Heterogeneous Processors: The Cell Broadband Engine (plus processor history & outlook)

IEEE SSC/CAS Joint Chapter, Jan. 18, 2011

H. Peter Hofstee

IBM Austin Research Laboratory

IBM Systems
Simplify your IT.

CMOS Microprocessor Trends, The First ~25 Years (Good old days)

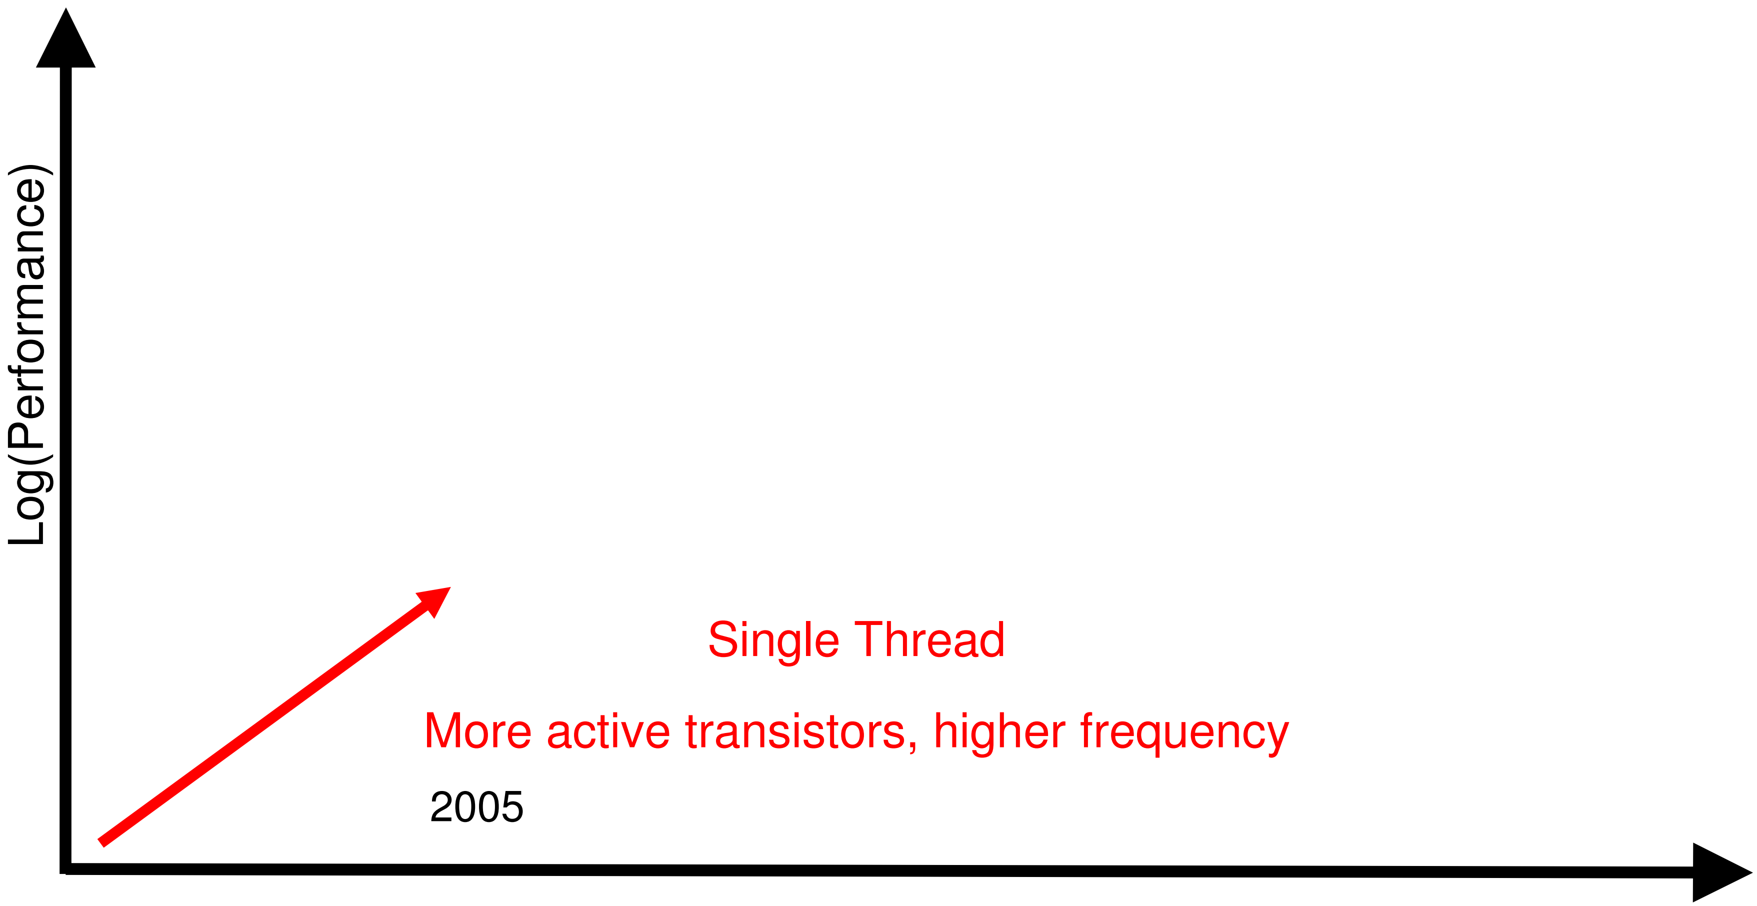

SPECINT

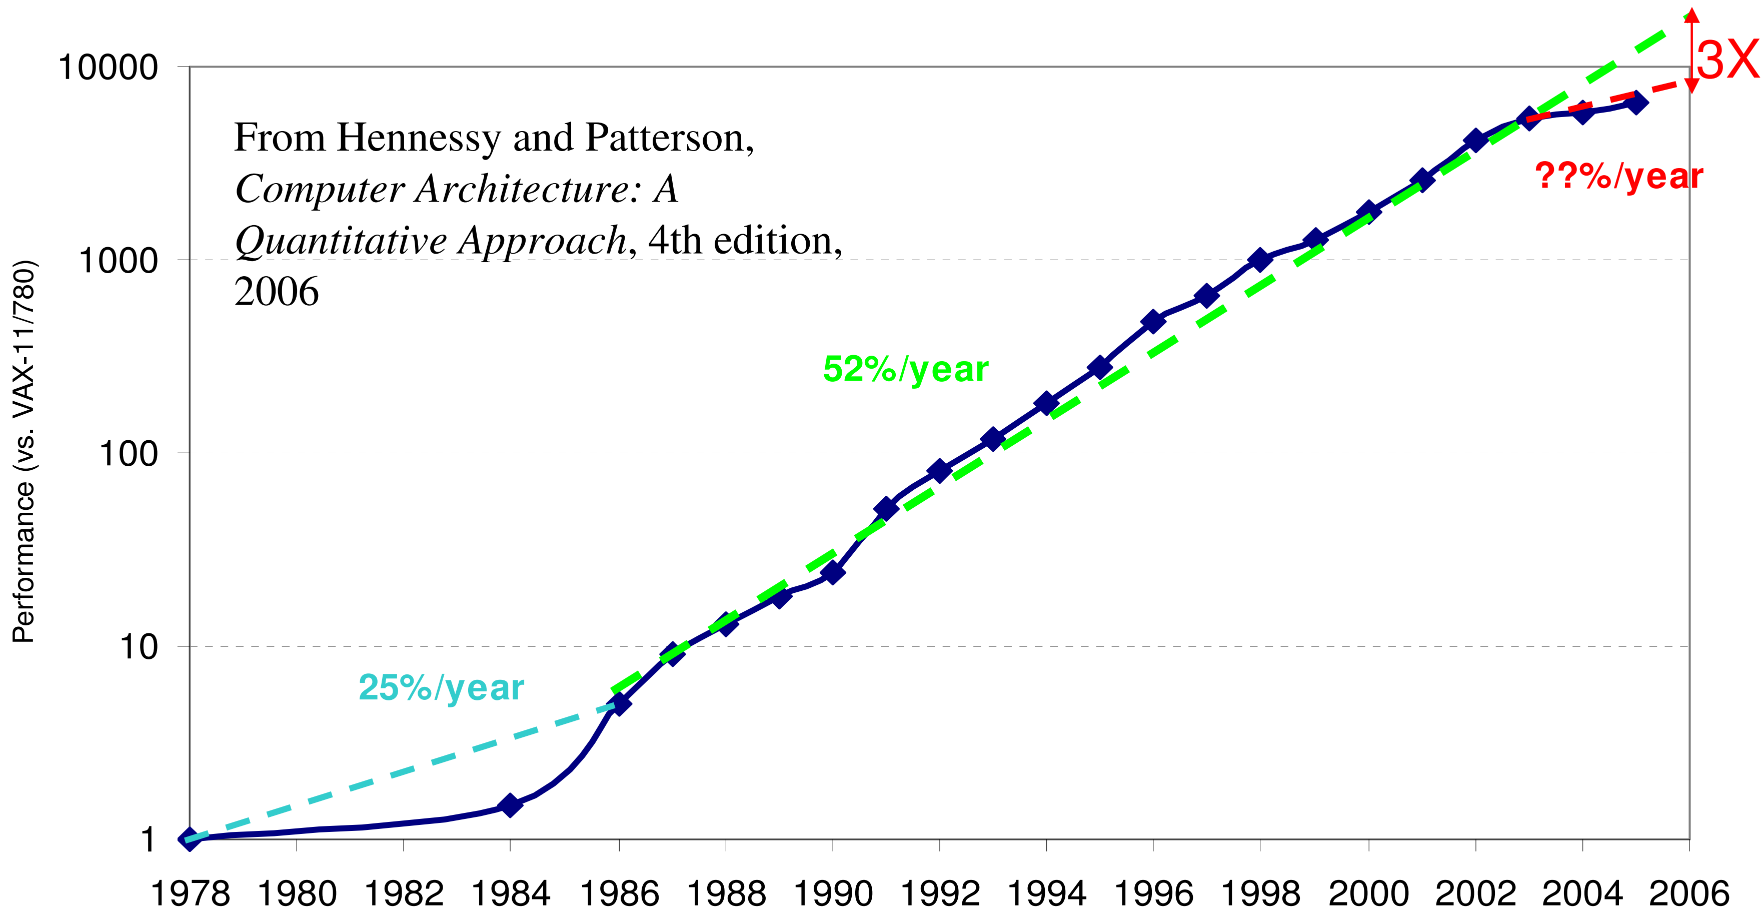

VAX : 25%/year 1978 to 1986
RISC + x86: 52%/year 1986 to 2002

1971 – 2000, A Comparison

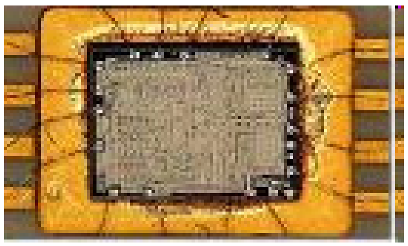

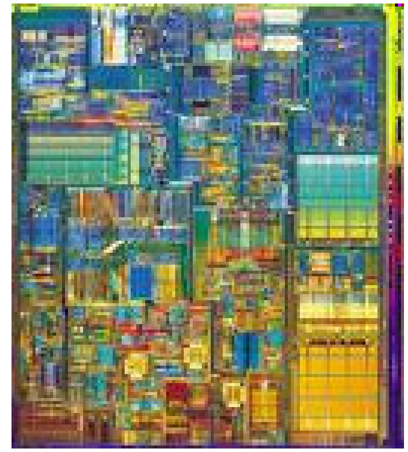

	1971(Intel 4004)	2000 (Intel Pentium III Xeon)	
Technology	10 micron (PMOS)	180 nm (CMOS)	
Voltage	15V	1.7V	(0.27V)
#Transistors	2,312	28M	(69M)
Frequency	740KHz	600MHz – 1GHz	(41MHz)
Cycles per Inst.	8	~1	
Chip size	11mm2	106mm2	
Power	0.45W	20.4 W(0.6GHz@1.7V)	
Power density	0.04W/mm2	0.18W/mm	
Inst/(Hz * #tr)	5.4 E-5	3.6 E-8	

