The Evolution of Power-Aware, High-Performance Clusters:
From the Datacenter to the Desktop

Wu-chun Feng
feng@lanl.gov

Research & Development in Advanced Network Technology (RADIANT)
Computer & Computational Sciences Division
Los Alamos National Laboratory
University of California

Keynote Address
IEEE Int'l Parallel & Distributed Processing Symposium
Workshop on High-Performance, Power-Aware Computing
4 April 2005

LA-UR-05-2850

Outline

- Motivation & Background
  - Where is High-Performance Computing (HPC)?
  - The Need for Efficiency, Reliability, and Availability
- Supercomputing in Small Spaces (http://sss.lanl.gov)
  - Past: Green Destiny (2001-2002)
    - Architecture & Experimental Results
  - Present: The Evolution of Green Destiny (2003-2004)
    - MegaScale Project: MegaProto
    - Orion Multisystems: DT-12 and DS-96
- Publications & Recognition
- Conclusion

Where is High-Performance Computing?

(Pictures: Thomas Sterling, Caltech & NASA JPL, and Wu Feng, LANL)

We have spent decades focusing on performance, performance, performance (and price/performance).

Where is High-Performance Computing? Top 500 Supercomputer List

- Benchmark
  - LINPACK: Solves a (random) dense system of linear equations in double-precision (64 bits) arithmetic.
    - Introduced by Prof. Jack Dongarra, U. Tennessee
- Evaluation Metric
  - Performance (i.e., Speed)
    - Floating-Operations Per Second (FLOPS)
- Web Site
  - http://www.top500.org

Where is High-Performance Computing? Gordon Bell Awards at SC

- Metrics for Evaluating Supercomputers (or HPC)
  - *Performance (i.e., Speed)*
    - Metric: Floating-Operations Per Second (FLOPS)
    - Example: Japanese Earth Simulator, ASCI Thunder & Q.
  - *Price/Performance → Cost Efficiency*
    - Metric: Acquisition Cost / FLOPS
    - Examples: LANL Space Simulator, VT System X cluster. (In general, Beowulf clusters.)

- Performance & price/performance are important metrics, but …

Where is High-Performance Computing? (Unfortunate) Assumptions

Adapted from David Patterson, UC-Berkeley

- Humans are infallible.
  - No mistakes made during integration, installation, configuration, maintenance, repair, or upgrade.
- Software will eventually be bug free.
- Hardware MTBF is already very large (~100 years between failures) and will continue to increase.
- Acquisition cost is what matters; maintenance costs are irrelevant.

- The above assumptions are even *more* problematic if one looks at current trends in HPC.

Reliability & Availability of Leading-Edge Supercomputers

Systems	CPUs	Reliability & Availability
ASCI Q	**8,192**	**MTBI: 6.5 hrs.** 114 unplanned outages/month. ◆ HW outage sources: storage, CPU, memory.
ASCI White	**8,192**	**MTBF: 5 hrs. (2001) and 40 hrs. (2003).** ◆ HW outage sources: storage, CPU, 3rd-party HW.
NERSC Seaborg	**6,656**	**MTBI: 14 days. MTTR: 3.3 hrs.** ◆ SW is the main outage source. **Availability: 98.74%.**
PSC Lemieux	**3,016**	**MTBI: 9.7 hrs.** **Availability: 98.33%.**
Google	**~15,000**	**20 reboots/day; 2-3% machines replaced/year.** ◆ HW outage sources: storage, memory. **Availability: ~100%.**

MTBI: mean time between interrupts; MTBF: mean time between failures; MTTR: mean time to restore

Source: Daniel A. Reed, UNC

Efficiency of Leading-Edge Supercomputers

- "Performance" and "Price/Performance" Metrics …
  - Lower efficiency, reliability, and availability.
  - Higher operational costs, e.g., admin, maintenance, etc.
- Examples
  - Computational Efficiency
    - Relative to Peak: Actual Performance/Peak Performance
    - Relative to Space: Performance/Sq. Ft.
    - Relative to Power: Performance/Watt
  - Performance: 2000-fold increase (since the Cray C90).
    - Performance/Sq. Ft.: Only 65-fold increase.
    - Performance/Watt: Only 300-fold increase.
  - Massive construction and operational costs associated with powering and cooling.

Ubiquitous Need for Efficiency, Reliability, and Availability

- Requirement: Near-100% *availability* with *efficient* and *reliable* resource usage.
  - E-commerce, enterprise apps, online services, ISPs, data and HPC centers supporting R&D.
- Problems

  Source: David Patterson, UC-Berkeley
  - Frequency of Service Outages
    - 65% of IT managers report that their websites were unavailable to customers over a 6-month period.
  - Cost of Service Outages
    - NYC stockbroker: **$ 6,500,000/hour**
    - Ebay (22 hours): **$ 225,000/hour**
    - Amazon.com: **$ 180,000/hour**
    - Social Effects: negative press, loss of customers who "click over" to competitor (e.g., Google vs. Ask Jeeves)

Where is High-Performance Computing?

(Pictures: Thomas Sterling, Caltech & NASA JPL and Wu Feng, LANL)

Efficiency, reliability, and availability will be *the* key issues of this decade.

Outline

- Motivation & Background
  - Where is High-Performance Computing (HPC)?
  - The Need for Efficiency, Reliability, and Availability
- Supercomputing in Small Spaces (http://sss.lanl.gov)
  - Past: Green Destiny (2001-2002)
    - Architecture & Experimental Results
  - Present: The Evolution of Green Destiny (2003-2004)
    - MegaScale Project: MegaProto
    - Orion Multisystems: DT-12 and DS-96
- Publications & Recognition
- Conclusion

Supercomputing in Small Spaces:
Efficiency, Reliability, and Availibility via Power-Aware HPC

- Goal
  - Improve efficiency, reliability, and availability (ERA) in large-scale computing systems.
    - Sacrifice a bit of raw performance (potentially).
    - Improve overall system throughput as the system will "always" be available, i.e., effectively no downtime, no HW failures, etc.
  - Reduce the total cost of ownership (TCO). Another talk …
- Crude Analogy
  - Formula One Race Car: Wins raw performance but reliability is so poor that it requires frequent maintenance. Throughput low.
  - Toyota Camry: Loses raw performance but high reliability results in high throughput (i.e., miles driven → answers/month).

COMPUTER & COMPUTATIONAL SCIENCES

How to Improve Efficiency, Reliability & Availability?

- Observation
  - High power α high temperature α low reliability.
  - Arrhenius' Equation*

    (circa 1890s in chemistry → circa 1980s in computer & defense industries)
    - As temperature increases by 10° C …
      - The failure rate of a system doubles.
    - Twenty years of unpublished empirical data .

* The time to failure is a function of $e^{-Ea/kT}$ where Ea = activation energy of the failure mechanism being accelerated, k = Boltzmann's constant, and T = absolute temperature

Processor	Clock Freq.	Voltage	Peak Temp.**
Intel Pentium III-M	500 MHz	1.6 V	252° F (122° C)
Transmeta Crusoe TM5600	600 MHz	1.6 V	147° F (64° C)

