Green Supercomputing Comes of Age

Prof. Wu-chun FENG, feng@cs.vt.edu, 540-231-1192
Depts. of Computer Science and Electrical & Computer Engineering
VirginiaTech
Invent the Future

Forecast

- Premise
  - *Supercomputing* consumes significant *electrical power*
- Consequences
  - Electrical power costs $$$$
  - "Too much" power affects efficiency, reliability, availability
- Contributions
  - *How green supercomputing has "come of age" …*
    - "Supercomputing in Small Spaces" Project (2002 – now)
      - Delivering green supercomputing to the masses, e.g., Green Destiny
    - The Green500 List (2006 – now)
      - Raising awareness in green supercomputing

Outline

- Motivation & Background
  - Status of Traditional High-Performance Computing (HPC)
  - The "Cost" of Achieving Performance
- Supercomputing in Small Spaces (http://sss.cs.vt.edu/)
  - Origin
    - Green Destiny
  - Evolution
    - Architectural: Orion Multisystems → Sun Microsystems
    - Software: ENERGYFIT™ Automated Power-Aware Run-Time OS Daemon
  - The Future?
- The **Green500** List: Past, Present, and Future
- Conclusion

Status of High-Performance Computing

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

The Focus on Performance:
The TOP500 List

A ranking of the fastest 500 supercomputers in the world

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

The Focus on Performance:

Gordon Bell Awards at Supercomputing (SC)

- Metrics for Evaluating Supercomputers
  - Performance (i.e., Speed)
    - Metric: Floating-Operations Per Second (FLOPS)
  - Price/Performance → Cost Efficiency
    - Metric: Acquisition Cost / FLOPS

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

The "Cost" of Achieving Performance

- **Electrical *power* costs $$$$.**

Worldwide Expense to Power and Cool the Server Installed Base, 1996–2010

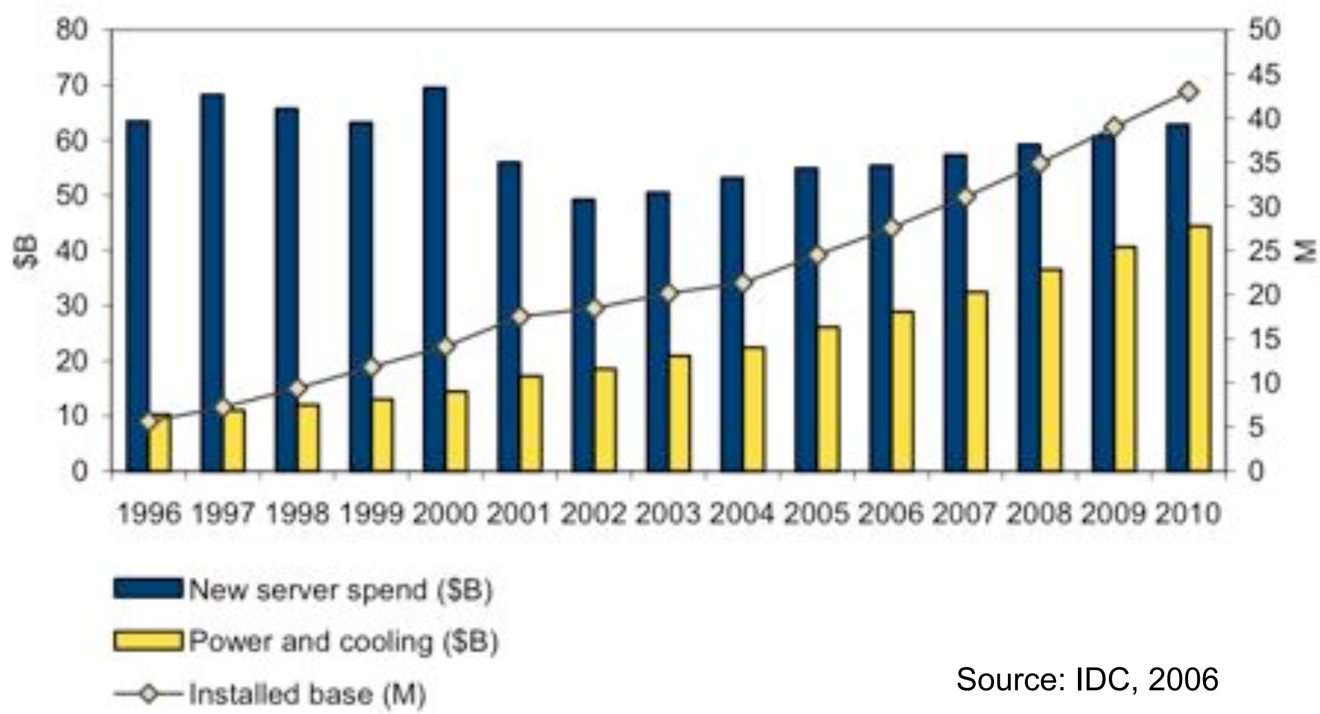

Source: IDC, 2006

The “Cost” of Achieving Performance

Examples: Power, Cooling, and Infrastructure $$$

- Japanese Earth Simulator
  - Power & Cooling: 12 MW/year → $10M/year

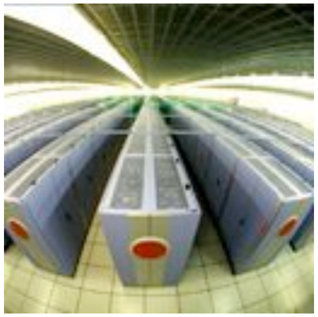

Virginia Tech
Invent the Future

- Google
  - "Hiding in Plain Sight, Google Seeks More Power," The New York Times, June 14, 2006.

The "Cost" of Achieving Performance

- ***Too much power* affects *efficiency, reliability, and availability.***

  - Anecdotal Evidence from a "Machine Room" in 2001 - 2002
    - **Winter**: "Machine Room" Temperature of **70-75° F**
      - Failure approximately ***once*** per week.
    - **Summer**: "Machine Room" Temperature of **85-90° F**
      - Failure approximately ***twice*** per week.

  - Arrenhius' Equation (applied to microelectronics)
    - ***For every 10° C (18° F) increase in temperature,***
      … the failure rate of a system doubles.*

* W. Feng, M. Warren, and E. Weigle, "The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs," *IEEE Cluster*, Sept. 2002.

The "Cost" of Achieving Performance

Systems	CPUs	Reliability & Availability
ASCI Q	8,192	**MTBF: 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.**
PSC Lemieux	**3,016**	**MTBF: 9.7 hrs.** **Availability: 98.33%.**
Google *(projected from 2003)*	***~450,000***	***~550 reboots/day; 2-3% machines replaced/yr.*** – **HW outage sources: storage, memory.** **Availability: ~100%.**

Source: Daniel A. Reed

The “Cost” of Achieving Performance

- Business Requirement of High-Performance Computing:

  E-commerce, enterprise apps, online services, ISPs, data and HPC centers supporting R&D

  – Near-100% *availability* with *efficient* and *reliable* resource usage.

- Problems
  – Frequency of Service Outages

    Source: David Patterson, UC-Berkeley

    - 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**
    - Amazon.com: **$ 1,200,000/hour**
    - Social Effects
      – Negative press, loss of customers who “click over” to competitor (e.g., Google vs. Ask Jeeves)

VirginiaTech Invent the Future

The Future of High-Performance Computing?

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

Outline

- Motivation & Background
  - Status of Traditional High-Performance Computing (HPC)
  - The "Cost" of Achieving Performance
- Supercomputing in Small Spaces (http://sss.cs.vt.edu/)
  - Origin
    - Green Destiny
  - Evolution
    - Architectural: Orion Multisystems → Sun Microsystems
    - Software: LOS ALAMOS NATIONAL LABORATORY ENERGYFIT COOL FAST RELIABLE COMPUTING ™ Automated Power-Aware Run-Time OS Daemon
  - The Future?
- The **Green500** List: Past, Present, and Future
- Conclusion

VirginiaTech Invent the Future

Supercomputing in Small Spaces

Efficiency, Reliability, and Availability via Green HPC
(Started in 2001 at Los Alamos Nat'l Lab. Now at Virginia Tech.)

- Goal: Improve efficiency, reliability, and availability (ERA) in large-scale supercomputing systems.
- Analogy

  Today's Supercomputer vs. Supercomputing in Small Spaces

  – Formula One Race Car: Wins raw performance but reliability is so poor that it requires frequent maintenance. Throughput low.

