



The Green500 List: Year One

W. Feng and T. Scogland
Department of Computer Science



<http://www.green500.org/>



W. Feng; The Green500 List: Year One; HPPAC at IPDPS; May 25, 2009



Motivation: Where is High-End Computing?



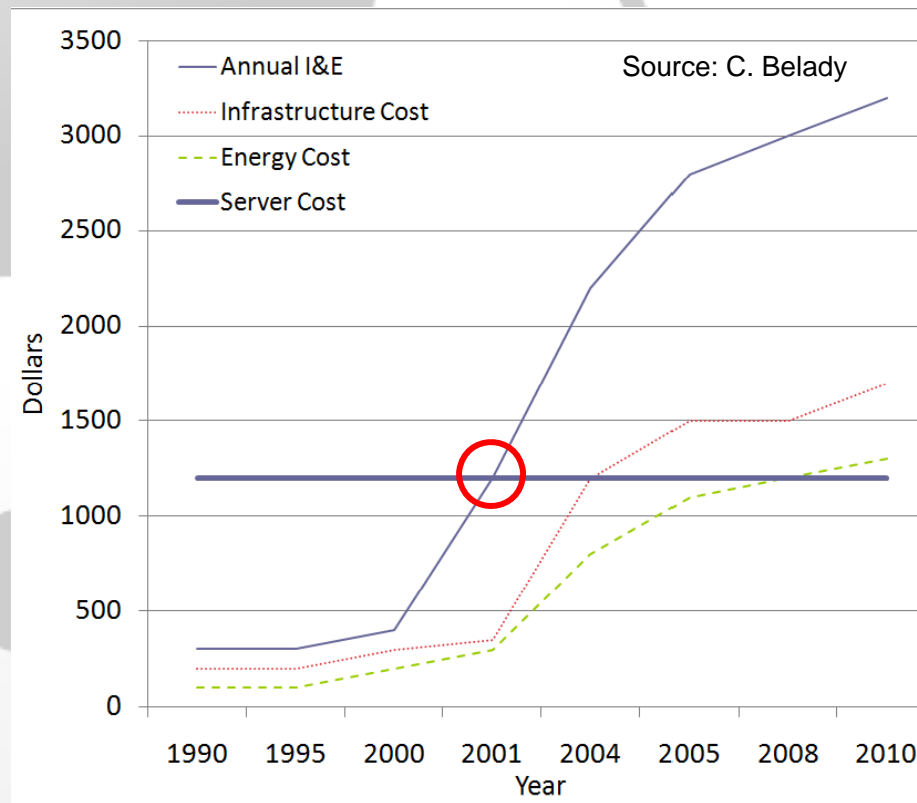
We have spent decades focusing on performance, performance, performance (and price/performance).*

Forecast

- The Green500 List: Origins
- The Green500 List: Year One
 - Track the progress of power-efficient supercomputing
 - Analyze trends
 - Reflect upon the implications of these trends
- The Green500 List: Evolution


Top 5 Reasons for Being Green

1. In 2001, the combined annual cost of “Infrastructure & Energy” for a 1U server surpassed its purchase cost.




Top 5 Reasons for Being Green

2. In 2004, the annual “Infrastructure” cost for a 1U server in a data center surpassed its purchase cost.



Top 5 Reasons for Being Green

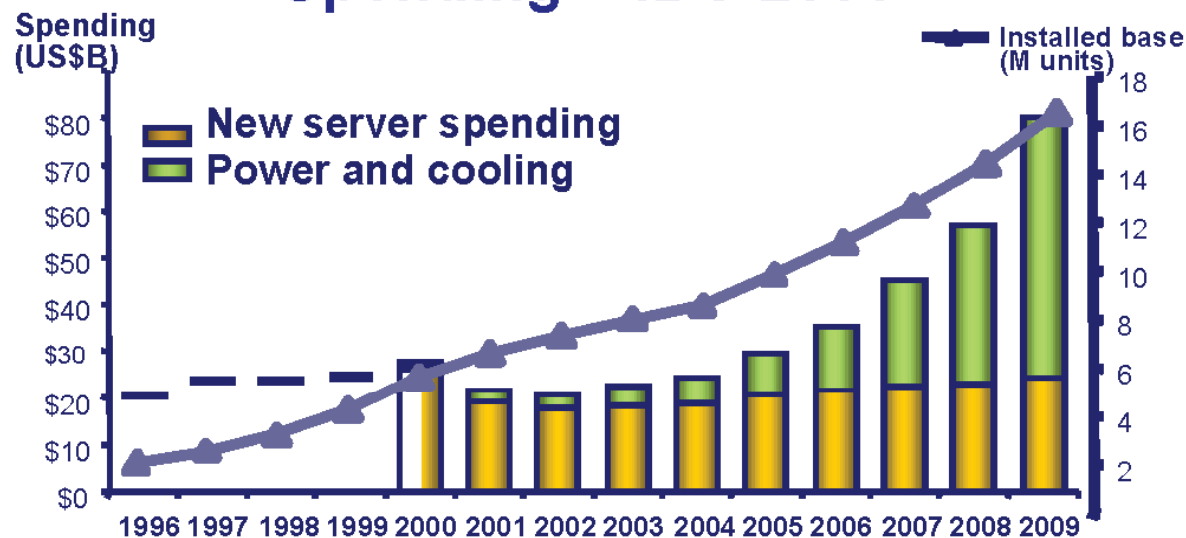
3. In 2008, the annual “Energy” cost for a 1U server in a data center surpassed its purchase cost.