Moore's Law for Power

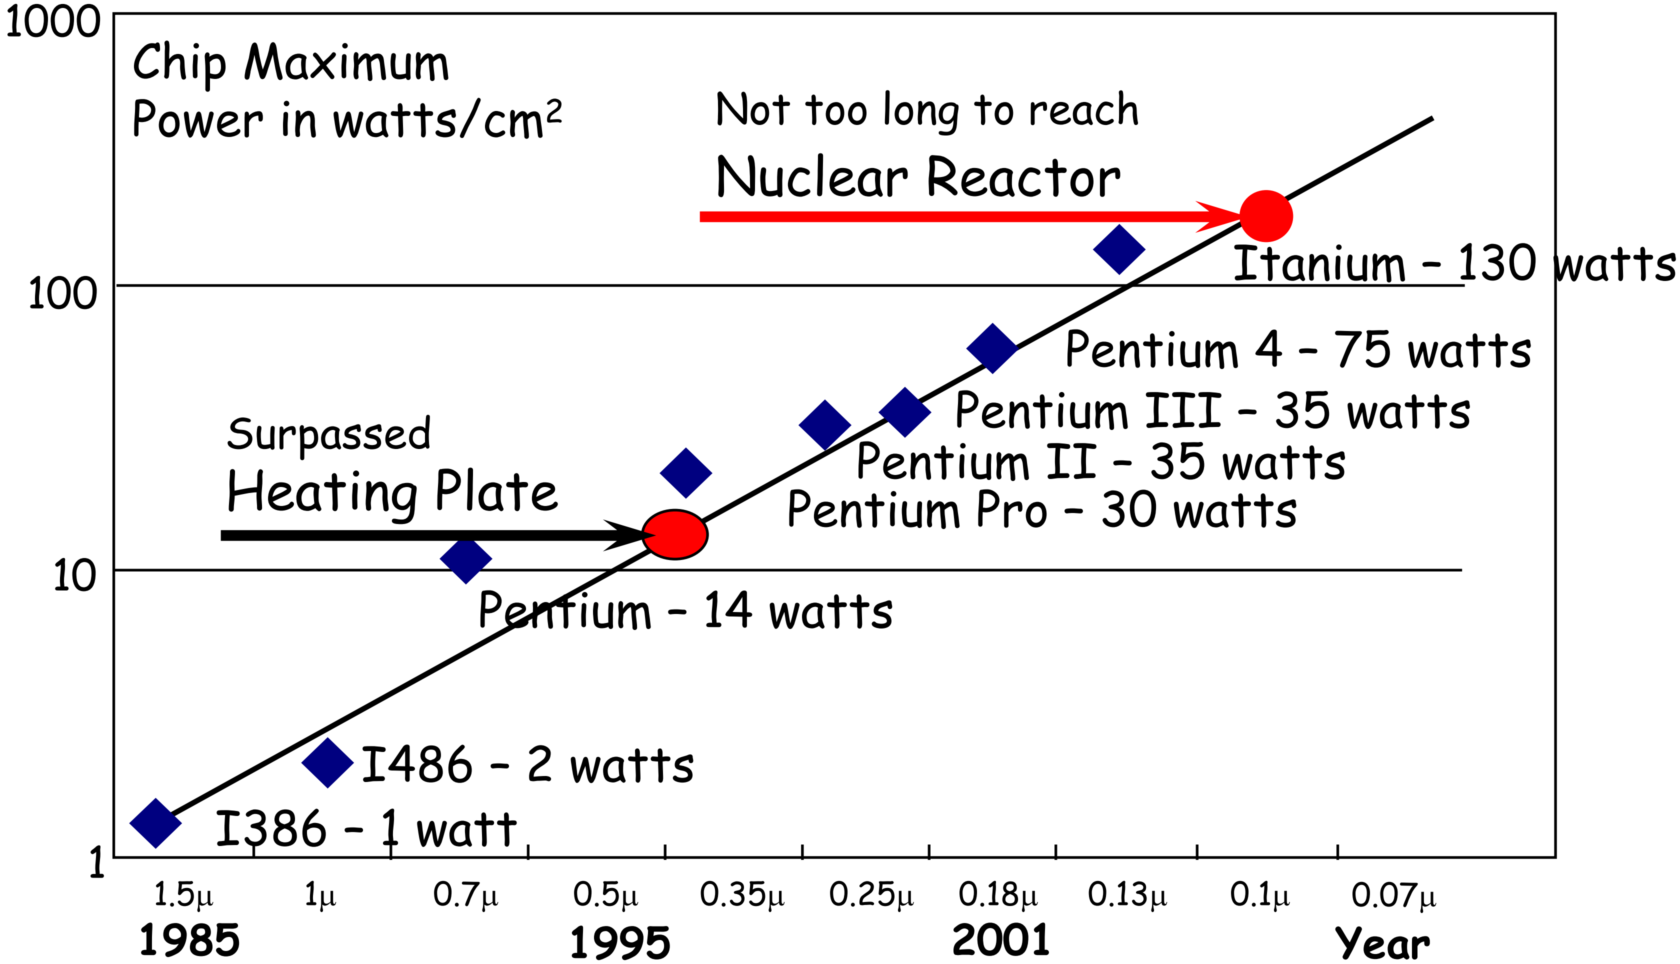

Source: Fred Pollack, Intel. New Microprocessor Challenges in the Coming Generations of CMOS Technologies, MICRO32 and Transmeta

Moore's Law for Power

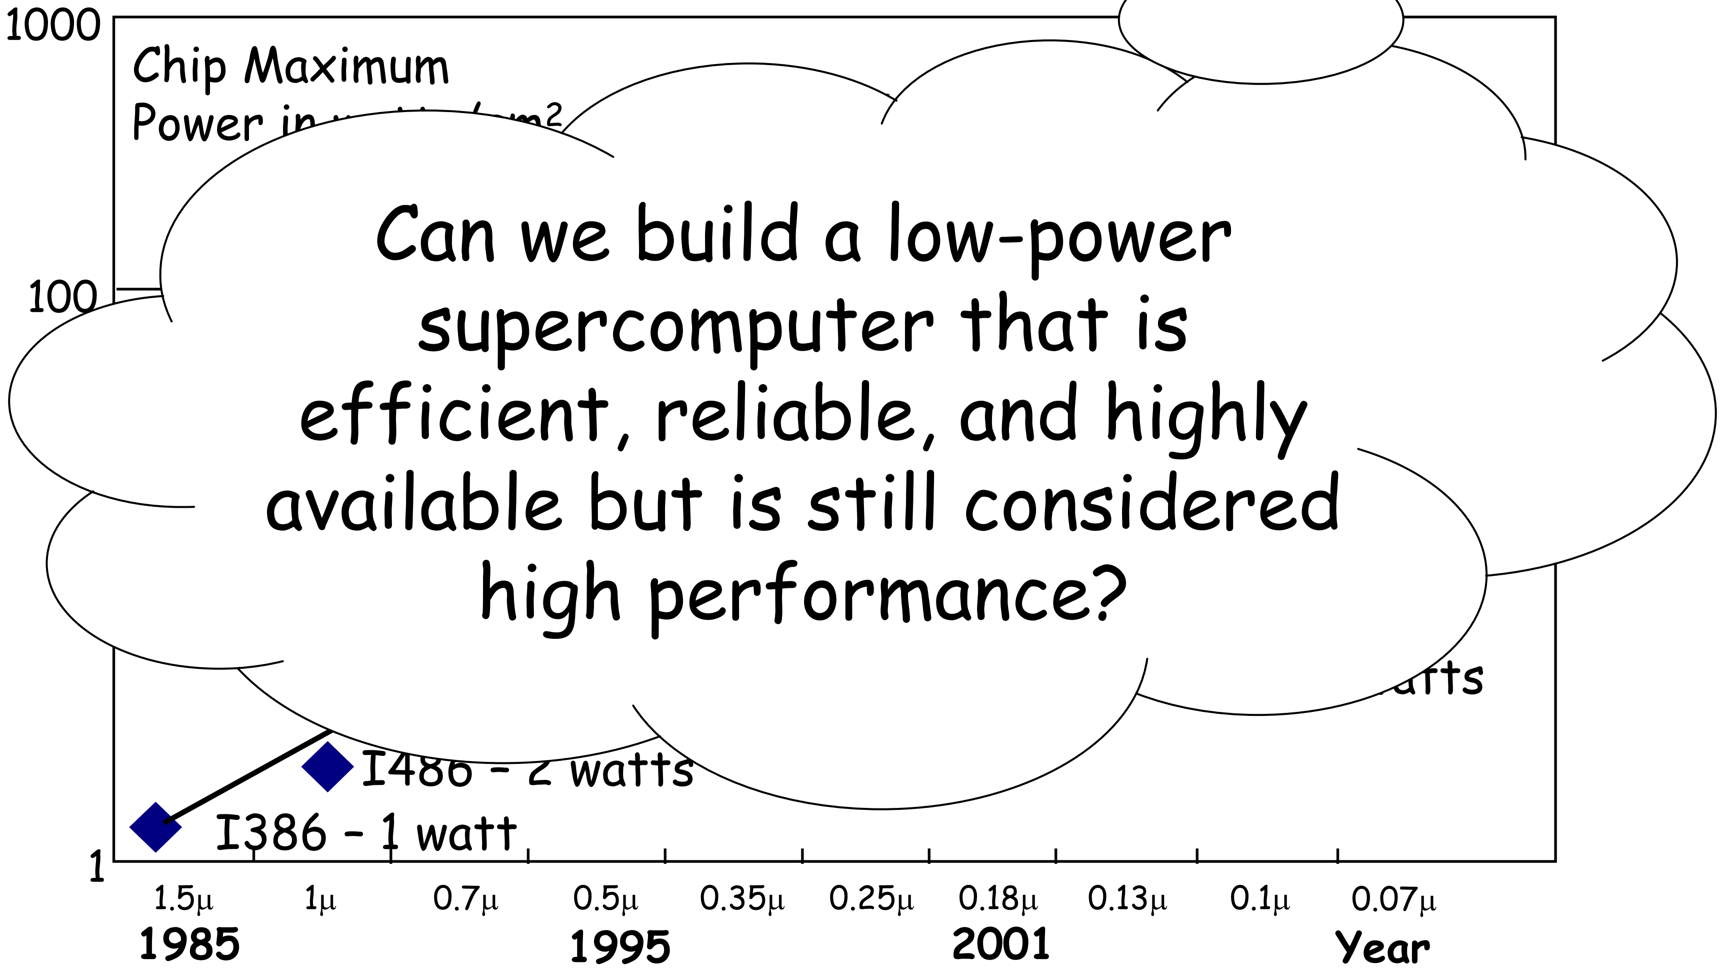

Can we build a low-power supercomputer that is efficient, reliable, and highly available but is still considered high performance?

Source: Fred Pollack, Intel. New Microprocessor Challenges in the Coming Generations of CMOS Technologies, MICRO32 and Transmeta