~2004: CMOS Devices hit a scaling wall

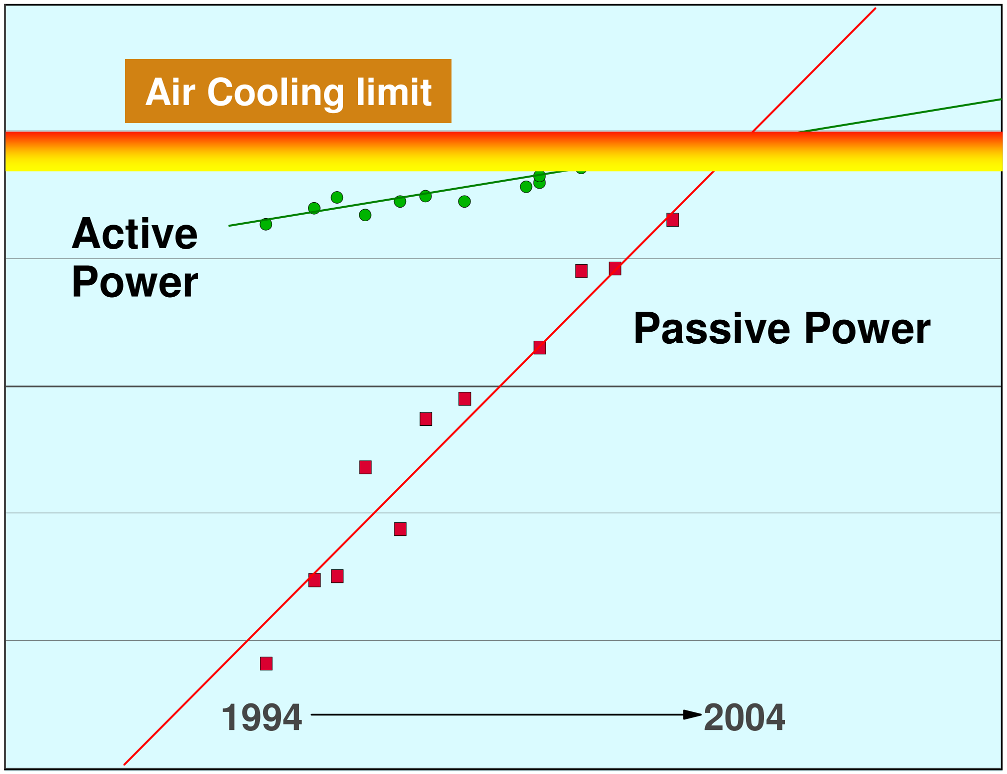

Isaac e.a. IBM

Microprocessor Trends

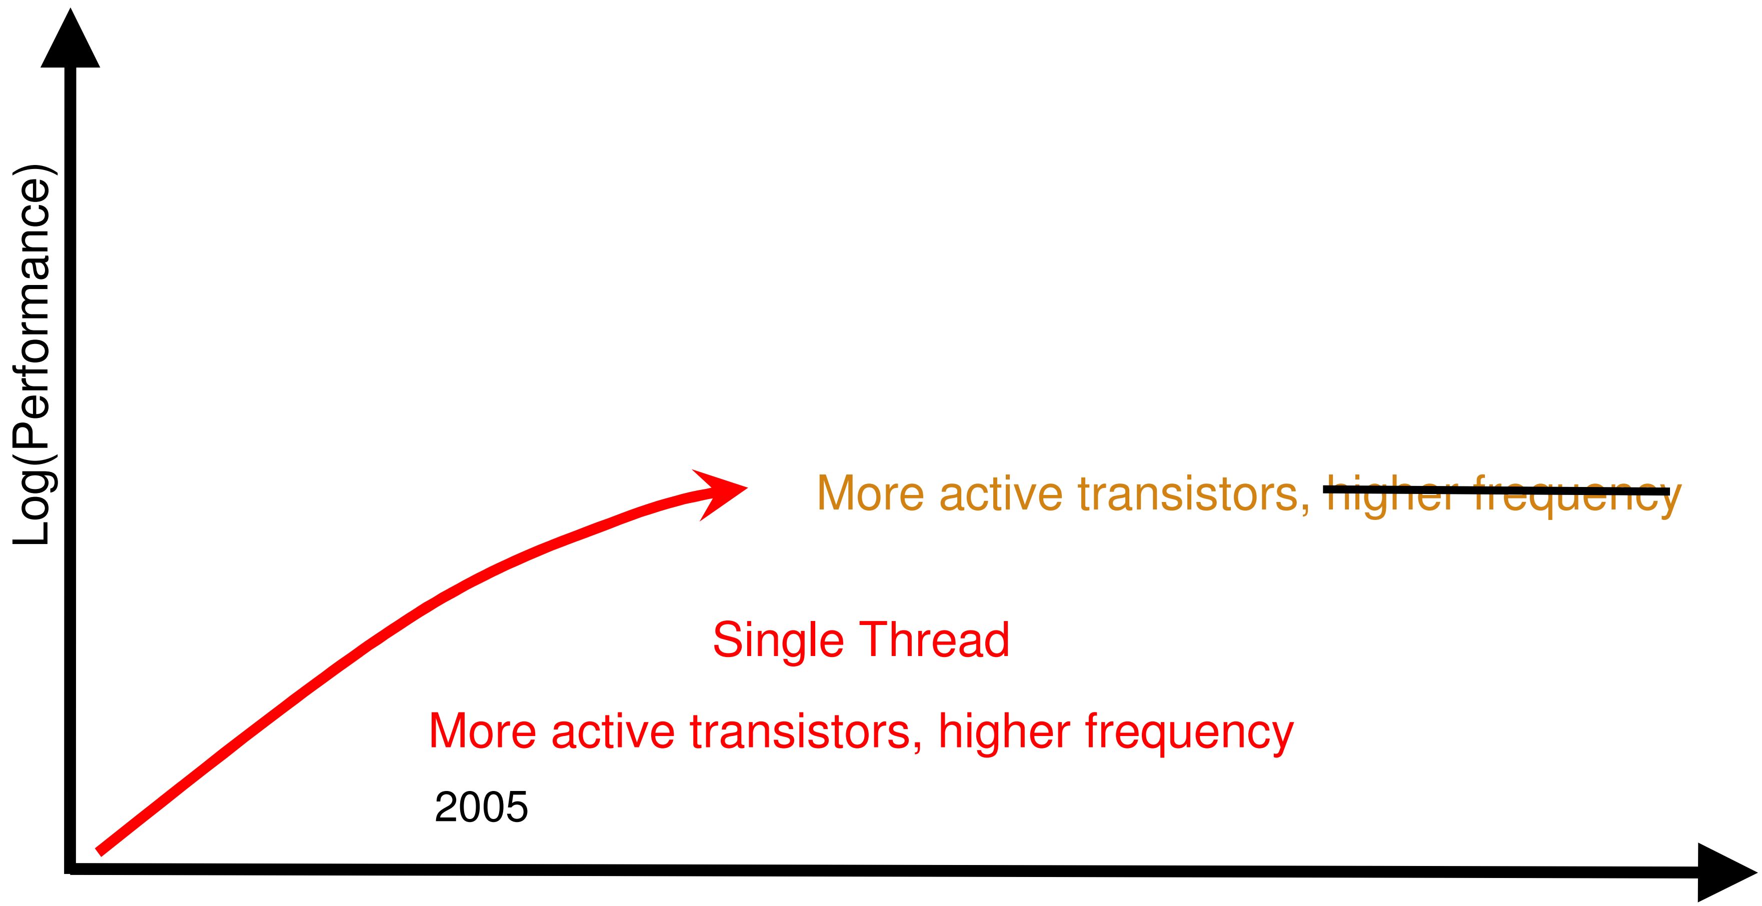

Microprocessor Trends

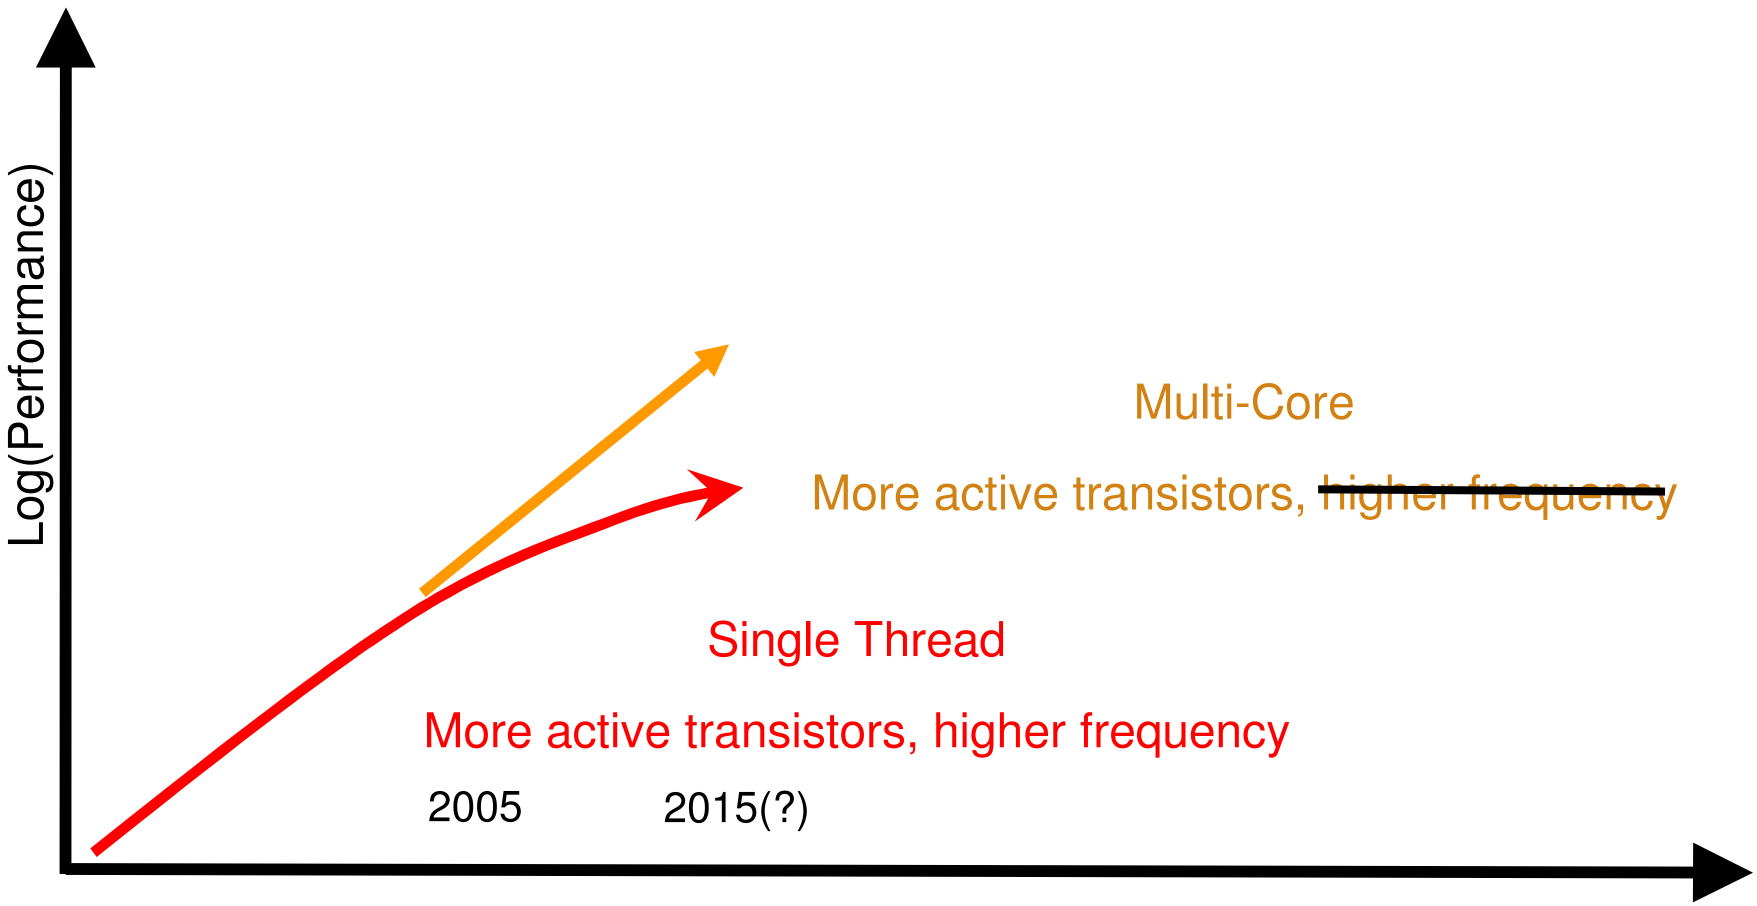

Multicore Power Server Processors

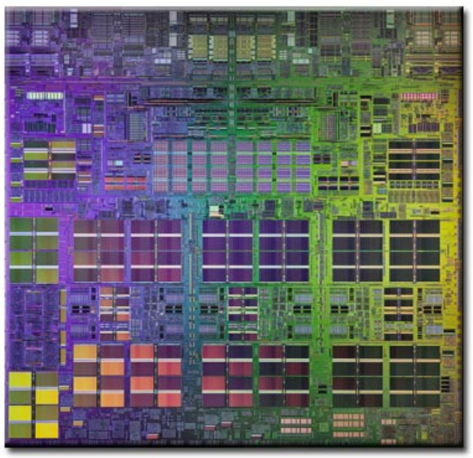

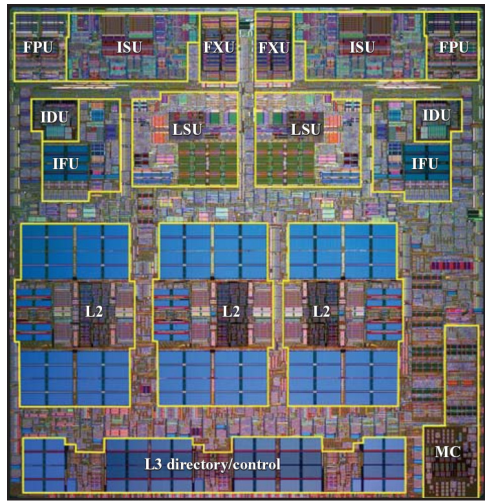

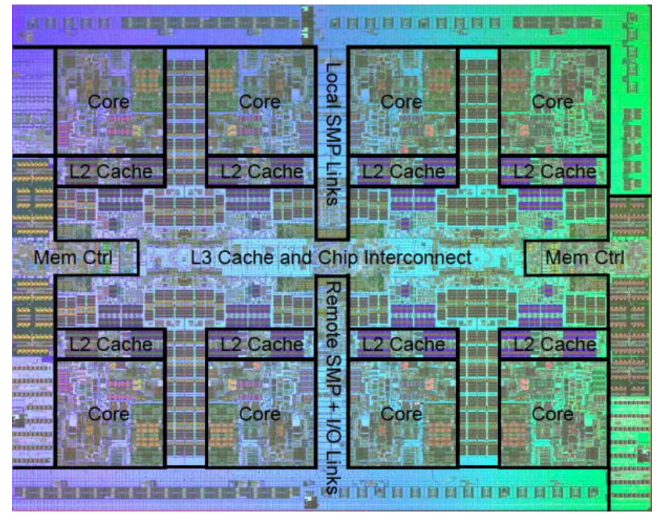

Power 4
2001
Introduces Dual core

Power 5
2004
Dual Core – 4 threads

Power 7
2009
8 cores – 32 threads

POWER7 Processor Chip : Threads & SIMD

- **Physical Design:**
  - 567mm² Technology: 45nm lithography, Cu, SOI, eDRAM
  - 1.2B transistors
    - Equivalent function of 2.7B
    - eDRAM efficiency
- **Features:**
  - Eight processor cores
    - 12 execution units per core
    - 4 Way SMT per core
    - 32 Threads per chip
    - 256KB L2 per core
  - 32MB on chip eDRAM shared L3
  - Dual DDR3 Memory Controllers
    - 100GB/s Memory bandwidth per chip sustained
  - Scalability up to 32 Sockets
    - 360GB/s SMP bandwidth/chip
    - 20,000 coherent operations in flight
- Two I/O Mezzanine (GX++) System Buses
  - Binary Compatibility with POWER6

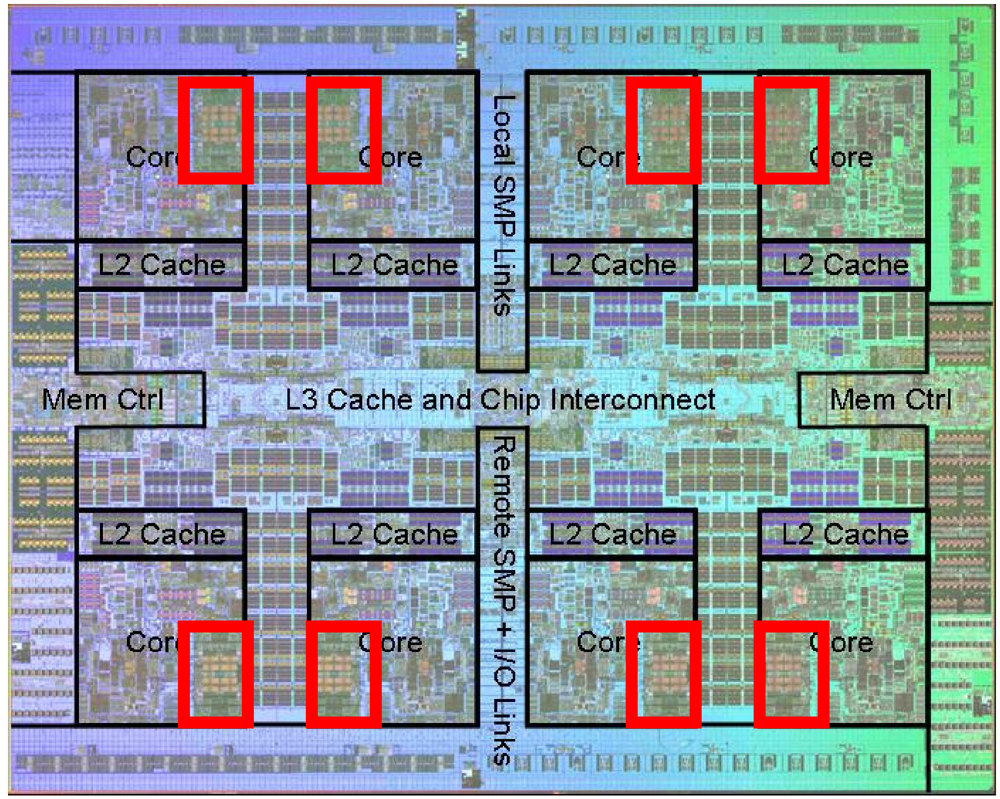

VSX Floating-Point / Media Units

P7 IH System Hardware – Node Front View

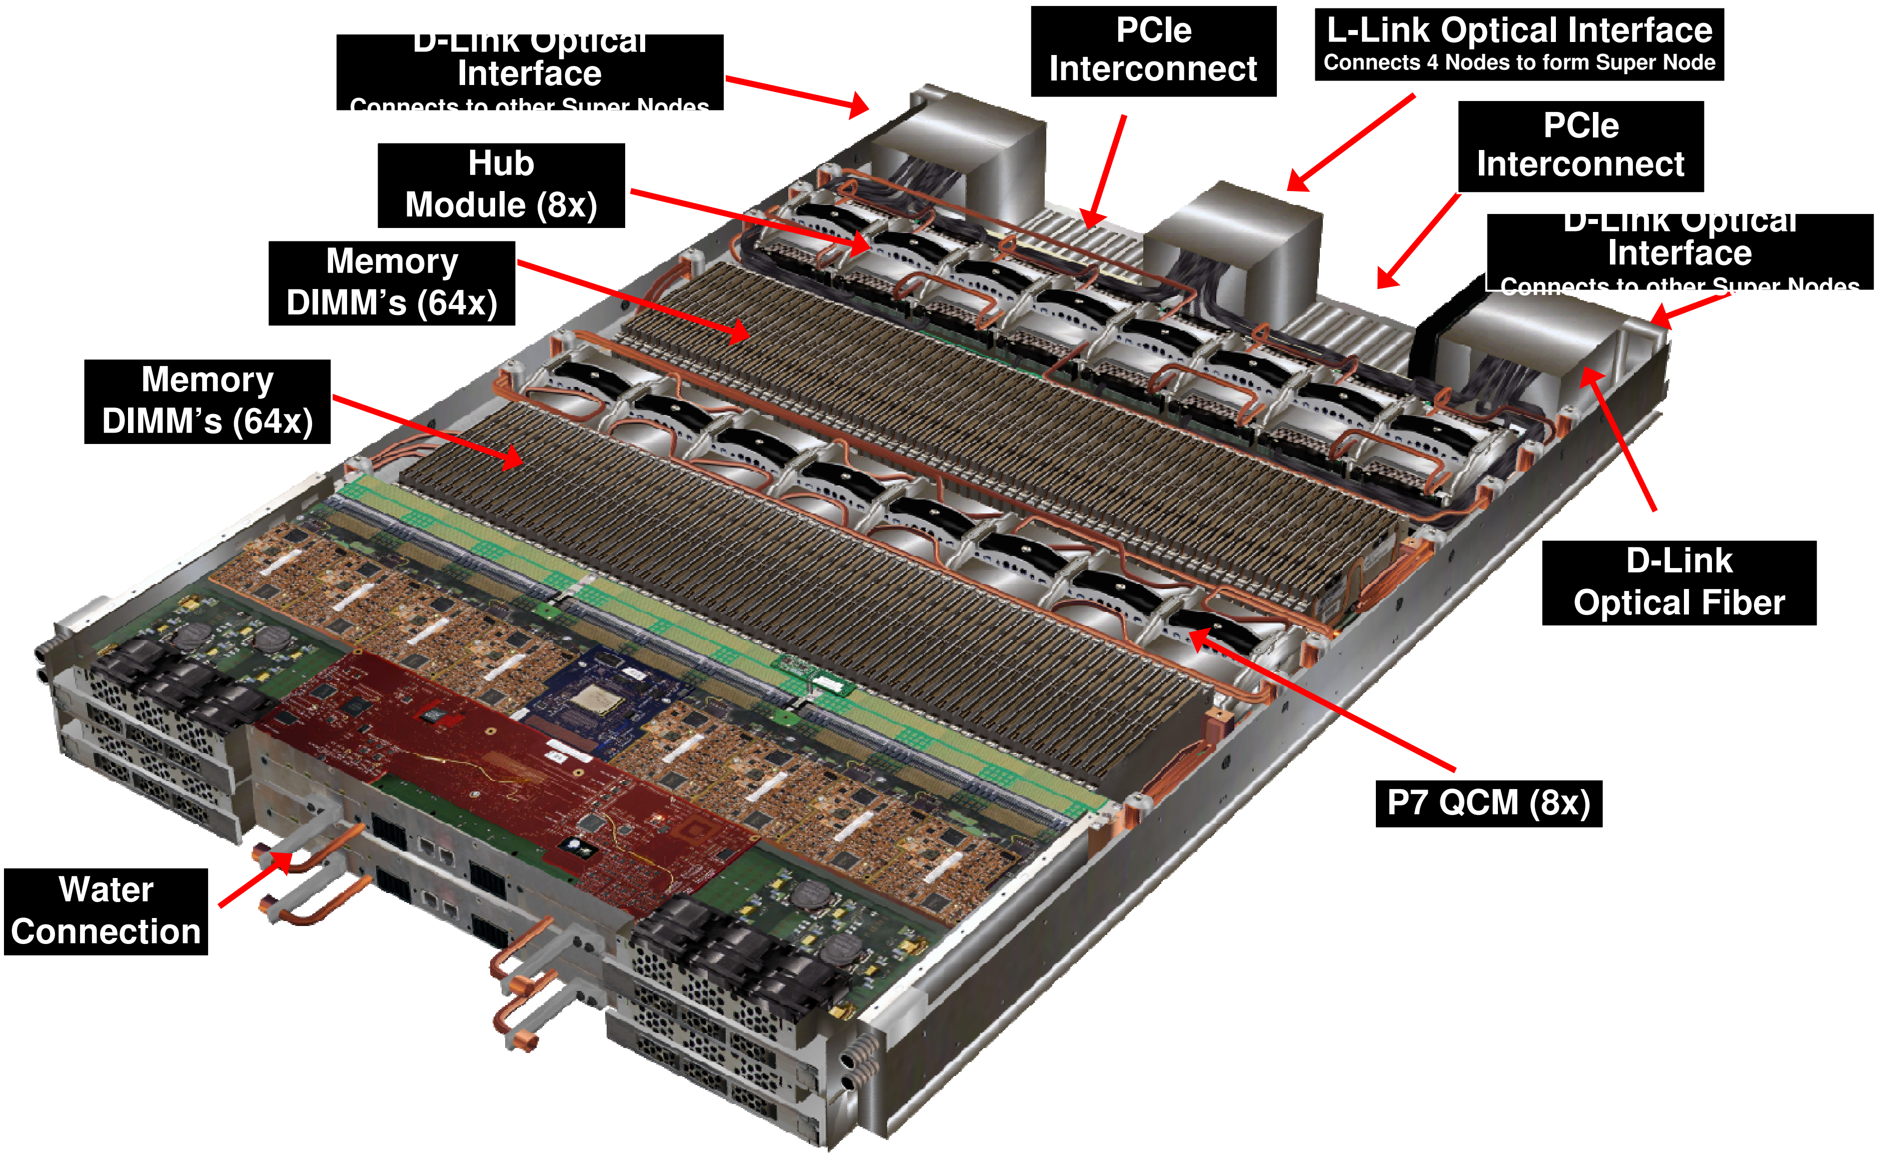

POWER7 IH Memory DIMM

Up to 128 DIMMs per Node

8 GB and 16 GB DIMMs

Dual Super Novas per DIMM

Water Cooled

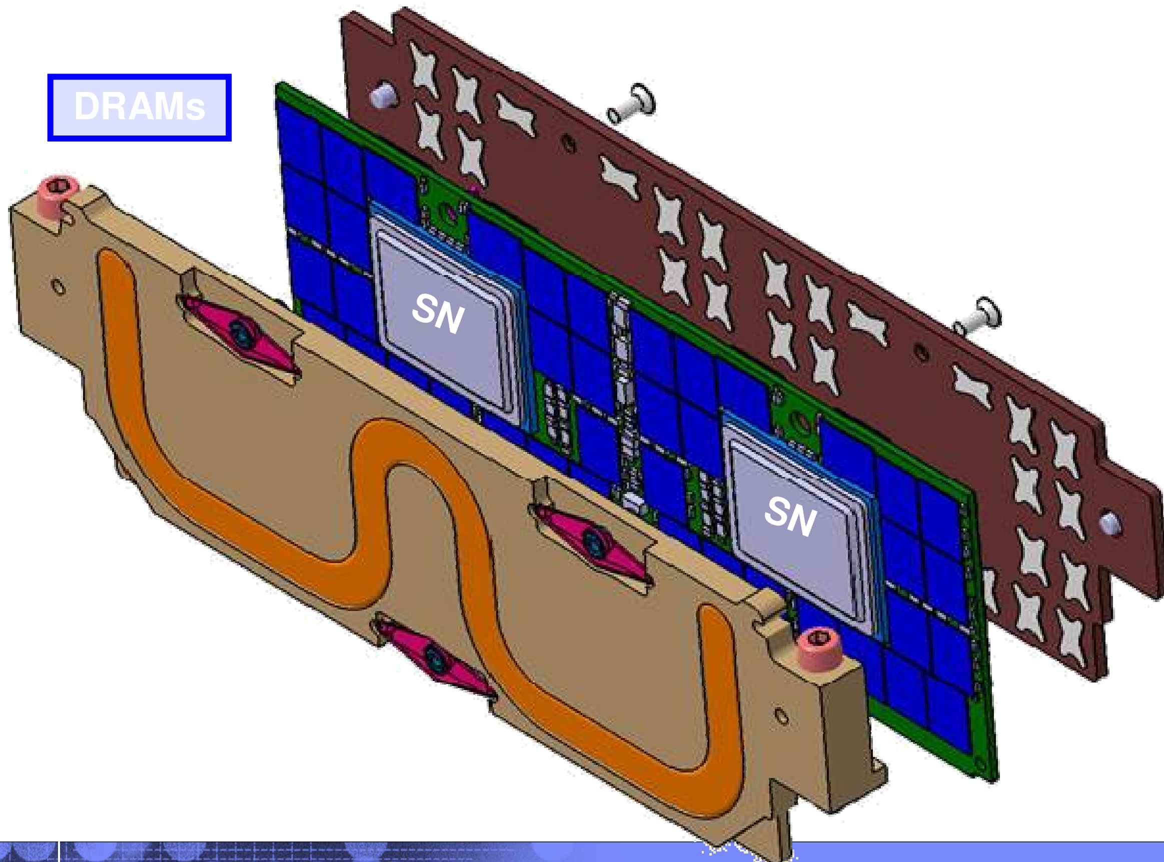

PERCS Node @ SC09

Why are (shared memory) CMPs dominant?

- A new system delivers nearly twice the throughput performance of the previous one without application-level changes.

- Applications do not degrade in performance when ported (to a next-generation processor).
  - This is an important factor in markets where it is not possible to rewrite all applications for a new system, a common case.

- Applications benefit from more memory capacity and more memory bandwidth when ported.
  - .. even if they do not (optimally) use all the available cores.