  – Nissan 370Z: Loses raw performance but high reliability results in high throughput (i.e., miles driven → answers/month).

VirginiaTech
Invent the Future

Green Destiny Supercomputer

(circa December 2001 – February 2002)

- A 240-Node Cluster in Five Sq. Ft.
- Each Node
  - 1-GHz Transmeta TM5800 CPU w/ *High-Performance Code-Morphing Software* running Linux 2.4.x
  - 640-MB RAM, 20-GB hard disk, 100-Mb/s Ethernet
- Total
  - 240 Gflops peak (Linpack: 101 Gflops in March 2002.)
  - 150 GB of RAM (expandable to 276 GB)
  - 4.8 TB of storage (expandable to 38.4 TB)

Transmeta TM5800 CPU: VLIW ASIC + CMS

A *Simpler* CPU

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

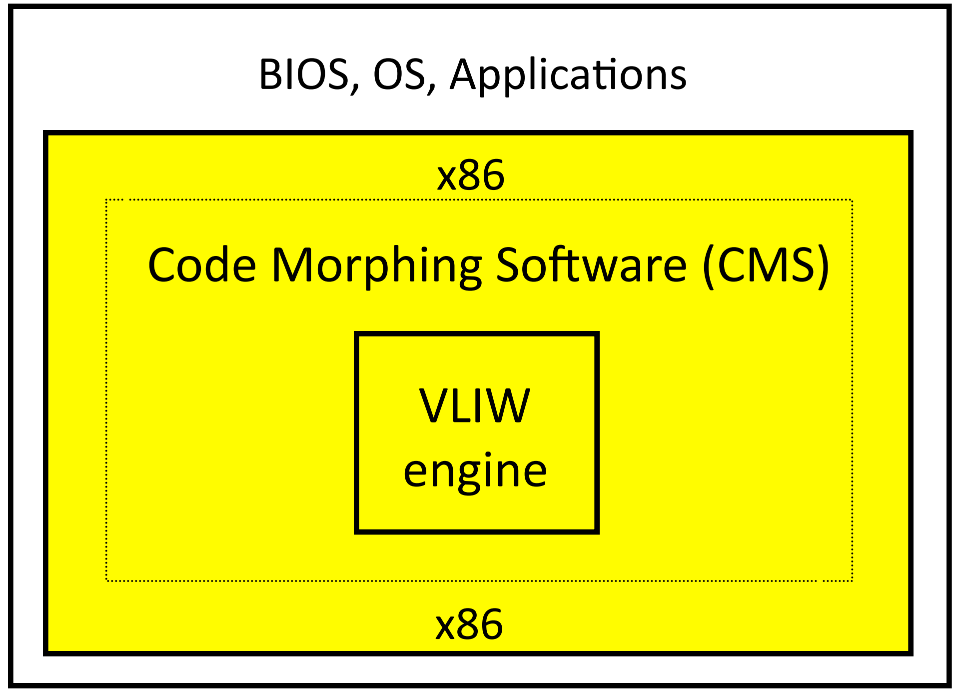

- VLIW Transistor Count ("Anti-Moore's Law")
  - ~¼ of Intel PIII → ~ 6x-7x less power dissipation
  - Less power → lower "on-die" temp. → better reliability & availability

Green Destiny Supercomputer

(circa December 2001 – February 2002)

- A 240-Node Cluster in Five Sq. Ft.
- Each Node
  - 1-GHz Transmeta TM5800 CPU w/ *High-Perfor* *Morphing Software* running Linux 2.4.x
  - 640-MB RAM, 20-GB hard disk, 100-Mb/s Ether
- Total
  - 240 Gflops peak (Linpack: 101 Gflops in March 2002.)
  - 150 GB of RAM (expandable to 276 GB)
  - 4.8 TB of storage (expandable to 38.4 TB)
  - *Power Consumption: Only 3.2 kW (diskless)*
- Reliability & Availability
  - *No unscheduled downtime in 24-month lifetime.*
    - Environment: A dusty 85°-90° F warehouse!

Equivalent Linpack to a
256-CPU SGI Origin 2000
(On TOP500 List at the time)