Top 5 Reasons for Being Green

4. In 2008, annual aggregate cost for “power & cooling” (\$26M) exceeds annual server spending (\$22M).


Power and cooling exceeds server Spending – IDC 2006



GREEN DESTINY: LOW-POWER SUPERCOMPUTER



Only Difference? The Processors



GREEN DESTINY "REPLICA": TRADITIONAL SUPERCOMPUTER

Top 5 Reasons for Being Green

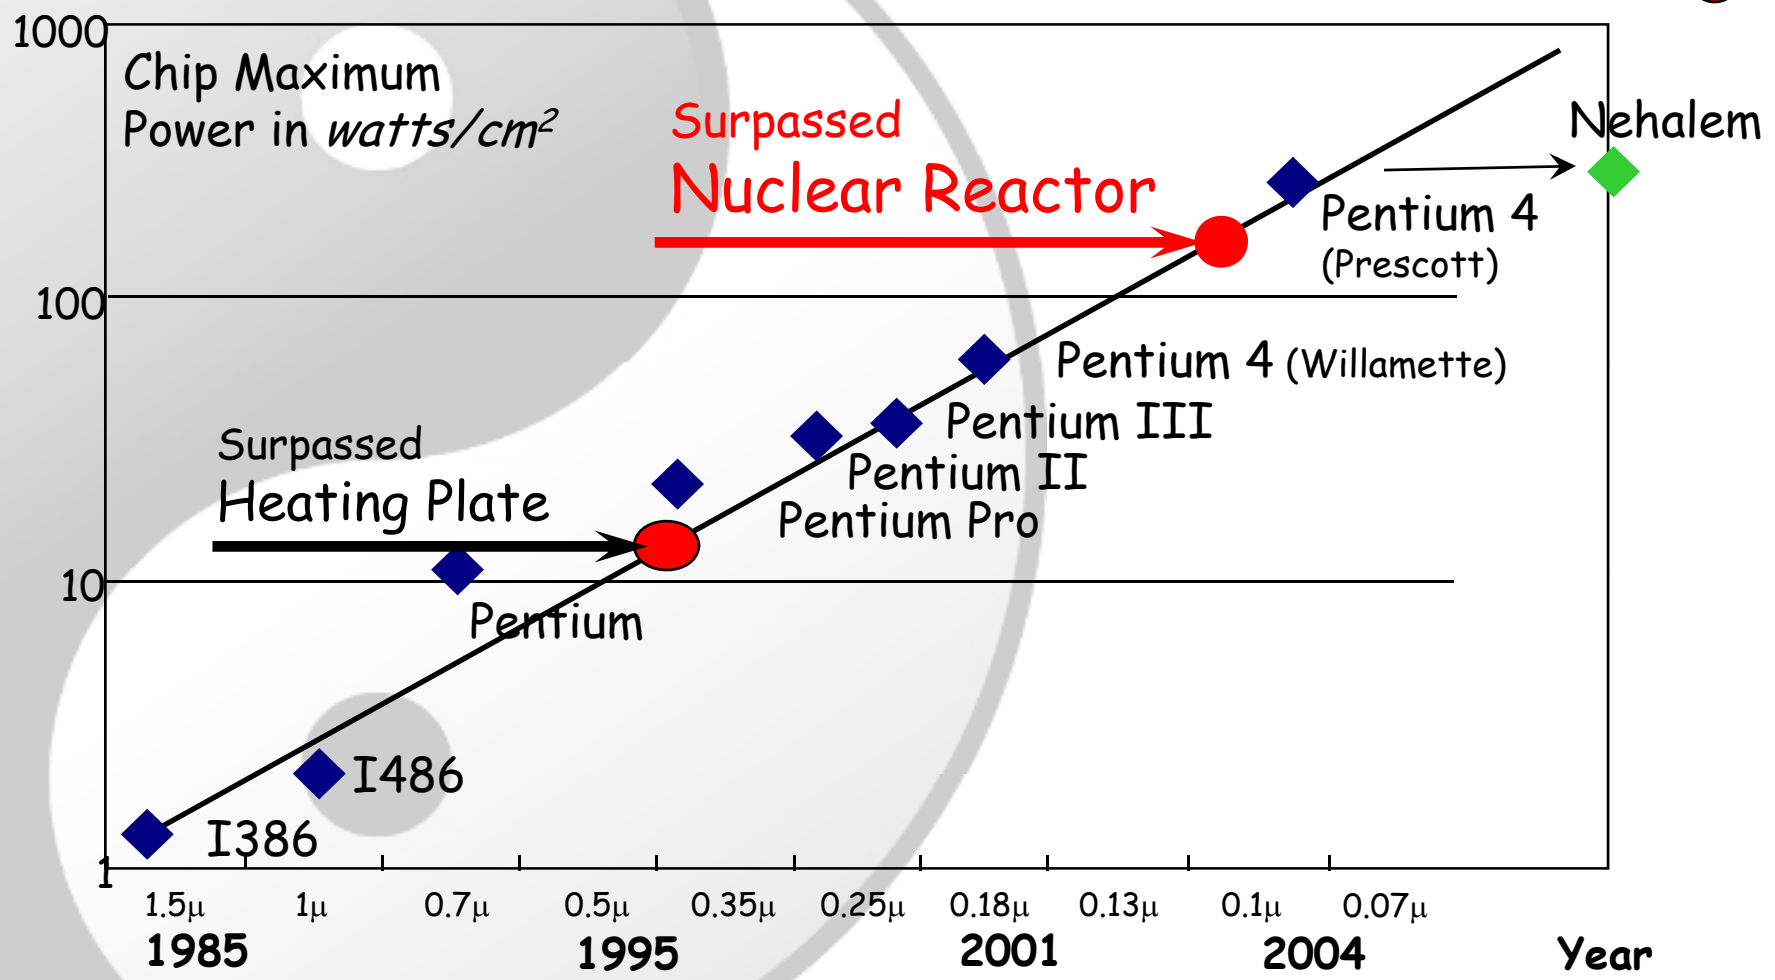
5. “Too much” power affects *efficiency, reliability, availability*.