- Even when a single application must be accelerated, large portions of code can be reused.

- Design cost is reduced, at least relative to the scenario where all available transistors are used to build a single processor.

PDSOI optimization results

Constant performance improvement, 20% per gen.

Optimizing for maximum performance for each core

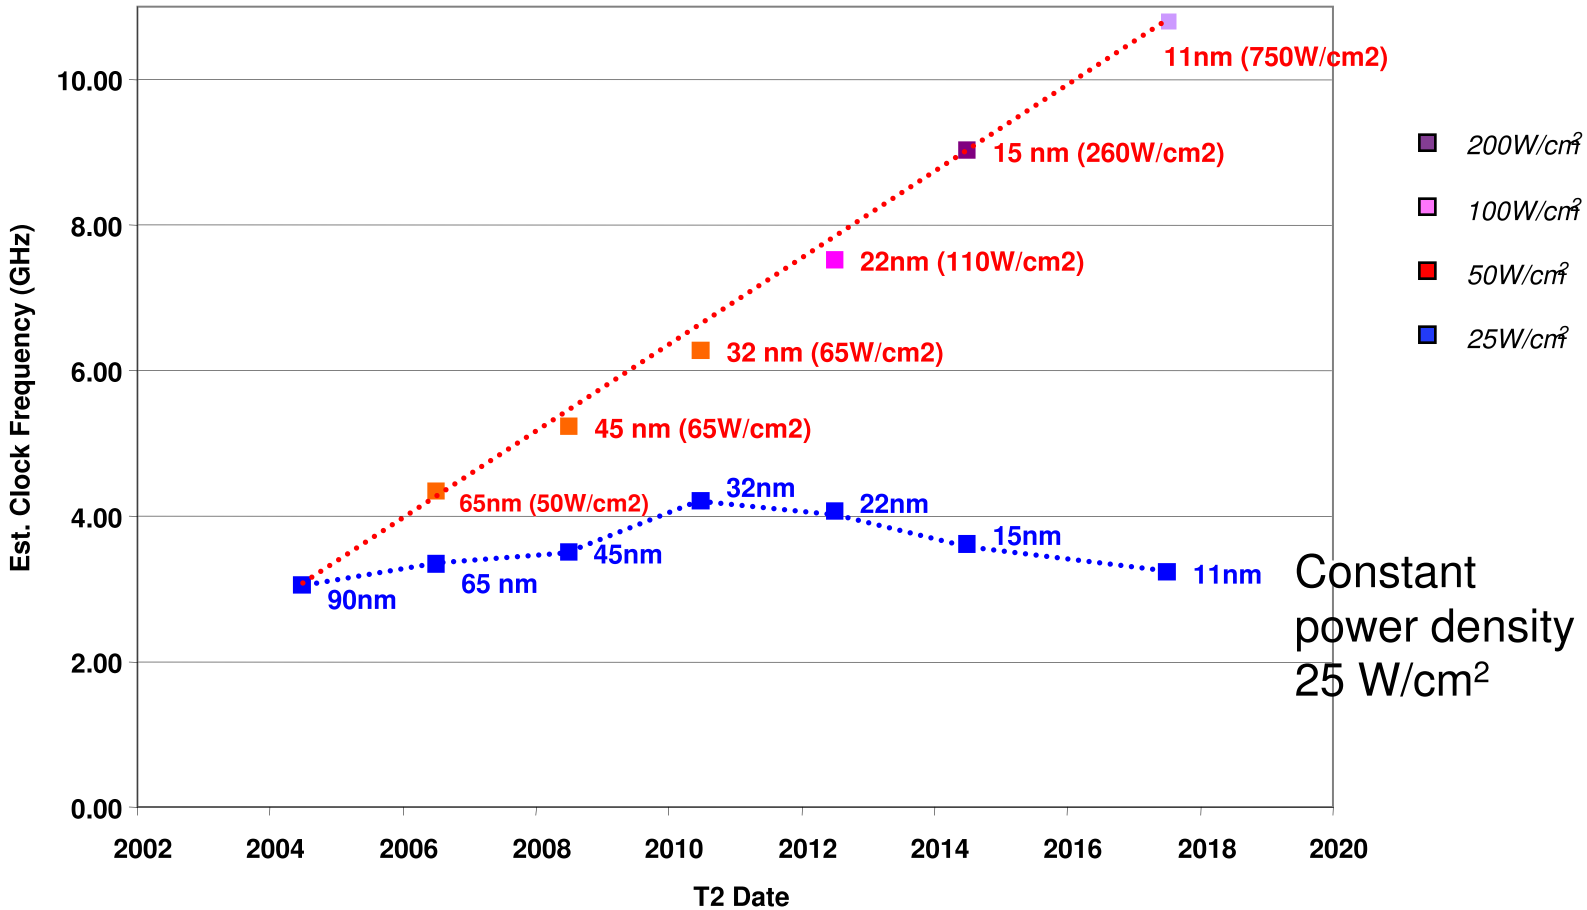

Constant power density 25 W/cm^2

D. Frank, C. Tyberg

Microprocessor Trends

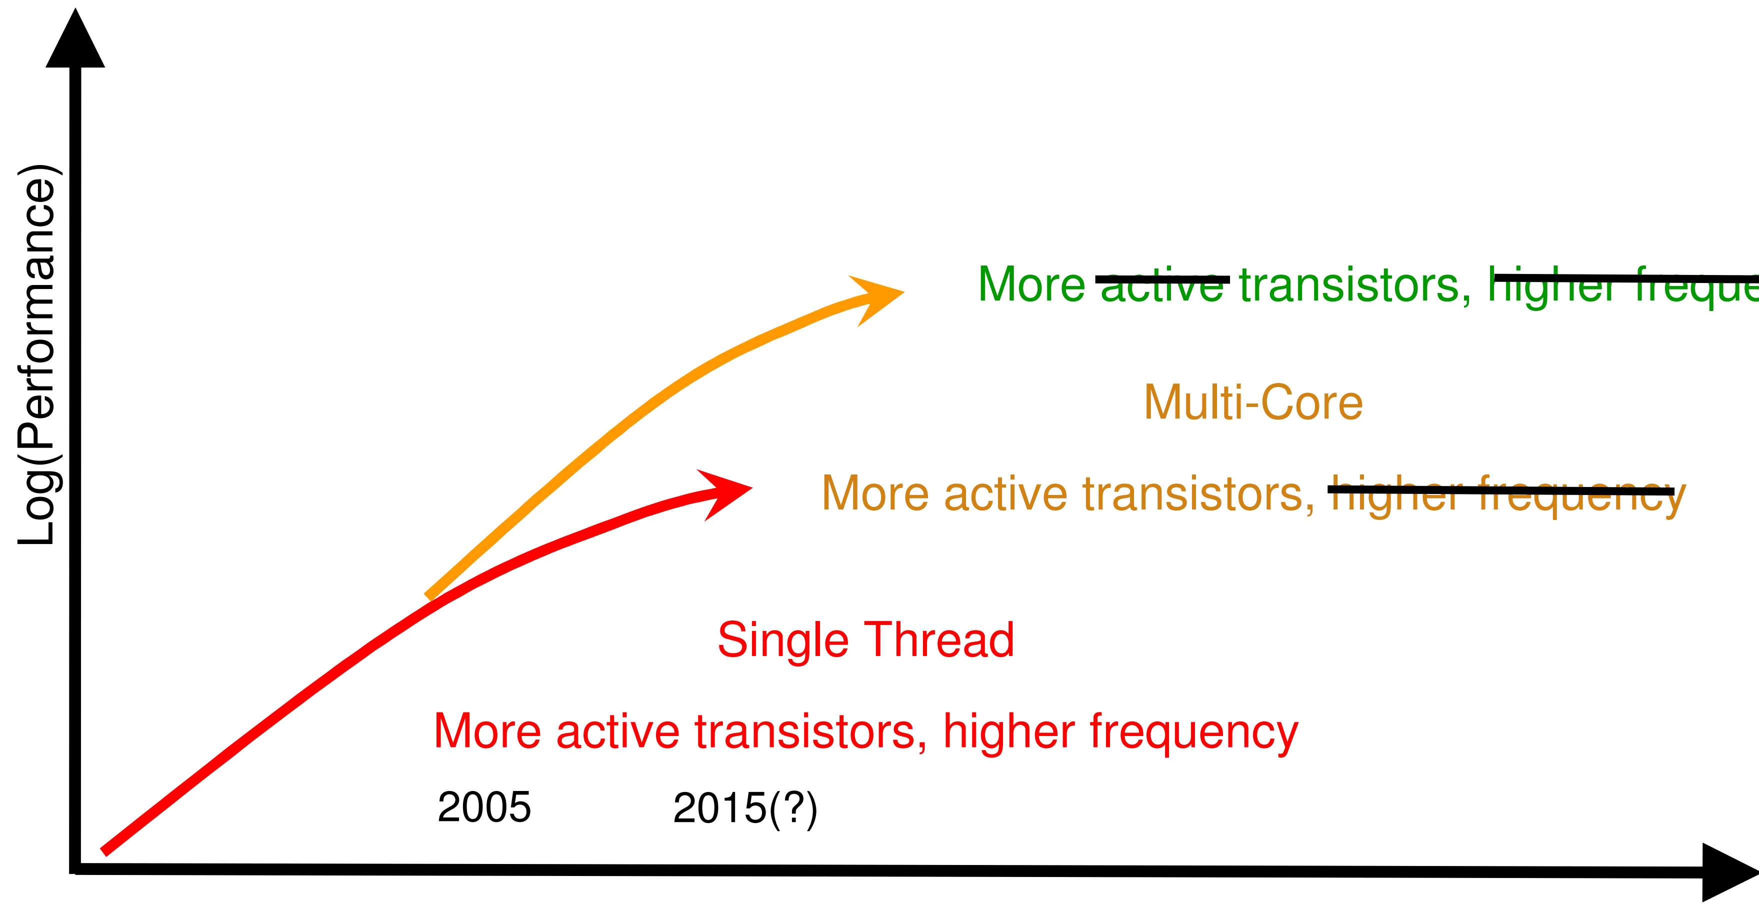

What are the options? What are we doing about it?

- Just live with it: accept slowdown
  - Use added transistors for bigger caches

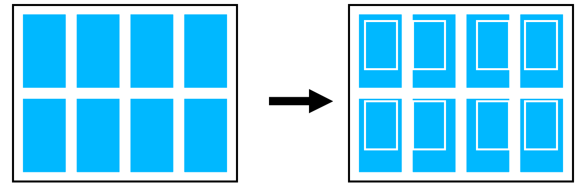

- Lighter weight threads only
  - More parallel=more efficient

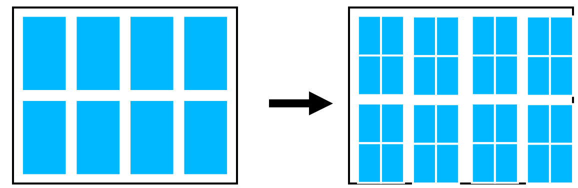

- Mix strong and light-weight cores
  - More parallel = more efficient

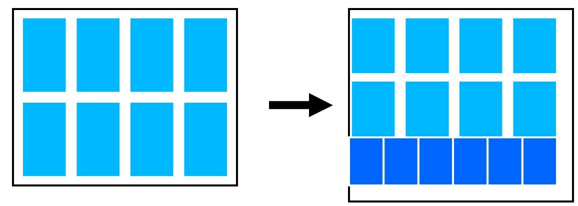

- Add accelerators
  - More specialized=more efficient

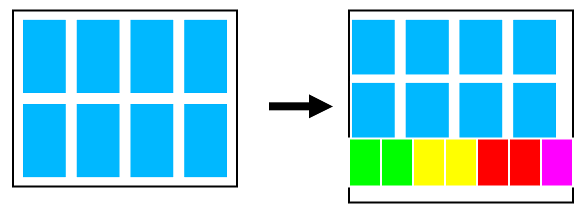

Microprocessor Trends

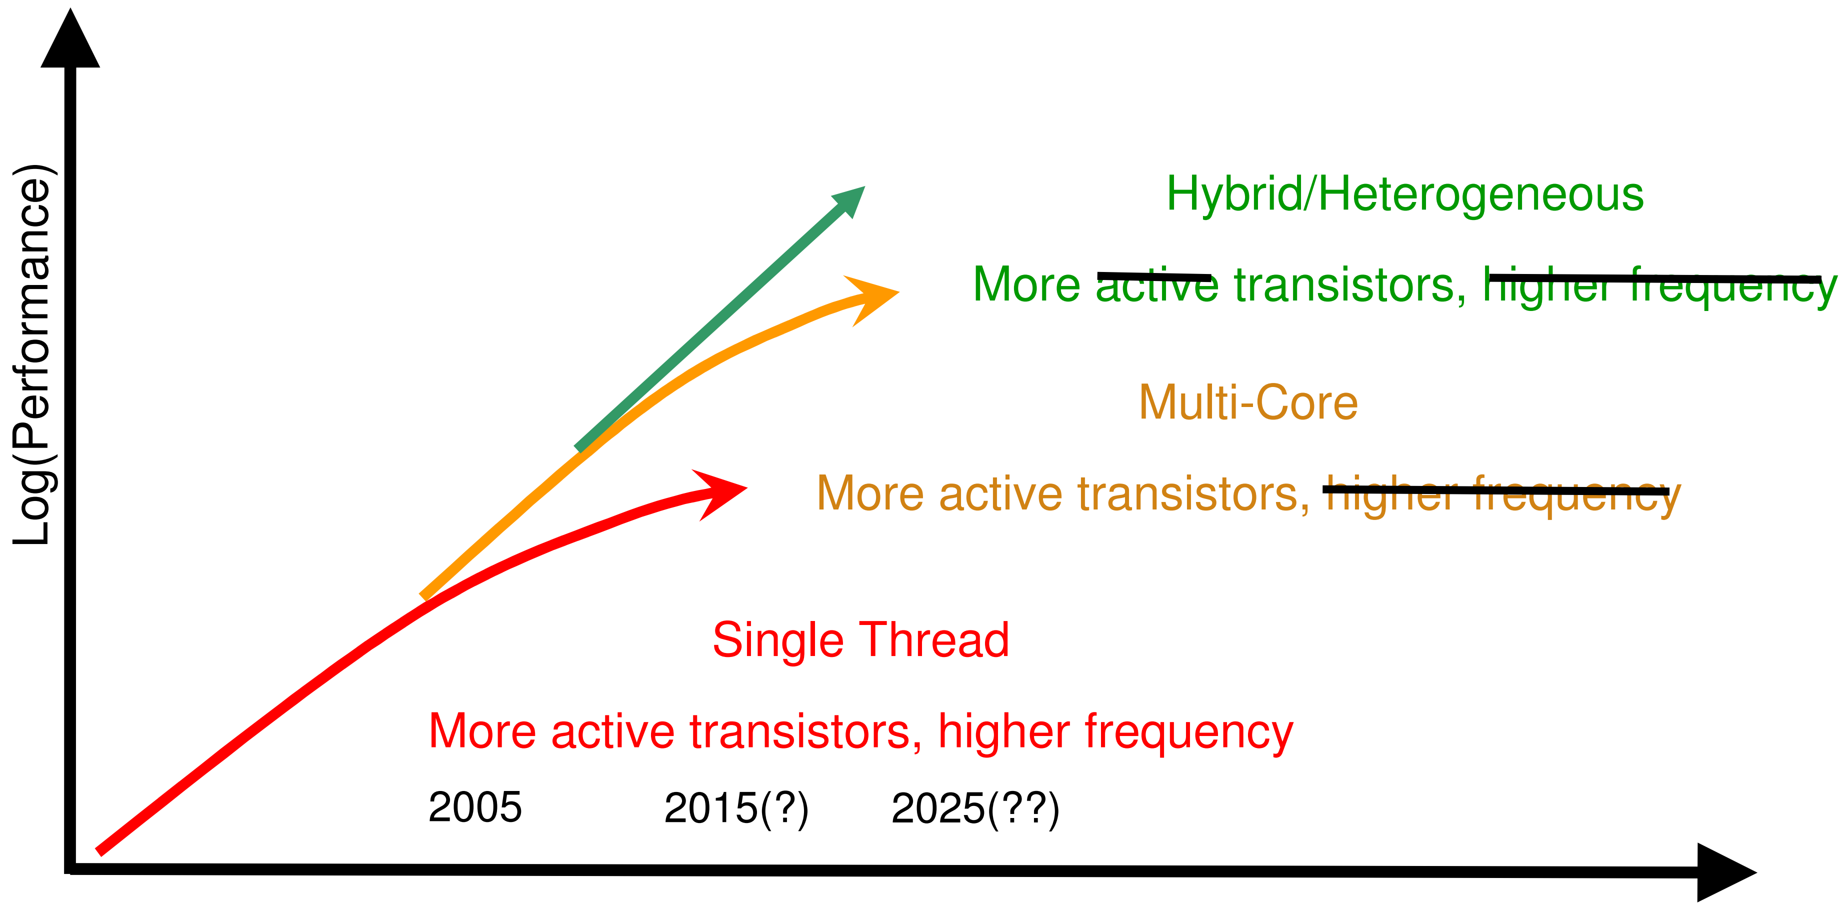

IBM Systems
Simplify your IT.

Heterogeneous Power Architecture Processors

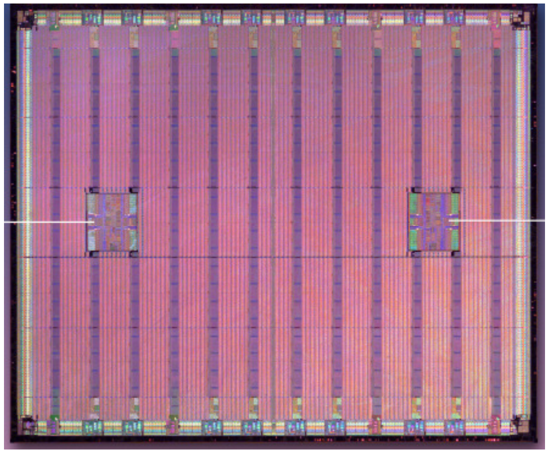

Xilinx Virtex II-Pro
2002
1-2 Power Cores
+ FPGA

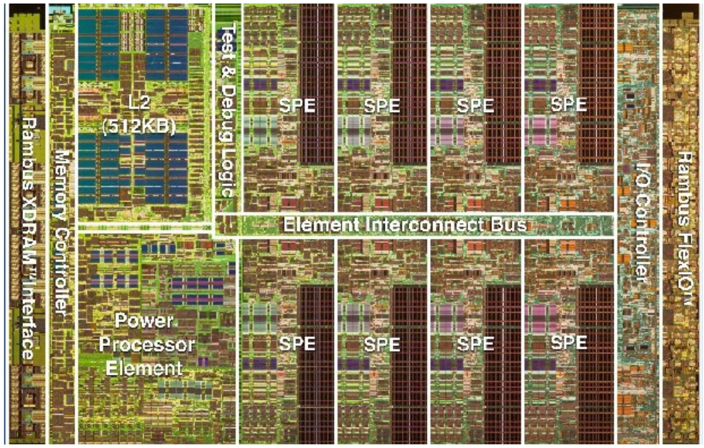

Cell Broadband Engine
2005
1 Power Core
+ 8 accelerator cores

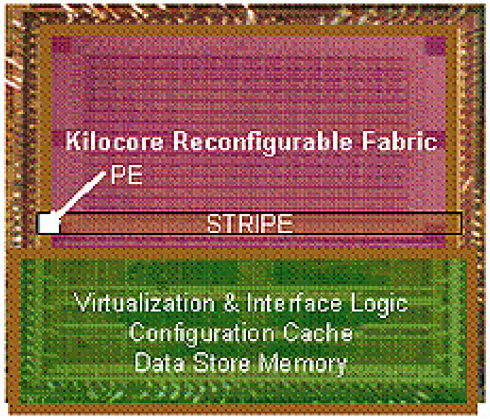

Rapport Kilocore
2006
1 Power Core
+ 256 accelerator cores

What Options do we Have? --LEVERAGING LOCALITY

- Ignoring locality is unbelievably wasteful
  - E.g. cache generally follows Pollack's rule sqrt(x) return on x investment
- Ample proof that enhancing locality can improve both efficiency and performance (dramatically)
  - PGAS / thread local vs global (Increasingly serious alternative to MPI)
  - Vector/GPU (HPC, TPC-H ..)
  - Cell local store / tasks (sorting, searching, mapreduce …)
  - Function placement (Haifa TCP-IP stack example)
- Where we have failed thus far:
  - Non-standard languages – rewrite required
  - Non-incremental changes to HW – big benefit but limited reach