Green Destiny: Low-Power Supercomputer

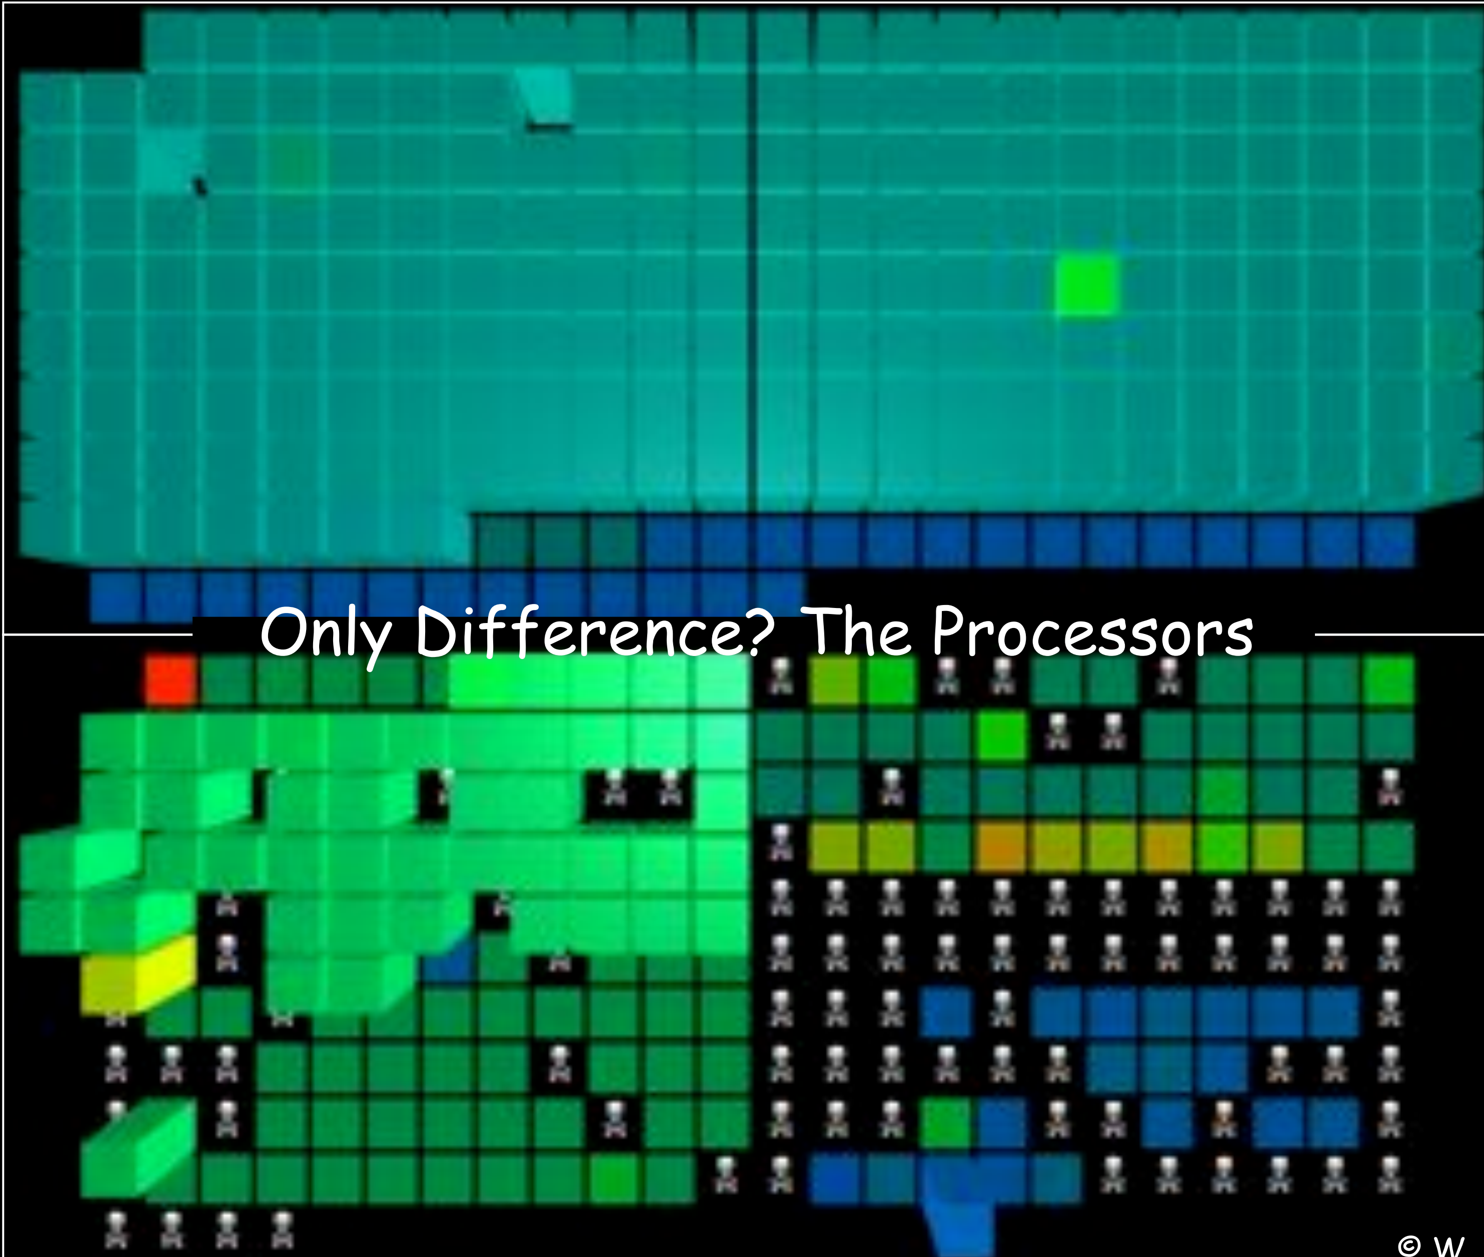

Only Difference? The Processors

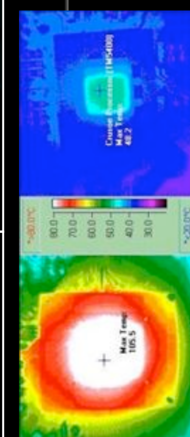

Green Destiny "Replica": Traditional Supercomputer

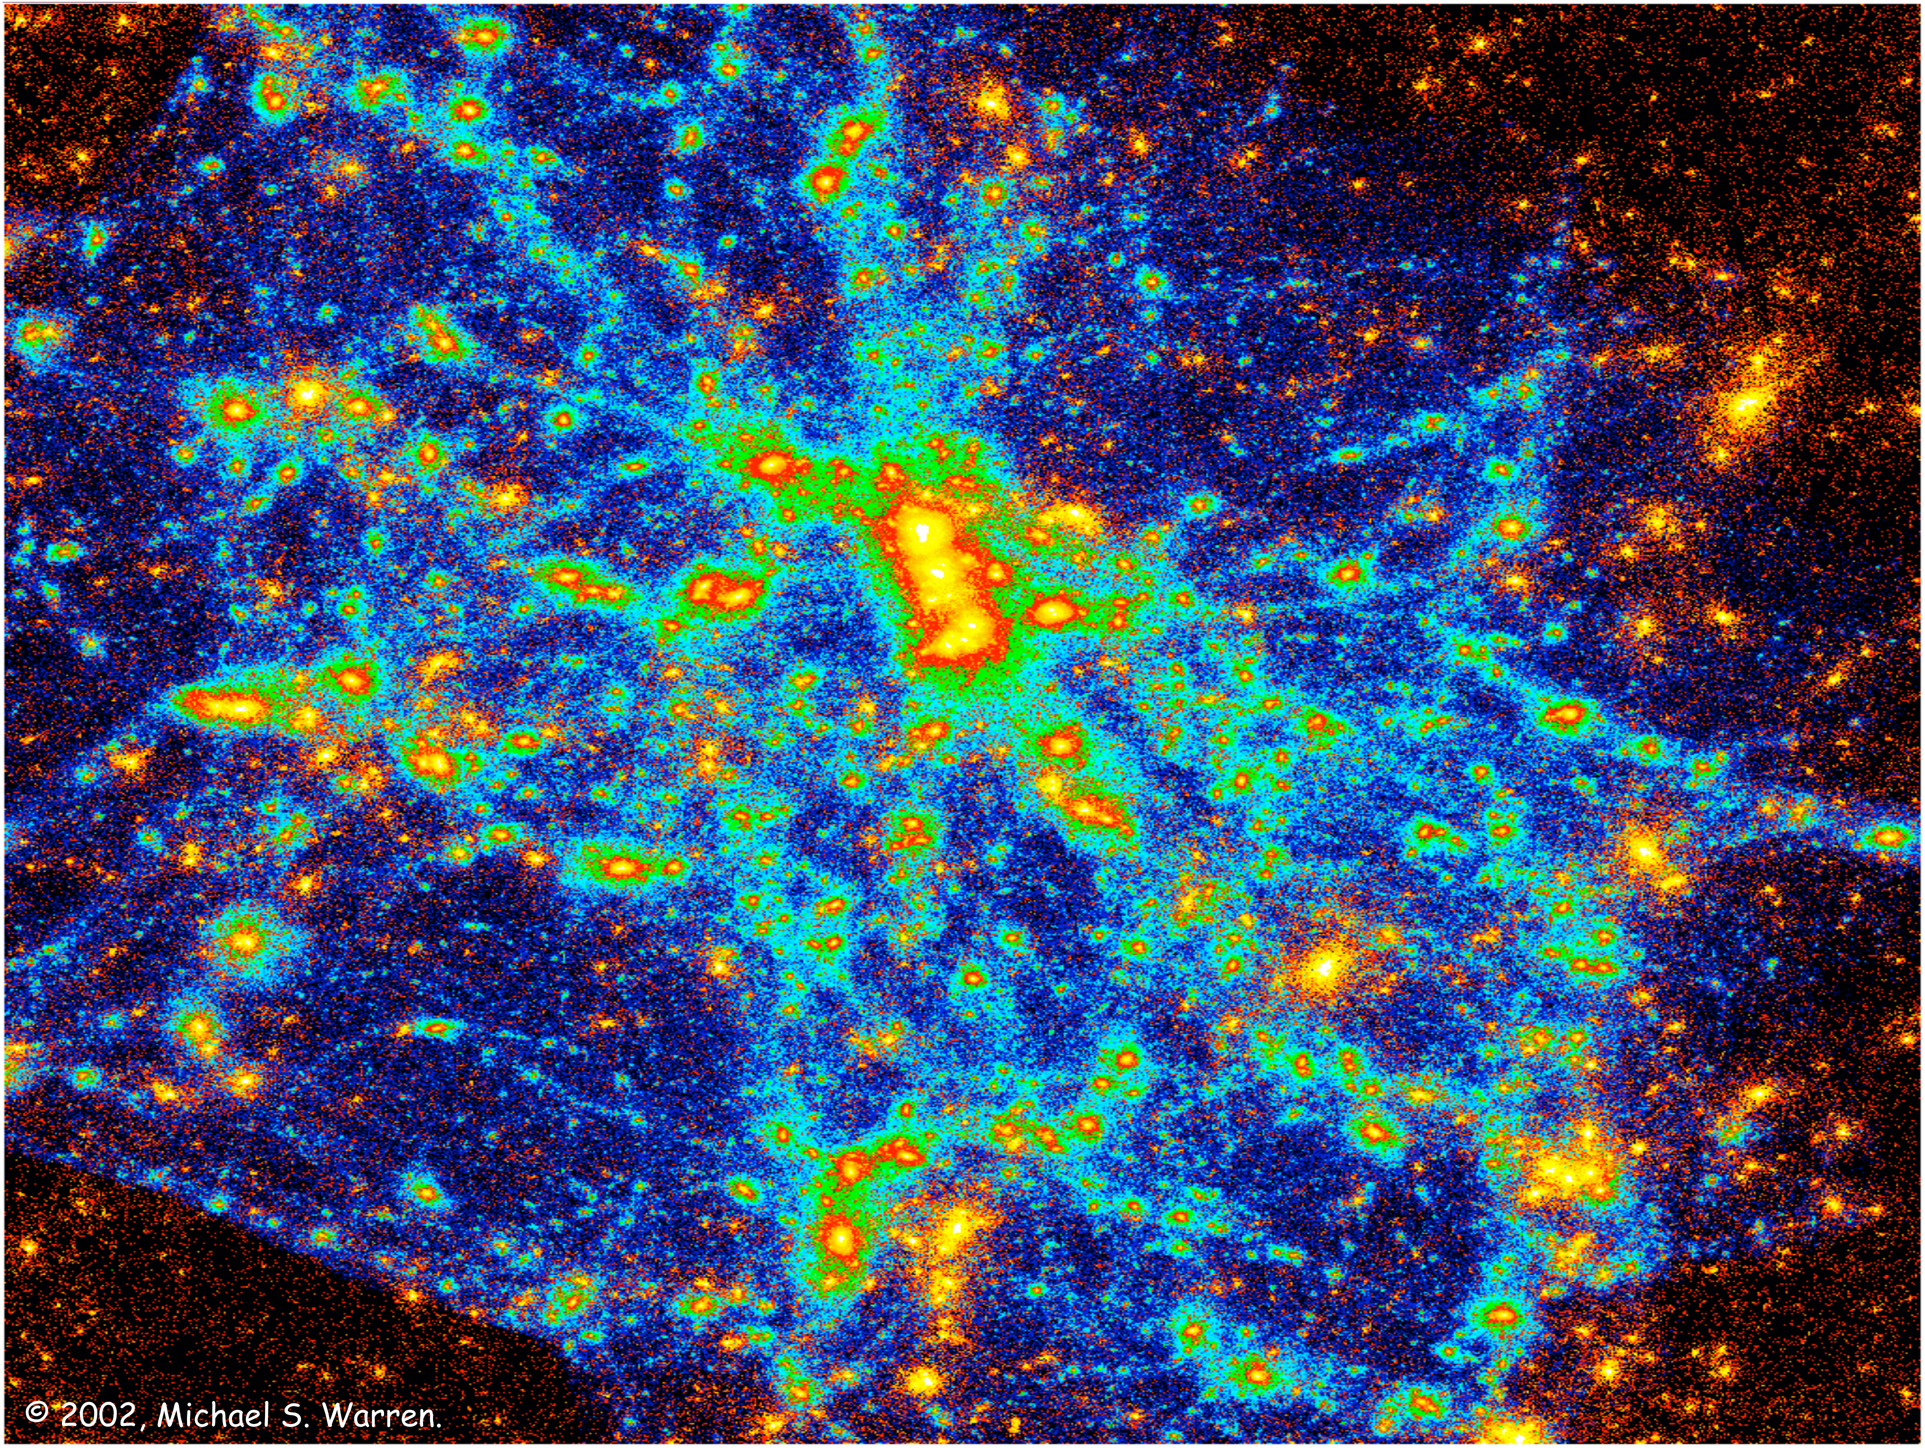

© 2002, Michael S. Warren.

Green Destiny Supercomputer

(circa December 2001 – February 2002)

- A 240-Node Cluster in Five Sq. Ft.
- Each Node
  - 1-GHz Transmeta TM5800 CPU w/ *High-Perfor* *Morphing Software* running Linux 2.4.x
  - 640-MB RAM, 20-GB hard disk, 100-Mb/s Ether
- Total
  - 240 Gflops peak (Linpack: 101 Gflops in March 2002.)
  - 150 GB of RAM (expandable to 276 GB)
  - 4.8 TB of storage (expandable to 38.4 TB)
  - *Power Consumption: Only 3.2 kW (diskless)*
- Reliability & Availability
  - *No unscheduled downtime in 24-month lifetime.*
    - Environment: A dusty 85°-90° F warehouse!

Equivalent Linpack to a
256-CPU SGI Origin 2000
(On TOP500 List at the time)

2003 WINNER

Featured in *The New York Times*, *BBC News*, and *CNN*.
Now in the *Computer History Museum*.

Yet in 2002 …

PCWorld

Supercomputing Coming to a Closet Near You?

- "In high-performance computing, no one cares about power & cooling, and no one ever will …"
- "Moore's Law for Power will stimulate the economy by creating a new market in cooling technologies."
- "Green Destiny is so low power that it runs just as fast when it is unplugged."

InfoWorld

HOME NEWS TEST CENTER

Green Destiny draws cheers and jeers

For many of the Los Alamos scientists, the unveiling of Green Destiny was their first introduction to blade servers -- never mind blade servers being used to build a supercomputer. The slew of expletives and exclamations that followed Feng's description of the system made it clear that the blades had captured the audience's attention. Some murmured, "Wow," while others let out multiple shouts of, "Jesus!" as their jaws dropped.

Several scientists here did not share the enthusiasm for Green Destiny, however. Los Alamos, after all, is the home to several massive supercomputers that take up entire floors of buildings and require several cooling systems shaped like mini-nuclear reactors to keep them running. These "real" supercomputers handle serious work, and some of the people running them consider Green Destiny a joke. One scientist walked out of Feng's presentation, making his feelings clear.