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, 3 rd -party HW.
Google <i>(projected from 2003)</i>	~450,000	~550 reboots/day; 2-3% machines replaced/yr. – HW outage sources: storage, memory. Availability: ~100%.


Source: Daniel A. Reed

Moore's Law for Power Consumption

Rocket
Nozzle

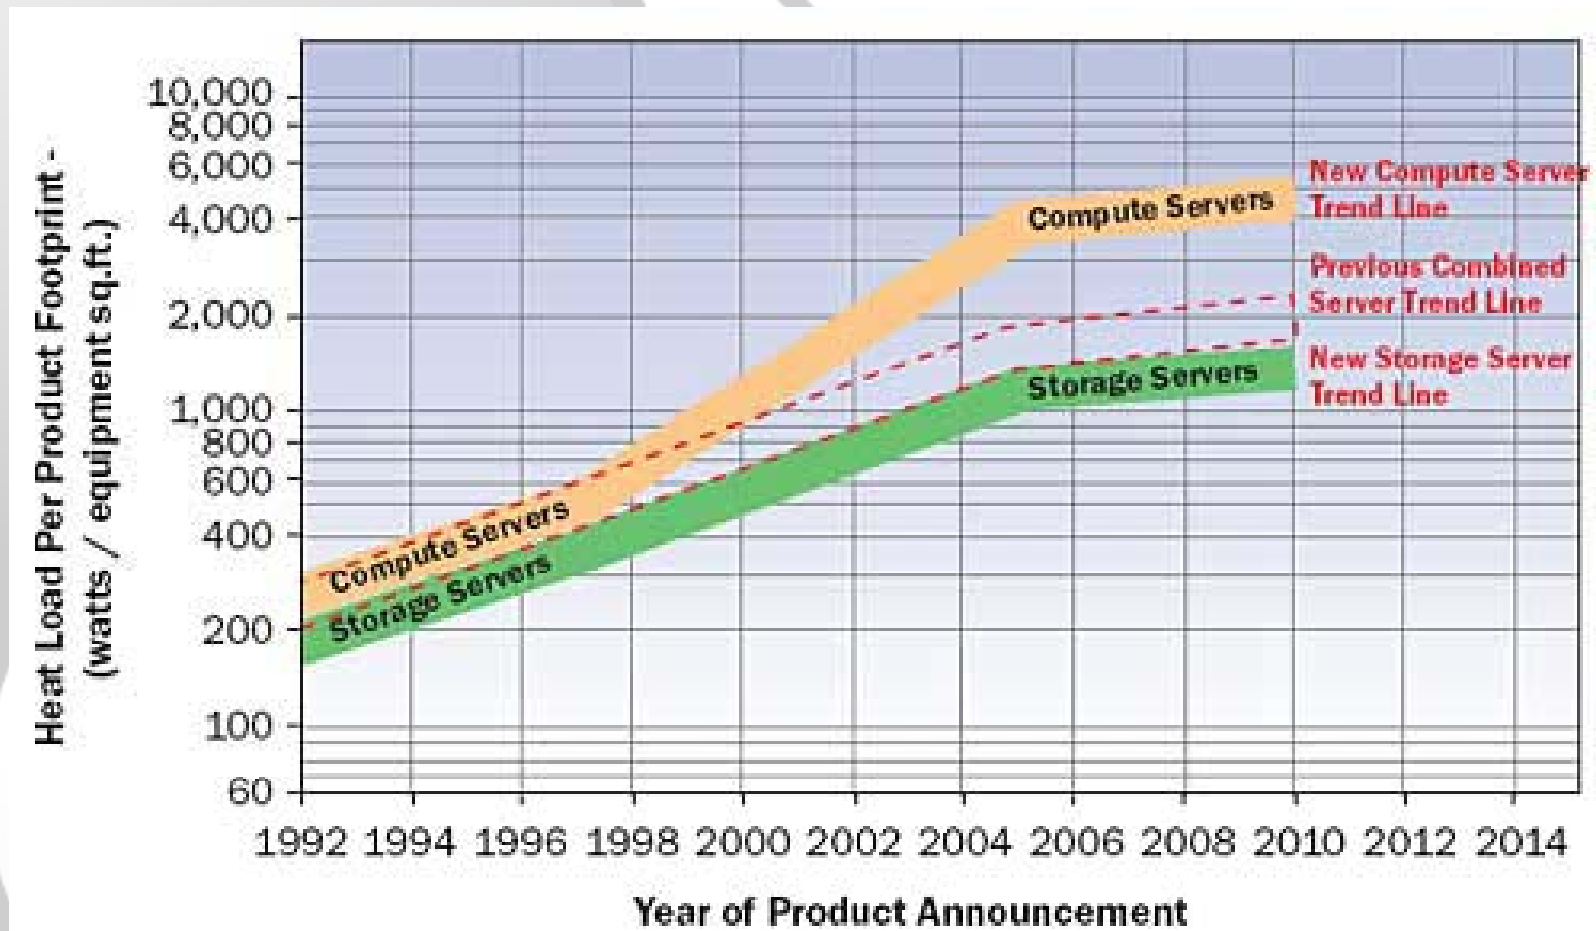


Average Power of an x86 Server



Source: IDC, 2006

Heat Load per Square Foot



Brief History:

From Green Destiny to The *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-Performance, Power-Aware Computing
 - 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”
- **9/2006:** Launch of *Green500* Web Site and RFC
 - <http://www.green500.org/> Generates feedback from hundreds

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)
 - *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 the TOP500 List at the time)

