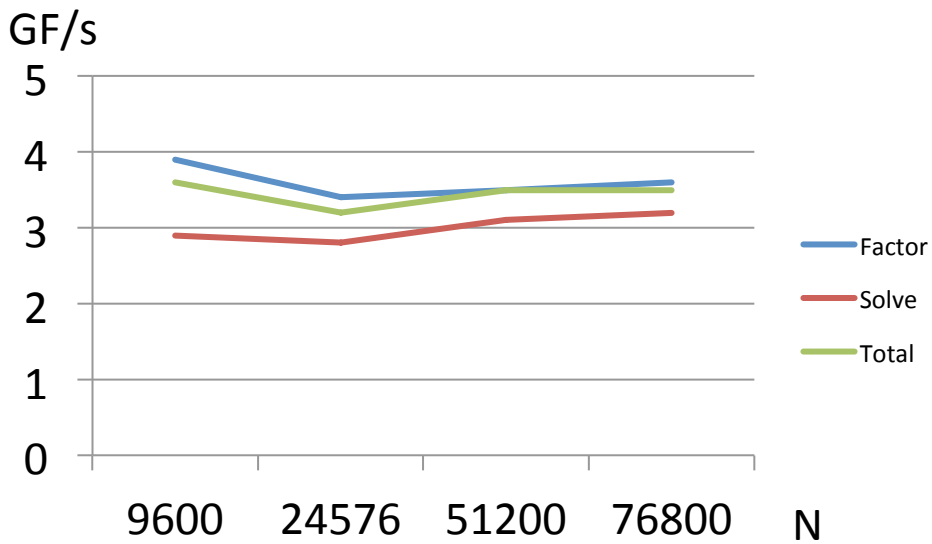

Efficiency:

Factor 48% w/o I/O
44% w I/O
Total 44% w/o I/O
43% w I/O

N	Disks	Factor		Solve (1024 RHS)		Total	
		Hrs	GF/s	Hrs	GF/s	Hrs	GF/s
9600	0	0.17	3.9	0.07	2.9	0.24	3.6
24576	32	3.2	3.4	0.5	2.8	3.7	3.2
51200	64	28.2	3.5	1.9	3.1	30.2	3.5
76800	118	94.6	3.6	4.2	3.2	98.8	3.5

Matrix-Matrix Multiplication

3D index space

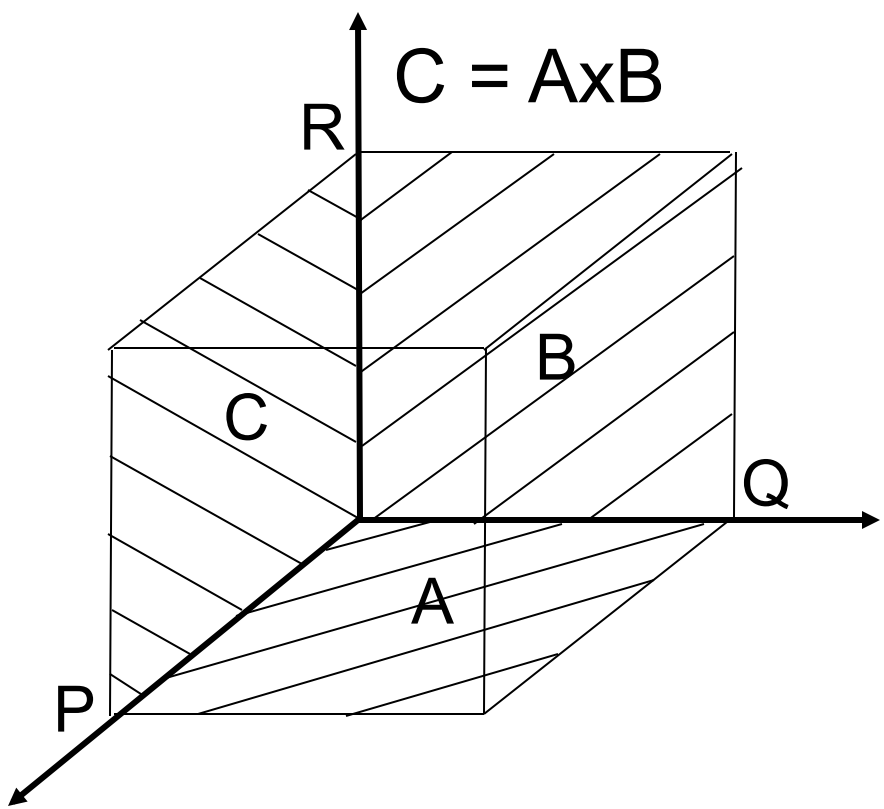

3-D matrix multiplication (in CMSSL):
comm. time ~$N^2/P^{2/3}$
compared to ~$N^2/P^{1/2}$
for a 2-D algorithms.

Communication efficient Linear Algebra has received renewed interest with interesting algorithms and analysis by Jim Demmel

Parallel Programming

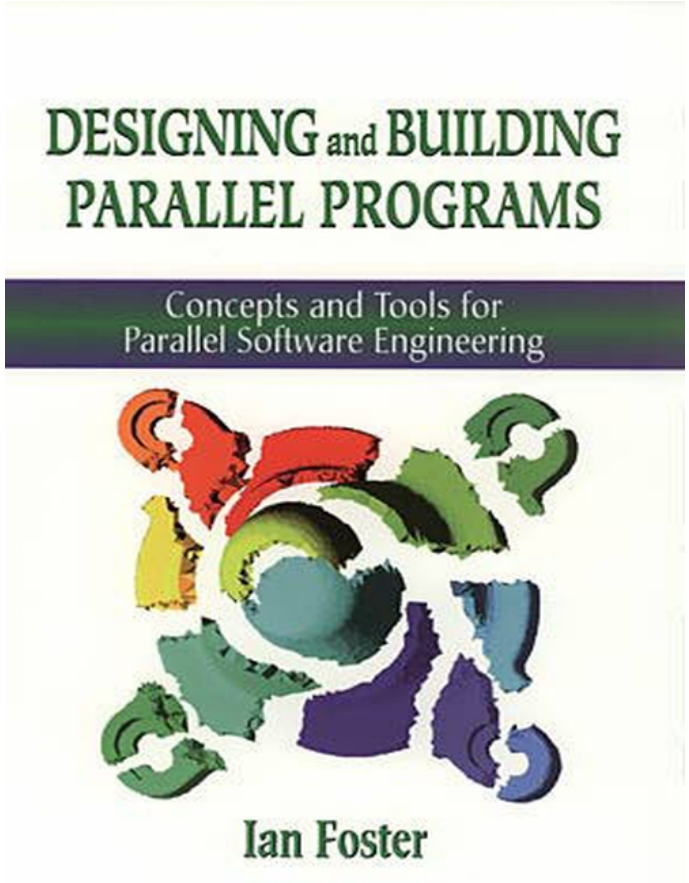

Partitioning Example (81,649 tetrahedra)

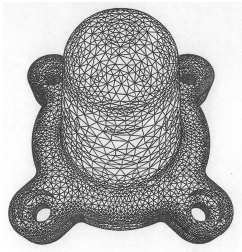

Mesh courtesy Mark Shepard, RPI

Parallel Spectral Bi-section

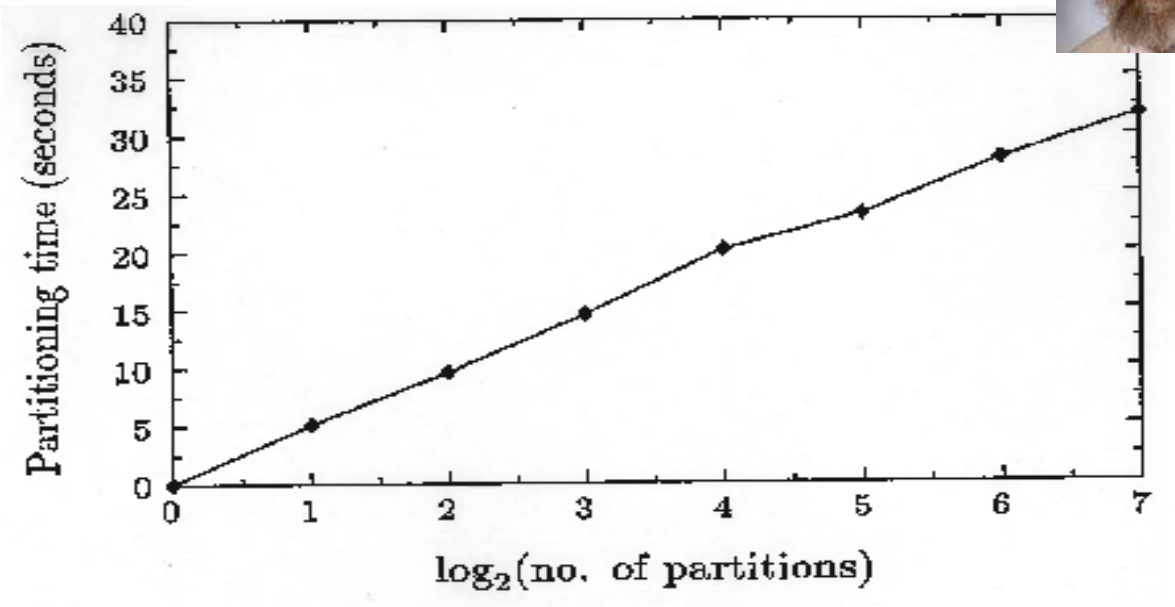

- Partitioning quality for 128 partitions
  - Graph edges: 152,878
  - Edges cut: 10,648
  - Percent cut: 7.0%