Moore's Law for Power

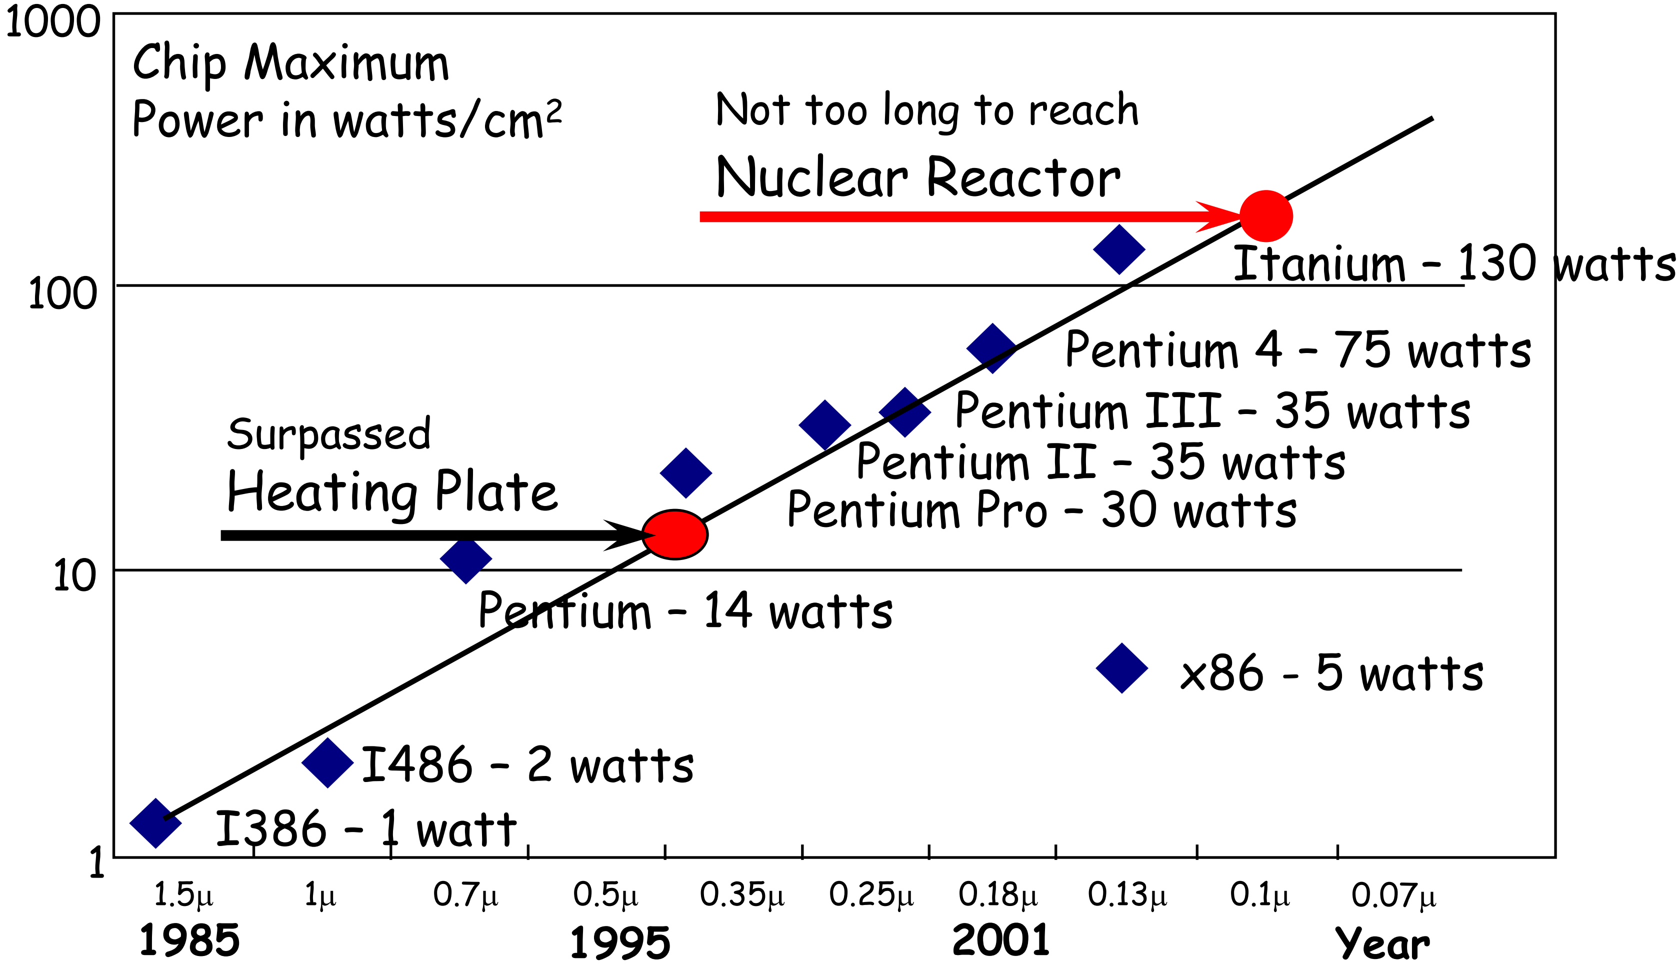

Source: Fred Pollack, Intel. New Microprocessor Challenges in the Coming Generations of CMOS Technologies, MICRO32 and Transmeta

Transmeta TM5600 CPU: VLIW + CMS

- VLIW Engine
  - Up to four-way issue
    - In-order execution only.
  - Two integer units
  - Floating-point unit
  - Memory unit
  - Branch unit

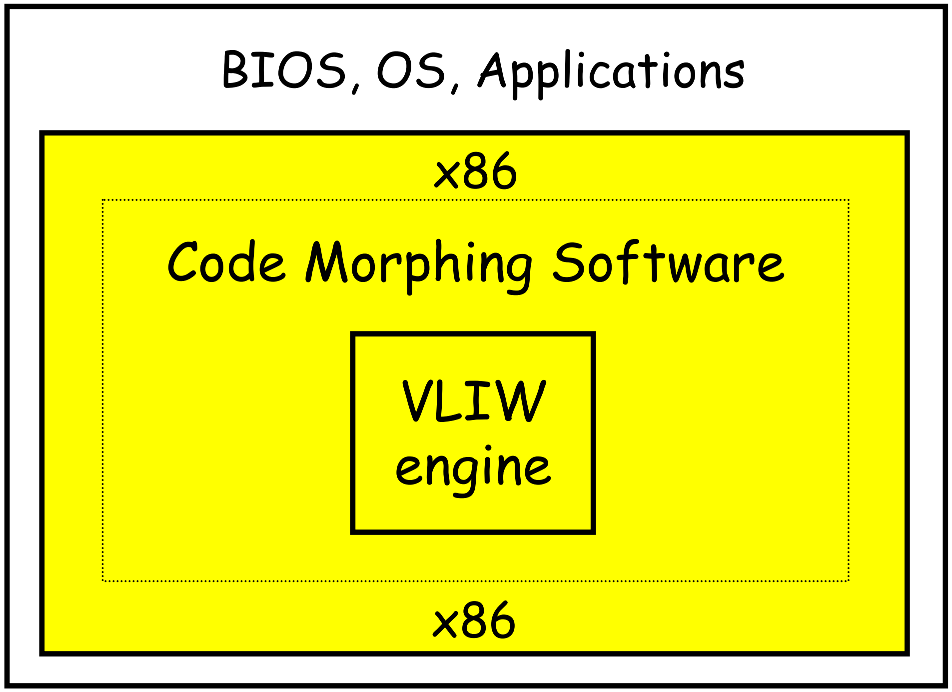

- VLIW Transistor Count ("Anti-Moore's Law")
  - ~ 25% of Intel PIII → ~ 7x less power consumption
  - Less power → lower "on-die" temp. → better reliability & availability

Transmeta TM5x00 CMS

- Code-Morphing Software (CMS)
  - Provides compatibility by dynamically "morphing" x86 instructions into simple VLIW instructions.
  - Learns and improves with time, i.e., iterative execution.
- High-Performance Code-Morphing Software (HP-CMS)
  - Results (circa 2001)
    - *Optimized to improve floating-pt. performance by 50%.*
    - *1-GHz Transmeta performs as well as a 1.2-GHz PIII-M.*
  - How?

Transmeta TM5x00 Comparison

Intel P4	MEM	MEM	ALU	ALU	FPU	SSE	SSE	Br
Transmeta TM5x00	MEM		ALU		FPU			Br

- Previous-generation Transmeta TM5800 + HP-CMS
  - Performs better than an Intel PIII over iterative scientific codes on a clock-for-clock-cycle basis.
  - Performs only *twice* as slow as the fastest CPU (at the time) rather than three times as slow.
- Efficeon, the current-generation CPU from Transmeta, rectifies the above mismatch in functional units.

Low-Power Network Switches

- WWP LE-410: 16 ports of Gigabit Ethernet
- WWP LE-210: 24 ports of Fast Ethernet via RJ-21s
- (Avg.) Power Dissipation / Port: A few watts.

"Green Destiny" Bladed Beowulf
(circa April 2002)

- A 240-Node Beowulf Cluster in Five Sq. Ft.
- Each Node
  - 667-MHz Transmeta TM5600 CPU w/ Linux 2.4.x
    - Upgraded to 1-GHz Transmeta TM5800 CPUs
  - 640-MB RAM, 20-GB HD, 100-Mb/s Ethernet (up to 3 interfaces)
- Total
  - 160 Gflops peak (240 Gflops with upgrade)
    - LINPACK: 101 Gflops in March 2003.
  - 150 GB of RAM (expandable to 276 GB)
  - 4.8 TB of storage (expandable to 38.4 TB)
  - *Power Consumption: Only 3.2 kW.*
- Reliability & Availability
  - *No unscheduled failures in 24 months.*