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

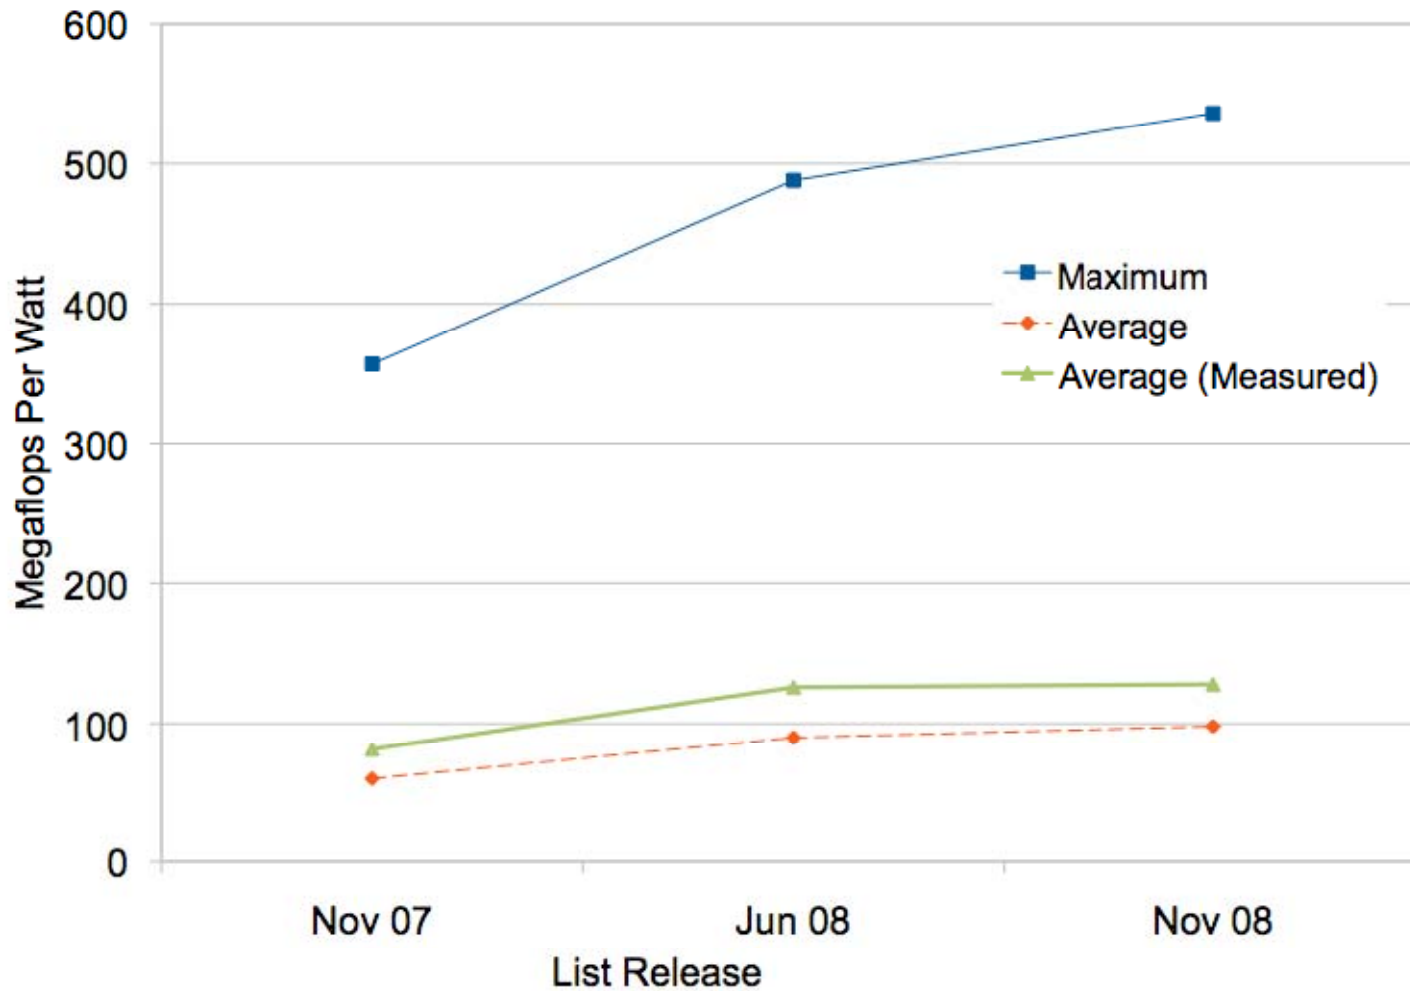
The Growth of The Green500 List

- From 0 to 10,000 page hits per month (on average)
- Number of reported sites increased by 20%
- The Green500 List BoF at SC|08
 - SRO: Standing Room Only
 - Its importance led to filming of the session