Source: J. Zdenek, K. Mathur, L. Johnsson, T.J. R. Hughes

Partitioning Examples - Scalability

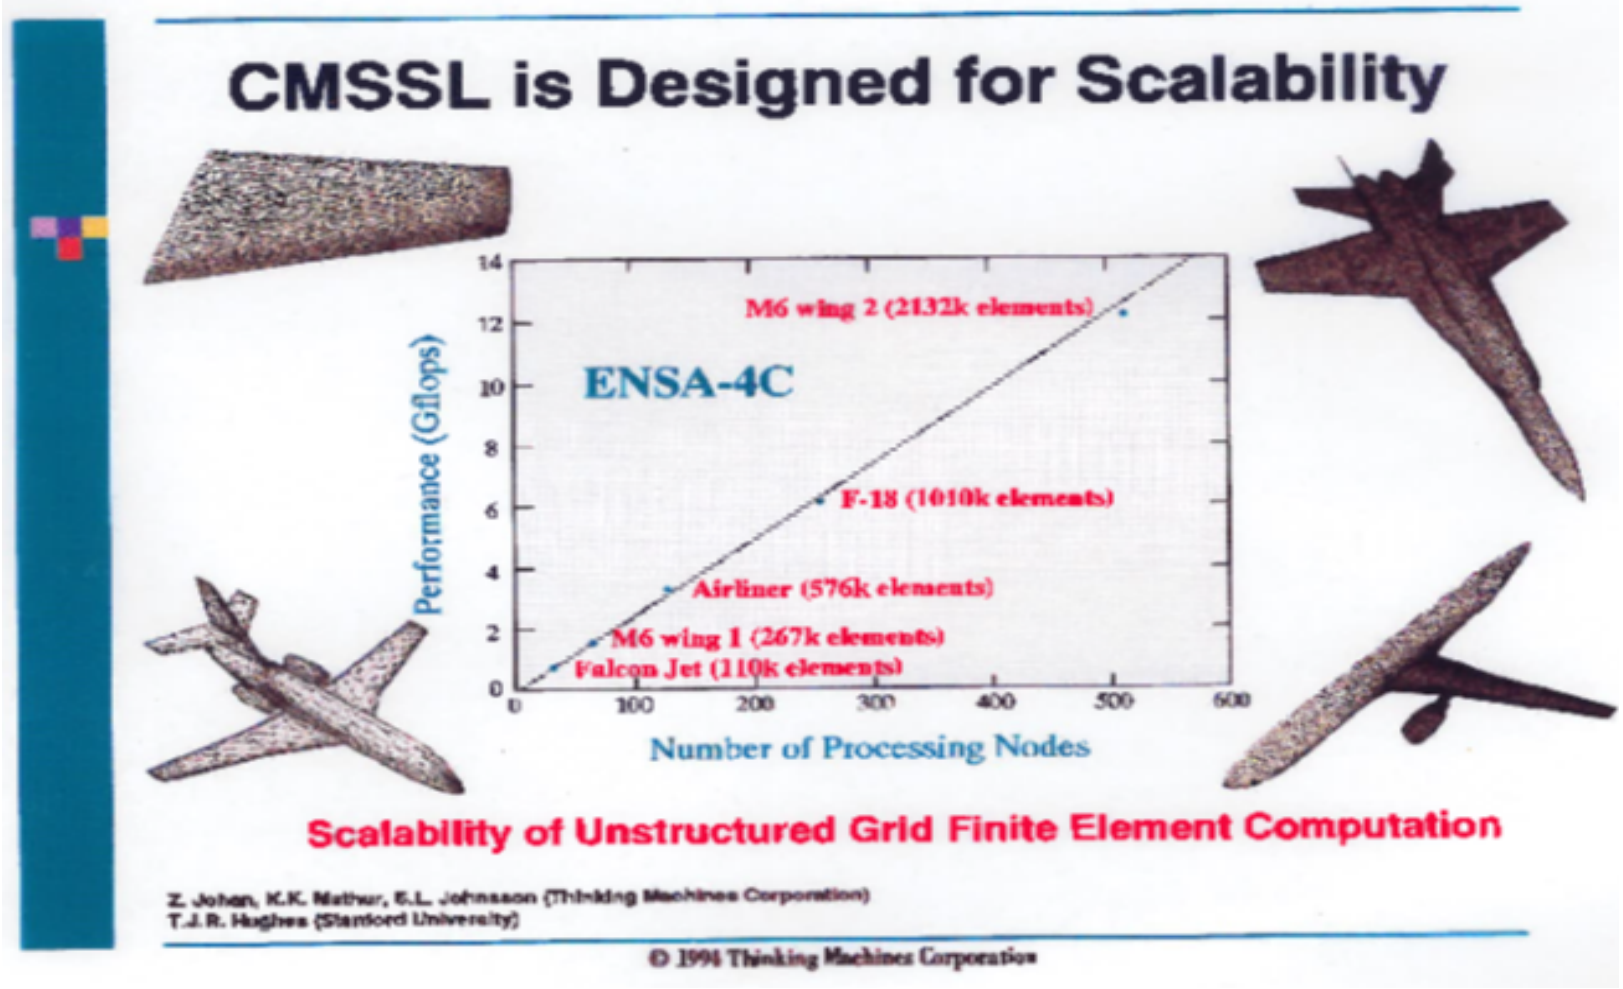

Source: J. Zdenek, K. Mathur, L. Johnsson, T.J. R. Hughes

Unstructured CFD

- M6 wing, 48,011 nodes, 256,566 tetrahedra, double precision

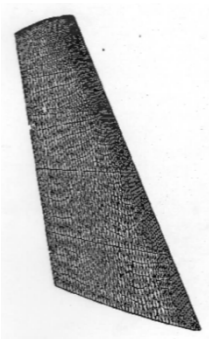

Operation	CM-5, 64-nodes	Cray C90
Total	14 min 15 sec 1.5 GF/s	42 min 12 sec 0.44 GF/s