2003 WINNER R&D 100

Wu FENG
feng@lanl.gov

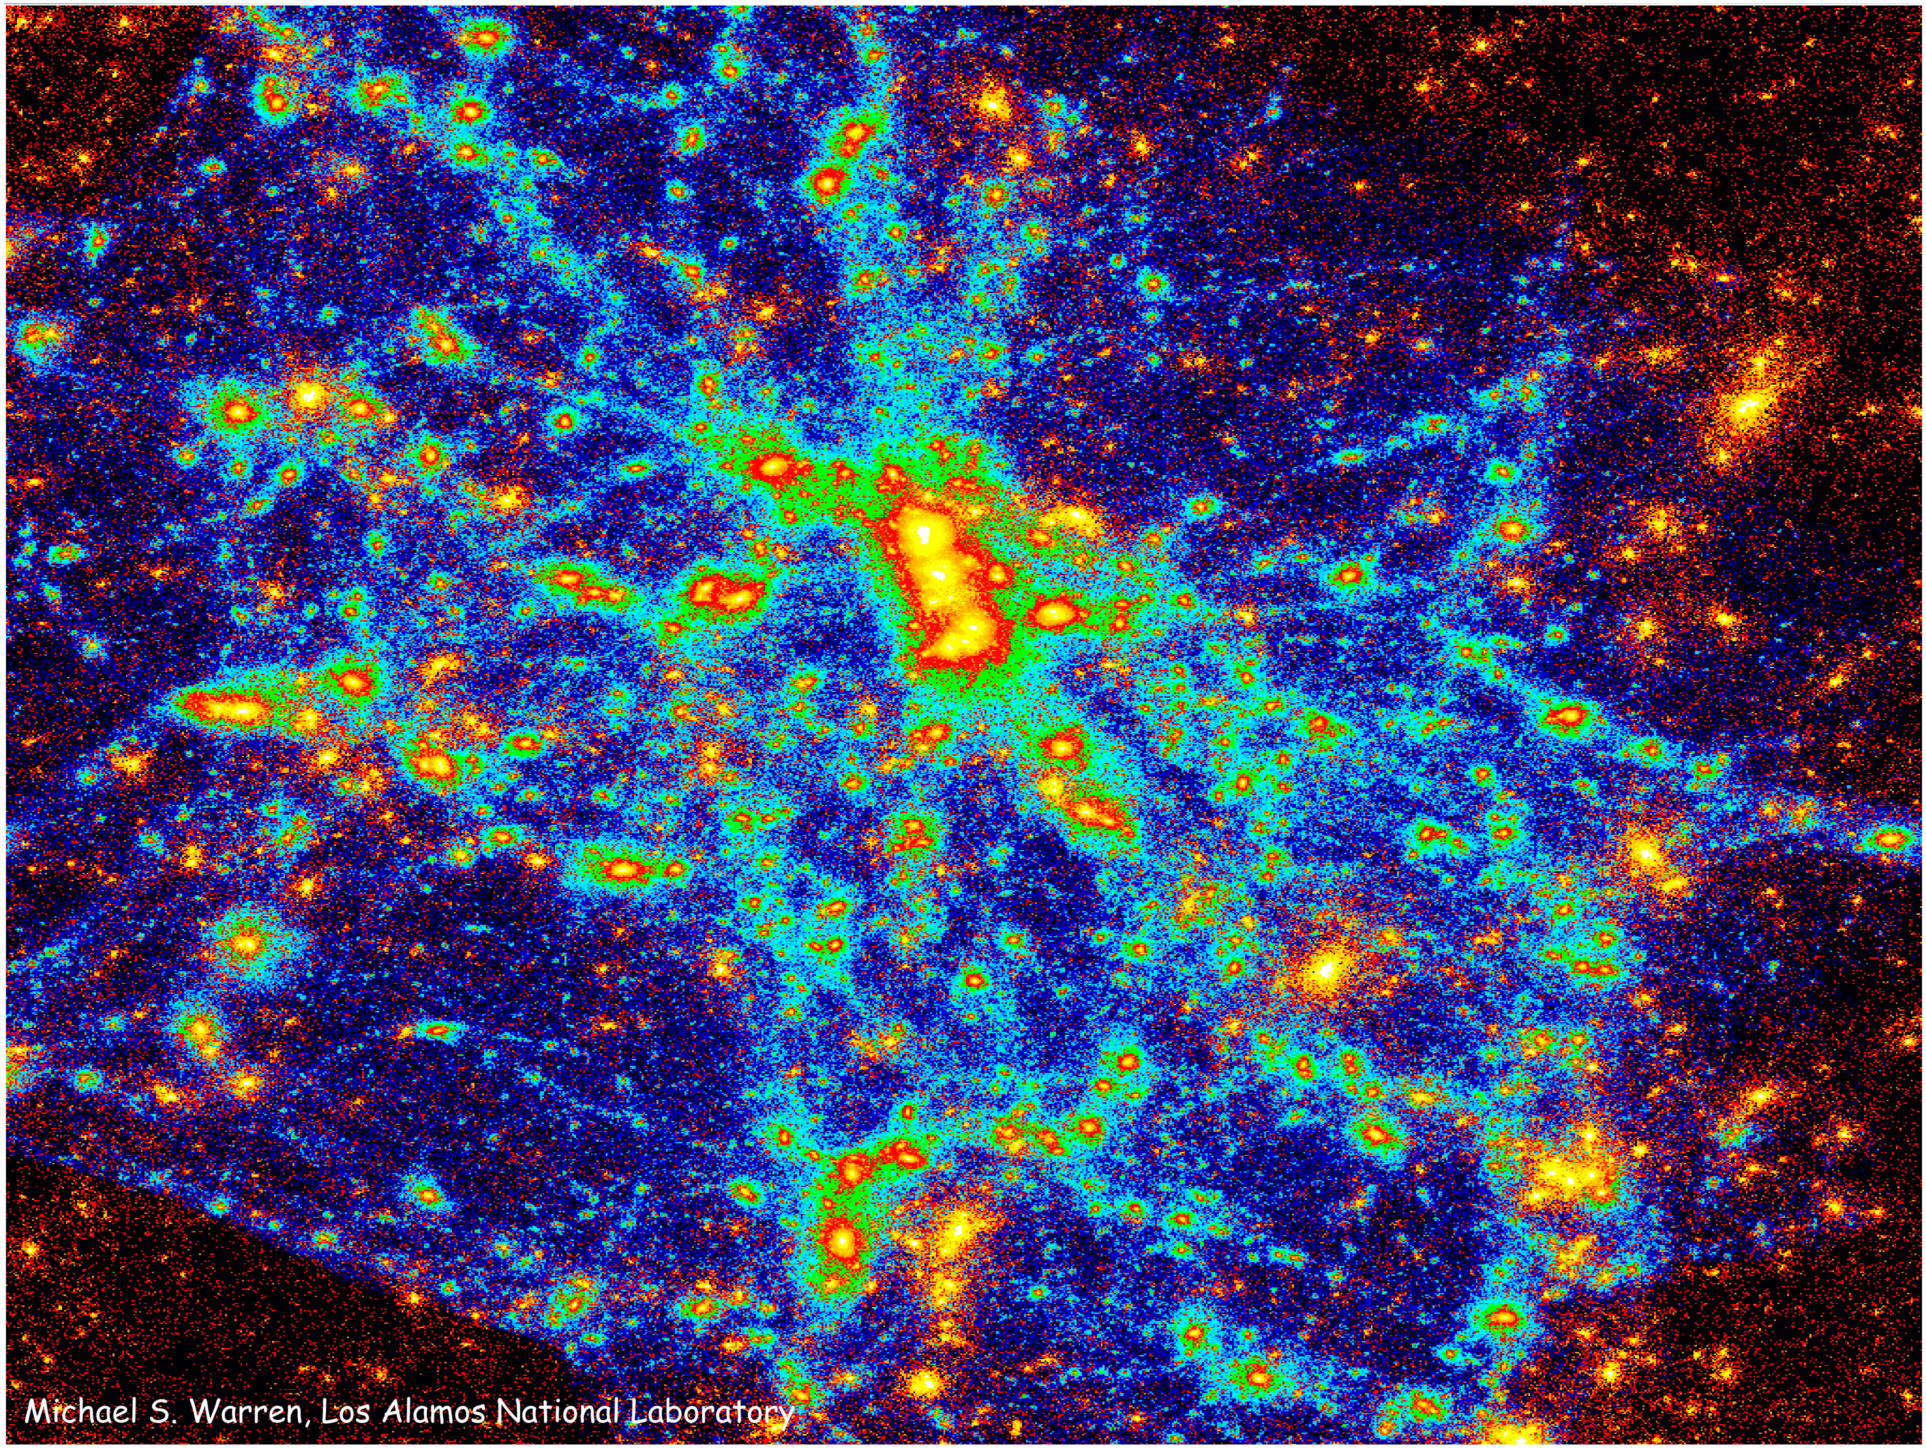
Michael S. Warren, Los Alamos National Laboratory

Parallel Computing Platforms ("Apples-to-Oranges" Comparison)

- Avalon (1996)
  - 140-CPU *Traditional Beowulf Cluster*
- ASCI Red (1996)
  - 9632-CPU *MPP*
- ASCI White (2000)
  - 512-Node (8192-CPU) *Cluster of SMPs*
- Green Destiny (2002)
  - 240-CPU *Bladed Beowulf Cluster*

Parallel Computing Platforms Running the N-body Code

Machine	Avalon Beowulf	ASCI Red	ASCI White	Green Destiny+
Year	1996	1996	2000	2002
Performance (Gflops)	18	600	2500	*58*
Area (ft^2)	120	1600	9920	5
Power (kW)	18	1200	2000	5
DRAM (GB)	36	585	6200	150
Disk (TB)	0.4	2.0	160.0	4.8
DRAM density (MB/ft^2)	300	366	625	30000
Disk density (GB/ft^2)	3.3	1.3	16.1	960.0
Perf/Space ($Mflops/ft^2$)	150	375	252	11600
Perf/Power (Mflops/watt)	1.0	0.5	1.3	11.6

Source: Michael S. Warren, Los Alamos National Laboratory

COMPUTER & COMPUTATIONAL SCIENCES

Parallel Computing Platforms Running the N-body Code

Machine	Avalon Beowulf	ASCI Red	ASCI White	Green Destiny+
Year	1996	1996	2000	2002
Performance (Gflops)	18	600	2500	*58*
Area (ft^2)	120	1600	9920	5
Power (kW)	18	1200	2000	5
DRAM (GB)	36	585	6200	150
Disk (TB)	0.4	2.0	160.0	4.8
DRAM density (MB/ft^2)	300	366	625	30000
Disk density (GB/ft^2)	3.3	1.3	16.1	960.0
Perf/Space ($Mflops/ft^2$)	150	375	252	11600
Perf/Power (Mflops/watt)	1.0	0.5	1.3	11.6

Source: Michael S. Warren, Los Alamos National Laboratory

Efficiency, Reliability, and Availability for …

- Green Destiny+
  - Computational Efficiency
    - Relative to Space: Performance/Sq. Ft.