Cell Broadband Engine

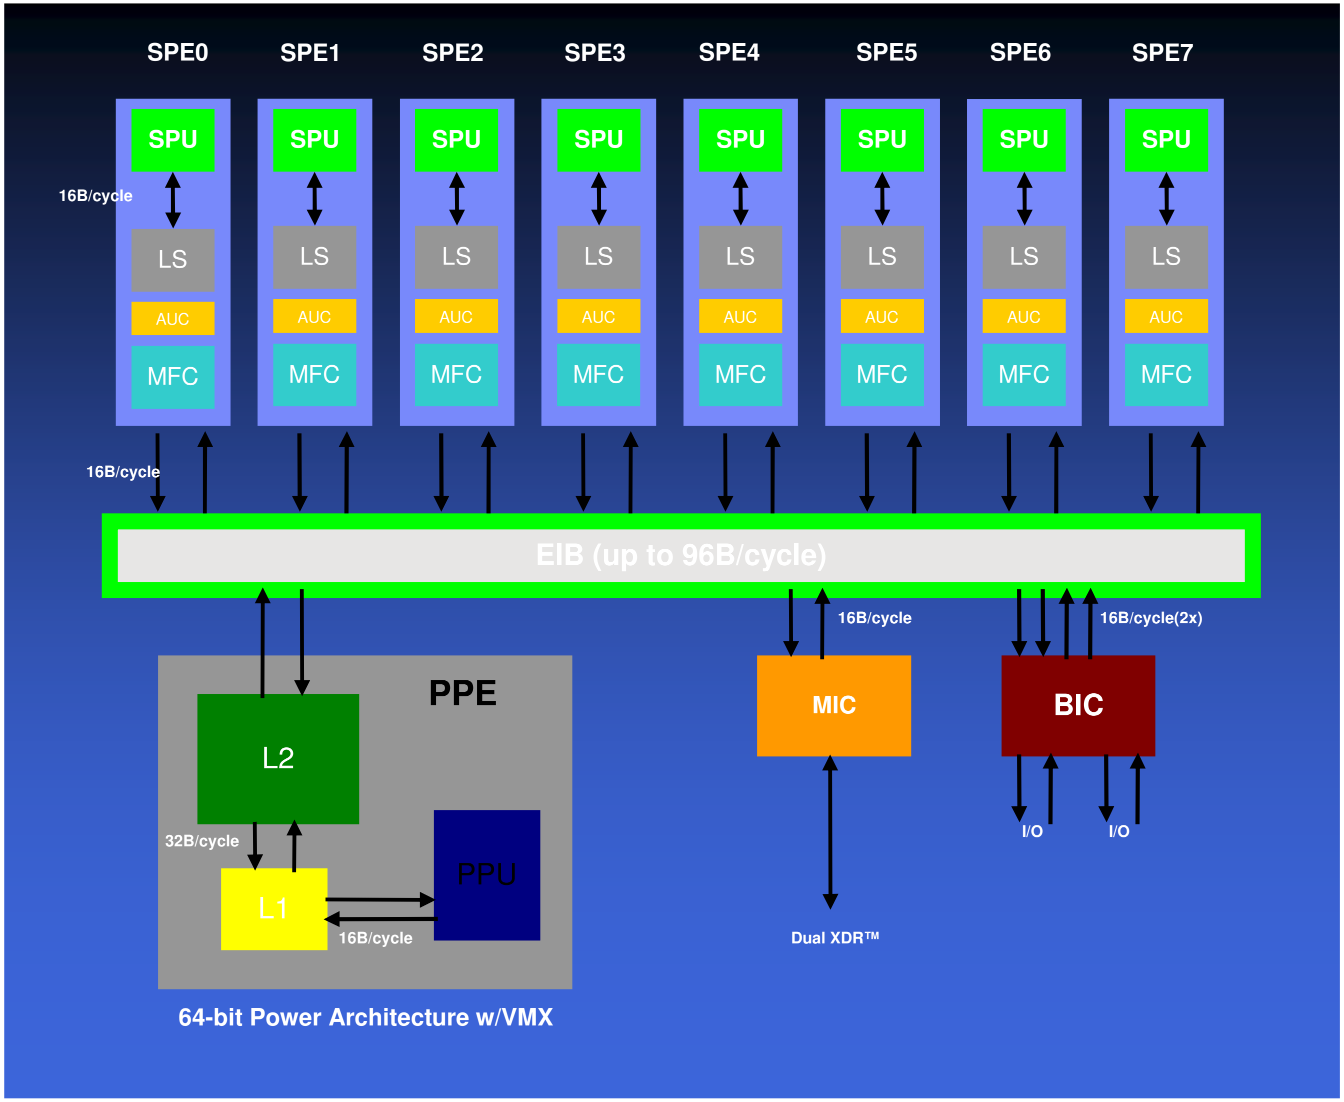

- Heterogeneous Multiprocessor
  - Power processor
  - Synergistic Processing Elements
- Power Processor Element (PPE)
  - general purpose
  - running full-fledged OSs
  - 2 levels of globally coherent cache
- Synergistic Proc. Element (SPE)
  - SPU optimized for computation density
  - 128 bit wide SIMD
  - Fast local memory
  - Globally coherent DMA

Memory Managing Processor vs. Traditional General Purpose Processor

Cell BE

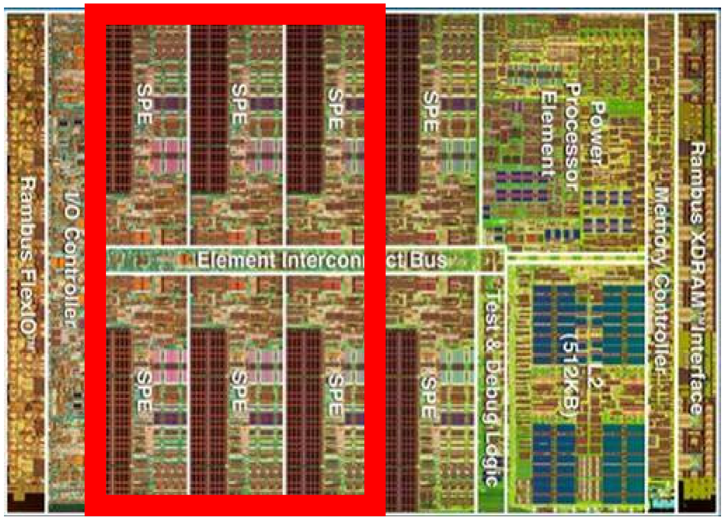

AMD

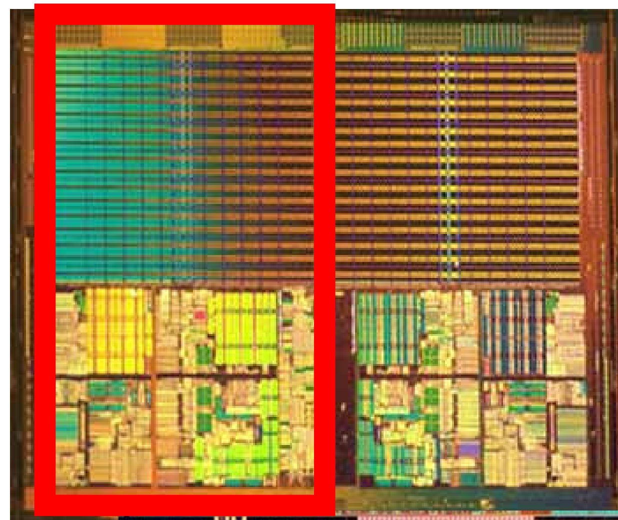

IBM

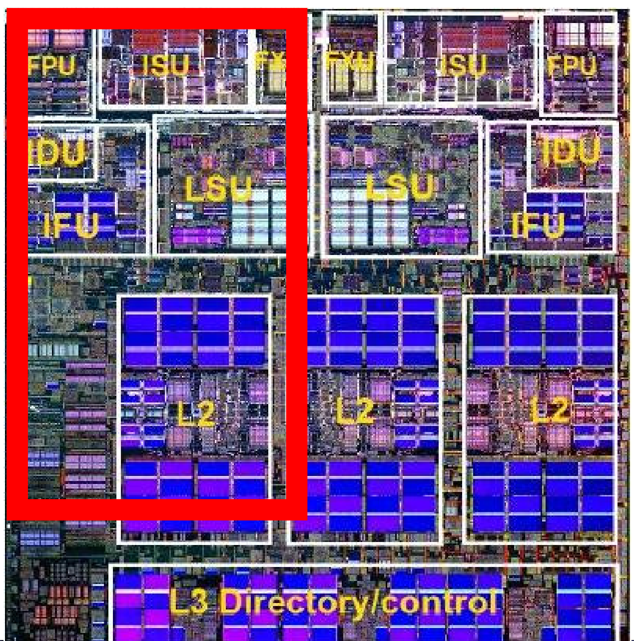

Intel

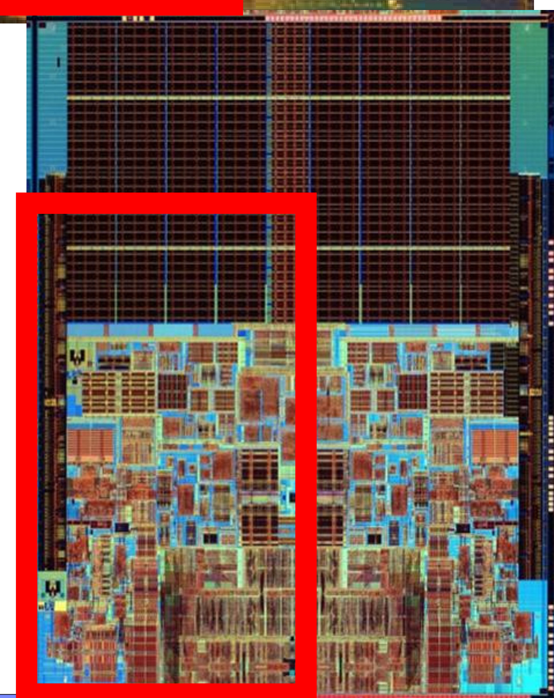

Three Generations of Cell/B.E.

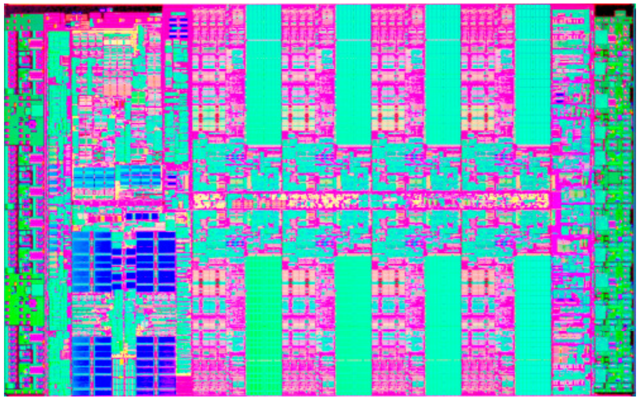

90nm SOI 236mm²

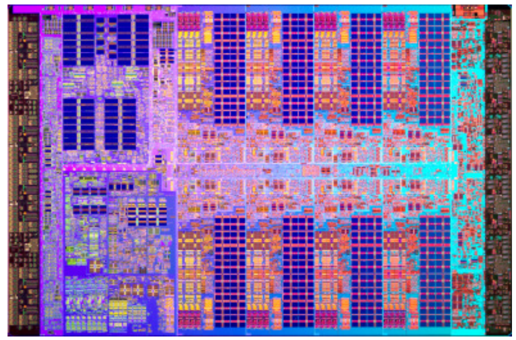

65nm SOI 175mm²

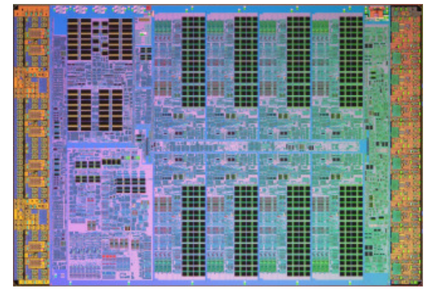

45nm SOI 116mm²

Cell/B.E.					
Generation	W (mm)	H (mm)	Area (mm^2)	Scaling from 90nm	Scaling from 65nm
90nm	19.17	12.29	235.48	100.0%	
65nm	15.59	11.20	174.61	74.2%	100.0%
45nm	12.75	9.06	115.46	49.0%	66.1%
Synergistic Processor Element (SPE)					
Generation	W (mm)	H (mm)	Area (mm^2)	Scaling from 90nm	Scaling from 65nm
90nm	2.54	5.81	14.76	100.0%	
65nm	2.09	5.30	11.08	75.0%	100.0%
45nm	1.59	4.09	6.47	43.9%	58.5%
Power Processor Element (PPE)					
Generation	W (mm)	H (mm)	Area (mm^2)	Scaling from 90nm	Scaling from 65nm
90nm	4.44	6.05	26.86	100.0%	
65nm	3.50	5.60	19.60	73.0%	100.0%
45nm	2.66	4.26	11.32	42.1%	57.7%

Takahashi e.a.

Cell Broadband Engine-based CE Products

Sony Playstation 3 and PS3

Toshiba Regza Cell

Image Processing on the Cell Broadband Engine

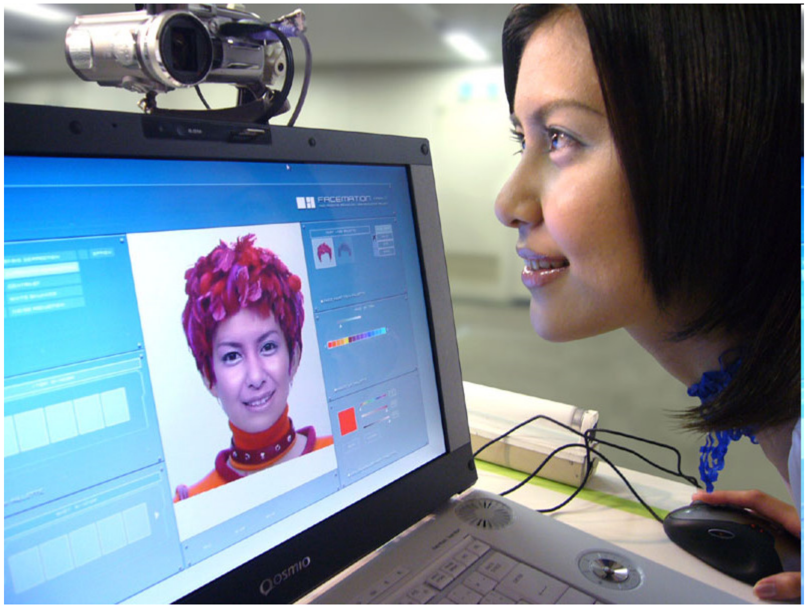

Toshiba Magic Mirror

Sony PlayStation®Move

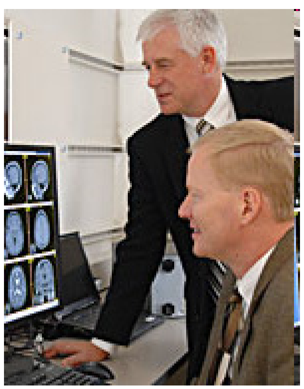

Mayo Clinic/IBM

Mercury Computer

Hitachi Medical

Axion Racing

Uses of Cell Technology beyond Consumer Electronics

- Three Generations of Server Blades Accompanied By 3 SDK Releases
  - IBM QS20
  - IBM QS21
  - IBM QS22
- Two Generations of PCIe Cell Accelerator Boards
  - CAB (Mercury)
  - PXCAB (Mercury/Fixstars/Matrix Vision)
- 1U Formfactor
  - Mercury Computer
  - TPlatforms
- Custom Boards
  - Hitachi Medical (Ultrasound)
  - Other Medical and Defense
- World's First 1 PFlop Computer
  - LANL Roadrunner
- Top 7 Green Systems
  - Green 500 list


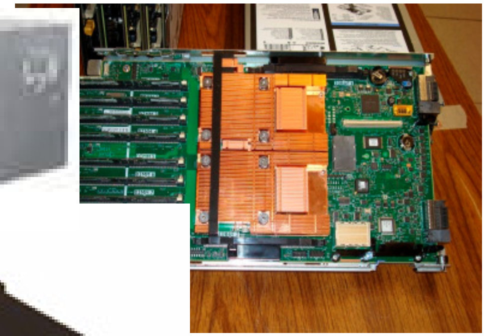



Optimization of Sparse Matrix-Vector Multiplication on Emerging Multicore Platforms

Samuel Williams[*,†], Leonid Oliker[*], Richard Vuduc[§], John Shalf[*], Katherine Yelick[*,†], James Demmel[†]
[*]CRD/NERSC, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
[†]Computer Science Division, University of California at Berkeley, Berkeley, CA 94720, USA
[§]CASC, Lawrence Livermore National Laboratory, Livermore, CA 94551, USA

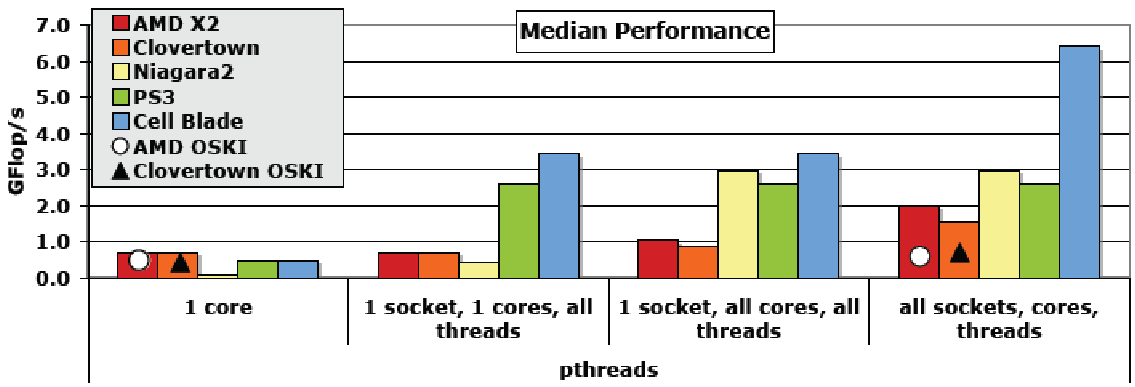

(a)

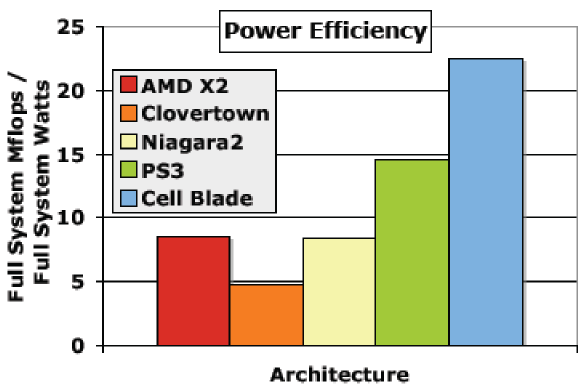

(b)

Figure 5: Architectural comparison of the median matrix performance showing (a) GFlop/s rates of OSKI and optimized SpMV on single-core, full socket, and full system and (b) relative power efficiency computed as total full system Mflop/s divided by sustained full system Watts (see Table 1).