- Falcon Jet, 19,417 nodes, 109,914 elements, double precision, 32-node CM-5

Operation	Default mapping	Random mapping
Total time	14min 6 sec	6 min 02 sec

Software

- MPI 1.0
  - Collective Communication
- Globus

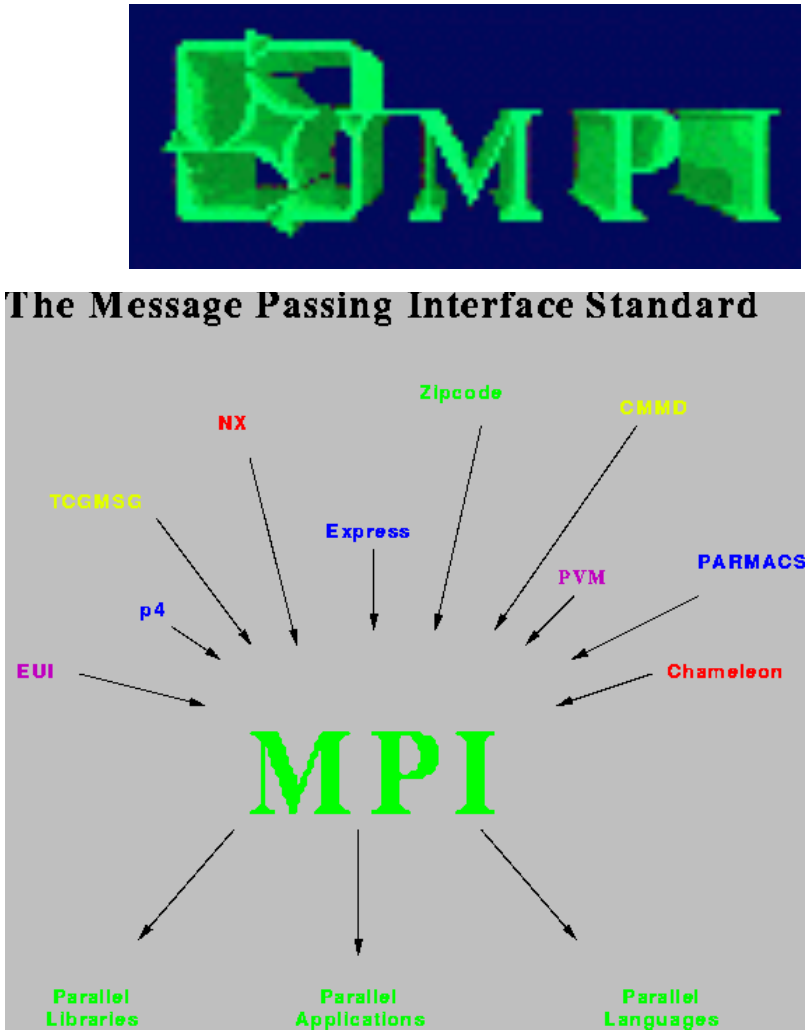

Globus

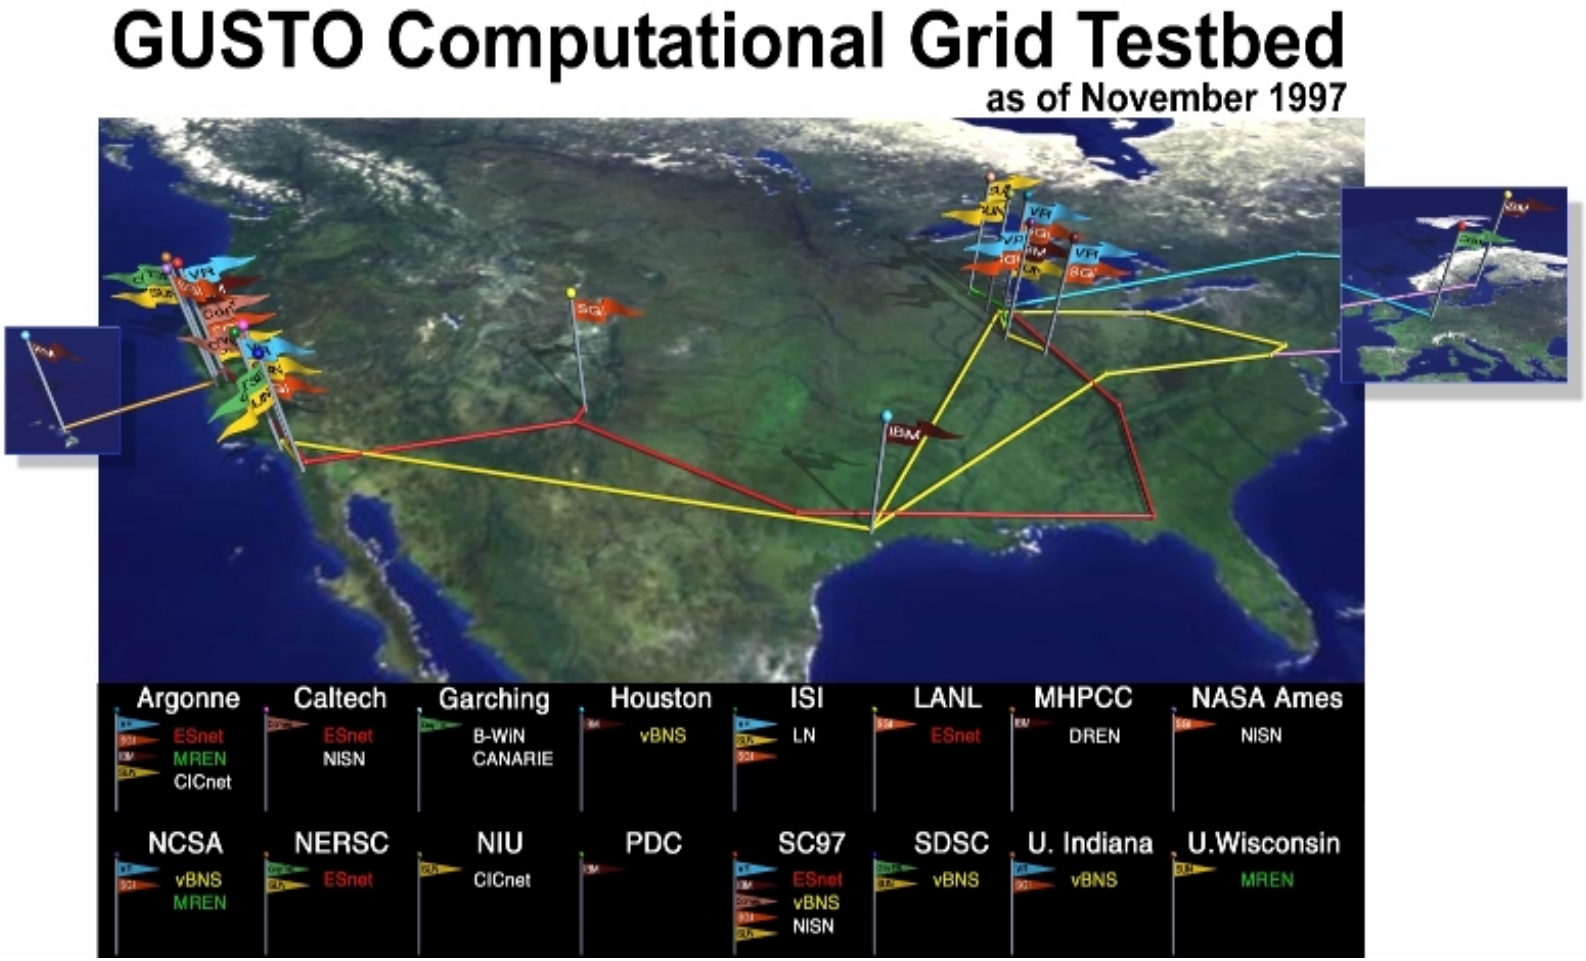

GUSTO tesbed for Grid applications demonstrated at Supercomputing97 exhibition

The Grid: Computational Steering

GEMSviz at iGRID 2000

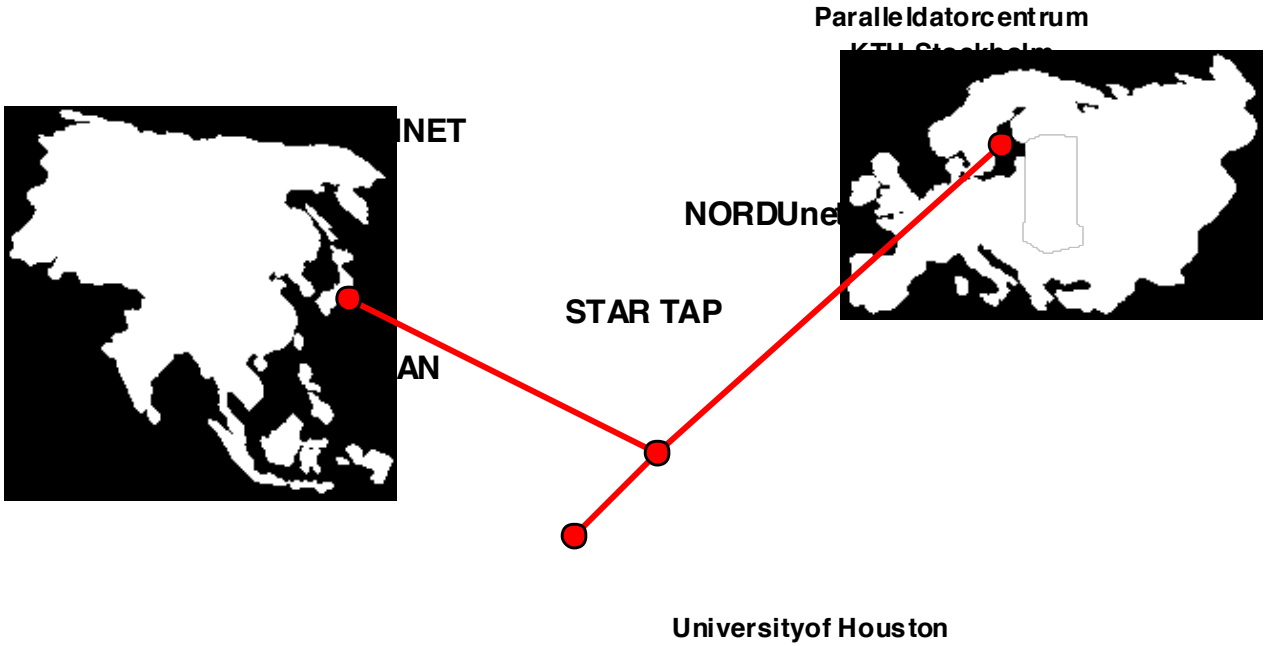

NORDUnet 2000

28 Sep 00 - #17

http://www.pdc.kth.se/projects/gemsviz

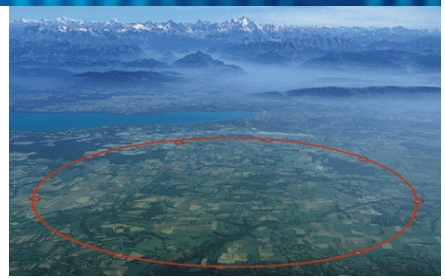

The LHC Detectors

ATLAS

CMS

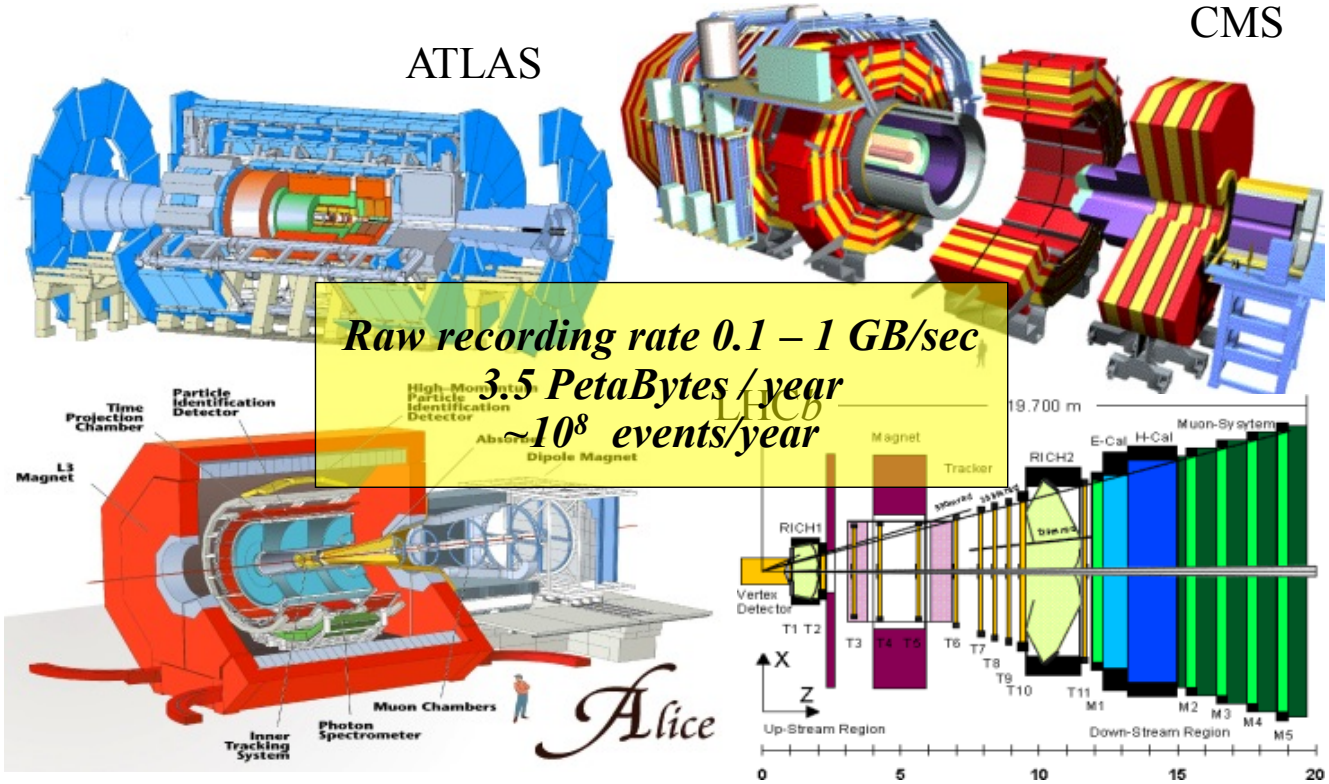

Raw recording rate 0.1 – 1 GB/sec
3.5 PetaBytes / year
~10^8 events/year

THIRTY YEARS OF PARALLEL COMPUTING AT ARGONNE: A Symposium
May 14-15, 2013
Argonne National Laboratory

UNIVERSITY of HOUSTON

GGF

- Globus Alliance
- European Grid Forum

Grid Application Development Software (GrADS)