      Up to 80x better.
    - Relative to Power: Performance/Watt

      Up to 25x better.
  - Reliability
    - MTBF: Mean Time Between Failures

      "Infinite"
  - Availability
    - Percentage of time that resources are available for HPC.

      Nearly 100%.

Outline

- Motivation & Background
  - Where is High-Performance Computing (HPC)?
  - The Need for Efficiency, Reliability, and Availability
- Supercomputing in Small Spaces (http://sss.lanl.gov)
  - Past: Green Destiny (2001-2002)
    - Architecture & Experimental Results
  - Present: The Evolution of Green Destiny (2003-2004)
    - MegaScale Project: MegaProto
    - Orion Multisystems: DT-12 and DS-96
- Publications & Recognition
- Conclusion

Inter-University Project: MegaScale

http://www.para.tutics.tut.ac.jp/megascale/

Univ. of Tsukuba Booth @ SC2004, Nov. 2004.

The Road from Green Destiny to Orion Multisystems

- Trends in High-Performance Computing
  - Rise of cluster-based high-performance computers.
    - Price/performance advantage of using "commodity PCs" as cluster nodes (Beowulf: 1993-1994.)
    - Different flavors: "homebrew" vs. "custom"

Architectures / Systems

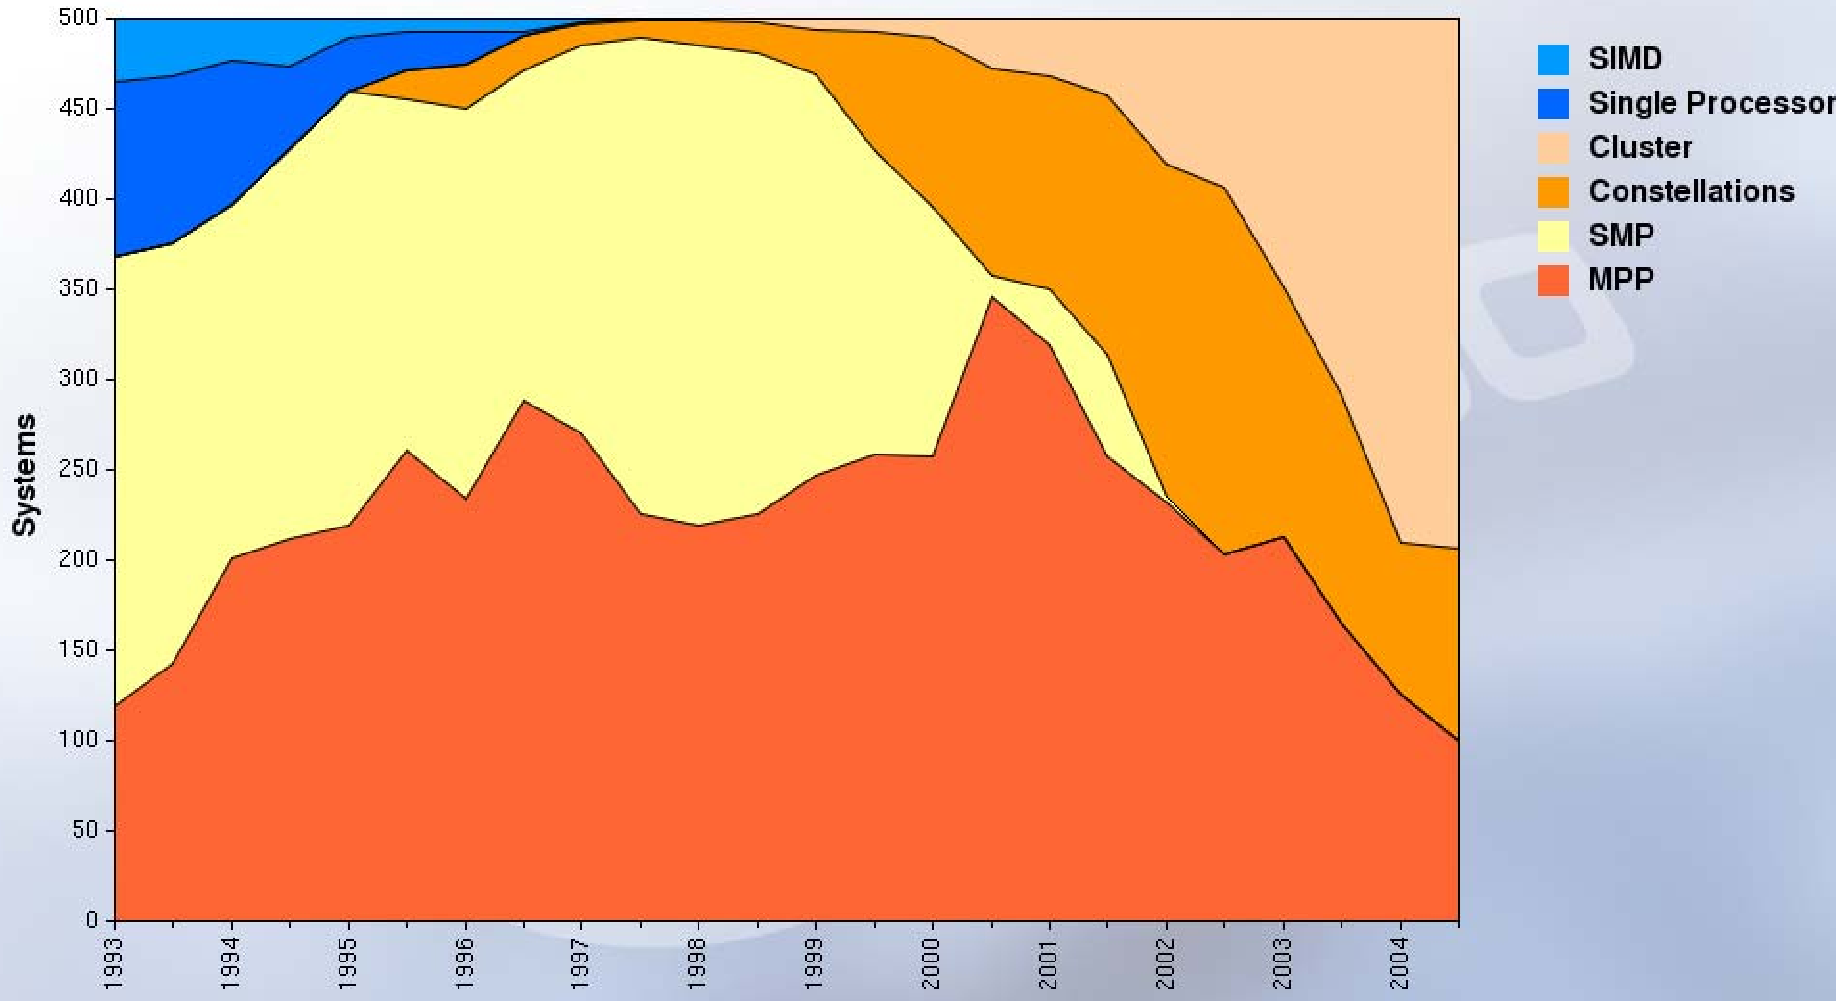

The Road from Green Destiny to Orion Multisystems

- Trends in High-Performance Computing
  - Rise of cluster-based high-performance computers.
    - Price/performance advantage of using "commodity PCs" as cluster nodes (Beowulf: 1993-1994.)
    - Different flavors: "homebrew" vs. "custom"
  - Maturity of open-source cluster software.
    - Emergence of Linux and MPI as parallel programming APIs.
  - Rapid decline of the traditional workstation.
    - Replacement of workstation with a PC.
    - 1000-fold (and increasing) performance gap with respect to the supercomputer.
    - Still a desperate need for HPC in workstation form.