Outline

- Motivation & Background
  - Status of Traditional High-Performance Computing (HPC)
  - The "Cost" of Achieving Performance
- Supercomputing in Small Spaces (http://sss.cs.vt.edu/)
  - Origin
    - Green Destiny
  - Evolution
    - Architectural: Orion Multisystems → Sun Microsystems
    - Software: ENERGYFIT™ Automated Power-Aware Run-Time OS Daemon
  - The Future?
- The **Green500** List: Past, Present, and Future
- Conclusion

VirginiaTech *Invent the Future*

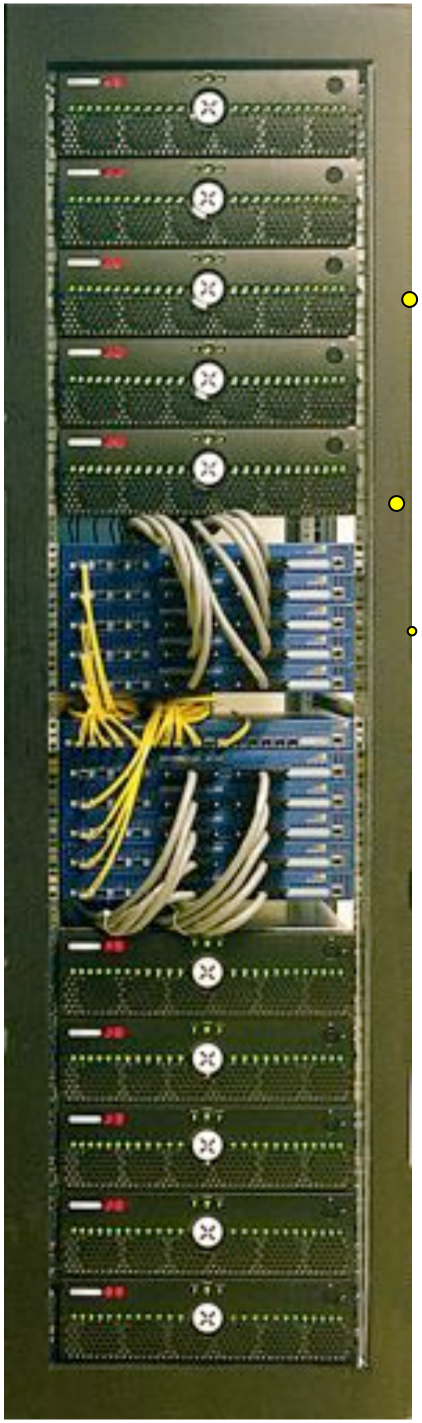

Cluster Technology

Low-Power Systems Design

Linux

But in the form factor of a workstation …

a *cluster workstation*

Orion DT-12 (circa 2004)

- LINPACK Performance
  - 14 Gflops
- Footprint
  - 3 sq. ft. (24" x 18")
  - 1 cu. ft. (24" x 4" x 18")
- Power Consumption
  - 170 watts at load

Outline

- Motivation & Background
  - Status of Traditional High-Performance Computing (HPC)
  - The "Cost" of Achieving Performance
- Supercomputing in Small Spaces (http://sss.cs.vt.edu/)
  - Origin
    - Green Destiny
  - Evolution
    - Architectural: Orion Multisystems → Sun Microsystems
    - Software: ENERGYFIT™ Automated Power-Aware Run-Time OS Daemon
  - The Future?
- The **Green500** List: Past, Present, and Future
- Conclusion

Automated Power-Aware Run-Time OS Daemon

- Self-Adapting Software for Energy Efficiency
  Conserve power & energy ***WHILE*** *maintaining performance.*

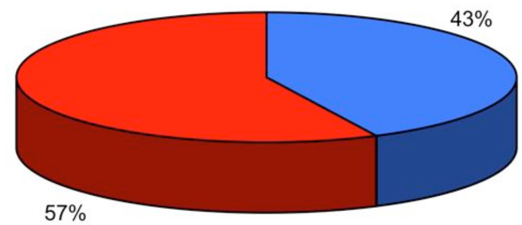

- Observations
  - Many commodity technologies support *dynamic voltage & frequency scaling (DVFS)*, which allows changes to the processor voltage and frequency at run time.
  - A computing system can trade off processor performance for power reduction.
    - *Power* α $V^2 \cdot f$ where V is the supply voltage of the processor and f is its frequency.
    - *Processor performance* α f

Software-Based Energy Efficiency via DVFS:

Key Observation

- Execution time of many programs is insensitive to CPU speed change.