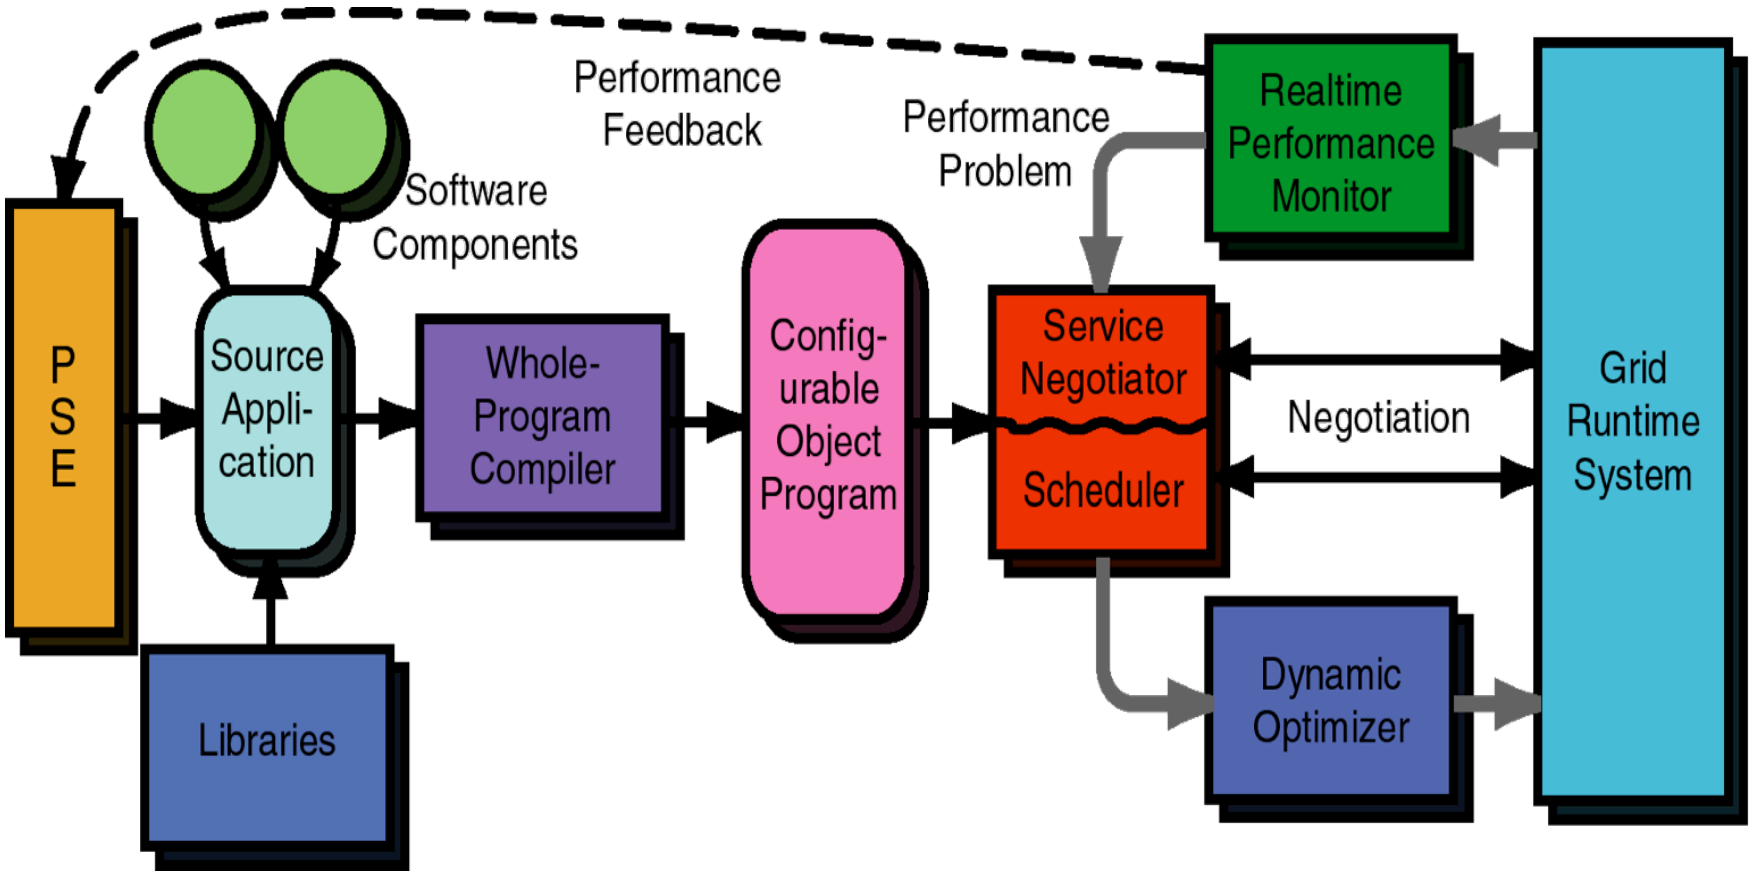

GrADS Scalapack Resource Selection

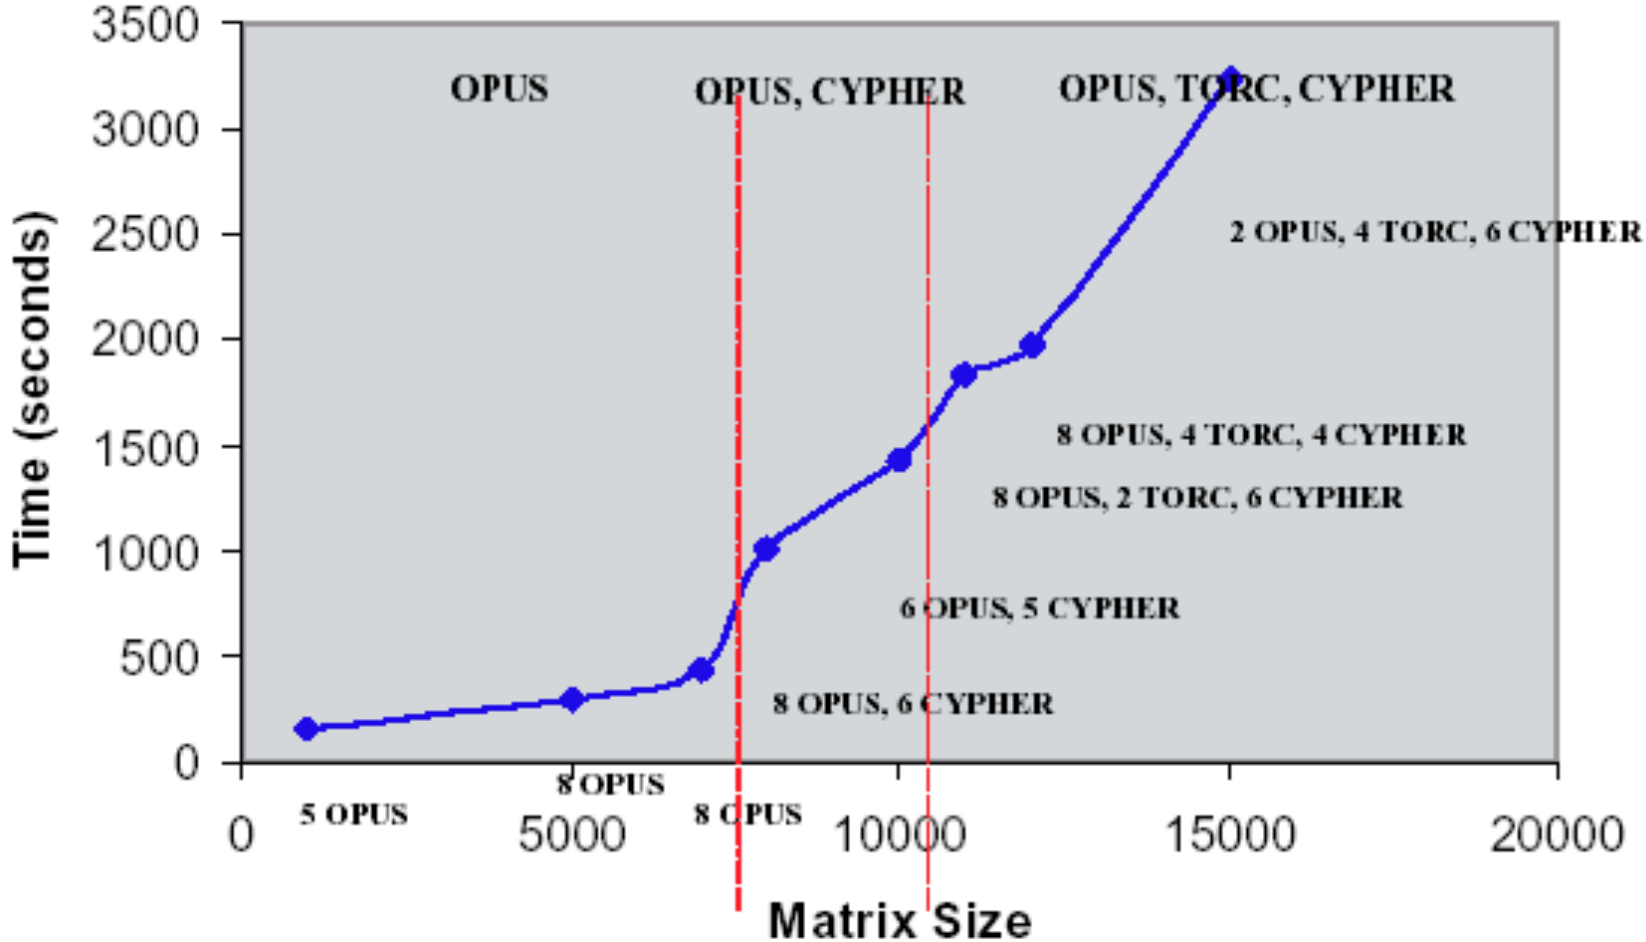

Energy Efficient Computing

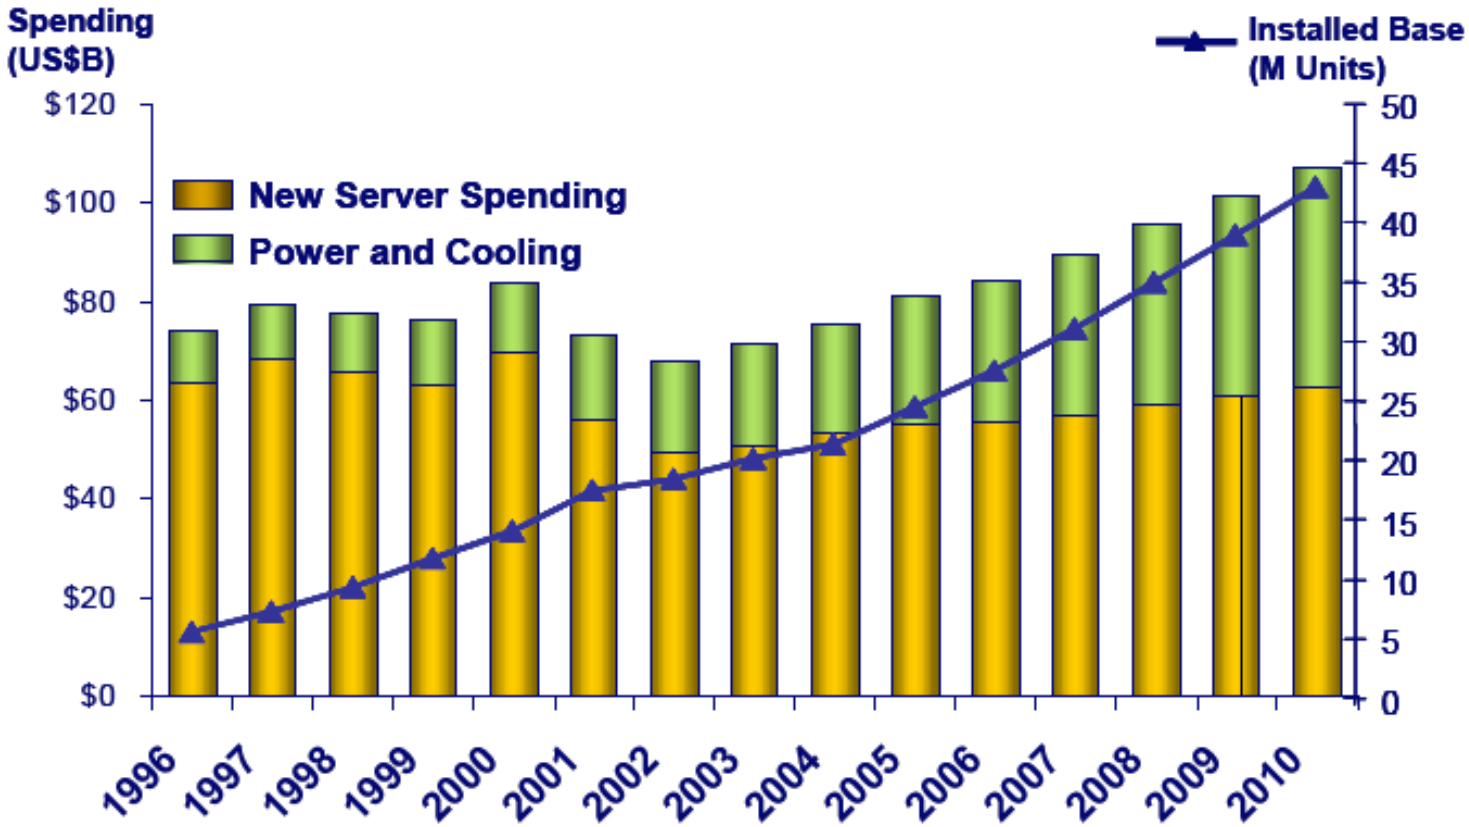

Estimate for 2008 purchase: 4 yr cooling cost ~1.5 times cluster cost

Source: IDC, 2006

Energy Efficiency Evolution

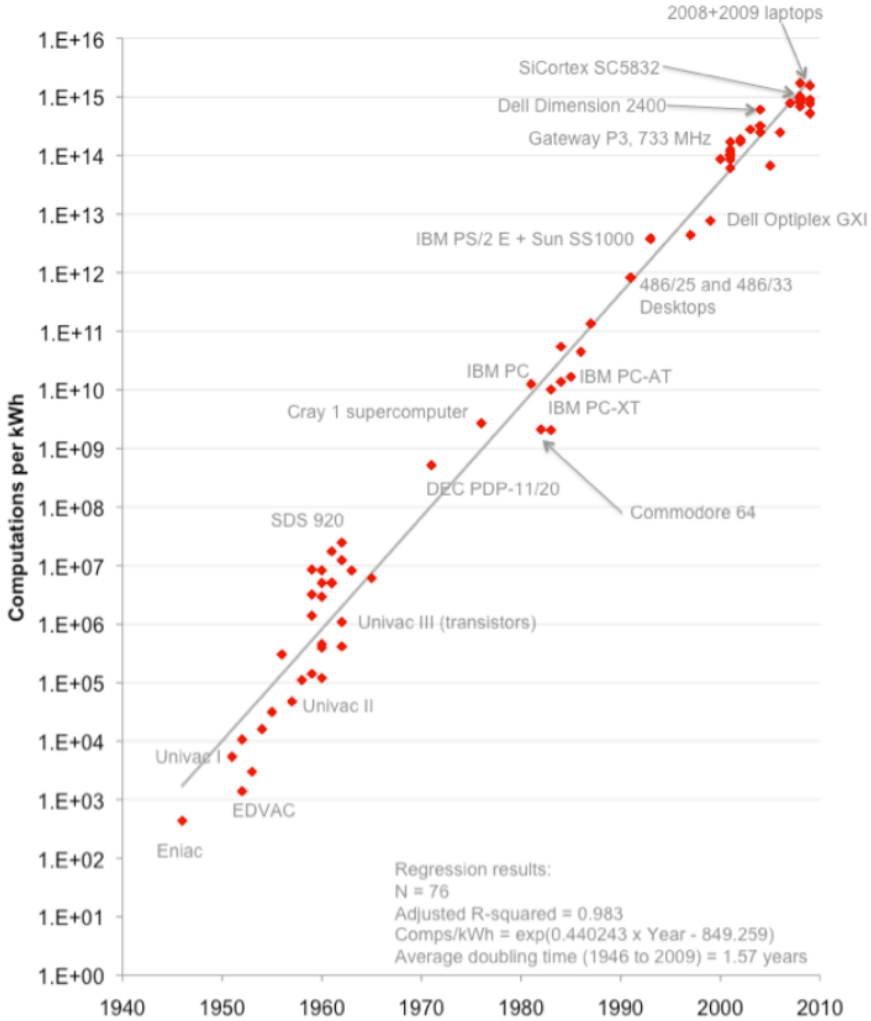

Energy efficiency doubling every 18.84 months on average measured as computation/kWh

Source: Assessing in the Trends in the Electrical Efficiency of Computation over Time, J.G. Koomey, S. Berard, M. Sanchez, H. Wong, Intel, August 17, 2009, http://download.intel.com/pressroom/pdf/computertrendsrelease.pdf

Top500 System Performance Evolution

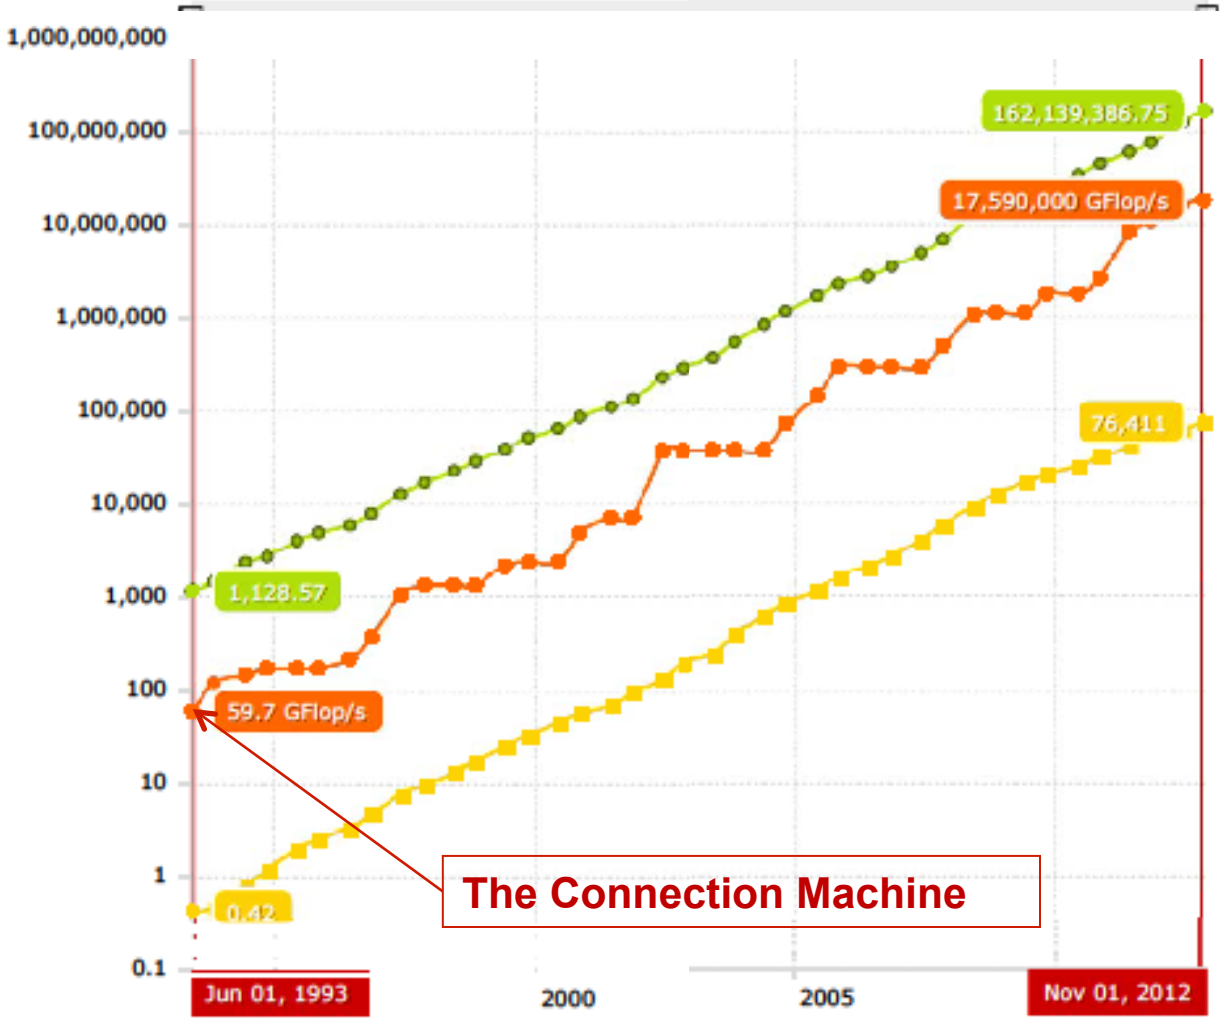

Performance doubling period on average:

No 1 – 13.64 months

No 500 – 12.90 months

www.top500.org

The Gap

The energy efficiency improvement as determined by Koomey does not match the performance growth of HPC systems as measured by the Top500 list

The Gap indicates a growth rate in energy consumption for HPC systems of about 20%/yr.