Current Cell: Integer Workloads

Breadth-First Search
Villa, Scarpazza, Petrini, Peinador
IPDPS 2007

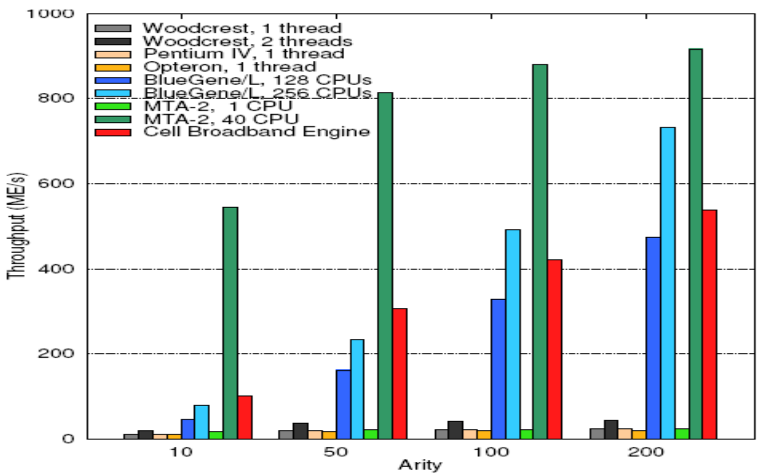

D.A. Bader et al. / Parallel Computing 33 (2007) 720–740

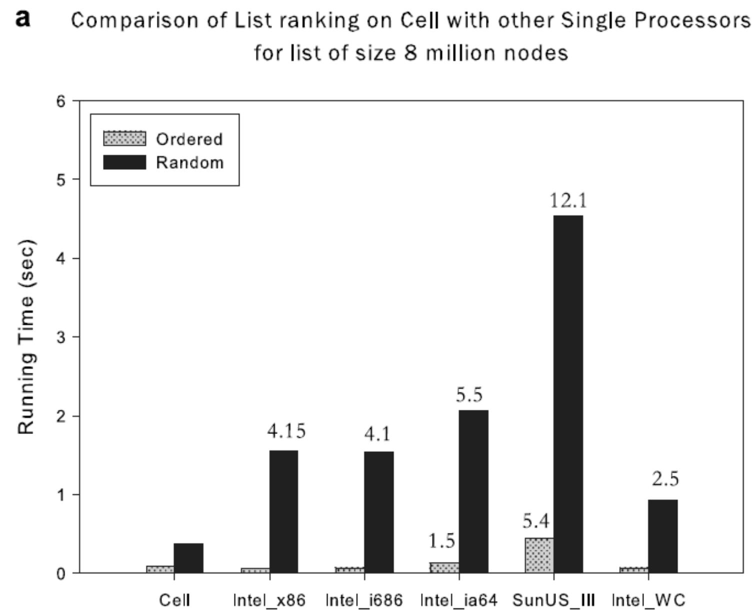

Mapreduce

Sangkaralingam, De Kruijf, Oct. 2007

Application Name	Application Type	Lines of Code		Speedup vs. Core2			BIPS		
		MapReduce	Serial	1-SPE	8-SPEs	8-SPE Ideal	1-SPE	8-SPEs	8-SPE Ideal
histogram	partition-dominated	345	216	0.16	0.15	2.44	1.56	1.51	24.49
kmeans	partition-dominated	324	318	0.91	3.00	6.92	2.08	7.35	17.01
linearRegression	map-dominated	279	114	0.34	2.59	2.67	1.47	11.32	11.70
wordCount	partition-dominated	226	324	0.87	0.96	10.26	1.52	1.74	18.64
NAS_EP	map-dominated	264	112	1.08	8.62	8.62	2.00	15.93	15.95
distributedSort	sort-dominated	171	93[c]	0.41	0.76	5.48	1.28	2.38	17.15

Sort:Gedik, Bordawekar, Yu (IBM)

Table 3: Out-of-core sort performance (in secs)

# items	16 SPEs bitonic	3.2GHz Xeon quick	3.2GHz Xeon quick 2-core	PPE quick
1M	0.0098	0.1813	0.098589	0.4333
2M	0.0234	0.3794	0.205728	0.9072
4M	0.0569	0.7941	0.429499	1.9574
8M	0.1372	1.6704	0.895168	4.0746
16M	0.3172	3.4673	1.863354	8.4577
32M	0.7461	7.1751	3.863495	18.3882
64M	1.7703	14.8731	7.946356	38.7473
128M	4.0991	30.0481	16.165578	79.9971

Playstation 3 high-level organization and PS3 cluster.

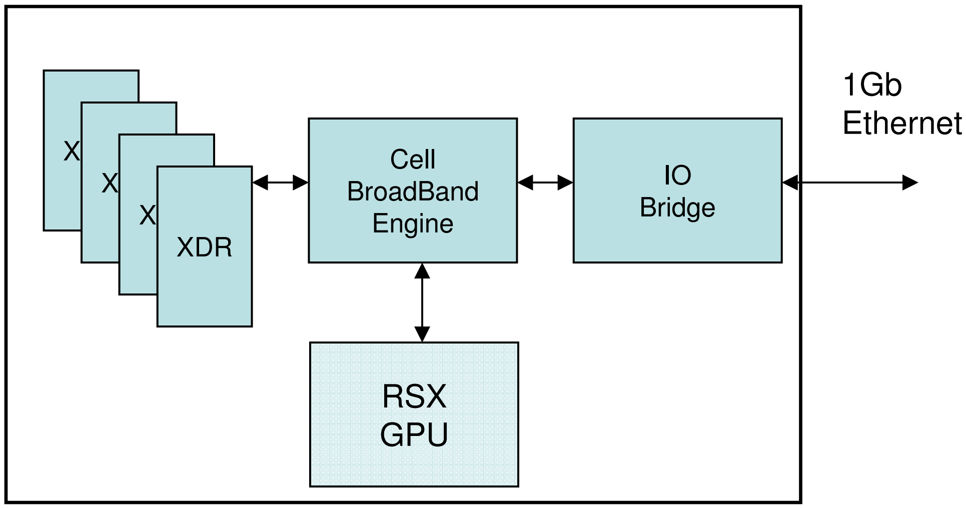

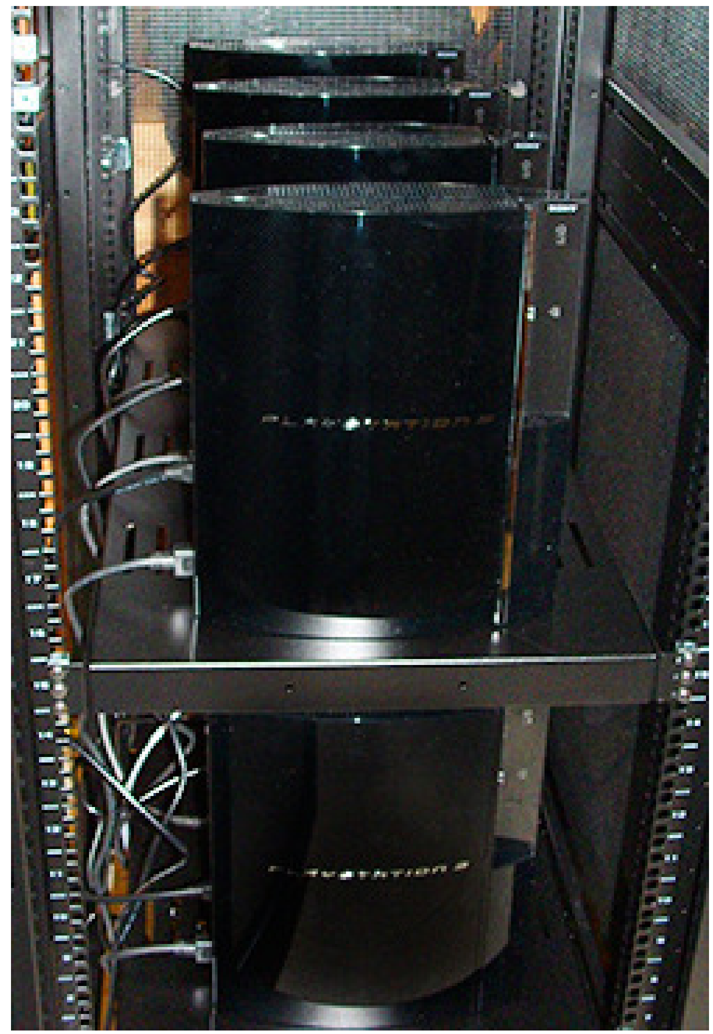

Roadrunner accelerated node and system.

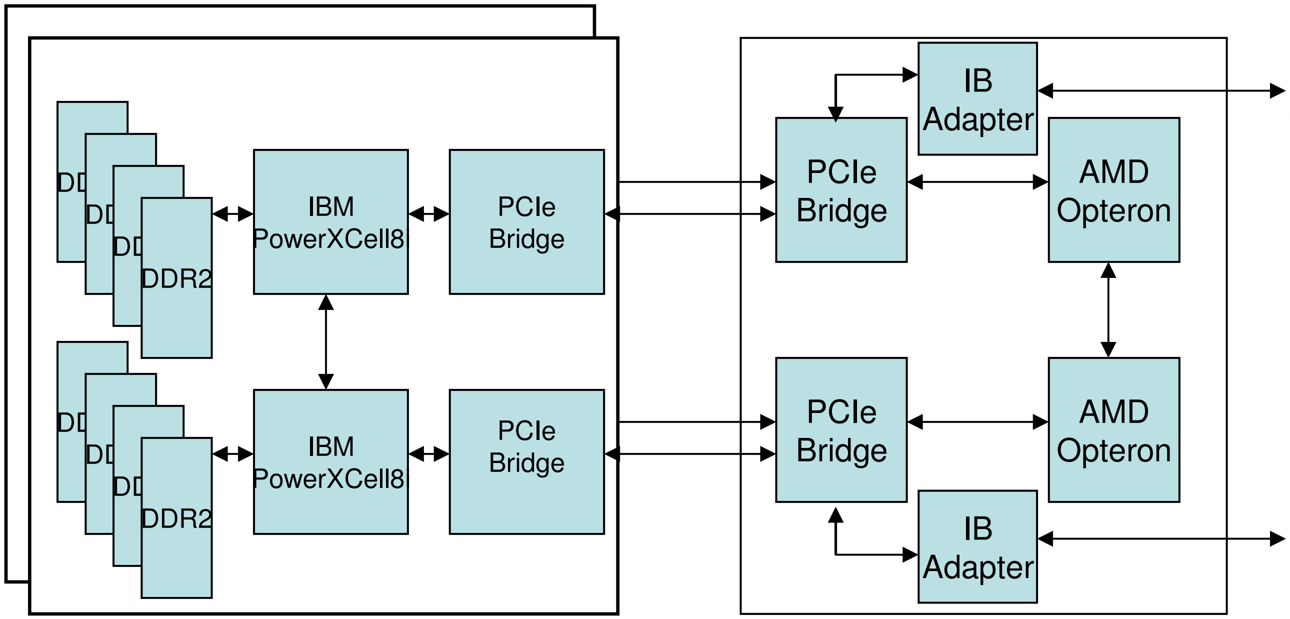

Roadrunner Accelerated Node

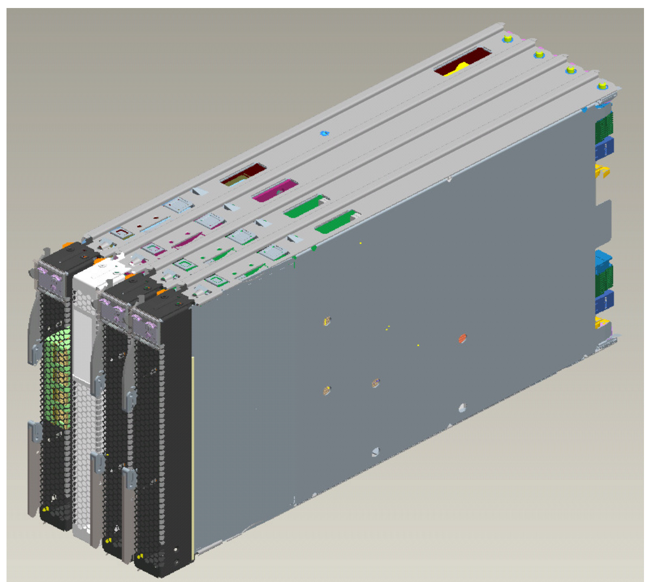

QPACE PowerXCell8i node card and system.

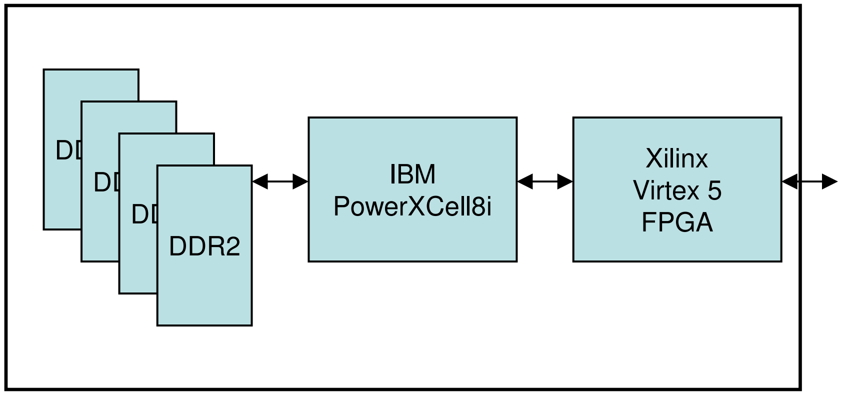

QPACE node card.

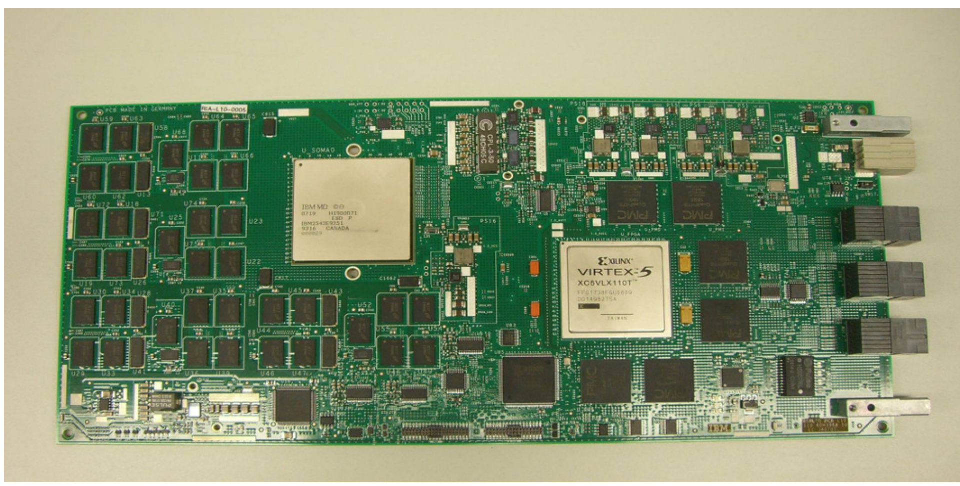

June 2010 Green500

Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
1	773.38	Forschungszentrum Juelich (FZJ)	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
1	773.38	Universitaet Regensburg	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
1	773.38	Universitaet Wuppertal	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
4	492.64	National Supercomputing Centre in Shenzhen (NSCS)	Dawning Nebulae, TC3600 blade CB60-G2 cluster, Intel Xeon 5650/ nVidia C2050, Infiniband	2580
5	458.33	DOE/NNSA/LANL	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Infiniband	276
5	458.33	IBM Poughkeepsie Benchmarking Center	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Infiniband	138
7	444.25	DOE/NNSA/LANL	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband	2345.5
8	431.88	Institute of Process Engineering, Chinese Academy of Sciences	Mole-8.5 Cluster Xeon L5520 2.26 Ghz, nVidia Tesla, Infiniband	480
9	418.47	Mississippi State University	iDataPlex, Xeon X56xx 6C 2.8 GHz, Infiniband	72
10	397.56	Banking (M)	iDataPlex, Xeon X56xx 6C 2.66 GHz, Infiniband	72

* Performance data obtained from publicly available sources including TOP500

Nov. 2010 Green 500

Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
1	1684.20	IBM Thomas J. Watson Research Center	NNSA/SC Blue Gene/Q Prototype	38.80
2	958.35	GSIC Center, Tokyo Institute of Technology	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows	1243.80
3	933.06	NCSA	Hybrid Cluster Core i3 2.93Ghz Dual Core, NVIDIA C2050, Infiniband	36.00
4	828.67	RIKEN Advanced Institute for Computational Science	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	57.96
5	773.38	Forschungszentrum Juelich (FZJ)	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
5	773.38	Universitaet Regensburg	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
5	773.38	Universitaet Wuppertal	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
8	740.78	Universitaet Frankfurt	Supermicro Cluster, QC Opteron 2.1 GHz, ATI Radeon GPU, Infiniband	385.00
9	677.12	Georgia Institute of Technology	HP ProLiant SL390s G7 Xeon 6C X5660 2.8Ghz, nVidia Fermi, Infiniband QDR	94.40
10	636.36	National Institute for Environmental Studies	GOSAT Research Computation Facility, nvidia	117.15

* Performance data obtained from publicly available sources including TOP500

Blue Gene/Q

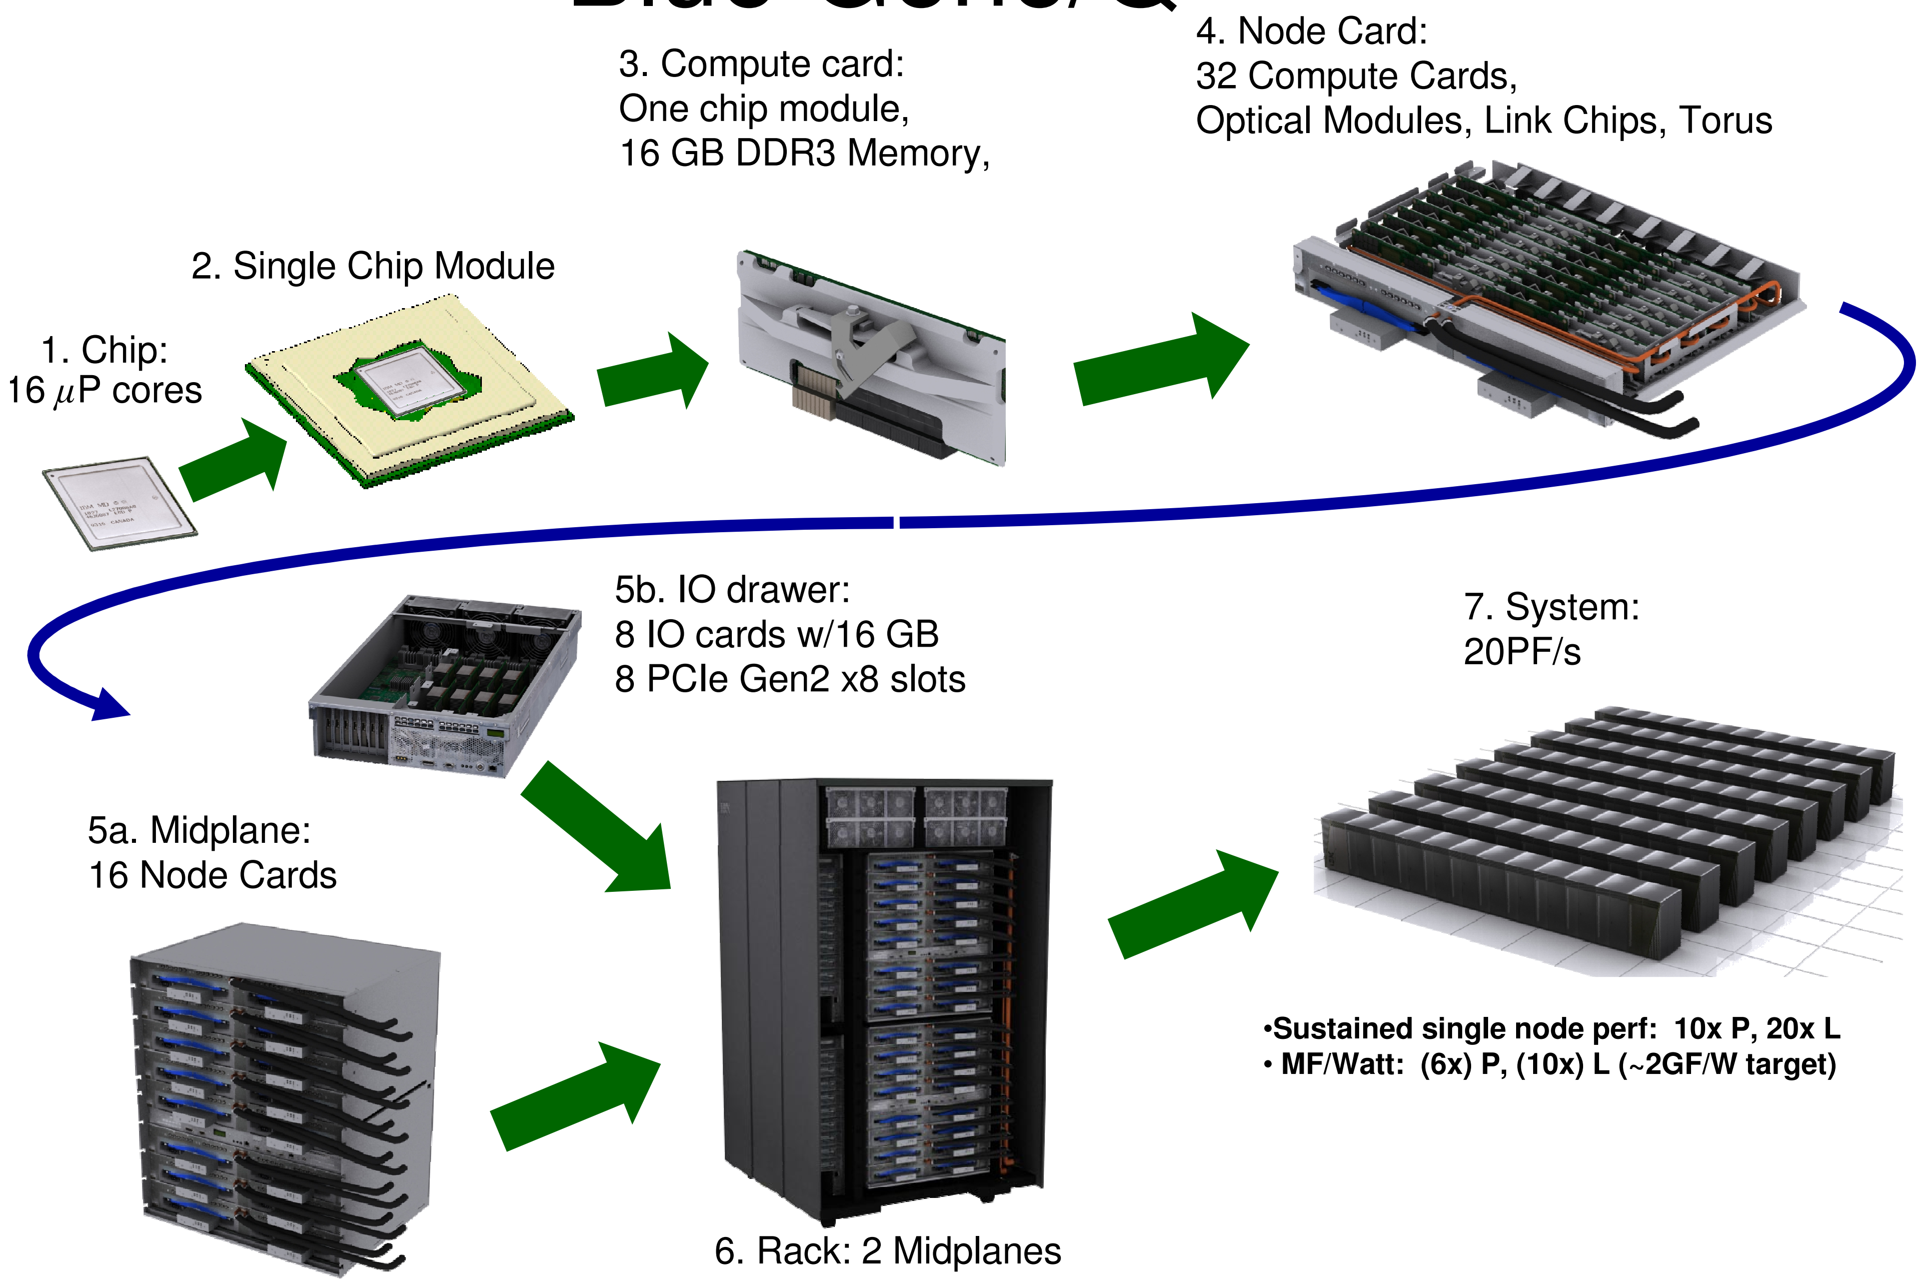
1. Chip:
16 μP cores
2. Single Chip Module
3. Compute card:
One chip module,
16 GB DDR3 Memory,
4. Node Card:
32 Compute Cards,
Optical Modules, Link Chips, Torus
5b. IO drawer:
8 IO cards w/16 GB
8 PCIe Gen2 x8 slots
5a. Midplane:
16 Node Cards
6. Rack: 2 Midplanes
7. System:
20PF/s
•Sustained single node perf: 10x P, 20x L
• MF/Watt: (6x) P, (10x) L (~2GF/W target)