Initial Observations: Year One

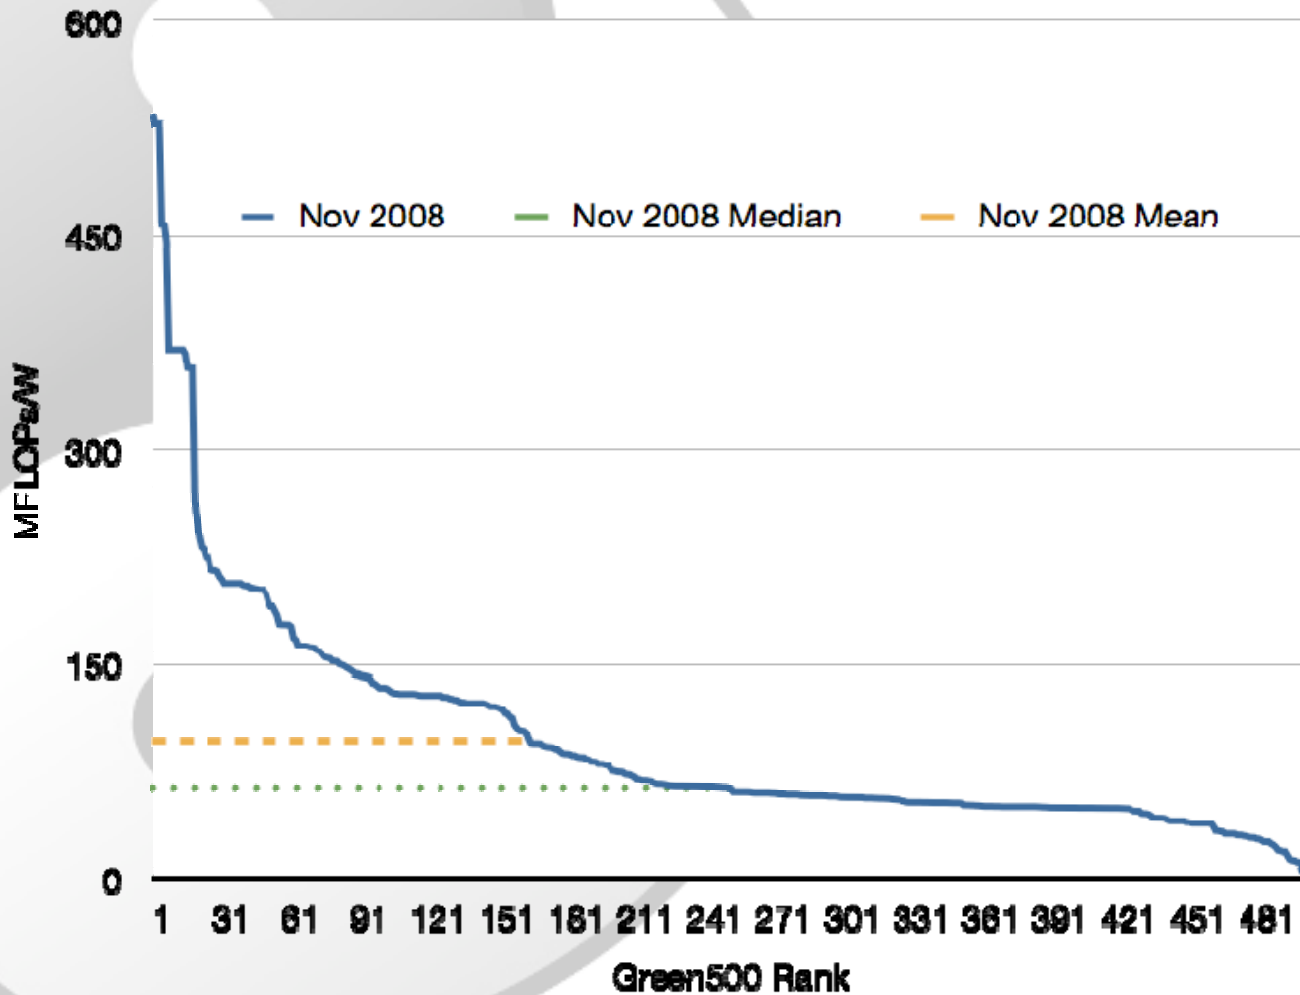
- Surpassed the 500 Mflops/W milestone
- Over 50% of the top supercomputers are now reporting their energy usage.
- 18% more hardware manufacturers participating!
- Many more to come!

Overall Power Efficiency (Mflops/Watt)