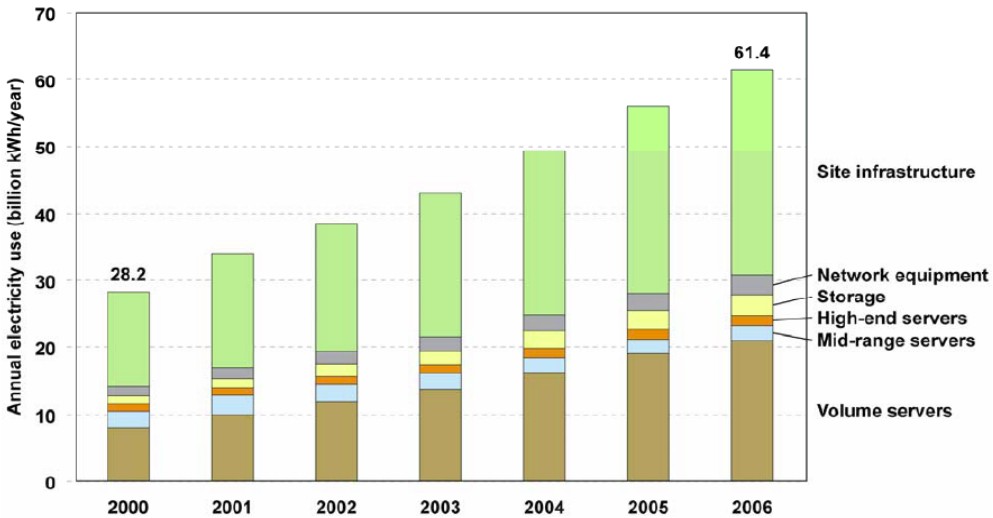

End use component	2000		2006		2000 – 2006 electricity use CAGR
	Electricity use (billion kWh)	% Total	Electricity use (billion kWh)	% Total	
Site infrastructure	14.1	50%	30.7	50%	14%
Network equipment	1.4	5%	3.0	5%	14%
Storage	1.1	4%	3.2	5%	20%
High-end servers	1.1	4%	1.5	2%	5%
Mid-range servers	2.5	9%	2.2	4%	-2%
Volume servers	8.0	29%	20.9	34%	17%
Total	28.2		61.4		14%

EPA study projections: 14% - 17%/yr

Uptime Institute projections: 20%/yr

PDC experience: 20%/yr

Report to Congress on Server and Data Center Energy Efficiency", Public Law 109-431, U.S Environmental Protection Agency, Energy Star Program, August 2, 2007, http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf

"Findings on Data Center Energy Consumption Growth May Already Exceed EPA's Prediction Through 2010!", K. G. Brill, The Uptime Institute, 2008, http://uptimeinstitute.org/content/view/155/147

What can be achieved with commodity technology and a focus on energy efficiency, with preserved programming model? (2009 – 2010)

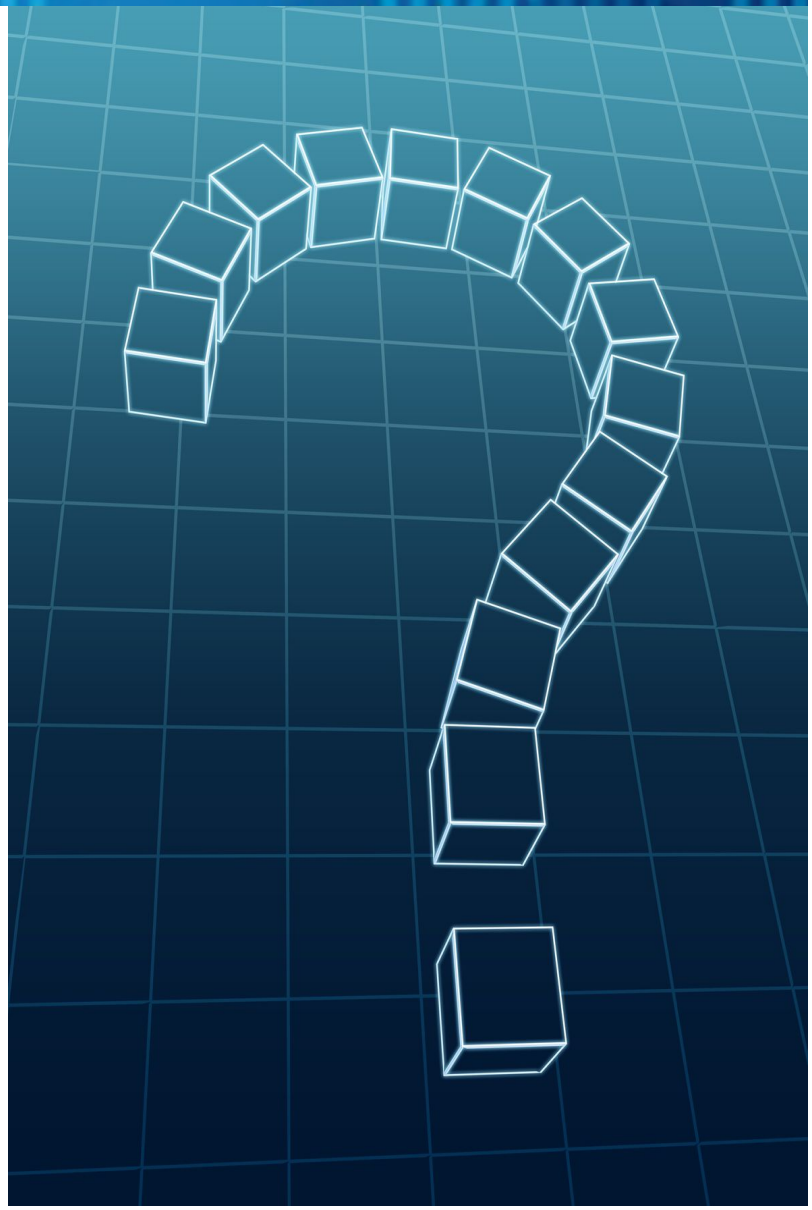

What Kind of Architecture (core)

How Small is Small

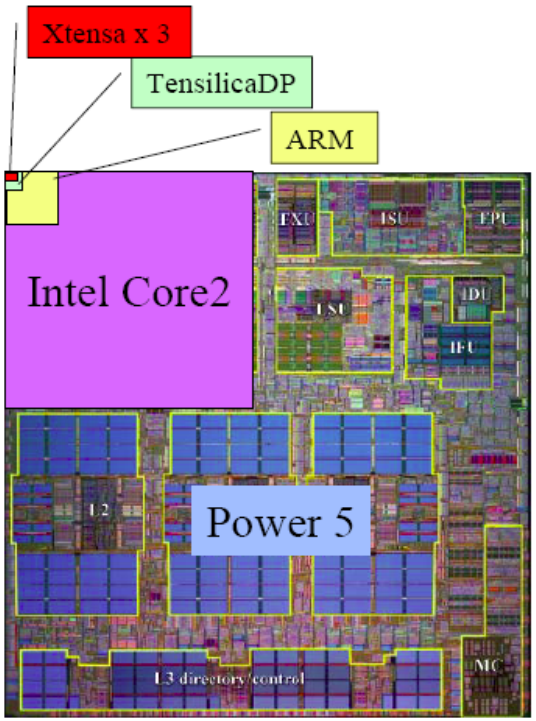

- Power5 (server)
  - 389mm^2
  - 120W@1900MHz
- Intel Core2 sc (laptop)
  - 130mm^2
  - 15W@1000MHz
- ARM Cortex A8 (toaster oven)
  - 5mm^2
  - 0.8W@800MHz
- Tensilica DP (cell phones)
  - 0.8mm^2
  - 0.09W@600MHz
- Tensilica Xtensa (Cisco Rtr)
  - 0.32mm^2 for 3!
  - 0.05W@600MHz

- Cubic power improvement with lower clock rate due to V^2F
- Slower clock rates enable use of simpler cores
- Simpler cores use less area (lower leakage) and reduce cost
- Tailor design to application to reduce waste

Each core operates at 1/3 to 1/10th efficiency of largest chip, but you can pack 100x more cores onto a chip and consume 1/20 the power

http://www.csm.ornl.gov/workshops/SOS11/presentations/j_shalf.pdf

What Type of Architecture?

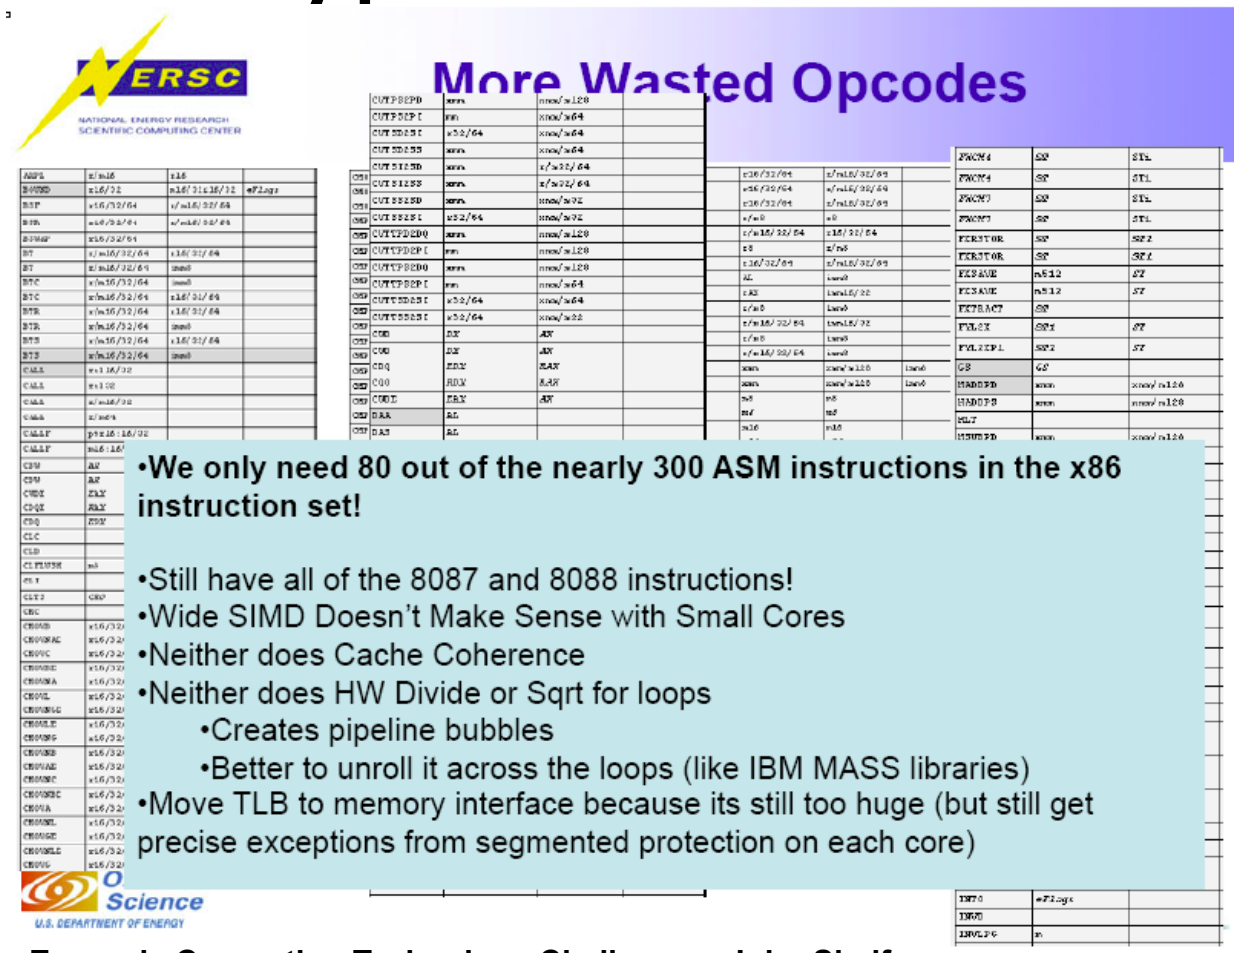

Exascale Computing Technology Challenges, John Shalf
National Energy Research Supercomputing Center, Lawrence Berkeley National Laboratory
ScicomP / SP-XXL 16, San Francisco, May 12, 2010