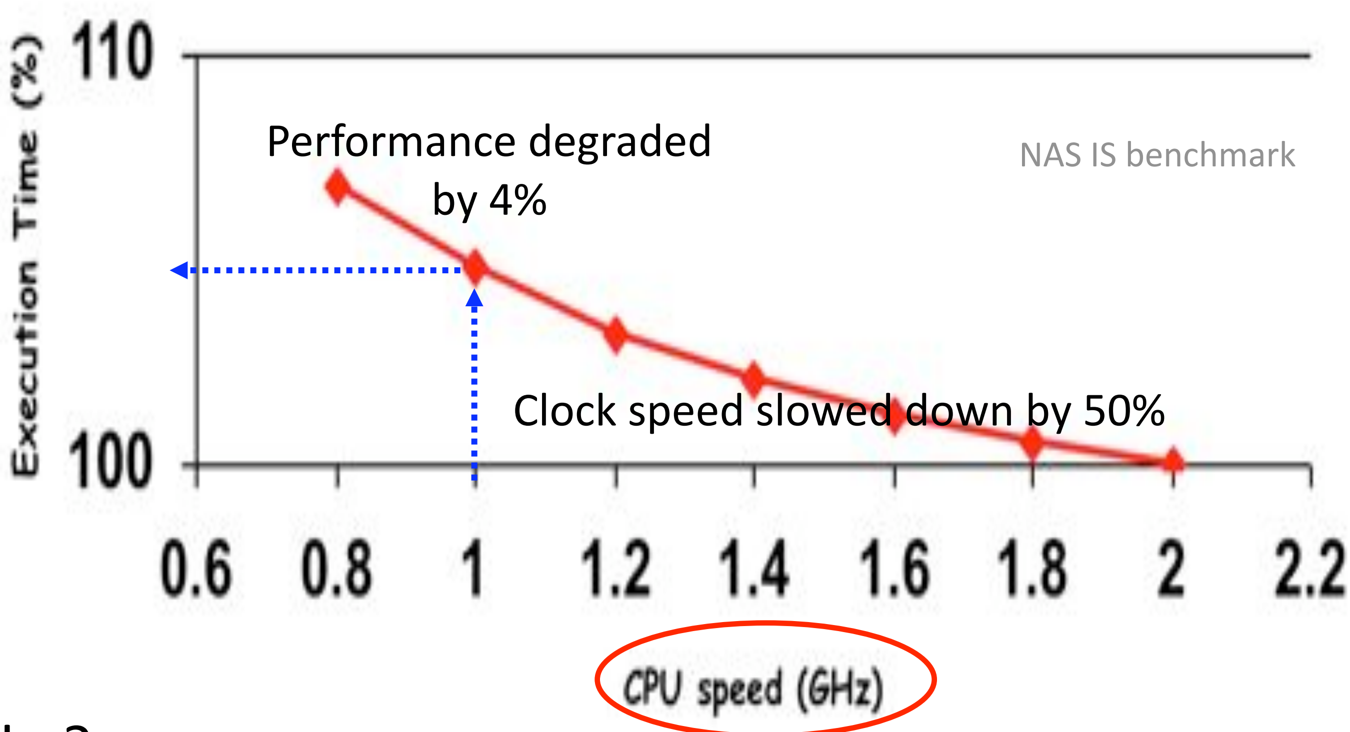

- Why?
  - The Memory & I/O Wall → *The Brick Wall*

Software-Based Energy Efficiency via DVFS:

Key Idea

- Applying DVFS to these programs will result in significant power & energy savings at a minimal performance impact.

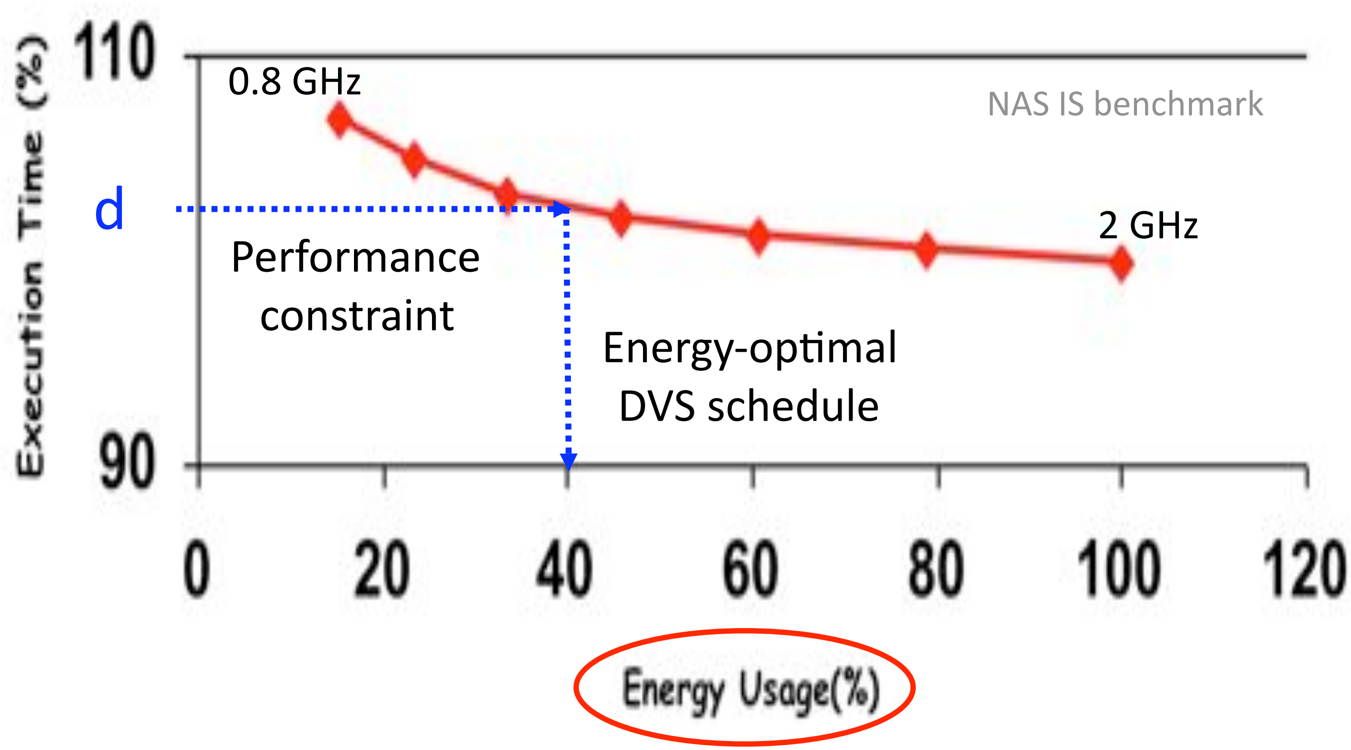

Why is Power Awareness via DVFS Hard?

- What is cycle time of a processor?
  - Frequency ≈ 2 GHz → Cycle Time ≈ 1 / (2 x 10^9) = 0.5 ns
- How long does the system take to scale voltage and frequency?