Evolution of Workstations: Performance Trends

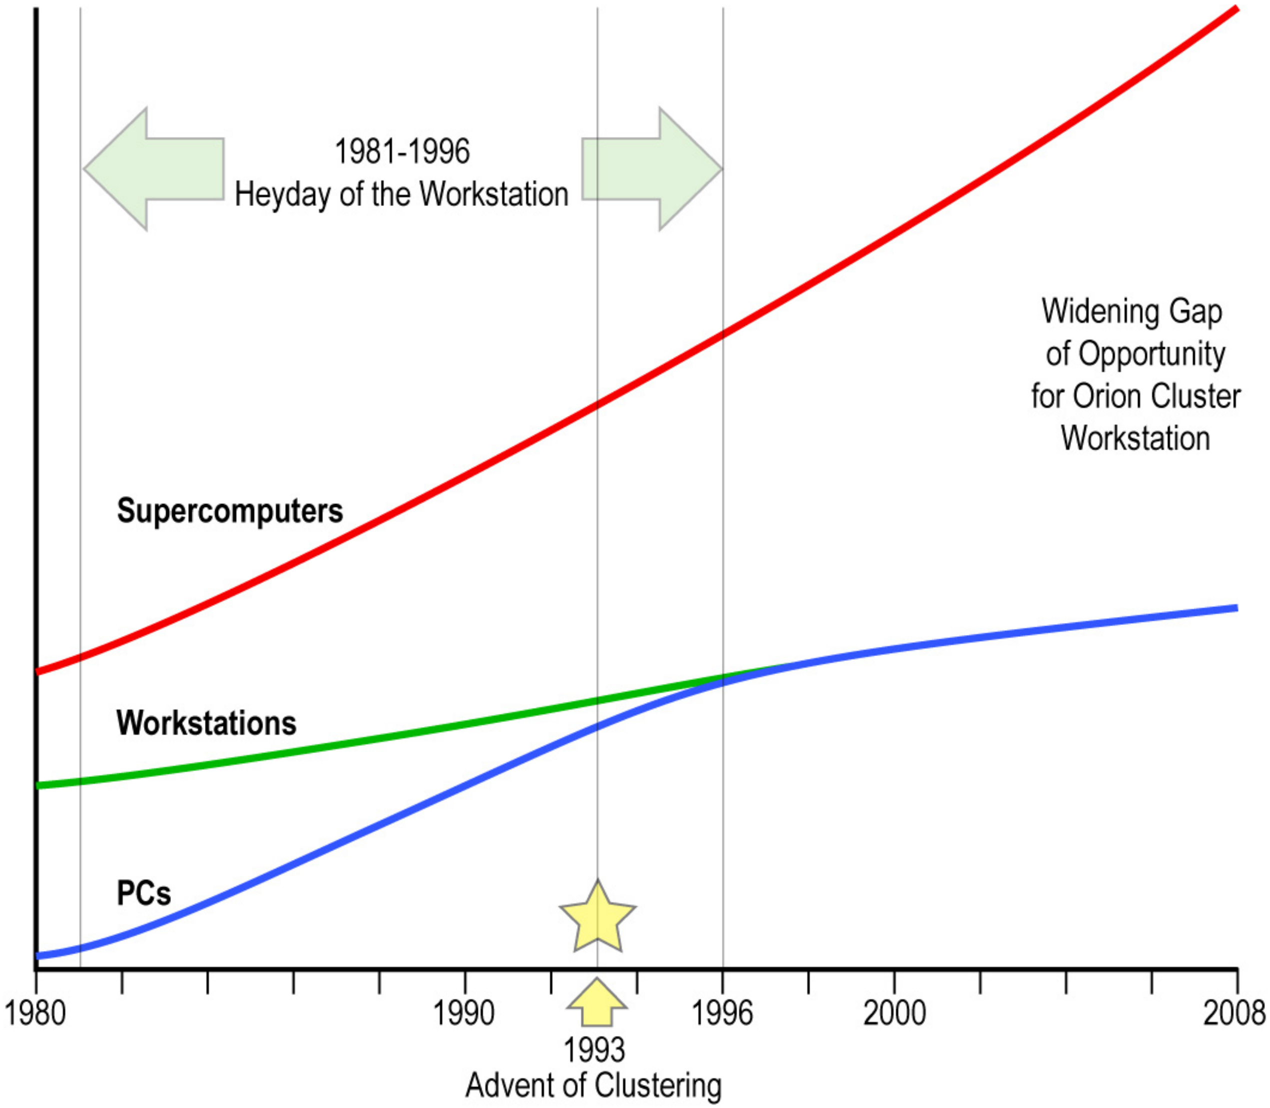

- PC performance caught up with workstations
  - PC OSes: NT and Linux
- A large gap has opened between PCs and super-computers
  - 3 Gflops vs. 3 Tflops

Source: Orion Multisystems, Inc.

Need: A Cluster Workstation

- Specifications
  - Desktop or deskside box with cluster inside
  - A cluster product – not an assembly
  - Scalable computation, graphics, and storage
  - Meets power limits of office or laboratory
- Reality of (Homebrew) Clusters
  - Ad-hoc, custom-built collections of boxes
  - Hard for an individual to get exclusive access (or even share access)
  - Power-, space-, and cooling-intensive
  - IT support required

Source: Orion Multisystems, Inc.

Wu FENG
feng@lanl.gov

Why a Cluster Workstation?

- Personal Resource
  - No scheduling conflicts or long queues.
  - Application debugging with scalability at your desktop
  - Redundancy possibilities (eliminate downtime)
- Improvement of Datacenter Efficiency
  - Off-load “repeat offender” jobs
  - Enable developers to debug their own code on their own system
  - Manage expectations
  - Reduce job turnaround time

Source: Orion Multisystems, Inc.

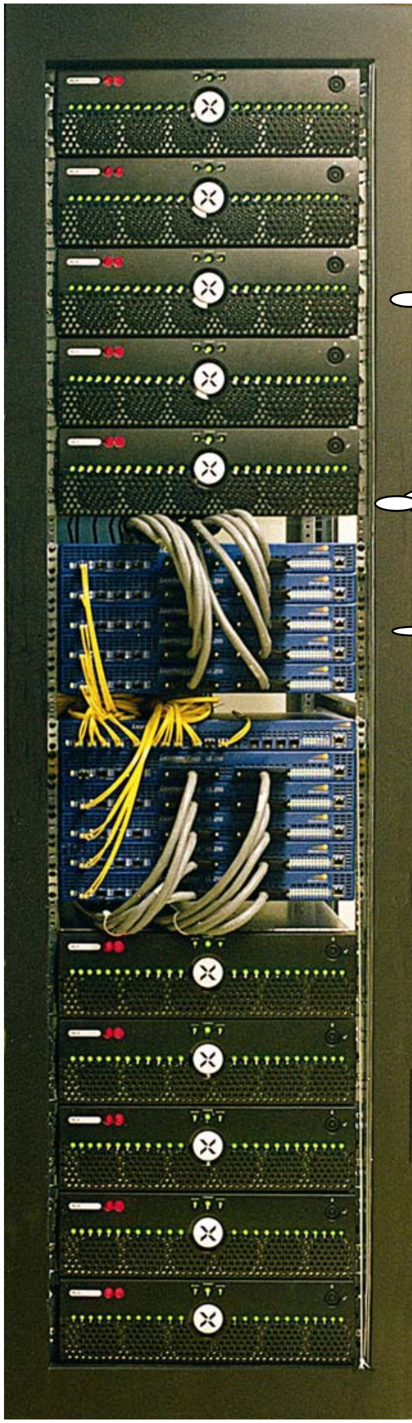

Cluster Technology

Low-Power Systems Design

Linux

But in the form factor of a workstation …
a *cluster workstation*

http://www.orionmultisystems.com

- LINPACK Performance
  - 13.80 Gflops
- Footprint
  - 3 sq. ft. (24" x 18")
  - 1 cu. ft. (24" x 4" x 18")
- Power Consumption
  - 170 watts at load
- How does this compare with a traditional desktop?

ORION DT-12 DESKTOP CLUSTER WORKSTATION

Imagine a 36 Gflop cluster **on your desk!**

12 Nodes
in a single computer

24 GBytes
memory capacity

36 Gflops
peak processing power

1 TByte
internal storage

DESIGNED FOR THE INDIVIDUAL

The Orion DT-12 cluster workstation is a fully integrated, completely self-contained, personal workstation based on the best of today's cluster technologies. Designed to be an affordable individual resource it is capable of 36 Gflops peak performance (18 Gflops sustained) with models starting at under $10k.

The Orion DT-12 cluster workstation provides supercomputer performance for the engineering, scientific, financial and creative professionals who need to solve computationally complex problems without waiting in the queue of the back-room cluster.

FASTER SOFTWARE DEVELOPMENT

The Orion DT-12 cluster workstation is the perfect platform for developers writing (and deploying) cluster software packages. It comes with cluster software development tools pre-installed, including libraries and a parallel compiler that allows you to spread one multiple-file compile to all the nodes in the system. Also included is a suite of system monitoring and management software.

NO ASSEMBLY REQUIRED

Orion workstations are designed from the ground up as a single computer. The entire system boots with the push of a button and has the ergonomics and ease of use of a personal computer. The modular, design allows for flexible configurations and scalability by stacking up to 4 systems as one 48 node cluster.

PRESERVE SOFTWARE INVESTMENTS

Orion workstations are built around industry standards for clustering: x86 processors, Ethernet, the Linux operating system and standard parallel programming libraries, including MPI, PVM and SGE. Existing Linux cluster applications run without modification.

PERFORMANCE AND FEATURES

The Orion DT-12 is a cluster of 12 x86-compatible nodes linked by a switched Gigabit Ethernet fabric. The cluster operates as a single computer with a single on-off switch and a single system image rapid boot sequence, which allows the entire system to boot in less than 90 seconds.

The Orion DT-12 cluster workstation is highly efficient, consuming a maximum of 220 Watts of power under peak load—about the same as an average desktop PC. It operates quietly, plugs into a standard 110V 15A wall socket and fits unobtrusively on a desk or lab bench.

ORION DS-96 DESKSIDE CLUSTER WORKSTATION

http://www.orionmultisystems.com

Imagine a 300 Gflop cluster...
under your desk.

96 Nodes
in a single computer

300 Gflops
peak processing power

192 GBytes
memory capacity

9.6 TBytes
internal storage

INCREASE YOUR PRODUCTIVITY
The Orion DS-96 cluster workstation is the highest performance general-purpose computing platform that can be plugged into a standard wall outlet and operated in an office or laboratory environment.

PRESERVE SOFTWARE INVESTMENTS
Orion workstations are built around industry standards for clustering: x86 processors, the Linux operating system and standard parallel programming libraries, including MPI, PVM and SGE. Your existing Linux cluster software applications can run without modification.

NO ASSEMBLY REQUIRED
Orion workstations are designed from the ground up as a single computer. The entire system boots with the push of a button and has the ergonomics and ease of use of a personal computer. Modular, solid state design allows for flexible configurations and scalability.

PERFORMANCE AND FEATURES
The Orion DS-96 cluster workstation is a fully integrated, completely self-contained personal workstation based on the best of today's cluster technologies and commodity components. Designed to be an individual or departmental resource, it is capable of 300 Gflops peak performance (150 Gflops sustained). The DS-96 is also highly efficient, consuming a maximum of 1500 Watts of power under peak load. It operates quietly, plugs into a standard 110V 15A wall socket, and fits unobtrusively beneath a desk or lab bench.