PUE – State-of-the-Art - Facebook

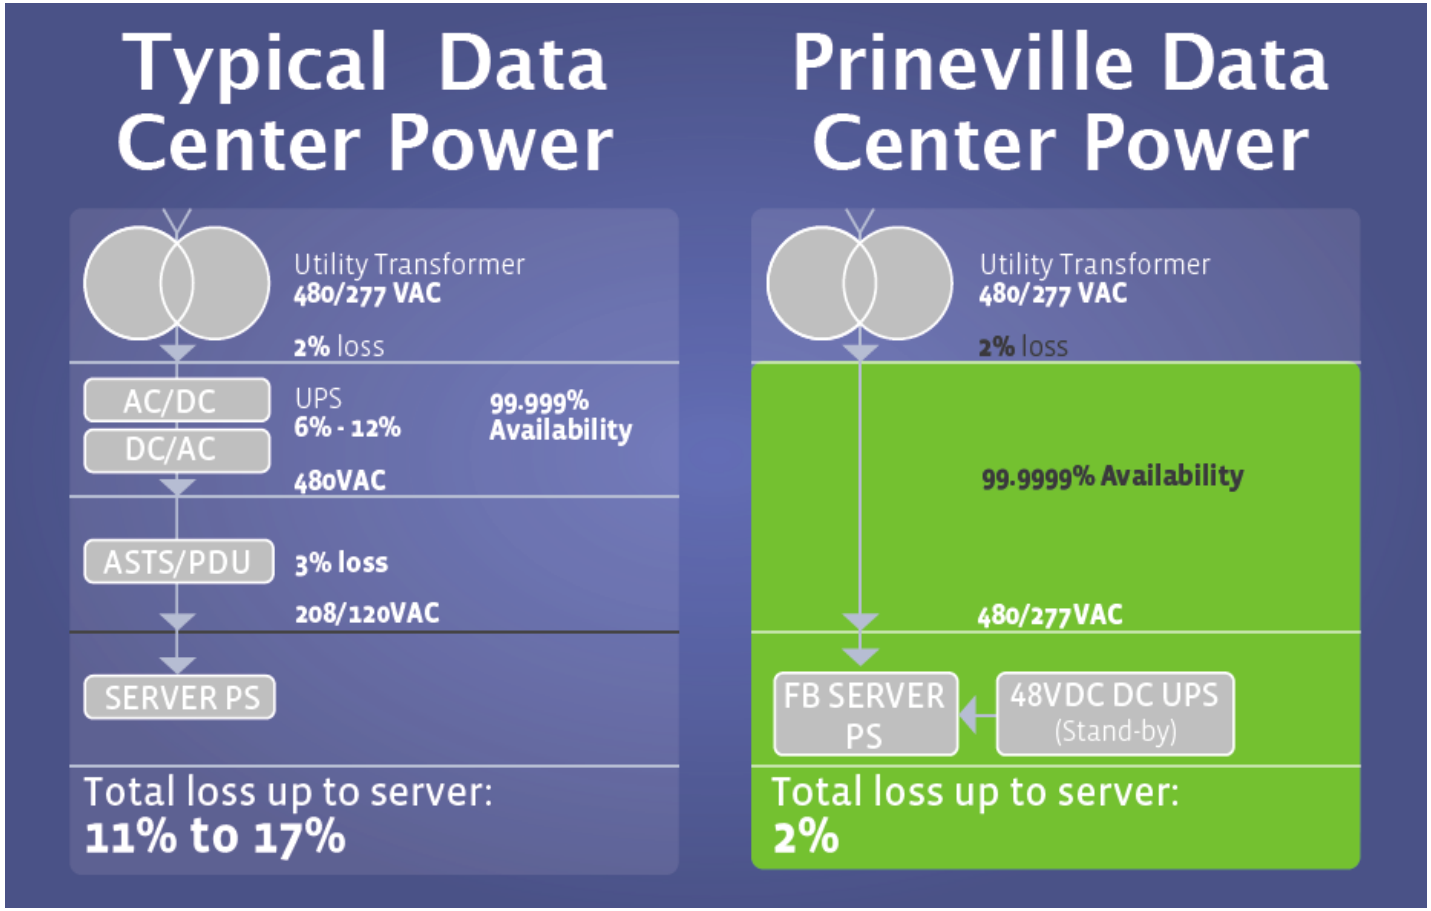

Google: UPS integrated with server PSU. UPS efficiency 99.9%

US Patent Office Application. June 1, 2007. Data Center Uninterruptible Power Distribution Architecture. http://appft1.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PG01&p=1&u=%2Fnetahtml%2FPTO%2Fsrchnum.html&r=1&f=G&l=50&s1=%2220080030078%22.PGNR.&OS=DN/20080030078&RS=DN/20080030078

Source: Amir Micahel, Facebook,August 17, 2011, http://www.hotchips.org/archives/hc23

Commodity Technology Prototype (2009)

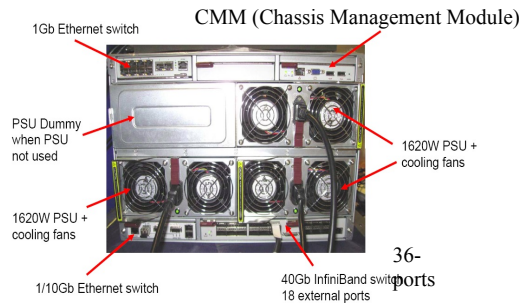

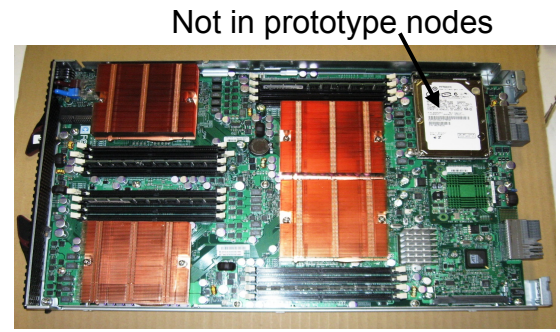

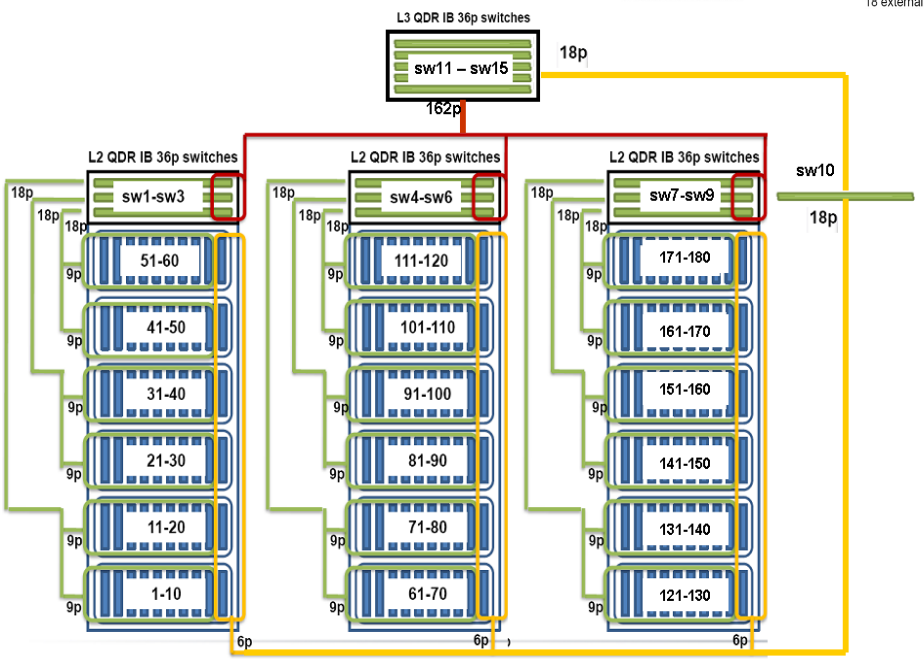

- New 4-socket blade with 4 DIMMs per socket supporting PCI-Express Gen 2 x16
- Four 6-core 2.1 GHz 55W ADP AMD Istanbul CPUs, 32GB/node
- 10-blade in a 7U chassis with 36-port QDR IB switch, new efficient power supplies.
- 2TF/chassis, 12 TF/rack, 30 kW (6 x 4.8)
- 180 nodes, 4320 cores, full bisection QDR IB interconnect

HPL Efficiency

• The prototype	79%
• BG/P	82%
• Reference platform (Dual socket HPTN)	91%
• Reference platform + GPU	53%

HPL Energy Efficiency

• The prototype	344 MF/W
• BG/P	357 MF/W
• Reference platform (Dual socket HPTN)	240 MF/W
• Reference platform + GPU	270 MF/W

Commodity Technology Prototype Results

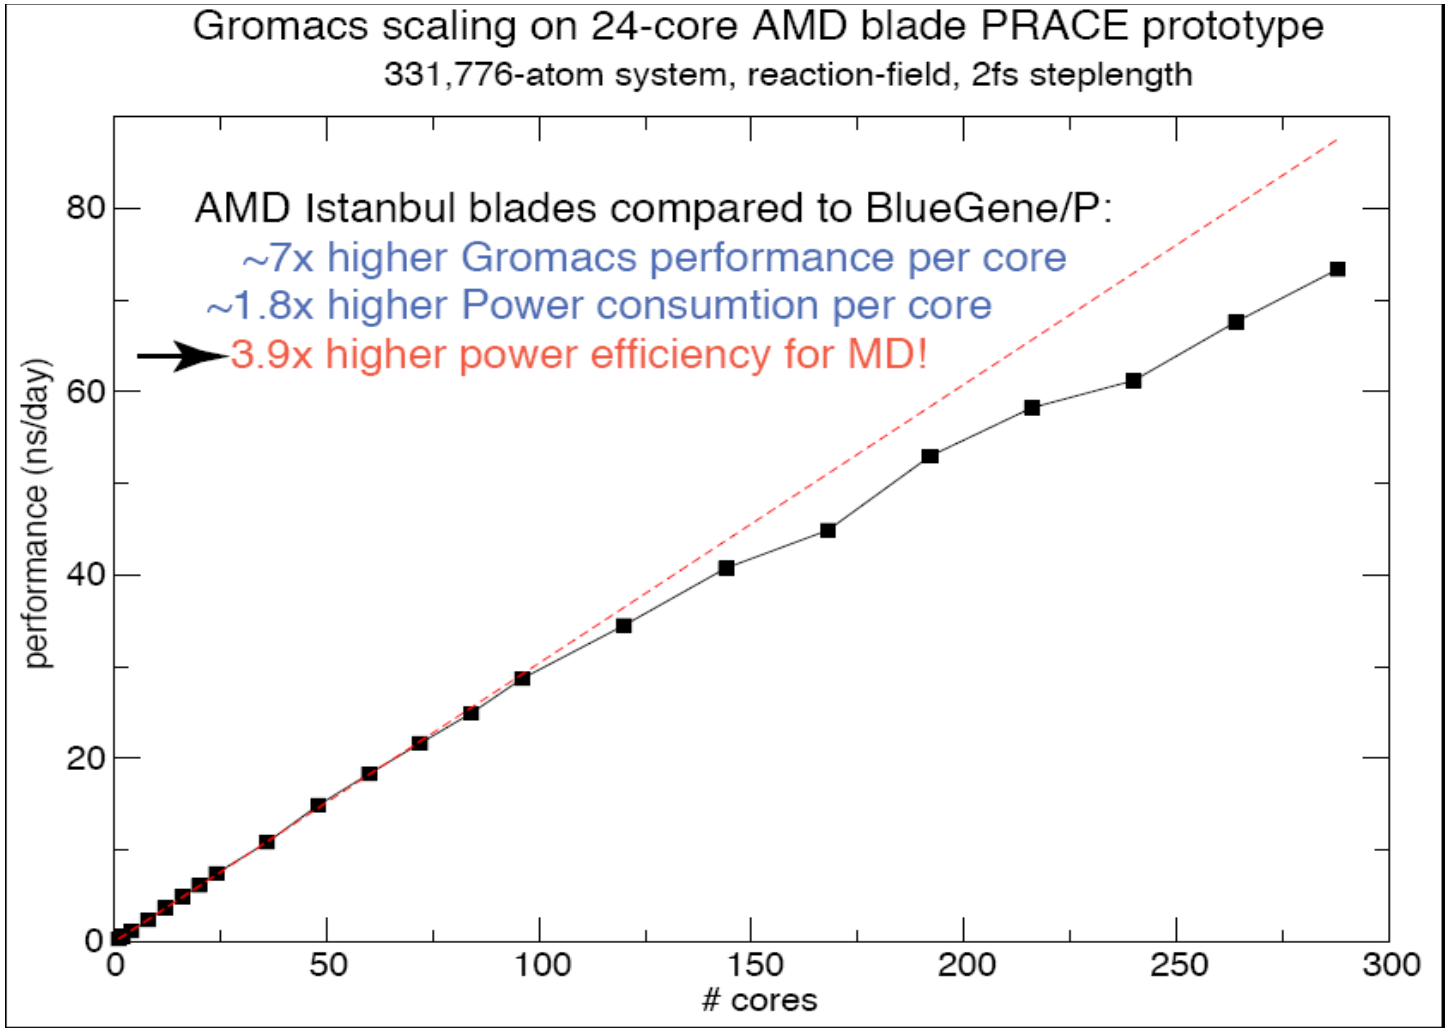

Lessons Learned

- Much can be achieved with focused engineering (without acceleration) at modest cost (which the internet companies have been pursuing for their workloads)
- Benchmark selection as important as ever (HPL not enough)

Nominal Energy Efficiency of Mobile CPUs, x86 CPUs and GPUs (2010)

ARM Cortex-9			ATOM			AMD 12-core			Intel 6-core			ATI 9370		
Cores	W	GF/W	Cores	W	GF/W	Cores	W	GF/W	Cores	W	GF/W	Cores	W	GF/W
4	~2	~0.5	2	2+	~0.5	12	115	~0.9	6	130	~0.6	1600	225	~2.3

nVidia Fermi			TMS320C6678			IBM BQC			ClearSpeed CX700		
Cores	W	GF/W	Cores	W	GF/W	Cores	W	GF/W	Cores	W	GF/W
512	225	~2.2	8	10	6	16	55	3.7	192	10	~10