O (10,000,000 cycles)

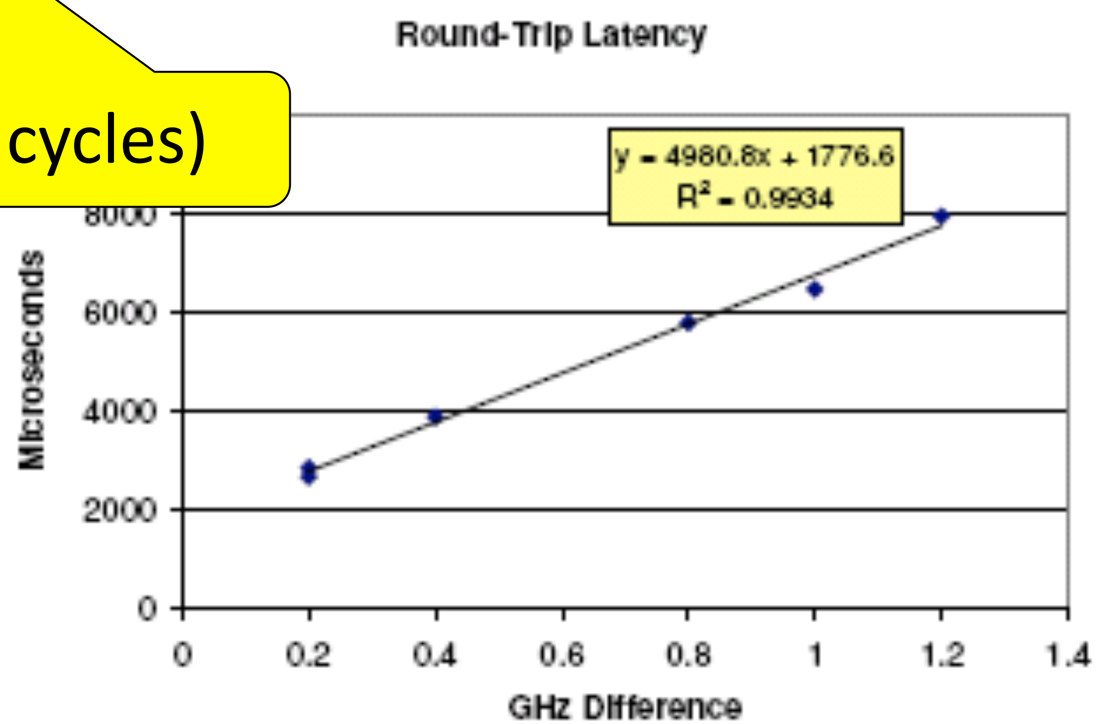

Problem Formulation:

LP-Based Energy-Optimal DVFS Schedule for

- Definitions
  - A DVFS system exports n $\{ (f_i, V_i) \}$ settings.
  - T_i : total execution time of a program running at setting i
- Given a program with deadline D, find a DVFS schedule (t_1*, …, t_n*) such that
  - If the program is executed for t_i seconds at setting i, the total energy usage E is minimized, the deadline D is met, and the required work is completed.

$$\min E = \sum_i P_i \cdot t_i$$

subject to

$$\sum_i t_i \leq D$$
$$\sum_i t_i / T_i = 1$$
$$t_i \geq 0$$

C. Hsu and W. Feng.
"A Power-Aware Run-Time System for High-Performance Computing," *SC|05,* Nov. 2005.

Experimental Set-Up

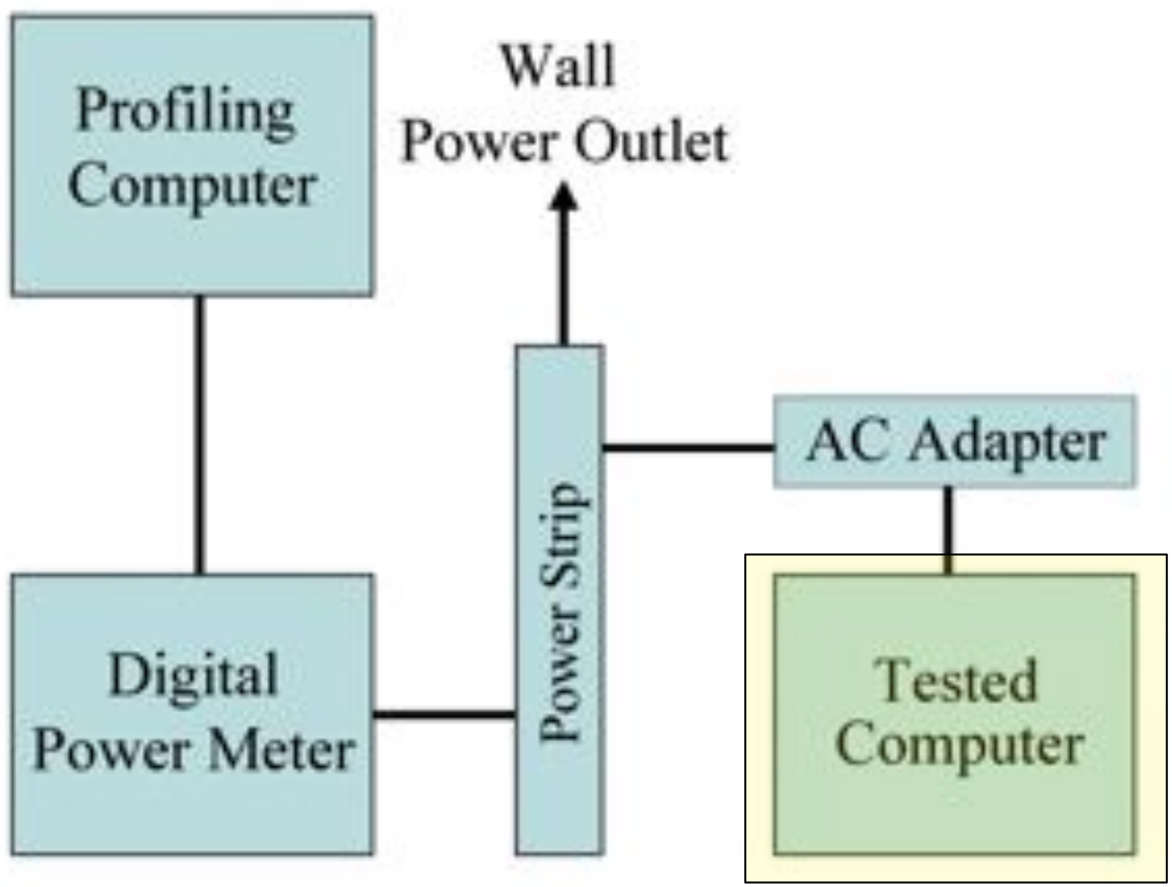

Experimental Specifics

- Tested Computer Platforms with DVFS Enabled
  - 64-bit AMD Opteron with *PowerNow!* enabled
    - 4-processor 4U server
    - 40-processor cluster (CAFfeine)
- Digital Power Meter
  - Yokogawa WT210: Continuously samples every 20 ms.
- Benchmarks Used
  - Uniprocessor: SPEC Benchmarks
  - Multiprocessor: NAS Parallel Benchmarks

VirginiaTech Invent the Future

ENERGYFIT on SPEC CPU Benchmarks

Type	SPEC Benchmarks	No Energy Software	State-of-the-Art Commercial: cpuspeed / Speedstep	State-of-the-Art Research: PID	ENERGYFIT δ = 5%
Memory Bound	swim	1.00 / 1.00	1.00 / 1.00	1.04 / 0.70	**1.04 / 0.61**
	tomcatv	1.00 / 1.00	1.00 / 1.00	1.03 / 0.69	1.00 / 0.85
	su2cor	1.00 / 1.00	0.99 / 0.99	1.05 / 0.70	1.03 / 0.85
	mgrid	1.00 / 1.00	1.00 / 1.00	1.18 / 0.77	1.03 / 0.89
	vortex	1.00 / 1.00	1.01 / 1.00	1.25 / 0.81	1.05 / 0.90
CPU Bound	go	1.00 / 1.00	1.00 / 1.00	1.37 / 0.88	1.05 / 0.96

normalized time / normalized energy

with respect to total execution time and system energy usage

δ = 5% means performance slowdown, if any, must be ≤ 5%

C. Hsu and W. Feng, "Effective Dynamic Voltage Scaling through CPU-Boundedness Detection," *Lecture Notes in Computer Science*, 2005.

VirginiaTech Invent the Future

ENERGYFIT on SPEC Java Business Benchmarks

Power Management	Watts	% Power Reduction
None	264	0%
cpuspeed	257	3%
On-demand	253	4%
ENERGYFIT	196	25%

Power Management	bops/watt
None	100.00%
cpuspeed	102.56%
On-demand	104.37%
ENERGYFIT	123.70%

C. Hsu and W. Feng, "Effective Dynamic Voltage Scaling through CPU-Boundedness Detection," *Lecture Notes in Computer Science*, 2005.