Performance and Productivity Challenges require a Multi-Dimensional Approach

Highly Productive Systems

POWER

Highly Scalable Multi-core Systems

Hybrid Systems

Comprehensive (Holistic) System Innovation & Optimization

HPC Cluster Directions

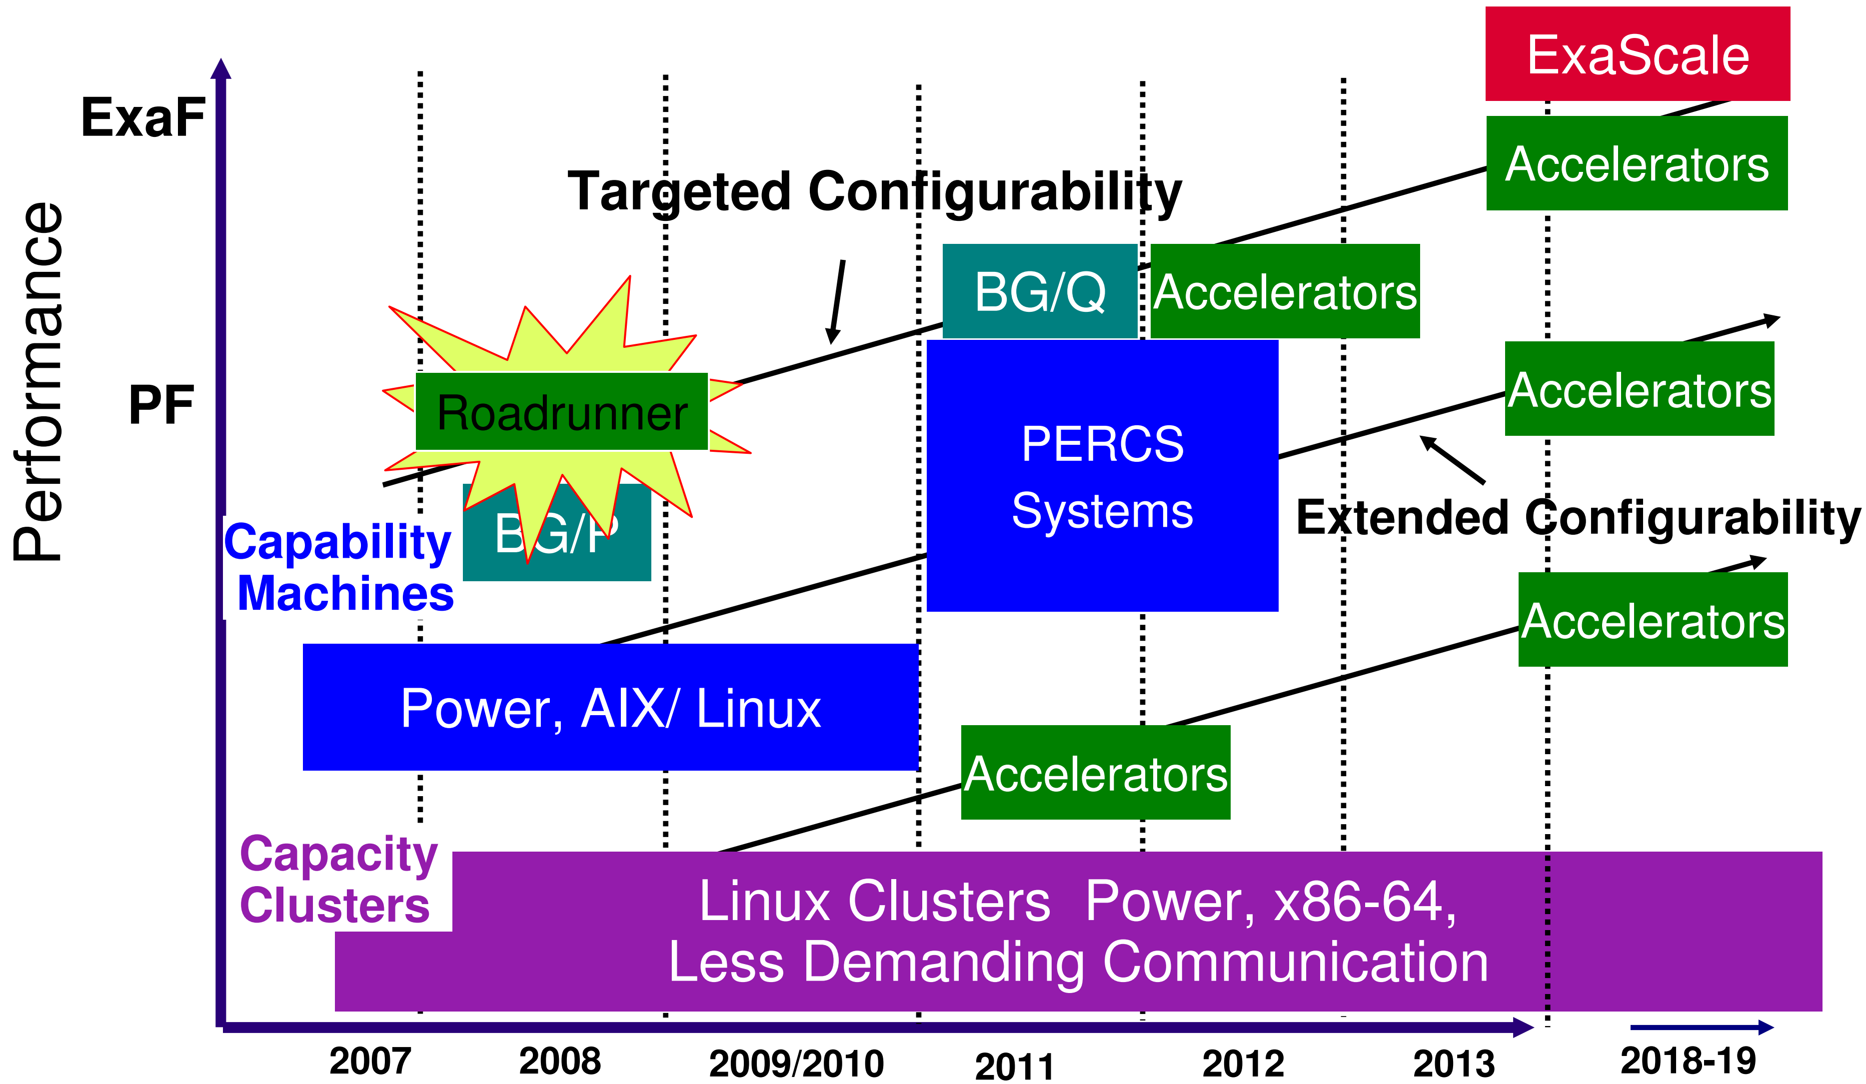

Top 500, Nov 2010

Site	Computer
National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT TH MPP, X5670 2.93Ghz 6C, NVIDIA GPU, FT-1000 8C NUDT
DOE/SC/Oak Ridge National Laboratory United States	Jaguar - Cray XT5-HE Opteron 6-core 2.6 GHz Cray Inc.
National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU Dawning
GSIC Center, Tokyo Institute of Technology Japan	TSUBAME 2.0 - HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows NEC/HP
DOE/SC/LBNL/NERSC United States	Hopper - Cray XE6 12-core 2.1 GHz Cray Inc.
Commissariat a l'Energie Atomique (CEA) France	Tera-100 - Bull bullx super-node S6010/S6030 Bull SA
DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband IBM
National Institute for Computational Sciences/University of Tennessee United States	Kraken XT5 - Cray XT5-HE Opteron 6-core 2.6 GHz Cray Inc.
Forschungszentrum Juelich (FZJ) Germany	JUGENE - Blue Gene/P Solution IBM
DOE/NNSA/LANL/SNL United States	Cielo - Cray XE6 8-core 2.4 GHz Cray Inc.

Next Era of Innovation – Hybrid Computing

Symmetric Multiprocessing Era

p6 → 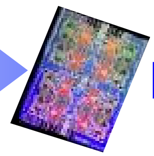 **p7**

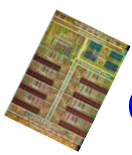

Cell →

BlueGene →

Technology Out

Driven by cores/threads

Hybrid Computing Era

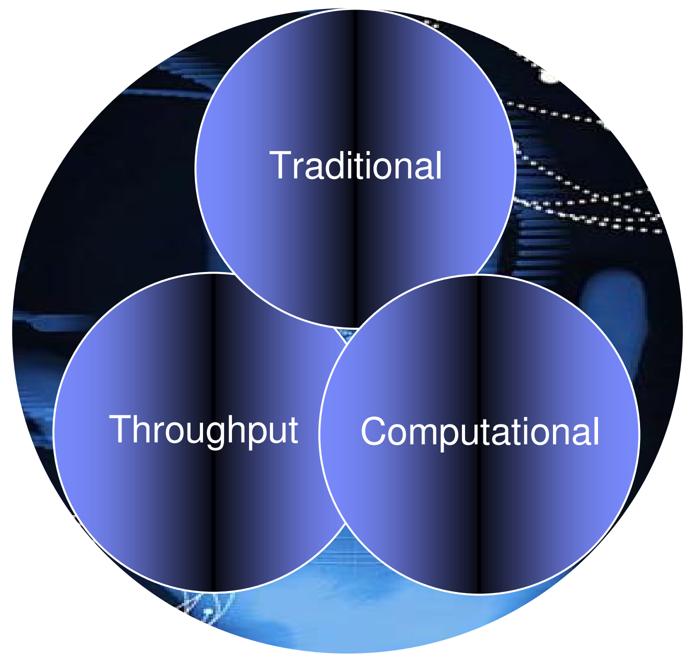

Market In

Driven by workload consolidation

Converging Software (Much Harder!)

- Software-Hardware Efficiency Driven by
  - Number of operations (we all learn this in school)
  - Degree of thread parallelism
  - Degree of data parallelism
  - Degree of locality (code and data, data more important)
  - Degree of predictability (code and data, data more important)

- Need a new Portable Framework
  - Allow portable format to retain enough information for run-time optimization.
  - Allow run-time optimization to heterogeneous hardware for each of the parameters above.

OpenCL vs CUDA ecosystem

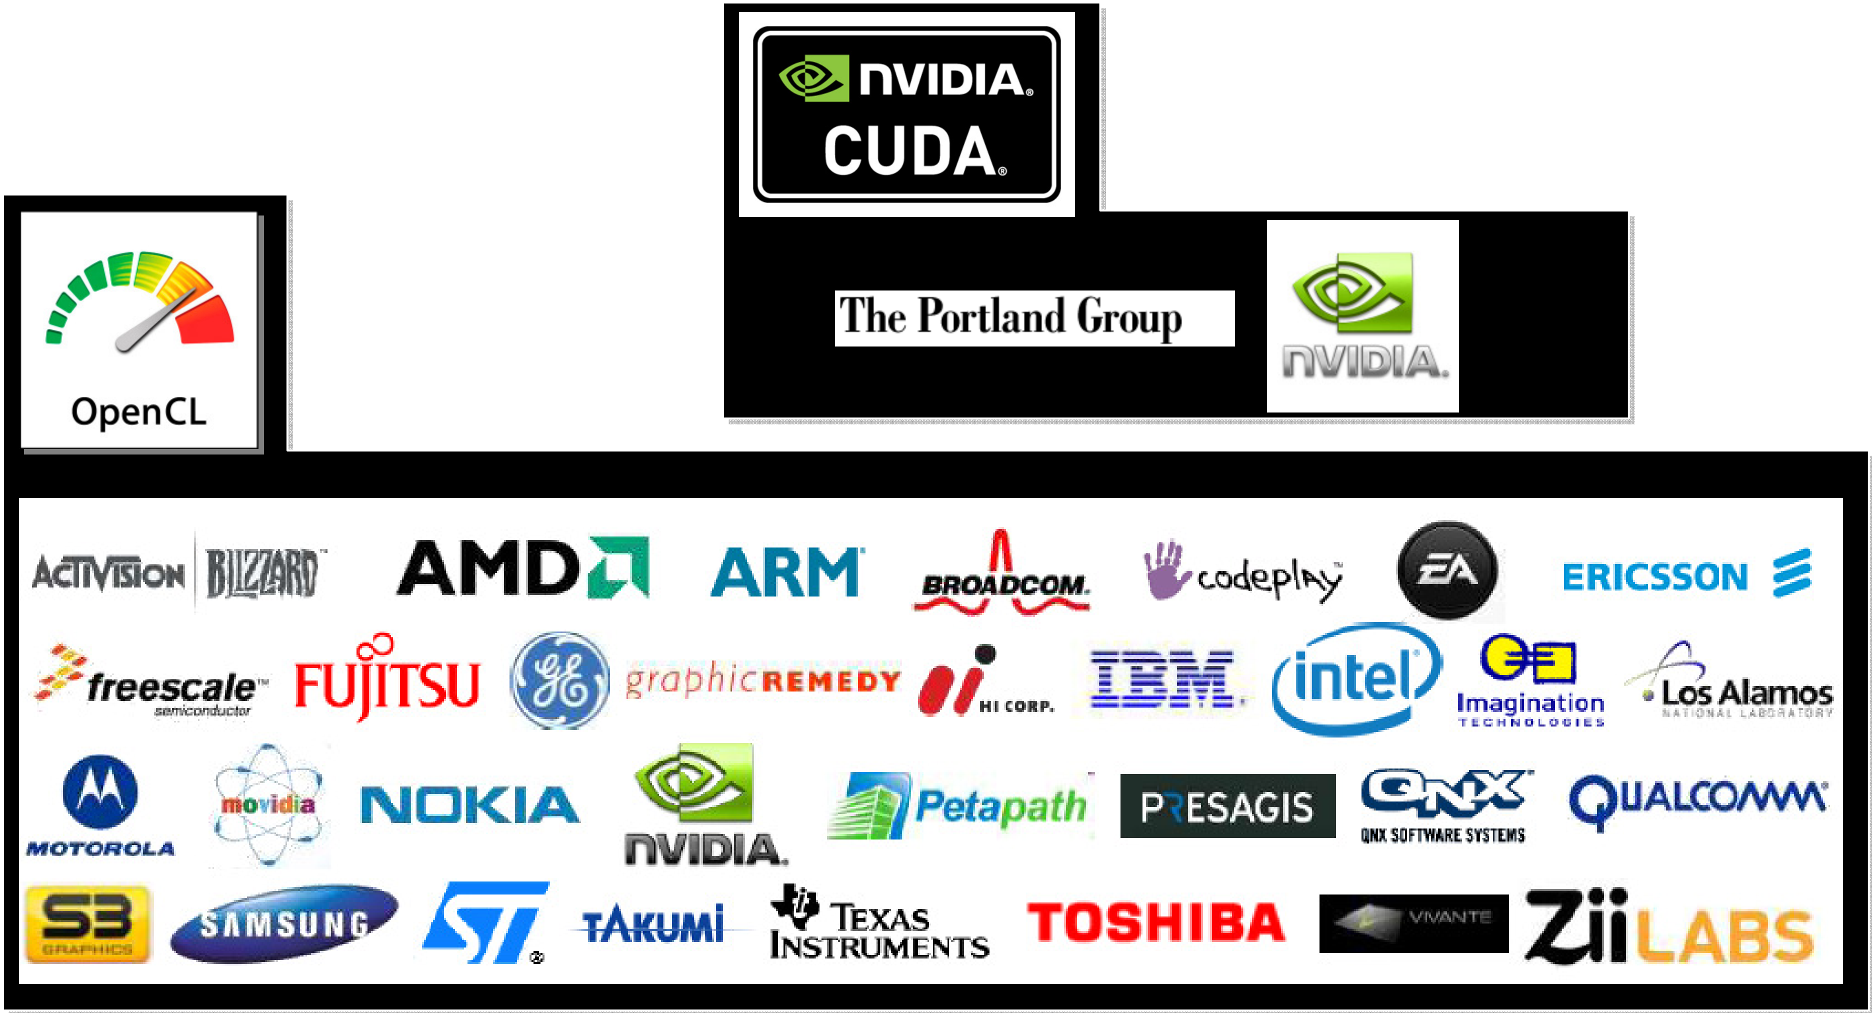

Two Standards for Programming the Node

- Two standards evolving from different sides of the market

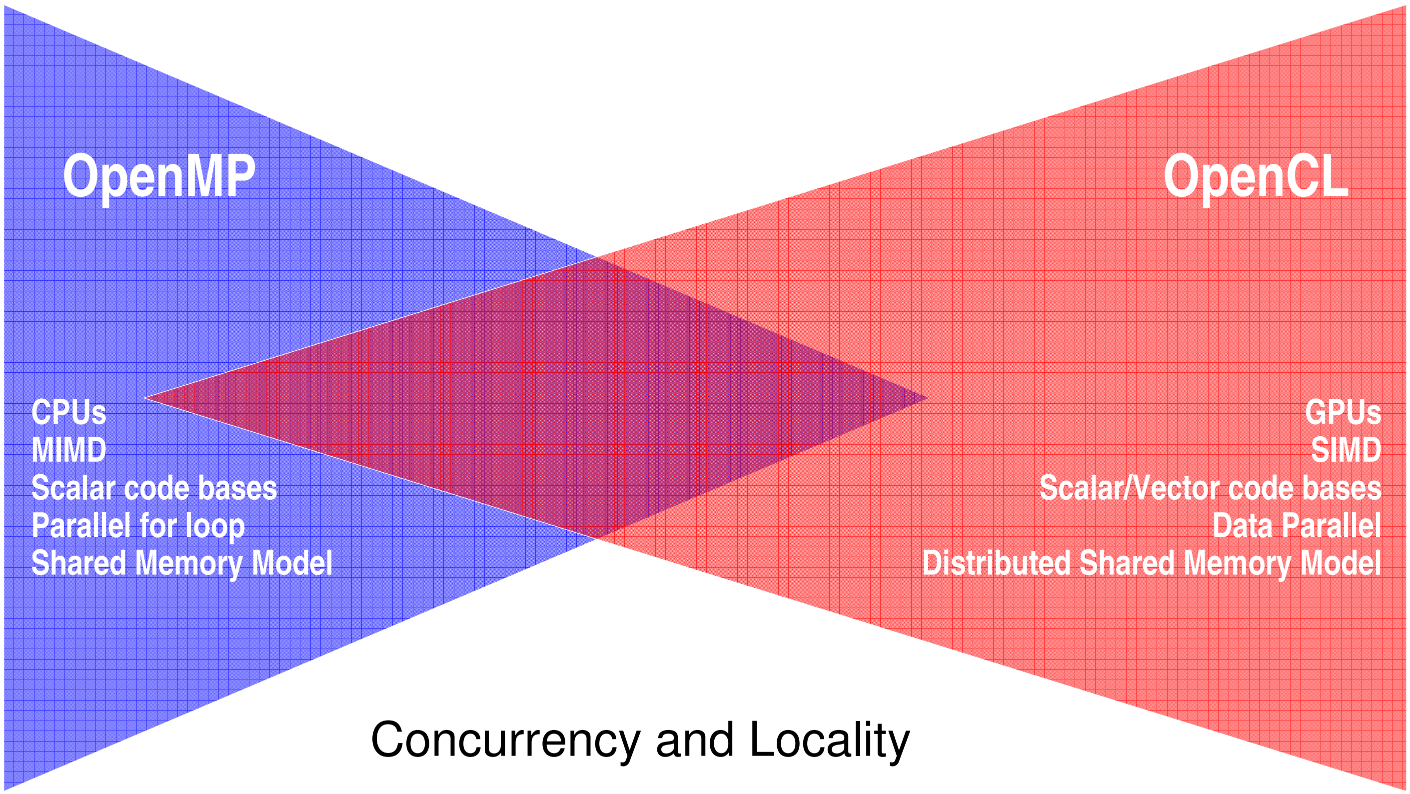

Hybrid Parallel Gas Dynamics in OpenCL

Los Alamos Booth (715) Configuration:

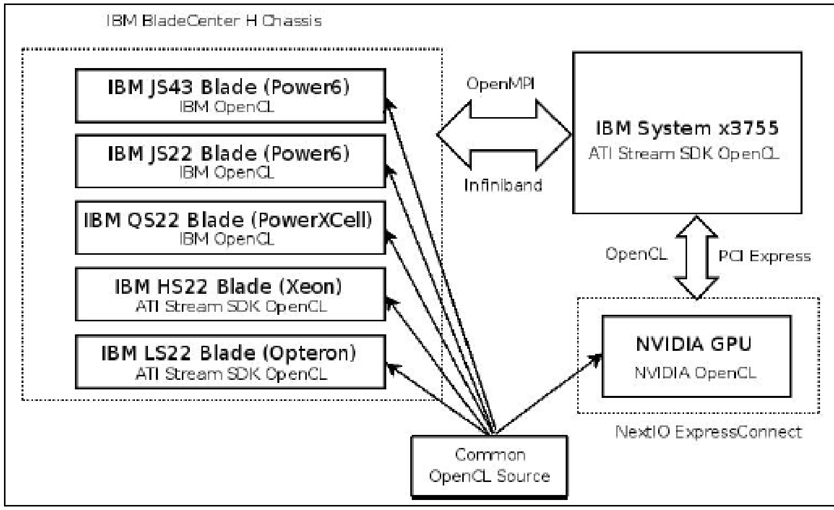

IBM Booth (1335) Configuration:

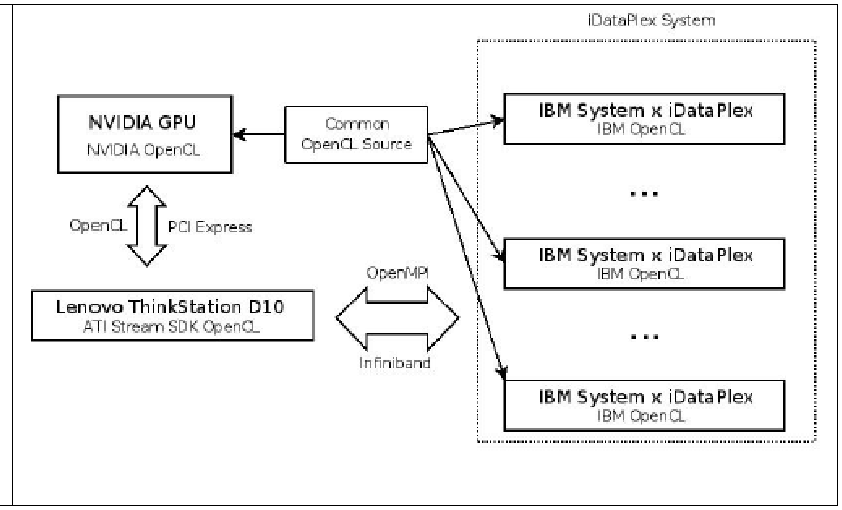

System	*Booth*	*Linux Distribution*	*IBM OpenCL Device*	*ATI Stream SDK OpenCL Device*	*NVIDIA OpenCL Device*	*NextIO ExpressConnect*
IBM BladeCenter QS22 Blade	Los Alamos (715)	Fedora Core	PowerXCell			
IBM BladeCenter JS22 & JS43 Express Blades	Los Alamos (715)	Red Hat	Power6			
IBM BladeCenter HS22 Blade	Los Alamos (715)	Fedora Core		Intel Xeon		
IBM BladeCenter LS22Blade	Los Alamos (715)	Fedora Core		AMD Opteron		
IBM System x3755	Los Alamos (715)	Fedora Core		AMD Opteron	GeForce 9800 GT	N2800
IBM System x iDataPlex	IBM (1335)	Fedora Core		AMD Opteron		
Lenovo ThinkStation D10	IBM (1335)	Fedora Core		AMD Opteron	Quadro FX 4600	

Hybrid Parallel Gas Dynamics Code: http://sourceforge.net/projects/hypgad/

IBM's OpenCL Development Kit for Linux on Power: http://www.alphaworks.ibm.com/tech/opencl

ATI's Stream Software Development Kit: http://developer.amd.com/gpu/ATIStreamSDKBetaProgram/Pages/default.aspx

NVIDIA's OpenCL for NVIDIA's CUDA Architecture GPUs: http://www.nvidia.com/object/cuda_opencl.html

Khronos OpenCL: http://www.khronos.org/opencl/

Finite difference, linear, hyperbolic, inhomogeneous PDE
EMRI Teukolsky

Khanna and McKennon

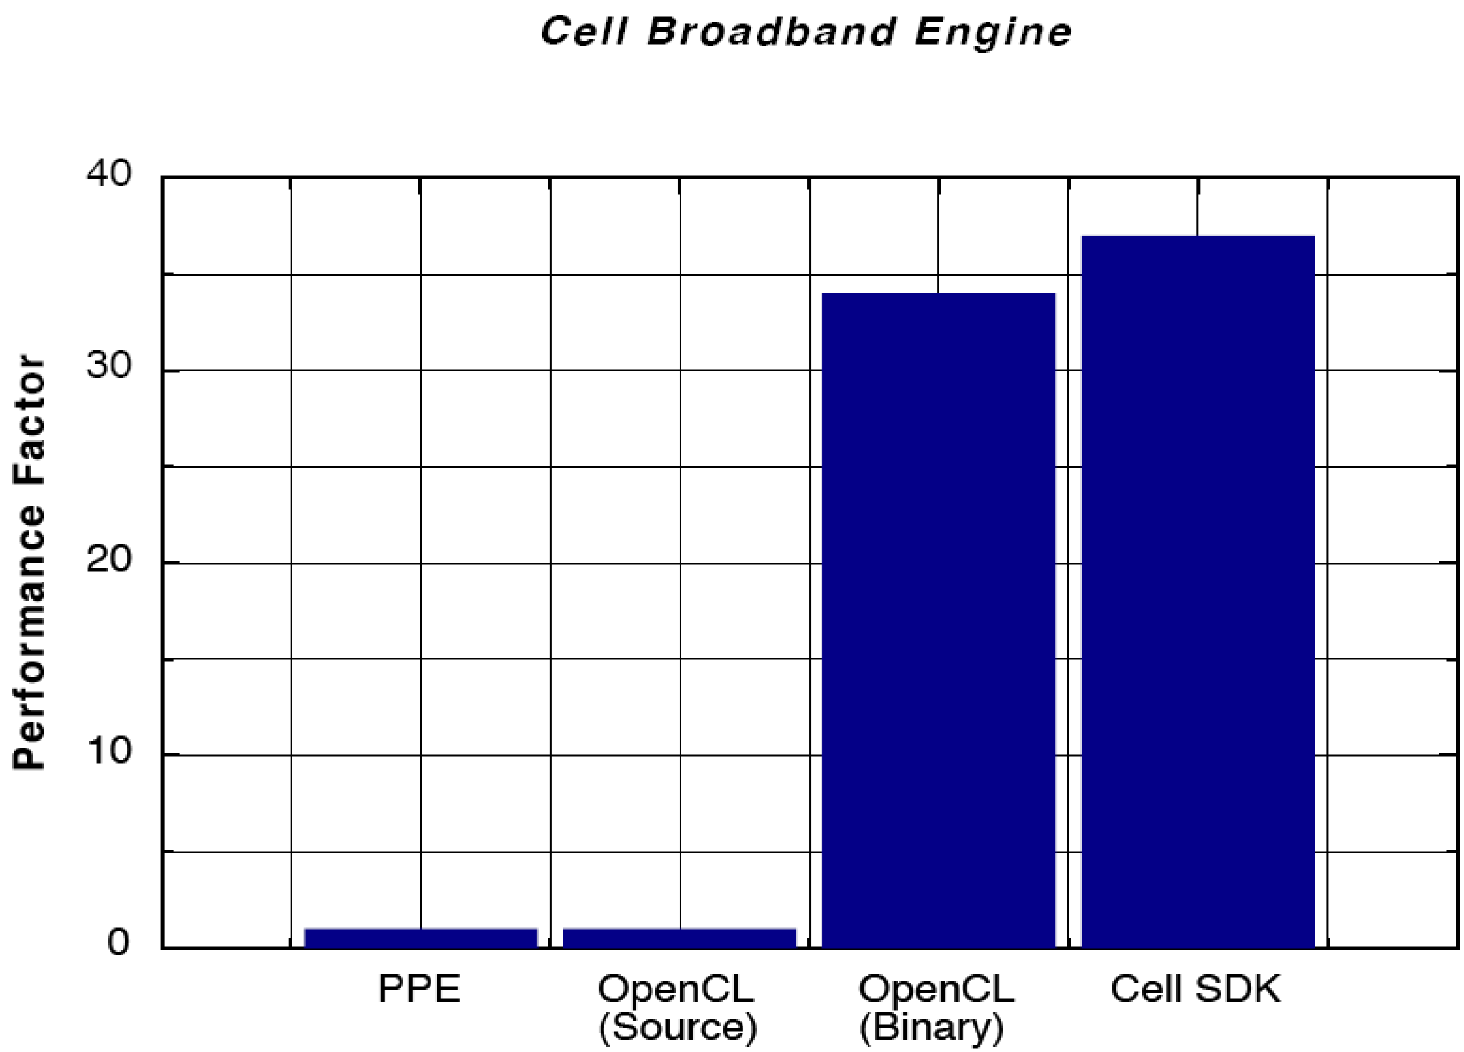

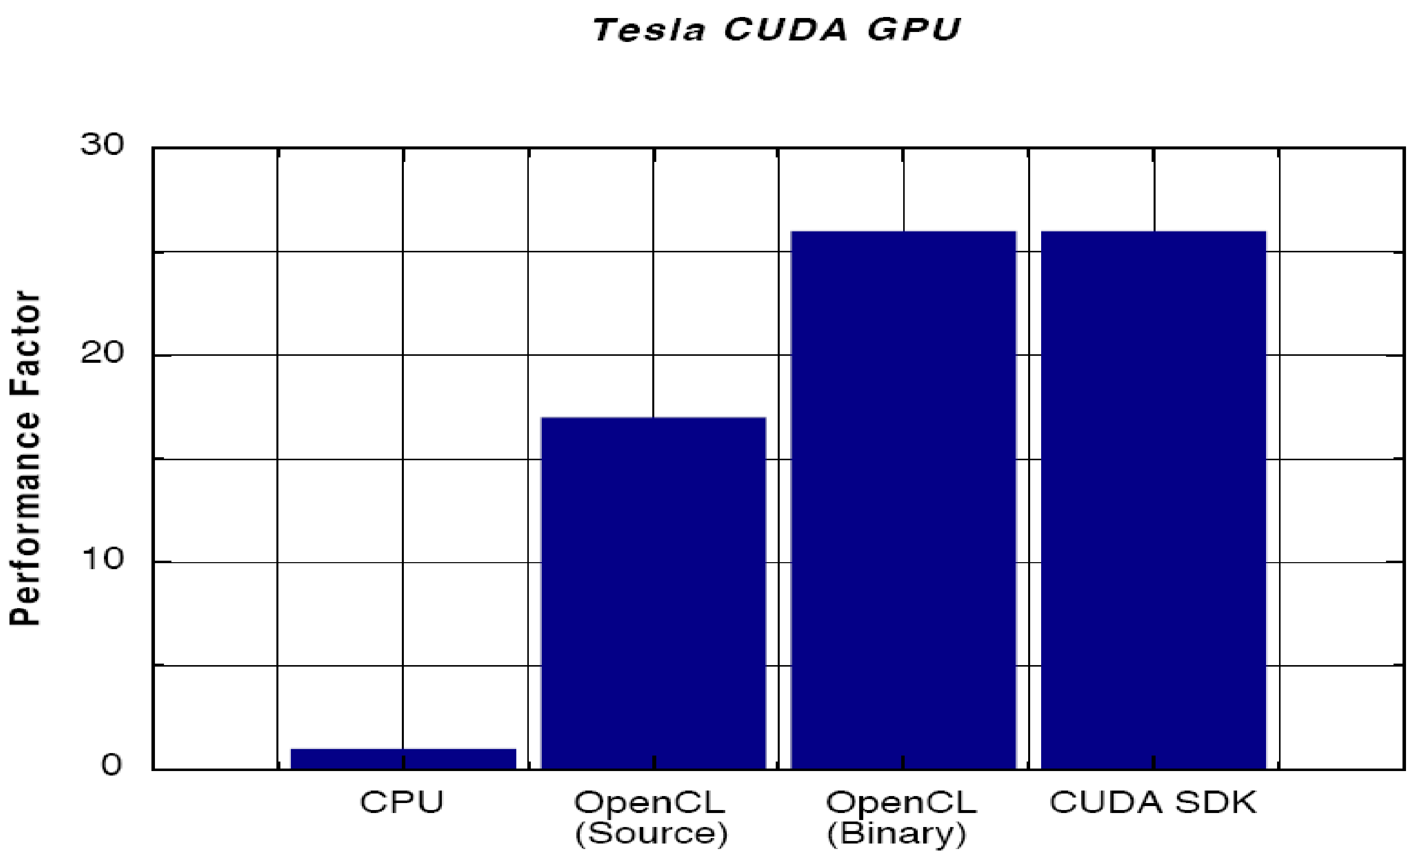

Figure 3. Overall performance of the EMRI Teukolsky Code accelerated by the Tesla CUDA GPU using OpenCL. The baseline here is the supporting system's CPU – an AMD Phenom 2.5 GHz processor.

Khanna & McKennon

IDENTICAL CODE BASE !

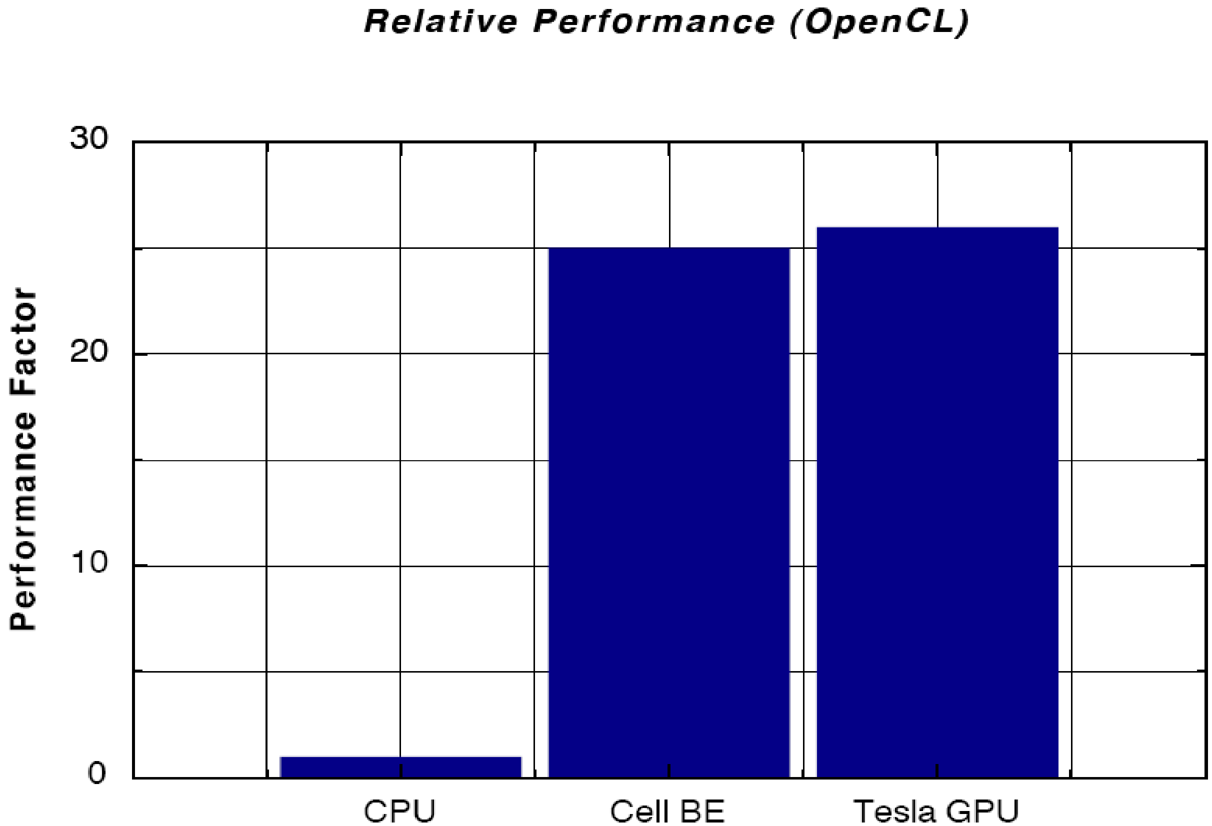

Figure 4. Relative performance of the OpenCL-based EMRI Teukolsky Code on all discussed architectures – CPU, CBE and GPU. The baseline here is the system CPU – an AMD Phenom 2.5 GHz processor.