The DS-96 is a cluster of 96 x86-compatible nodes linked by an integrated Gigabit Ethernet fabric. The cluster operates as a single computer, with a single on-off switch, and a single-system-image rapid boot sequence which allows the entire system to boot in less than 2 minutes. The DS-96 comes with standard Linux and drivers pre-installed, including an optimized MPI message-passing library. Also included is a suite of cluster software development tools, system monitoring and system management software.

Recall ....
GD: 101 Gflops

- LINPACK Performance
  - 110.8 Gflops
- Footprint
  - 3 sq. ft. (17" x 25")
  - 6 cu. ft. (17" x 25" x 25")
- Power Consumption
  - 1580 watts at load
- Road to Tflop?
  - 10 DS-96s → ~ 1 Tflop LINPACK

Parallel Computing Platforms Running LINPACK

Machine	ASCI Red	ASCI White	Green Destiny+
Year	1996	2000	2002
Performance (Gflops)	2379	7226	101.0
Area (ft^2)	1600	9920	5
Power (kW)	1200	2000	5
DRAM (GB)	585	6200	150
Disk (TB)	2.0	160.0	4.8
DRAM density (MB/ft^2)	366	625	30000
Disk density (GB/ft^2)	1	16	960
Perf/Space ($Mflops/ft^2$)	1487	728	20202
Perf/Power (Mflops/watt)	2	4	20

Parallel Computing Platforms Running LINPACK

Machine	ASCI Red	ASCI White	Green Destiny+	Orion DS-96
Year	1996	2000	2002	2005
Performance (Gflops)	2379	7226	101.0	110.8
Area (ft^2)	1600	9920	5	2.95
Power (kW)	1200	2000	5	1.58
DRAM (GB)	585	6200	150	96
Disk (TB)	2.0	160.0	4.8	7.68
DRAM density (MB/ft^2)	366	625	30000	32542
Disk density (GB/ft^2)	1	16	960	2603
Perf/Space ($Mflops/ft^2$)	1487	728	20202	37559
Perf/Power (Mflops/watt)	2	4	20	70

What's Inside?

Orion Multisystems' Workstation Architecture

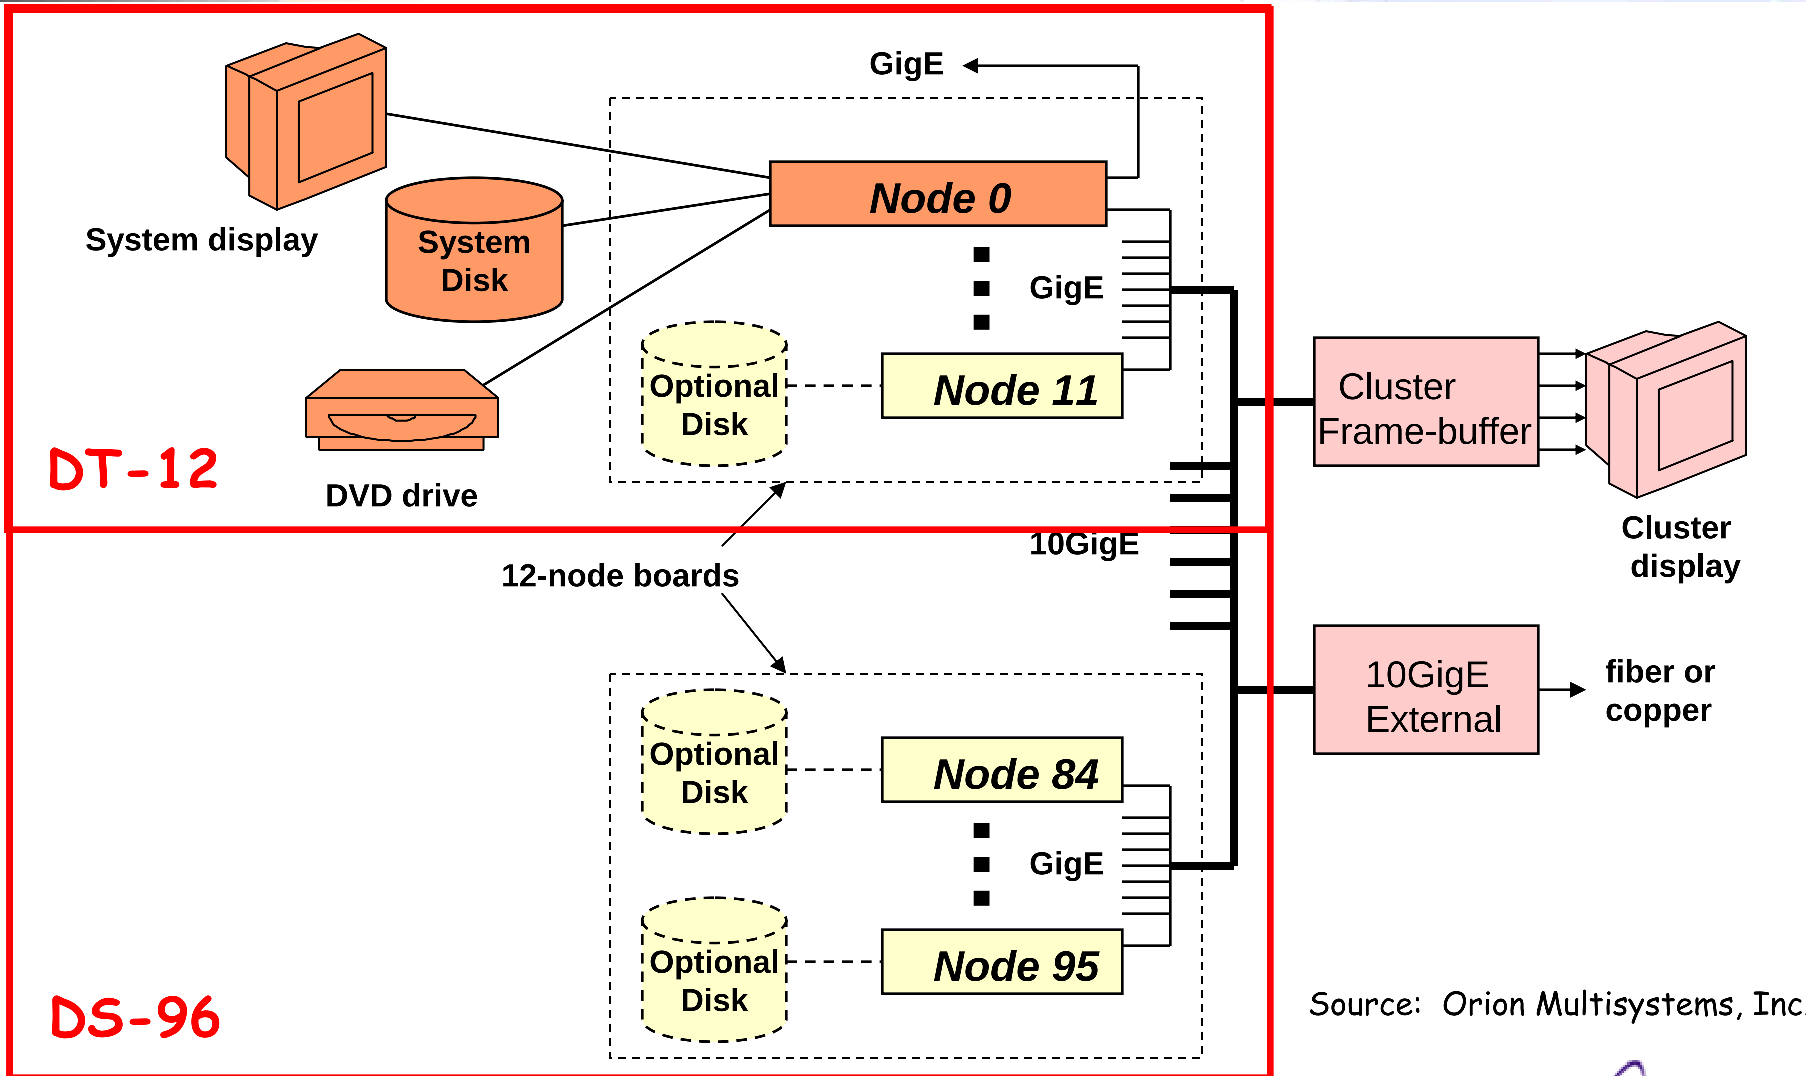

Source: Orion Multisystems, Inc.

Orion Multisystems: Key Hardware Features

- Standard Low-Power Components
  - Low-power, high performance/watt x86 processor
  - Low-power mobile disk-drive per node for distrib. storage
- Low-Power Ethernet Fabric
  - Low-voltage signaling, no PHYs, no cables
  - On-board GigE between nodes
  - On-midplane 10GigE between boards
- Private Monitoring and Control Bus
  - Out-of-band cluster-wide management
  - Synchronized high-resolution hardware clock
- Integrated Video
  - ATi Mobility Radeon 9000 – 64MB integrated DDR

Orion Multisystems: Software Package

- Basic Operating System
  - Fedora Linux Core 2, Version 2.6.6 kernel with Orion drivers; 2.4 kernel support
  - Low-latency Ethernet protocol stack
  - Single system image (SSI) clustering based on NFS
  - Head-node hard drive NFS-mounted on all nodes
  - Head-node video system with GNOME/X-windows
- Development Tools and Libraries
  - MPI over TCP/IP
  - `gcc` reference compiler; `distcc` distributed compiler; `mpicc`
  - ATLAS math library; FFTW fast fourier transform library
  - Sun Grid Engine (SGE) job scheduler

Orion Multisystems: Software Package (continued)

- Ease-of-Use Features
  - Power On: Single-system rapid boot sequence.
  - Monitoring & Control Software: Ganglia cluster monitor with Orion-specific extensions
- Demo Applications
  - LINPACK HPL
    - Solves a (random) dense system of linear equations in double-precision (64 bits) arithmetic.
  - mpiBLAST
    - Performs pairwise genome search alignment, parallelized with MPI.
    - Delivers super-linear speed-up, relative to NCBI's BLAST on a uniprocessor.
- Overall: Make a Linux cluster look like a standalone.