ENERGYFIT on NAS Parallel Benchmarks

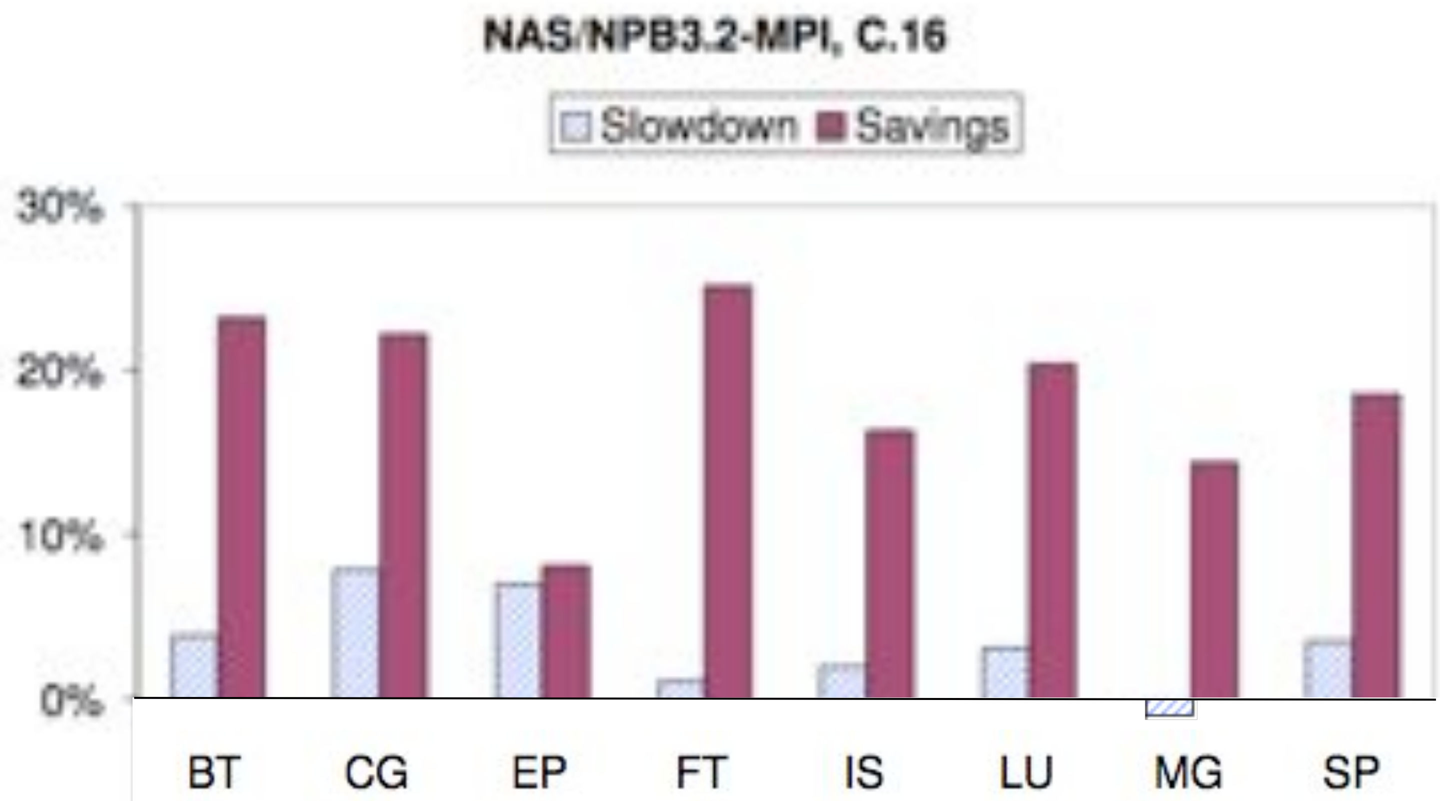

Energy reduction (15%)
Performance improvement

C. Hsu and W. Feng.
"A Power-Aware Run-Time System for High-Performance Computing,"
SC|05, Nov. 2005.

Outline

- Motivation & Background
  - Status of Traditional High-Performance Computing (HEC)
  - The "Cost" of Achieving Performance
- Supercomputing in Small Spaces (http://sss.cs.vt.edu/)
  - Origin
    - Green Destiny
  - Evolution
    - Architectural: Orion Multisystems → Sun Microsystems
    - Software: ENERGYFIT™ Automated Power-Aware Run-Time OS Daemon
  - The Future?
- The **Green500** List: Past, Present, and Future
- Conclusion

The Future of Energy-Efficient Supercomputing?

- Heterogeneity & Simplicity
  - Use the right processor (i.e., tool) for the job
    - Remember the "floating-point unit" (FPU) co-processor?
  - General-purpose CPU + GPU + FPGA (?)
- The GPU?
  - AMD/ATI Radeon 4000 Series → AMD/ATI Radeon 5000 Series
    - Programming Model: Brook+ to OpenCL
  - NVIDIA Tesla → NVIDIA Fermi
    - Programming Model: CUDA and OpenCL
- A Fused CPU+GPU?
  - AMD/ATI Fusion
  - Intel Larrabee

Electrostatic Potential for Molecular Dynamics

Viral Capsid

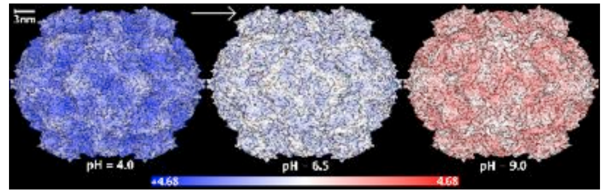

Processor + Optimization	Execution Time	Speed-Up
CPU	80690.20	-
GPU	355.64	227x
GPU + Kernel Split	219.20	368x
GPU + Multi-Level 1	51.43	1569x
GPU + Kernel Split + Multi-Level 1	35.04	2303x

VirginiaTech Invent the Future

Electrostatic Potential for Molecular Dynamics

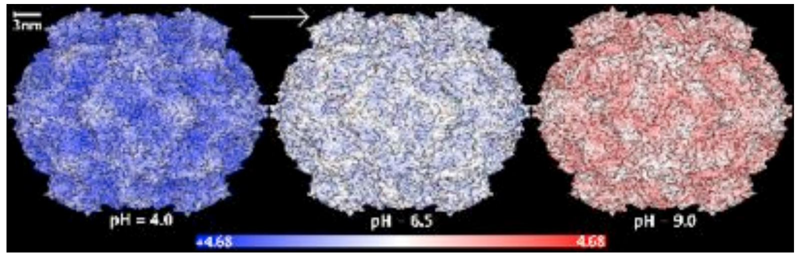

Processor + Optimization	Power (W)	Execution Time (s)	Energy (J)	Energy-Delay Product (J • s)
CPU serial	98.53	2,037.53	200,757.83	409,050,101
CPU multi-core	98.53	1,135.61	111,891.65	127,065,277
GPU + kernel split	228.07	9.50	2,166.67	20,583

Outline

- Motivation & Background
  - Status of Traditional High-Performance Computing (HPC)
  - The "Cost" of Achieving Performance
- Supercomputing in Small Spaces (http://sss.cs.vt.edu/)
  - Origin
    - Green Destiny
  - Evolution
    - Architectural: Orion Multisystems → Sun Microsystems
    - Software: ENERGYFIT™ Automated Power-Aware Run-Time OS Daemon
  - The Future?
- The **Green500** List: Past, Present, and Future
- Conclusion

Brief History:

From Green Destiny to Green500 List

- 2/2002: Green Destiny (http://sss.lanl.gov/ → http://sss.cs.vt.edu/)
  - "Honey, I Shrunk the Beowulf!" *31st Int'l Conf. on Parallel Processing*, August 2002.
- 4/2005: Workshop on High-Perf., Power-Aware Computing (at the IEEE International Parallel & Distributed Processing Symposium)
  - Keynote address generates initial discussion for Green500 List
- 4/2006 and 9/2006: Making a Case for a Green500 List
  - *Workshop on High-Performance, Power-Aware Computing*
  - Jack Dongarra's *CCGSC Workshop* → "The Final Push" (TOP500 List)
- 9/2006: Launch of Green500 Web Site and RFC
  - http://www.green500.org/ → Generates feedback from hundreds
- 11/2007: Debut of the first Green500 List

Creating The *Green500* List

Ranking the most energy-efficient supercomputers in the world

- Challenges
  - Find a metric*
    - Performance / Measured Power
      - For now, Linpack FLOPS from TOP500 / Measured Power
        - » Use Peak Power, if Measured Power is unavailable.
  - Decide what to measure
    - The largest possible "unit" (excludes facility cooling)
  - How to measure it
    - In-line power measurement
  - Example
    - A Green500 supercomputer using measured power (instead of peak power) jumps up 60-80 spots on the Green500 List.

* Towards Efficient Supercomputing: Choosing the Right Efficiency Metric, IEEE IPDPS/HPPAC, April 2005.