Nominal approximate estimates

2nd Prototype

DSP Prototype

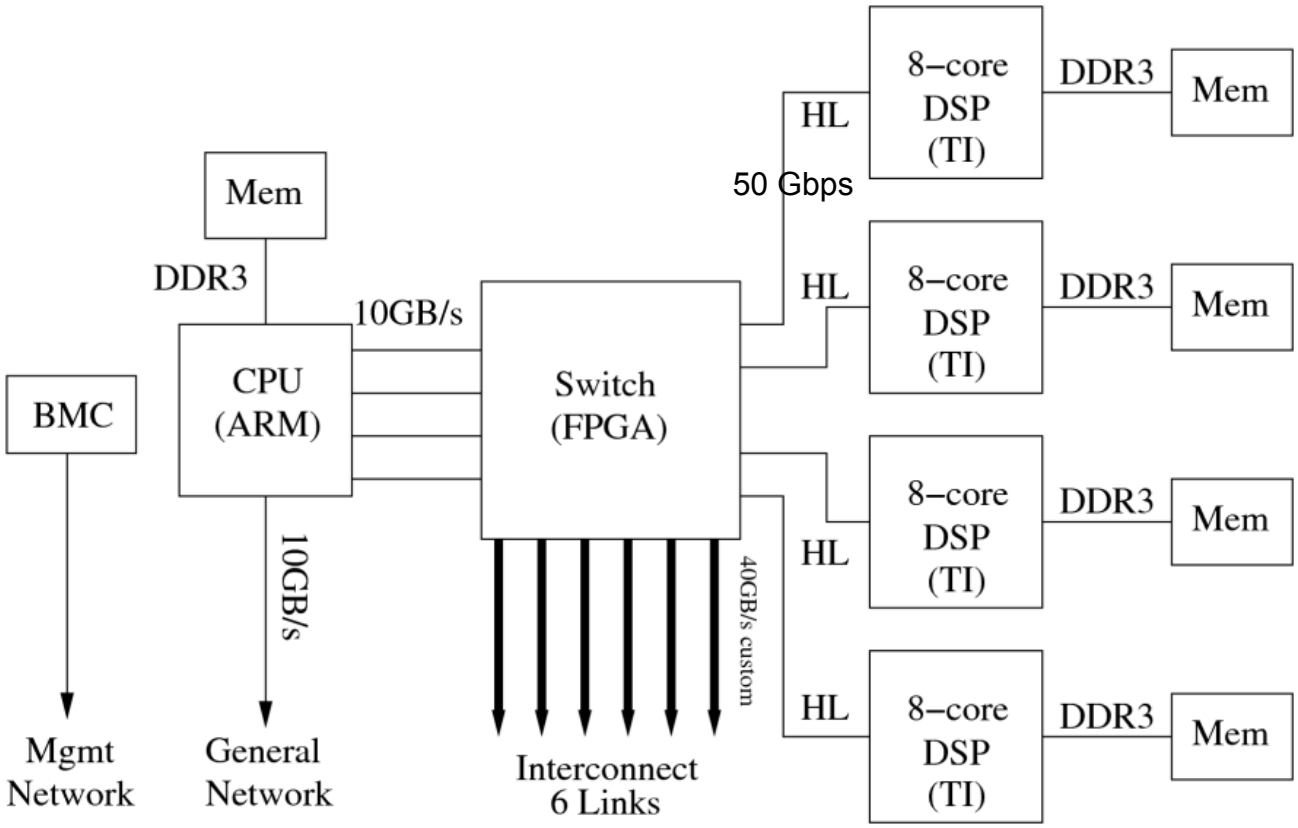

Instrumentation of the C6678 Module

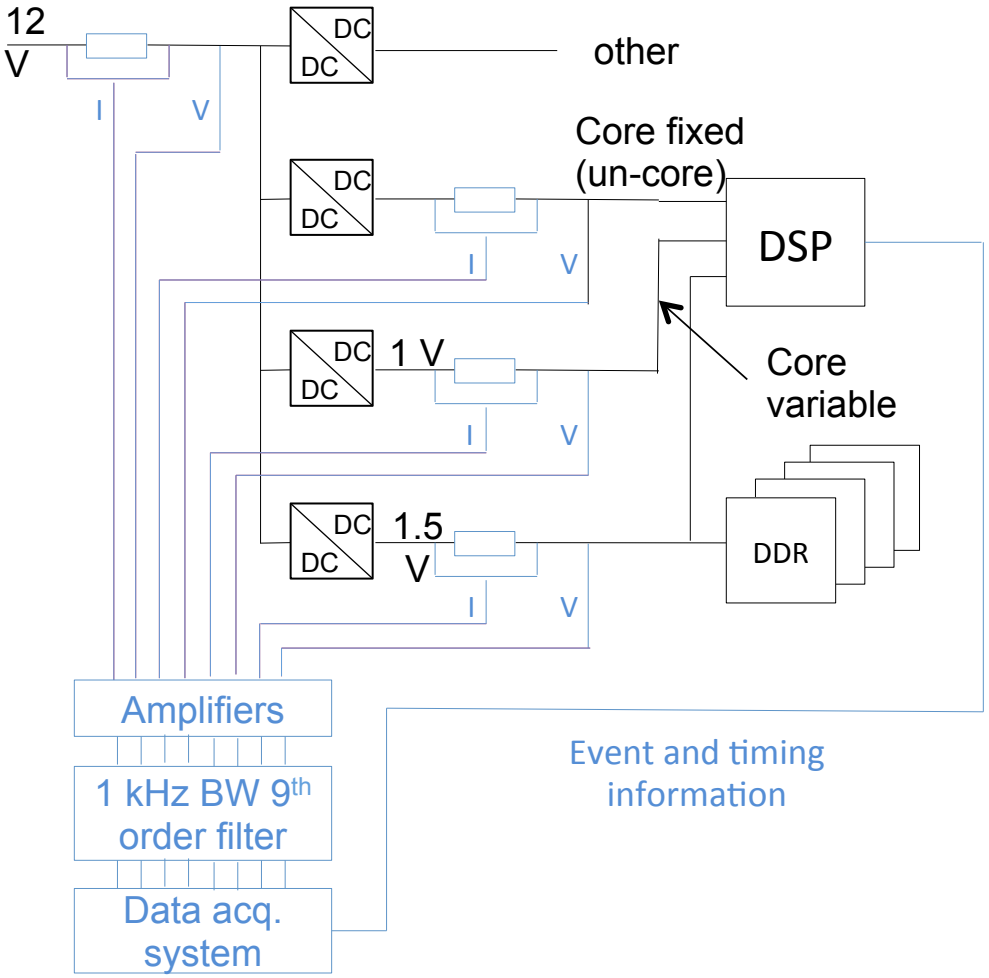

- Four differential channels for Current
- Four differential channels for Voltage
- Sampling rate 125 kHz, 125/8 kHz per channel
- Accuracy better than 1%

STREAM (Unoptimized)

Copy	BW (GB/s)	Core (W)	Memory (W)	Total (W)	Core (GB/J)	DSP (GB/J)
L1	124.2	6.81	0.42	14.20	18.24	17.19
L2	47.6	6.54	0.42	13.85	7.27	6.84
MCSM	38.7	6.65	0.42	13.99	5.81	5.47
DDR3-1333	3.0	5.59	1.95	14.47	0.54	0.40

Scale	BW (GB/s)	Core (W)	Memory (W)	Total (W)	Core (GB/J)	DSP (GB/J)
L1	123.7	7.17	0.42	14.69	17.23	16.28
L2	41.8	6.70	0.42	14.05	6.24	5.50
MCSM	21.9	6.44	0.42	13.71	3.40	3.19
DDR3-1333	2.9	5.60	1.91	14.44	0.52	0.87

L1: 98 % of peak
L2: 75 % of peak
DDR3: 90 % of peak best of served (for optimized HPL)

C6678 Linpack – 2.64 GF/J

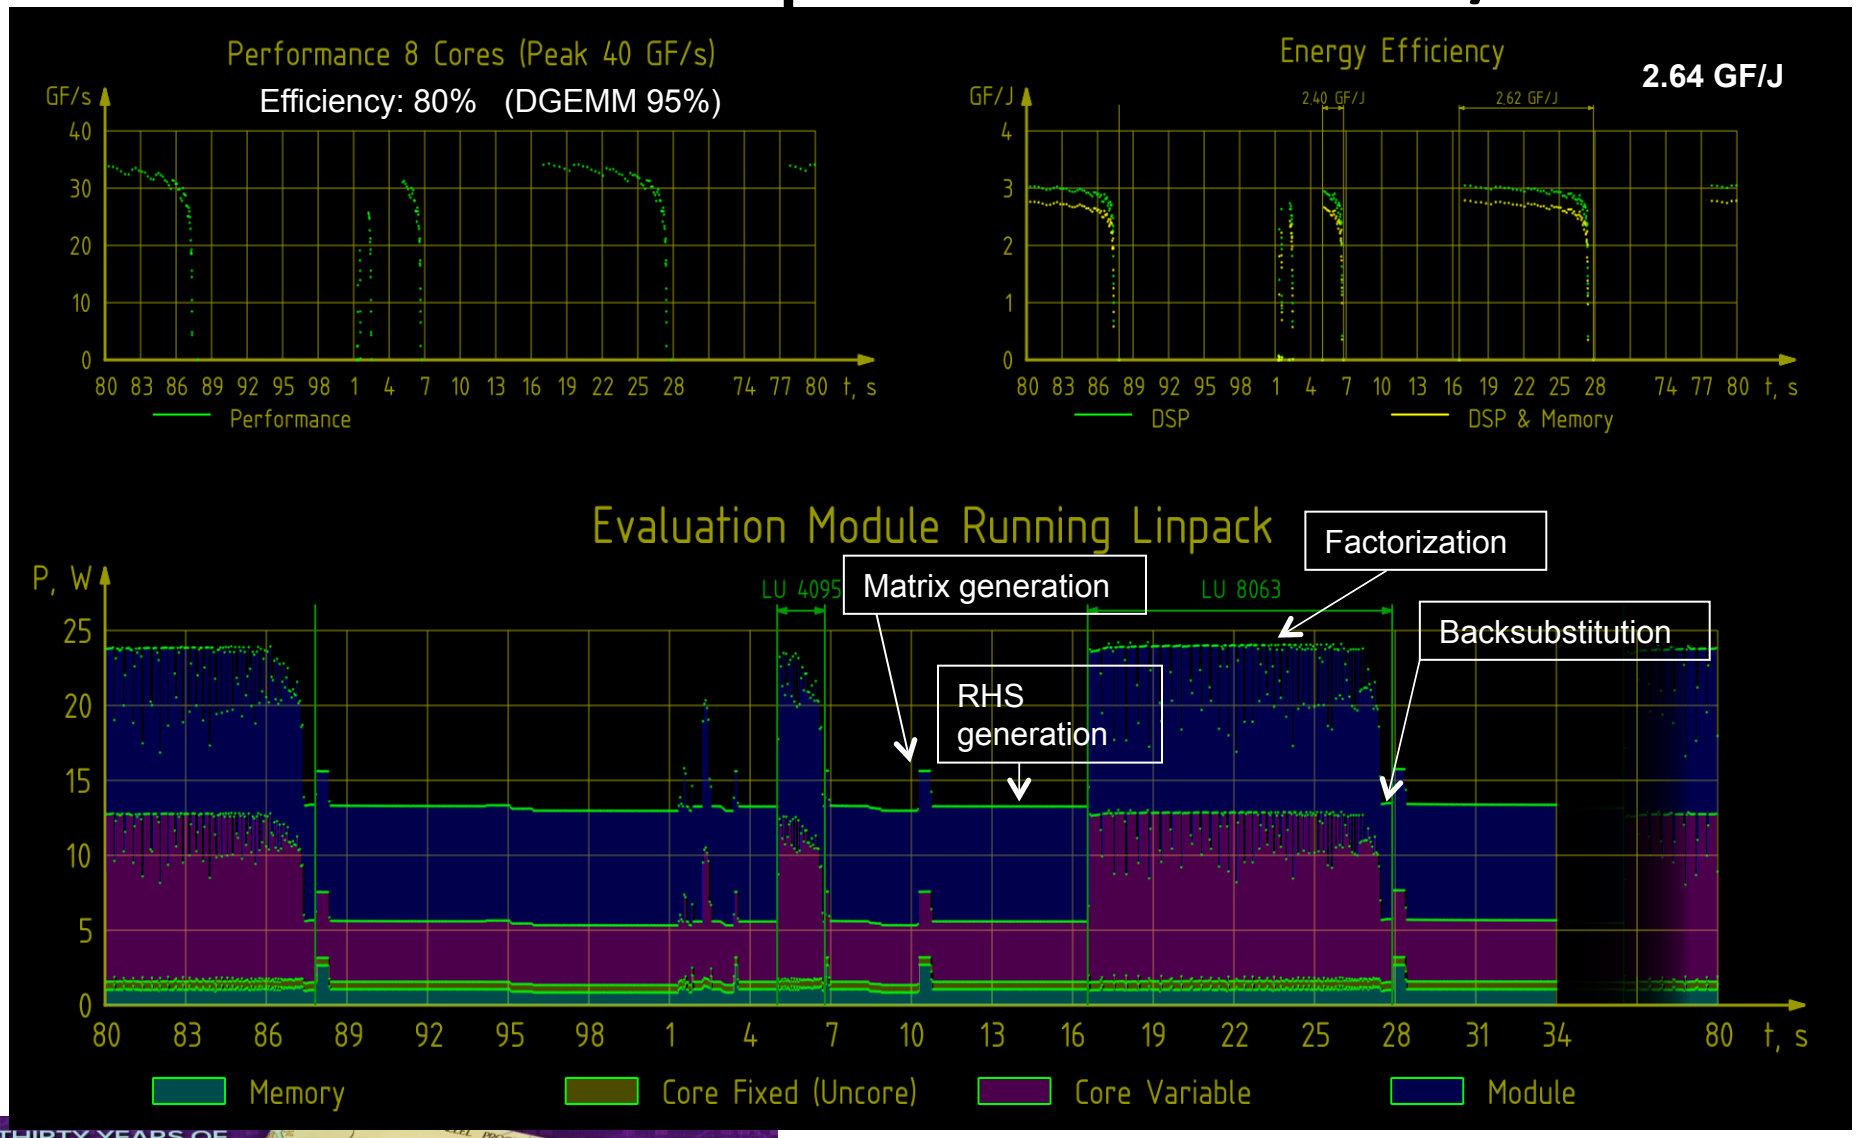

Size	Perf. (GF/s) Peak 40 GF/s	Efficiency (%)	Power (W) Cores+Memory	Energy Efficiency (F/J)
127	0.7	1.7	4.8	143
255	3.0	7.6	5.6	541
511	6.8	16.9	7.0	961
1023	13.0	32.4	8.4	1548
2047	19.2	48.1	9.5	2035
4095	25.9	64.8	10.7	2421
8063	30.9	77.2	11.7	2644