Microprocessor Trends

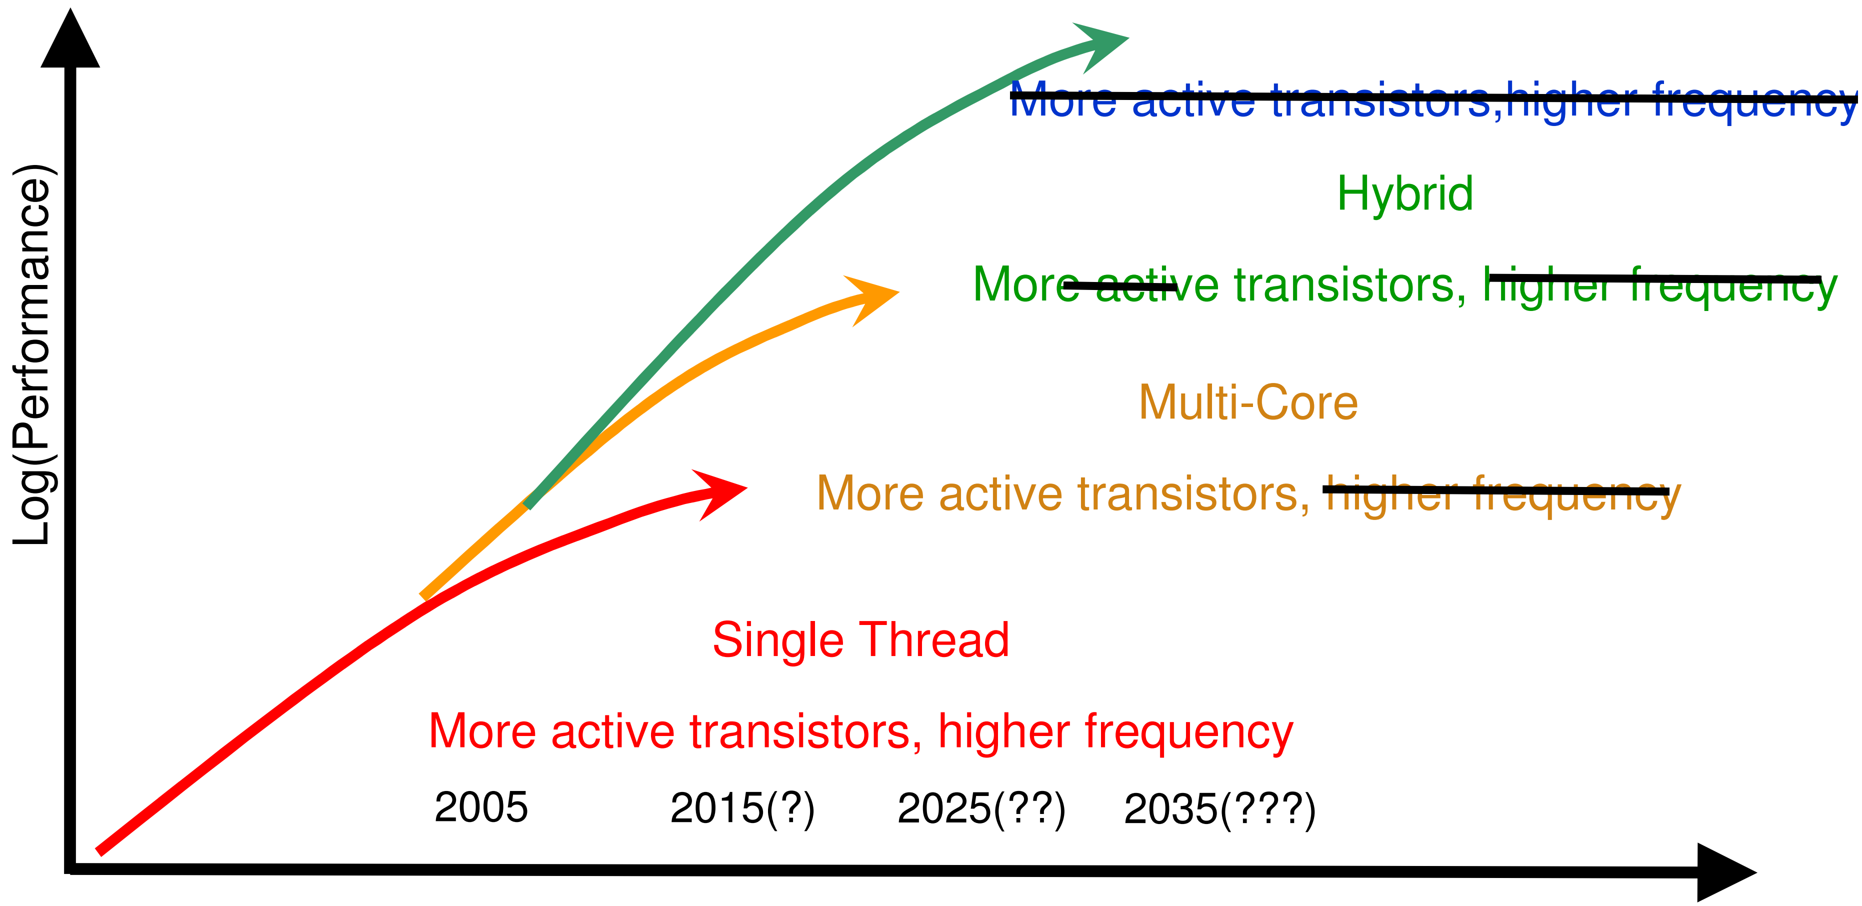

Microprocessor Trends

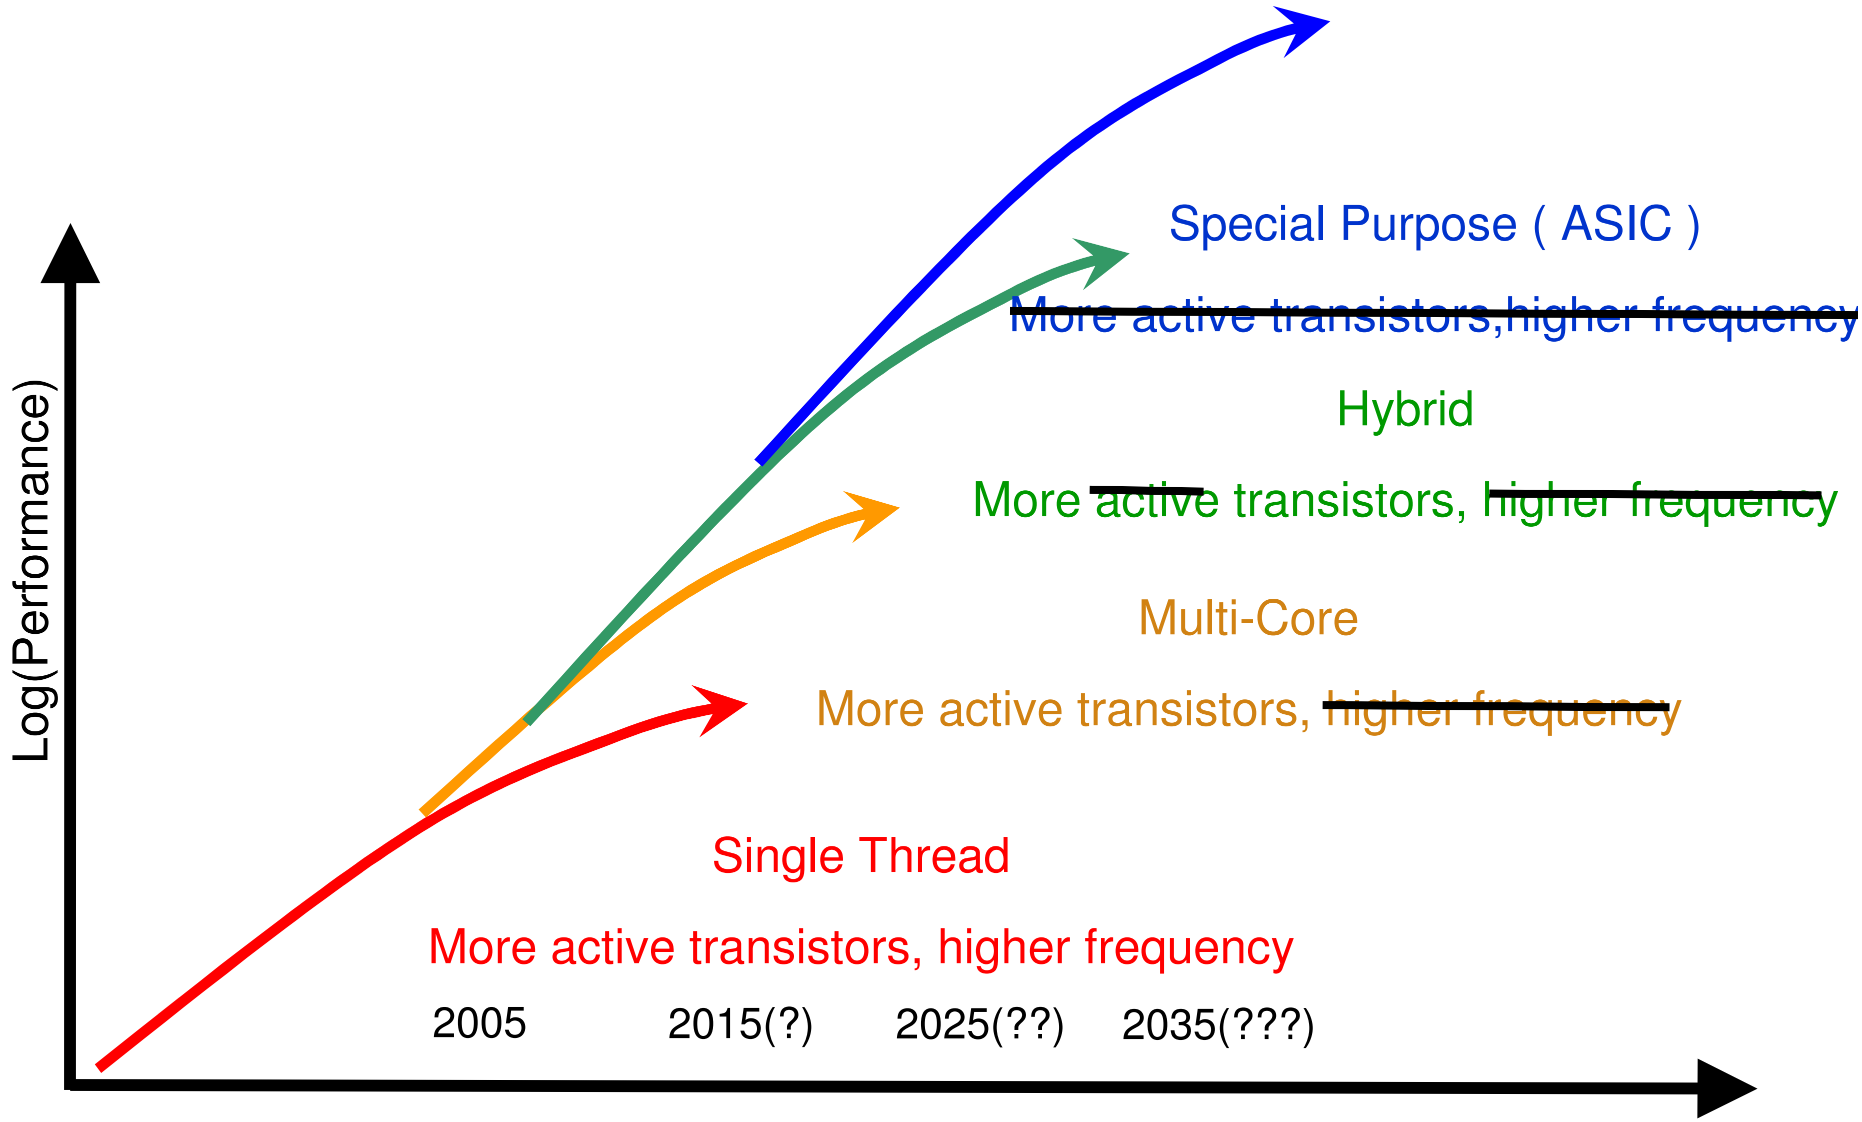

Microprocessor Trends

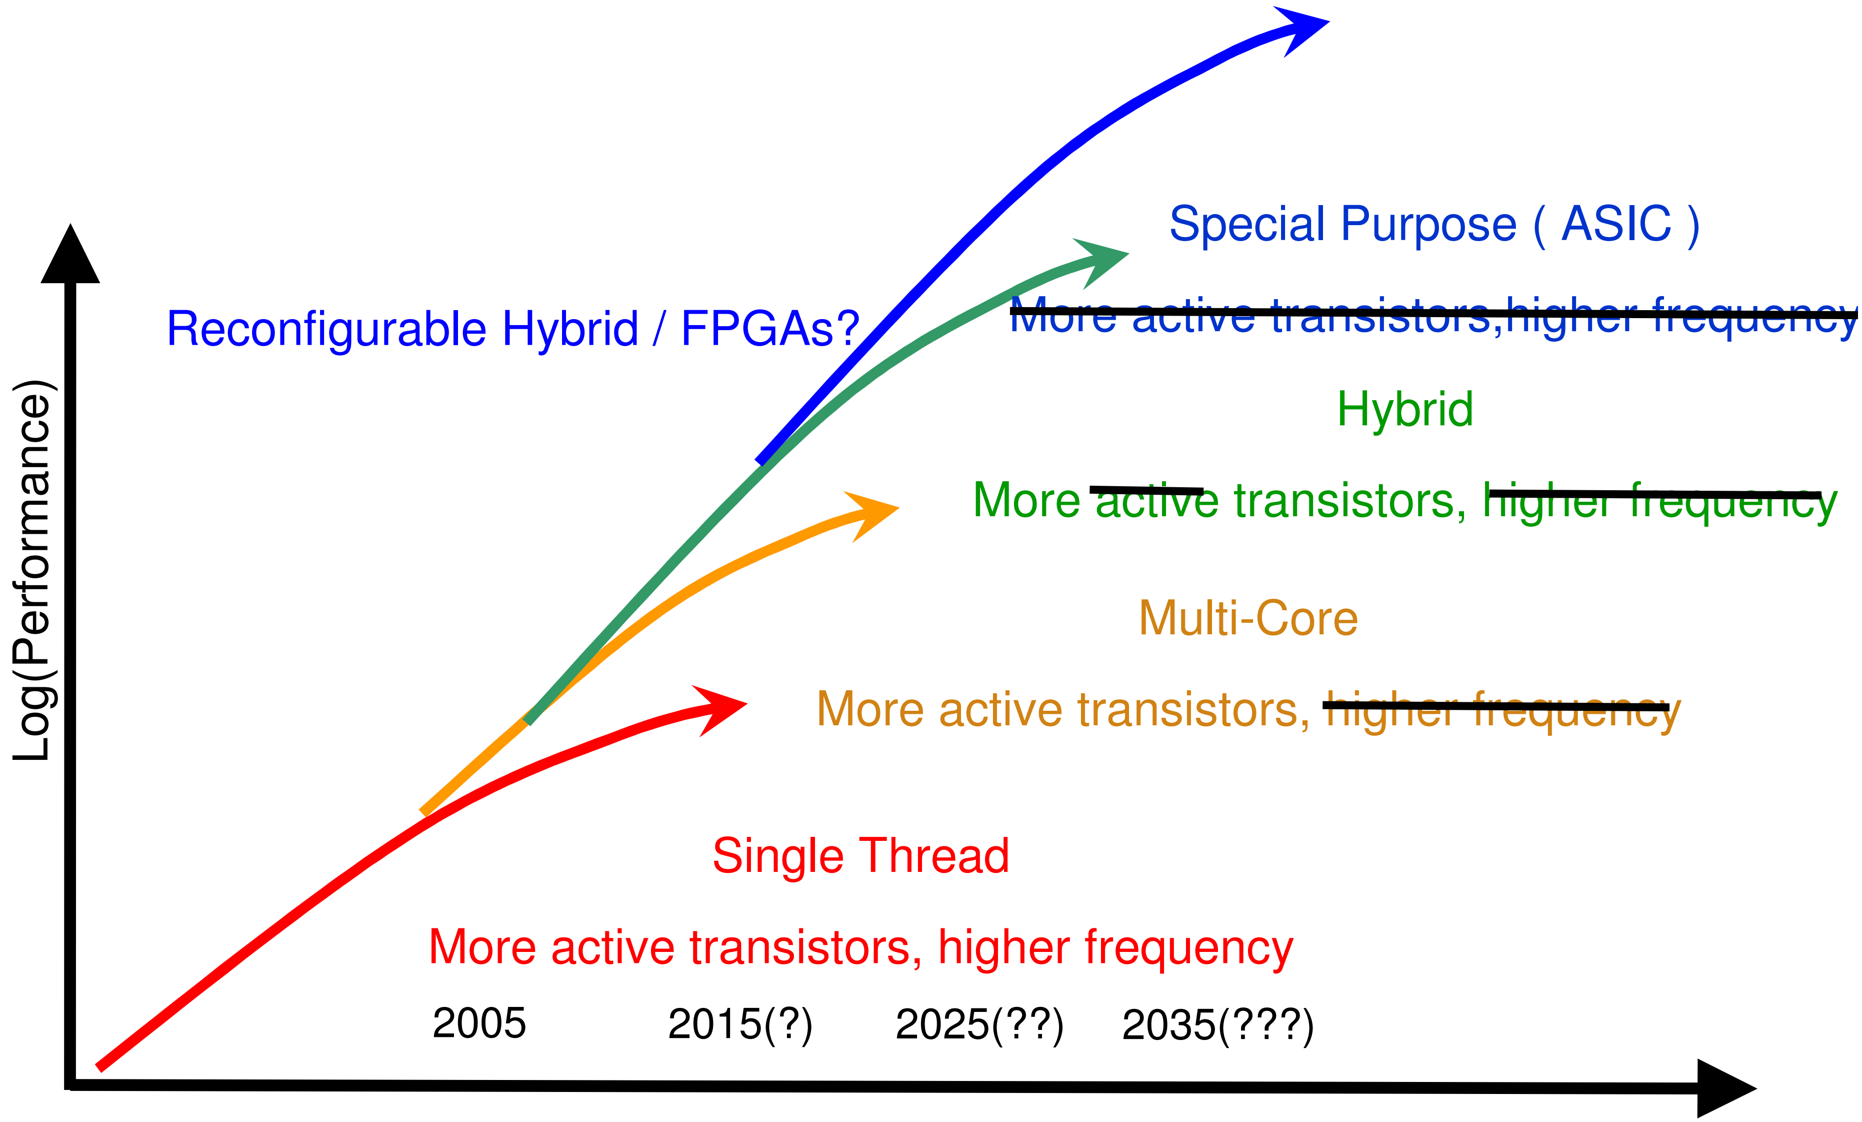

Five Decades of Innovations (Beyond Scaling)

1960s -1970s

S/360 Model 67 first virtualized machine

Air / liquid hybrid cooling technology

1980s

Thermal conduction cooling technology

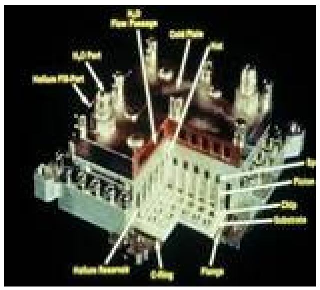

VM virtualization

mid-1990s

CMOS processors

Flat plate conduction cooling technology

late-1990s

Modular refrigeration cooling technology

Copper chip

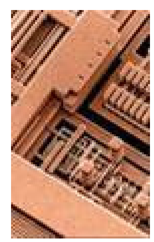

2000s

IBM Energy Efficiency Institute, Austin, TX

High-k metal gates

Airgap

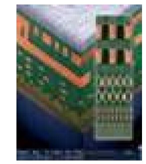

Cell BE processor

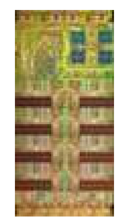

eDRAM

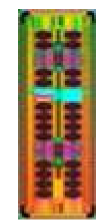

3D chip stacking

Summary

- Technology limits drive fundamental change:
  - First multi-core, then hybrid and eventually special-purpose again?
  - Cell an early example of hybrid
- What is next:
  - Continued Focus on Efficiency
    - Technology developments require it
  - Increasing Focus on Ease of Use
    - Focus on efficiency fundamentals in code
    - Make accelerators "invisible" for most customers
    - Commercial and CE applications, not just HPC
    - Not an easy thing to do
  - Increasing Focus on Standards-Based Programming
    - OpenMP & OpenCL
    - …
  - Continue to Broaden Application Reach for Hybrid Systems

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Special notices (cont.)

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, AIX/L, AIX/L(logo), alphaWorks, AS/400, BladeCenter, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, ClusterProven, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DirectTalk, Domino, DYNIX, DYNIX/ptx, e business(logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, i5/OS, IBM, IBM(logo), ibm.com, IBM Business Partner (logo), Informix, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, OpenPower, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, Passport Advantage, POWERparallel, Power PC 603, Power PC 604, PowerPC, PowerPC(logo), Predictive Failure Analysis, pSeries, PTX, ptx/ADMIN, RETAIN, RISC System/6000, RS/6000, RT Personal Computer, S/390, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SpaceBall, System/390, The Engines of e-business, THINK, Tivoli, Tivoli(logo), Tivoli Management Environment, Tivoli Ready(logo), TME, TotalStorage, TURBOWAYS, VisualAge, WebSphere, xSeries, z/OS, zSeries.

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: Advanced Micro-Partitioning, AIX 5L, AIX PVMe, AS/400e, Chiphopper, Chipkill, Cloudscape, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, DS4000, DS6000, DS8000, e-business(logo), e-business on demand, eServer, Express Middleware, Express Portfolio, Express Servers, Express Servers and Storage, General Purpose File System, GigaProcessor, GPFS, HACMP, HACMP/6000, IBM TotalStorage Proven, IBMLink, IMS, Intelligent Miner, iSeries, Micro-Partitioning, NUMACenter, On Demand Business logo, POWER, PowerExecutive, Power Architecture, Power Everywhere, Power Family, Power PC, PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, pure XML, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, System i, System i5, System p, System p5, System Storage, System z, System z9, S/390 Parallel Enterprise Server, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Virtualization Engine, Visualization Data Explorer, X-Architecture, z/Architecture, z/9.

A full list of U.S. trademarks owned by IBM may be found at: http://www.ibm.com/legal/copytrade.shtml.
The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
UNIX is a registered trademark of The Open Group in the United States, other countries or both.
Linux is a trademark of Linus Torvalds in the United States, other countries or both.
Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.
Intel, Itanium, Pentium are registered tradas and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States, other countries or both.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.
AltiVec is a trademark of Freescale Semiconductor, Inc.
Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.
Other company, product and service names may be trademarks or service marks of others.