Selected Publications

http://sss.lanl.gov (… about three years out of date …)

- C. Hsu and W. Feng, "Effective Dynamic Voltage Scaling through CPU-Boundedness Detection," *IEEE/ACM MICRO Workshop on Power-Aware Computer Systems*, Dec. 2004.
- W. Feng and C. Hsu, "The Origin and Evolution of Green Destiny," *IEEE Cool Chips VII*, Apr. 2004.
- W. Feng, "Making a Case for Efficient Supercomputing," *ACM Queue*, Oct. 2003.
- W. Feng, "Green Destiny + mpiBLAST = Bioinfomagic," *10th Int'l Conf. on Parallel Computing (ParCo'03)*, Sept. 2003.
- M. Warren, E. Weigle, and W. Feng, "High-Density Computing: A 240-Processor Beowulf in One Cubic Meter," *SC 2002,* Nov. 2002.
- W. Feng, M. Warren, and E. Weigle, "Honey, I Shrunk the Beowulf!," *Int'l Conference on Parallel Processing*, Aug. 2002.

Sampling of Media Overexposure

- "Parallel BLAST: Chopping the Database," *Genome Technology*, Feb. 2005.
- "Start-Up Introduces a Technology First: The Personal Supercomputer," *LinuxWorld*, Sept. 2004.
- "New Workstations Deliver Computational Muscle," *Bio-IT World*, August 30, 2004.
- "Efficient Supercomputing with Green Destiny," *slashdot.org*, Nov. 2003.
- "Green Destiny: A 'Cool' 240-Node Supercomputer in a Telephone Booth," *BBC News*, Aug. 2003.
- "Los Alamos Lends Open-Source Hand to Life Sciences," *The Register*, June 29, 2003.
- "Servers on the Edge: Blades Promise Efficiency and Cost Savings," *CIO Magazine*, Mar. 2003.
- "Developments to Watch: Innovations," *BusinessWeek*, Dec. 2002.
- "Craig Venter Goes Shopping for Bioinformatics …," *GenomeWeb*, Oct. 2002.
- "Not Your Average Supercomputer," *Communications of the ACM*, Aug. 2002.
- "At Los Alamos, Two Visions of Supercomputing," *The New York Times*, Jun. 25, 2002.
- "Supercomputing Coming to a Closet Near You?" *PCWorld.com*, May 2002.
- "Bell, Torvalds Usher Next Wave of Supercomputing," *CNN*, May 2002.

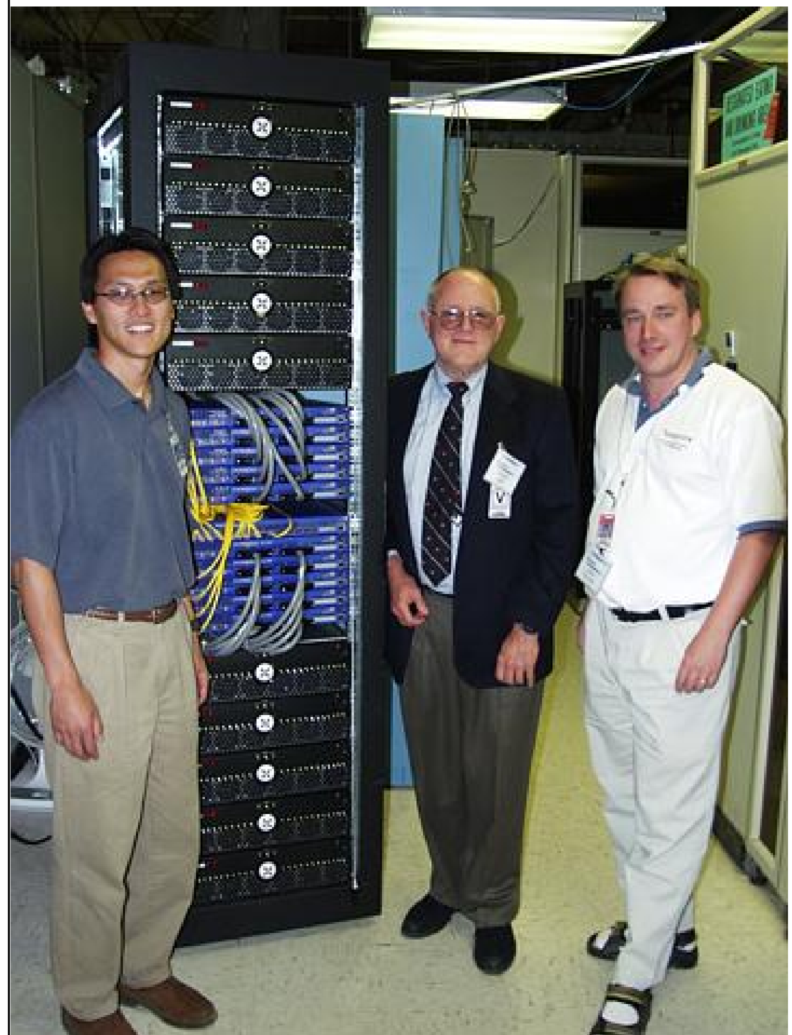

Adding to the Media Hype …

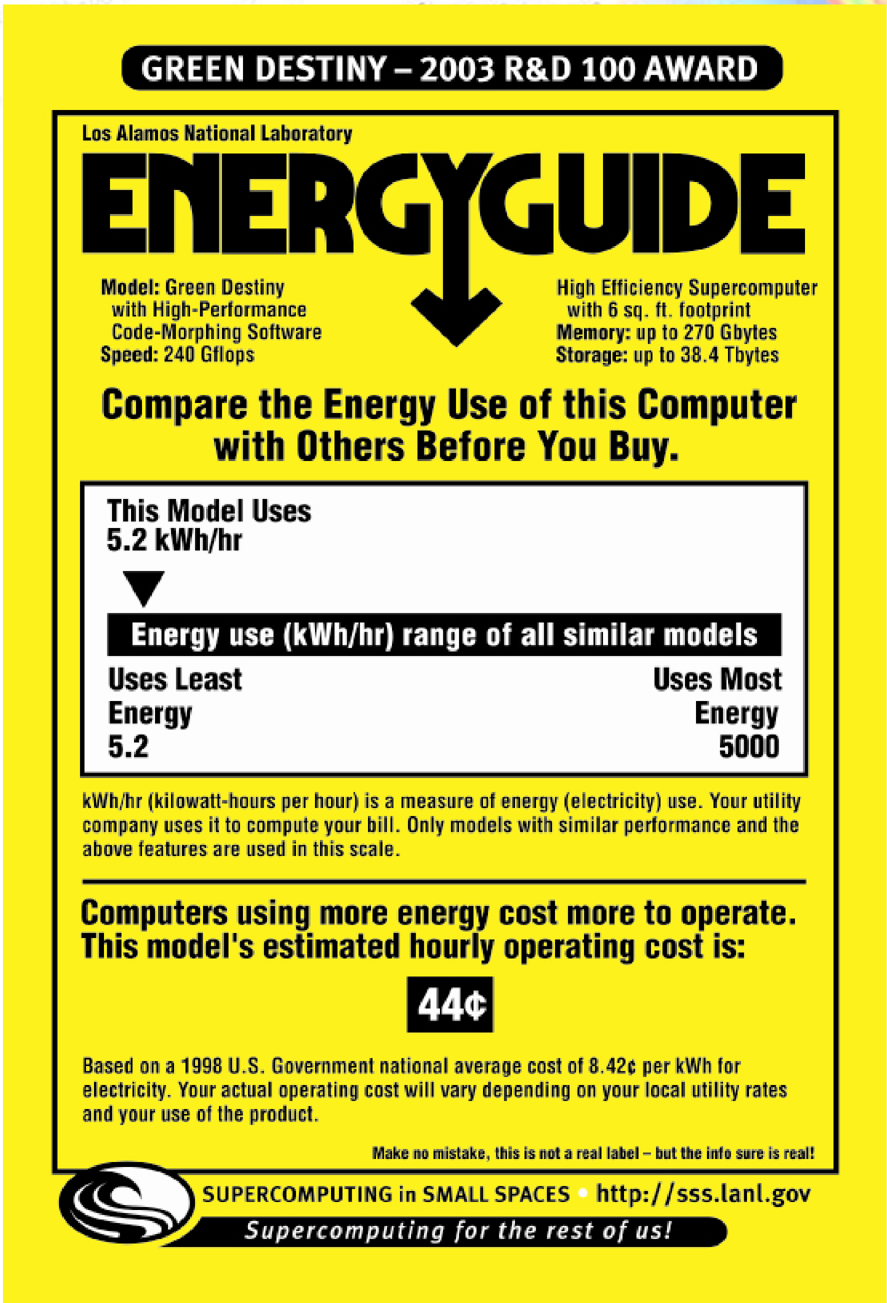

Conclusion

- Efficiency, reliability, and availability will be *the* key issues of this decade.
- Approach: Reduce power consumption via HW or SW.
- Cheesy Sound Bite for the DS-96 Personal Deskside Cluster (PDC):

" … the horsepower of a 134-CPU Cray T3E 1200 on LINPACK but in the power envelope of a hairdryer …"

Los Alamos
NATIONAL LABORATORY
UNIVERSITY OF CALIFORNIA

SUPERCOMPUTING
in SMALL SPACES

http://sss.lanl.gov

RADIANT

Research And Development In Advanced Network Technology

http://www.lanl.gov/radiant

Wu-chun (Wu) Feng
feng@lanl.gov