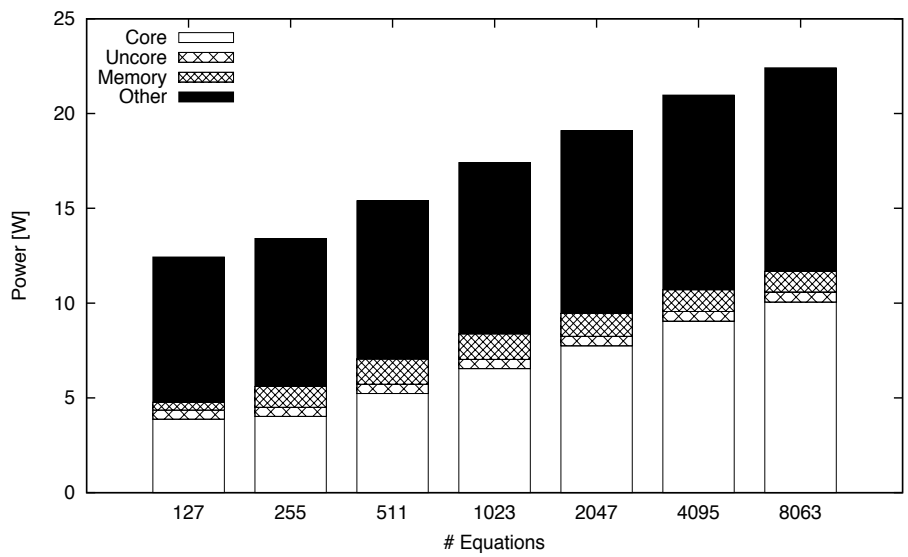

Results in Technology Perspective

CPU	"Current"			"Next Generation"		
	Feature size	GF/J	GF/mm^2	Feature size	GF/J	GF/mm^2
AMD Interlagos (16C, 2.7 GHz)	32	1.3	0.55			
AMD FirePro 7900	40	3.0	1.19			
AMD S9000				28	3.6	2.21
AMD Brazos	40					
AMD Llano	32					
AMD Trinity	32	1.0	0.40			
IBM Blue Gene/Q	45	3.7	0.57			
IBM Power7	45	1.4	0.48			
Intel Sandy Bridge (8C, 3.1 GHz)	32	1.3	0.46			
Intel Ivy Bridge (4C,3.5 GHz)				22	1.45	0.70
Nvidia Fermi	40	2.66	1.26			
Nvidia Kepler 20x				28	5.2	2.36
Nvidia Tegra2	40	~1	0.04			
Nvidia Tegra3	40					
TI TMS320C6678	40	4	~3			
TI 66AK2Hx				28		
Xilinx Vertex-6	40	5-10				
Xilinx Vertex-7				28	~13	0.236

Lessons Learned *so far…..*

- Efficiencies comparable to that for x86 CPUs is achievable for simple HPC benchmarks such as STREAM and HPL
- Achieving comparable levels of efficiency is hard; most of existing software eco-system proprietary
- Existing SoC has a fair number of functional units of limited or no interest for HPC (and that do consume significant power)
- DSPs, as judged by the TI DSP, viable starting point for energy efficient HPC node design

- HP Moonshot
  - TI DSP cartridge
  - Calxeda ARM cartridge
  - Intel ATOM cartridge
  - .....

HP Moonshot

If the public cloud were a country, it would rank fifth in electricity consumption.

Reducing that number by even 50% would save the equivalent of the electricity consumption of the United Kingdom.

Next Prototype?

Prototype III – Enhanced Mobile Video CPU

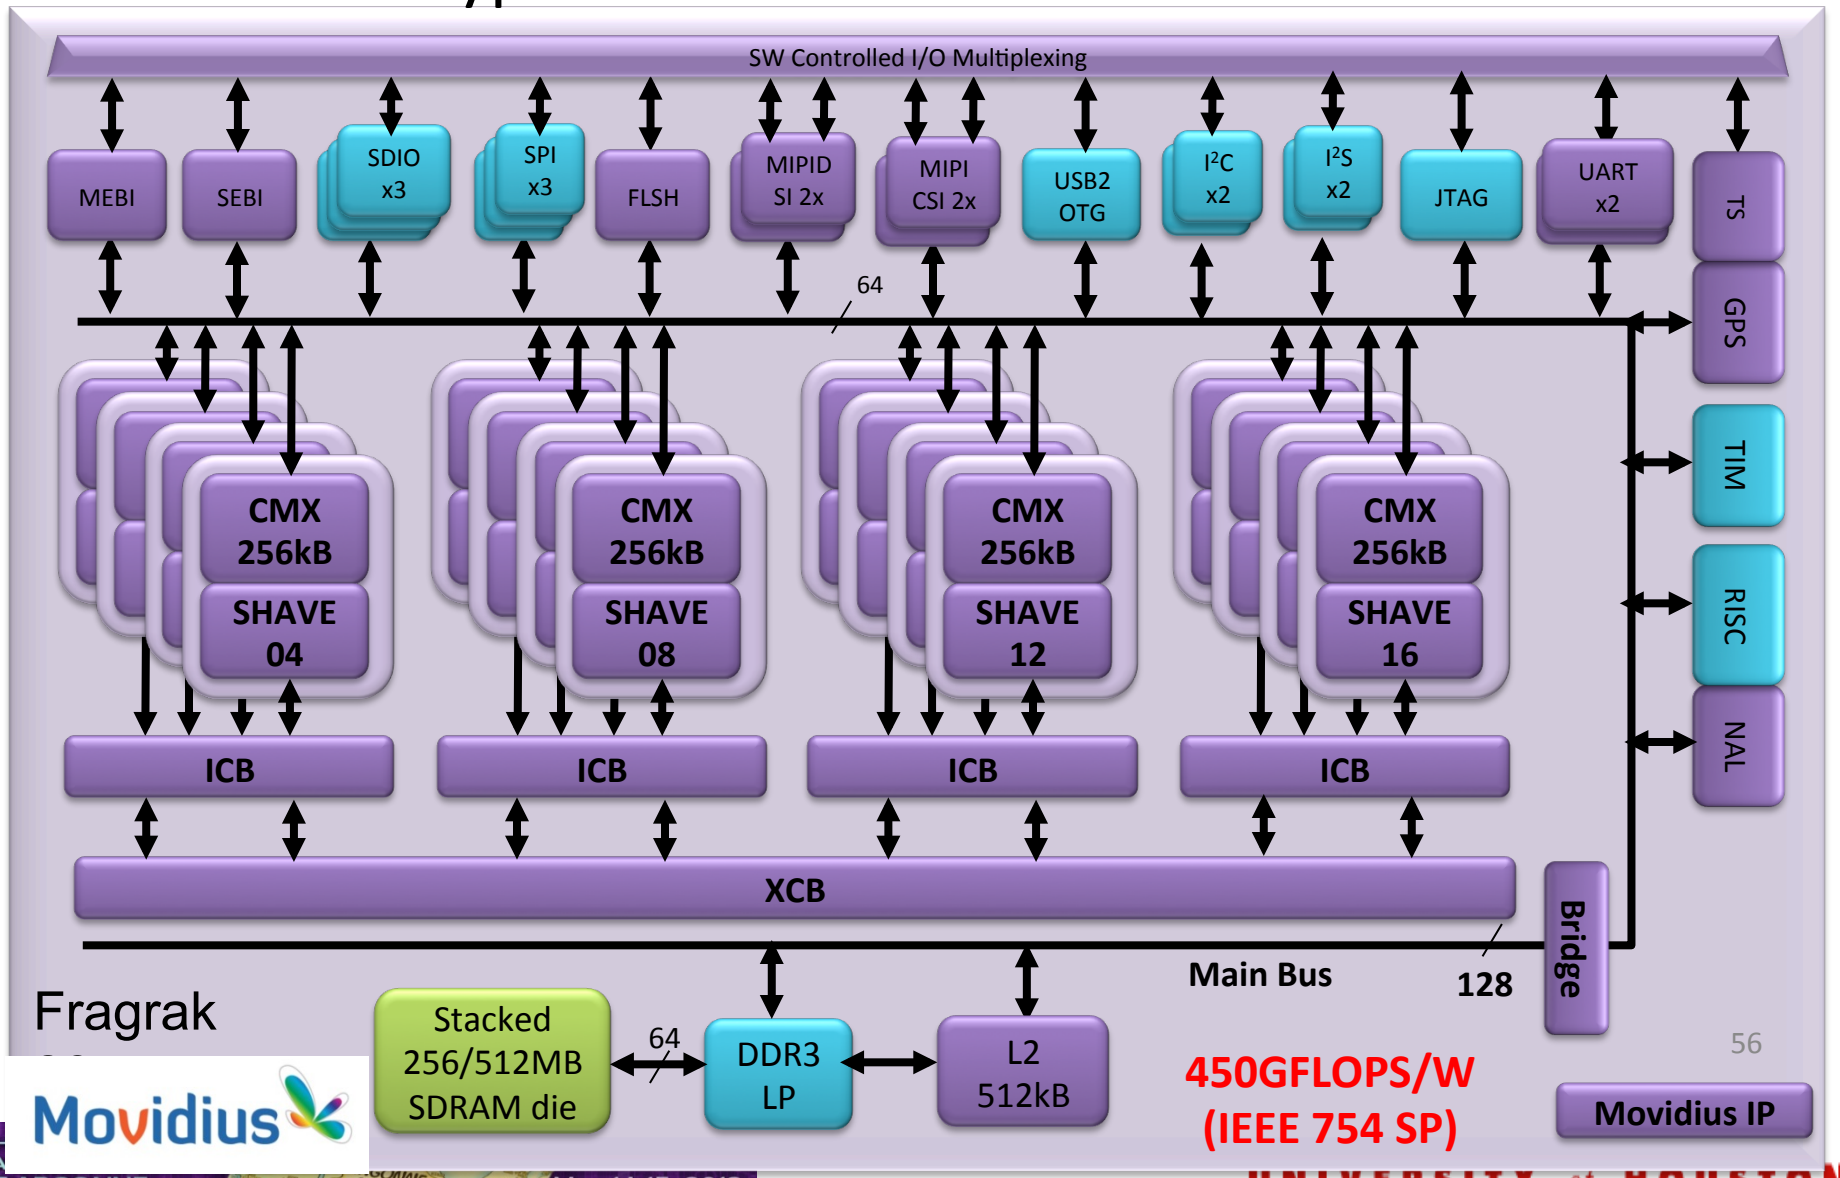

Fragrak

Movidius
AT ARGONNE: A Symposium
May 14-15, 2013
Argonne National Laboratory
UNIVERSITY of HOUSTON

SHAVE in Comparison

	SHAVE Fragrak	BlueGene/Q	NVIDIA Kepler
Clock Frequency	800 MHz	1600 MHz	1006 MHz
Cores/Threads	16	16	1536
FP Performance	51.2 GF/s	204.8 GF/s	~ 1000 GF/s
Power	0.35 W	55 W	195 W
Memory	1 GB	16 GB	2 GB
Memory Bandwidth	12.8 GB/s	42.7 GB/s	192.2 GB/s
Network Bandwidth	3 GB/s	22 GB/s	12.8 GB/s
Energy Efficiency	146 GFLOP/J	3.7 GFLOP/J	5.1 GFLOP/J
FLOP/Memory Cap	51 FLOP/B	12 FLOP/B	466 FLOP/B
FLOP/Memory BW	4 FLOP/B	2.5 FLOP/B	5.2 FLOP/B
FLOP/Network BW	17 FLOP/B	9.3 FLOP/B	78 FLOP/B

Summary

- Trickle-up effect (Justin Rattner) could continue to make HPC exiting/dynamic

- Compliments to Argonne for its influence and contributions to the HPC and other communities

Thanks to many for many interesting formative discussions ….

Thank
you!
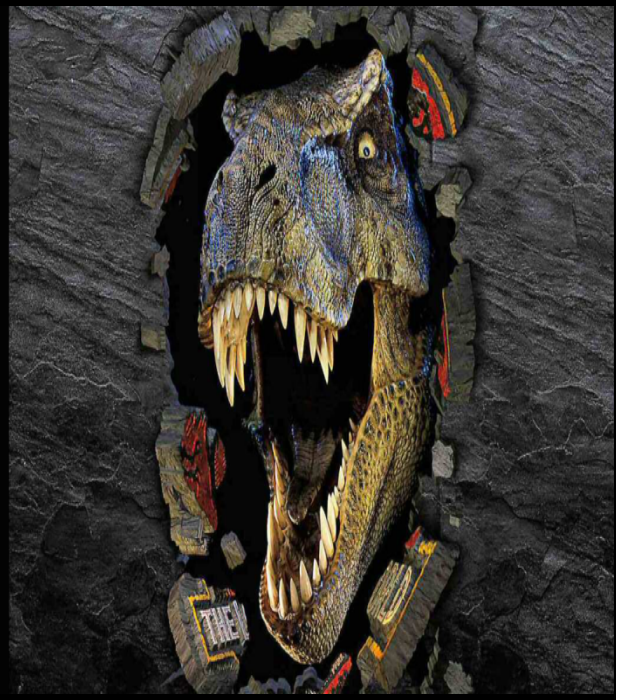
Any Questions?