The *Green500* List … Today…

- What is “The Green500 List”?
  - Effectively, a re-ordering of the TOP500 List but based on energy efficiency, specifically FLOPS/Watt.
- How often is “The Green500 List” released?
  - Every June and November
- How does the Green500 engage the HPC community?
  - “Birds of a Feather” (BoF) Sessions: Typically ISC and SC
  - Workshop on High-Performance, Power-Aware Computing
- How important is “The Green500 List”?
  - The Green500 List BoF at IEEE/ACM SC 2008
    - “Standing room only” (SRO)
    - Its importance led to filming of the session

5th Edition of The Green500 List

Green500 Rank	MFLOPS/W	Site*		Total Power (kW)	TOP500 Rank*
1	536.24	Interdisciplinary Centre for Mathematical and Computational Modeling, University of Warsaw	BladeCenter QS22 Clu	34.63	422
2	458.33	DOE/NNSA/LANL	BladeCenter QS22/LS2 DC 1.8 GHz, Infiniband	138.00	61
2	458.33	IBM Poughkeepsie Benchmarking Center	BladeCenter QS22/LS2 DC 1.8 GHz, Infiniband	138.00	62
4	444.94	DOE/NNSA/LANL	BladeCenter QS22/LS2 DC 1.8 GHz, Voltaire In	2483.47	1
5	428.91	National Astronomical Observatory of Japan	GRAPE-DR accelerator	51.20*	277
6	371.67	ASTRON/University Groningen	Blue Gene/P Solution	94.50	124
7	371.67	IBM - Rochester	Blue Gene/P Solution	126.00	84
7	371.67	IBM Thomas J. Watson Research Center	Blue Gene/P Solution	126.00	85
7	371.67	Max-Planck-Gesellschaft MPI/IPP	Blue Gene/P Solution	126.00	86
7	371.67	Bulgarian State Agency for Information Technology and Communications (SAITC)	Blue Gene/P Solution	63.00	245

Trends from the 5th Green500 List

- *Overall energy efficiency continues to improve …*
  - Average efficiency increased by 10% (98 MFlops/Watt → 108 MFlops/Watt) despite aggregate power of the list increasing by 15% (200 MW → 230 MW).
- *A self-made accelerator-based supercomputer catapults into the #5 spot on the Green500.*
  - The first Green500 supercomputer with arguably more than a million processing elements: 2,097,152.
  - Will this approach of aggregating many more less powerful processors for better overall performance be a trend to keep an eye on?
    - Tesla S1070: 4 GPUs → 960 processing cores "in a box"

Trends from the 5th Green500 List

- *Commodity processors continue to nip at the heels of previous-generation custom processors*.
  - 20 of the top 50 energy-efficient supercomputers utilize commodity processors.
- *For the first time, maximum energy efficiency remains the same, but three 500-Mflops/Watt supercomputers drop out of the Green500*.

The Growth of The Green500 List

- From 0 to *20,000* page hits per month (on average)
- Number of reporting sites increased by 20%
- The Green500 List BoF
  - "Standing Room Only" (SRO)
  - Its importance led to filming of the session

Issues …

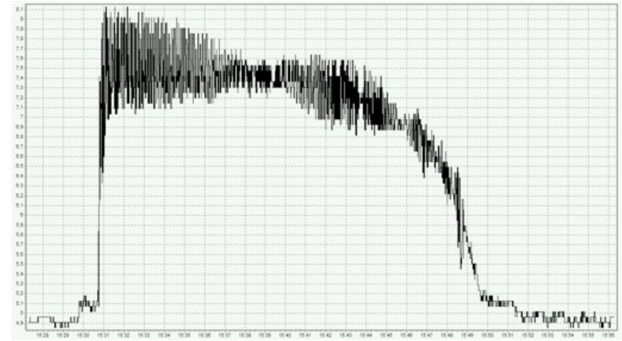

Discussion at the Green500 BoF at SC'08

- *Eligibility:* Where is the entry bar?
- *Benchmark:* LINPACK, HPCC, NAS? But then what is the metric?
  - Should mixed-precision arithmetic be allowed?
- *Metric:* Performance/Watt favors smaller machines but others favor large machines
- *What to Measure & How to Measure:* Peak vs. Average. In-line vs. facility level.
- *Classification:* List of "classes of machine" rather than individual machines? Likely not.
  - See *Code Morphing Software* in Green Destiny.

VirginiaTech Invent the Future

The *Green500* List … Tomorrow?

The Green500 is about *energy-efficient* supercomputing. Why limit the list only to the top *performing* 500 supercomputers in the world?

- Eligibility
  - Place the bar at the last year's #500 Top500
    - Put focus on energy efficiency rather than "re-ordered performance."
- Measurement: Continue with status quo …
  - Measure only machine, not cooling facilities. Average power.
- Metric
  - FLOPS/watt
- Benchmark
  - Stay with LINPACK … for now …
- Presentation
  - Enable different views for the data.

For More Information

- The Green500 List
  http://www.green500.org/
- Acknowledgements and Thanks
  - The TOP500 Team and the Green500 Team


Coming to IEEE Computer Magazine, December 2007

The Green500 List: Encouraging Sustainable Supercomputing

The performance-at-any-cost design mentality ignores supercomputers' excessive power consumption and need for heat dissipation and will ultimately limit their performance. Without fundamental change in the design of supercomputing systems, the performance advances common over the past two decades won't continue.

VirginiaTech
Invent the Future

Selected References: Available at http://sss.cs.vt.edu/

Green Supercomputing

- **"Green Supercomputing Comes of Age," *IT Professional*, Jan.-Feb. 2008.**
- **"The Green500 List: Encouraging Sustainable Supercomputing," *IEEE Computer*, Dec. 2007.**
- "CPU MISER: A Performance-Directed, Run-Time System for Power-Aware Clusters," *36th International Conference on Parallel Processing*, Sept. 2007.
- **"Green Supercomputing in a Desktop Box," *IEEE Workshop on High-Performance, Power-Aware Computing* at *IEEE IPDPS,* Mar. 2007.**
- "Making a Case for a Green500 List," *IEEE Workshop on High-Performance, Power-Aware Computing* at *IEEE IPDPS,* Apr. 2006.
- **"A Power-Aware Run-Time System for High-Performance Computing," *SC|05*, Nov. 2005.**
- "The Importance of Being Low Power in High-Performance Computing," *Cyberinfrastructure Technology Watch* (NSF), Aug. 2005.
- "Green Destiny & Its Evolving Parts," *International Supercomputer Conference,* Innovative Supercomputer Architecture Award,., Apr. 2004.
- "Making a Case for Efficient Supercomputing," ACM Queue, Oct. 2003.
- "Green Destiny + mpiBLAST = Bioinfomagic," 10th Int'l Conf. on Parallel Computing (ParCo'03), Sept. 2003.
- **"Honey, I Shrunk the Beowulf!," Int'l Conf. on Parallel Processing, Aug. 2002.**

Conclusion

- Recall that …
  - *Supercomputing* consumes significant *electrical power*
- Consequences
  - Electrical power costs $$$$
  - "Too much" power affects efficiency, reliability, availability
- Contributions
  - *How green supercomputing has "come of age" …*
    - "Supercomputing in Small Spaces" Project
      - Delivering green supercomputing to the masses, e.g., Green Destiny
    - The Green500 List
      - Raising awareness in green supercomputing

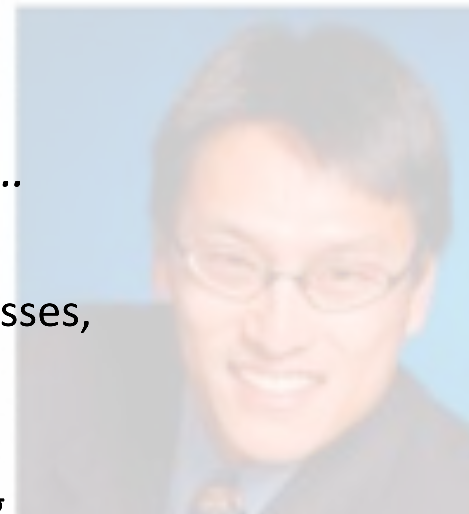

Wu-chun FENG, Ph.D.

feng@cs.vt.edu

SyNeRGy Laboratory

http://synergy.cs.vt.edu/

mpiBLAST

http://www.mpiblast.org/

Green500

http://www.green500.org/

SUPERCOMPUTING
in SMALL SPACES

http://sss.cs.vt.edu/