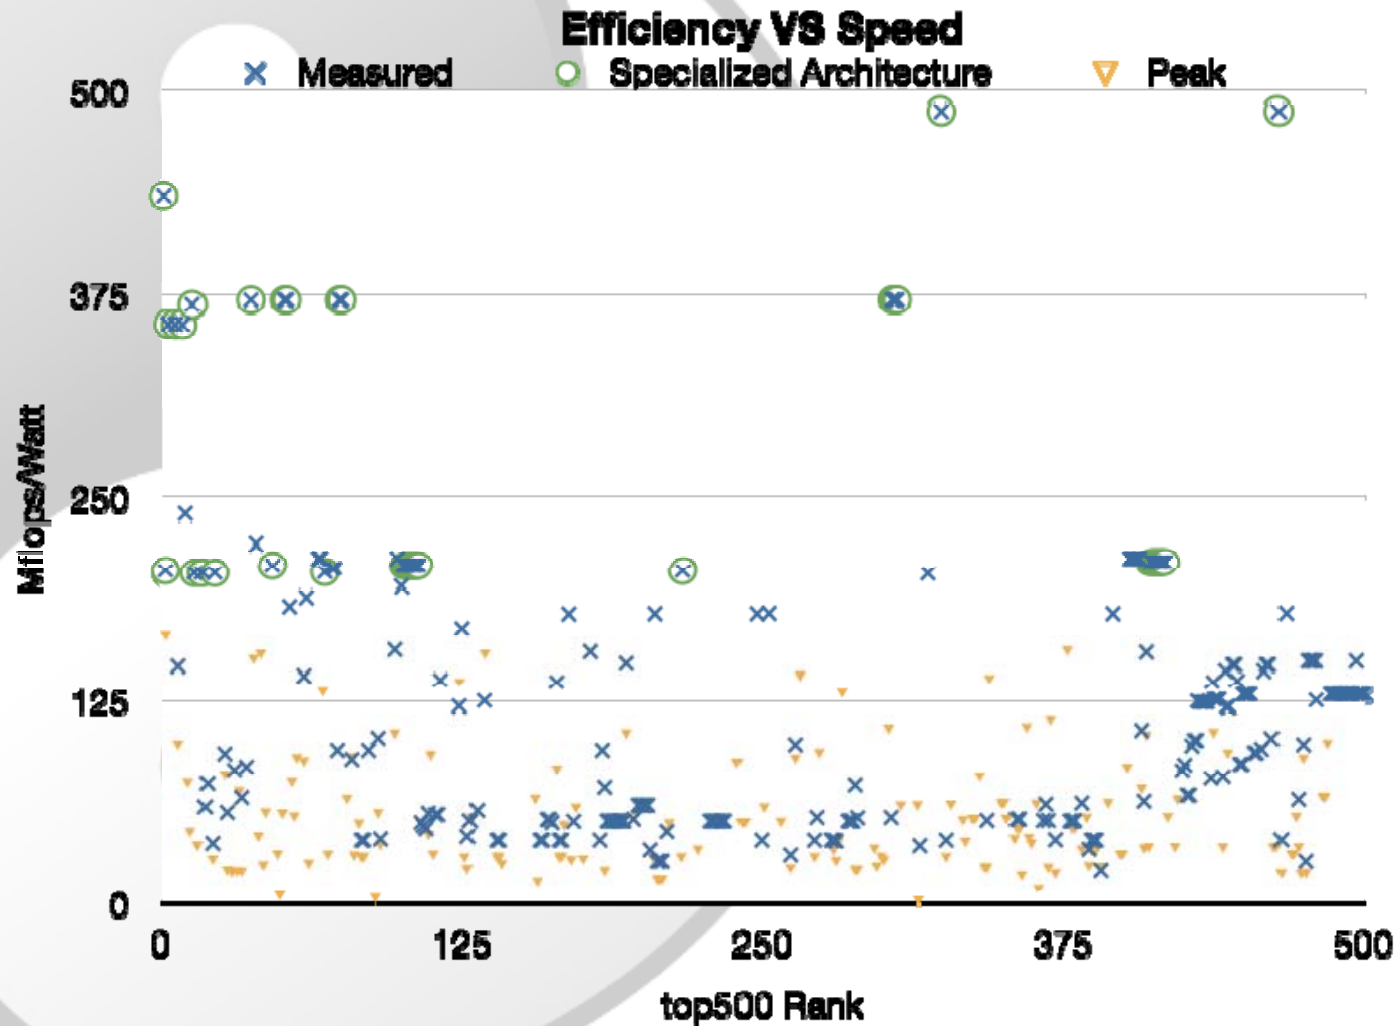


Overall Distribution for Power Efficiency


Efficiency Distribution Nov 2008



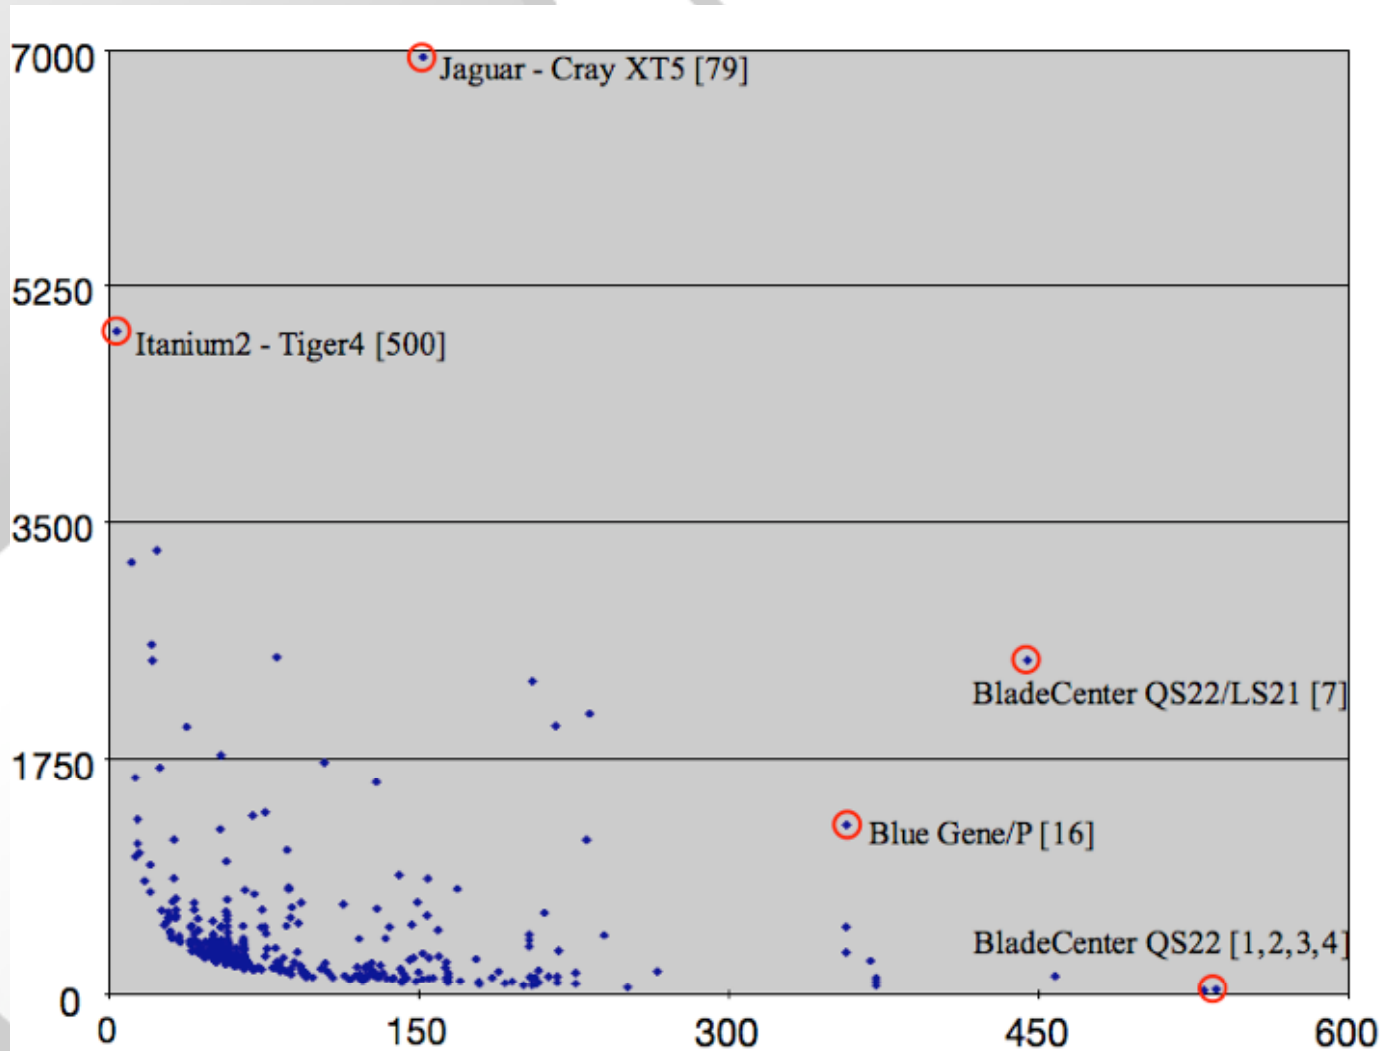
Power Efficiency vs. Performance (i.e., Speed)



Power Efficiency vs. Performance (i.e., Speed)



Power Consumption vs. Mflops/Watt




Commodity Catches Custom ...

- During “The Green500 List: Year One” ...
 - Commodity supercomputers approach “custom” supercomputers in power efficiency
- Custom
 - Blue Gene/L: 210 Mflops/Watt
 - Blue Gene/P: 370 Mflops/Watt
- Commodity
 - Intel quad-core systems achieve *265 Mflops/Watt*

The Green500 List by Country

- Top 10 Countries Ranked by Aggregate Power



The Green500 List at Year One

- The Green500 BoF

1. *What to Measure?* Continue to focus on the power efficiency of the machine and do not consider the cooling facilities or the accompanying infrastructure.
2. *What Machines to Include?* Expand the Green500 List to be more than just a re-ordering of the TOP500 List ranked by power efficiency in FLOPS per watt.
3. *What Benchmark(s) to Use?* Investigate a different benchmark (or benchmark suite) that more appropriately stresses and captures the performance of the entire machine.
4. *What Metric to Use?* Investigate a different performance metric to evaluate energy efficiency.
5. *How to Present?* Enable different ways to view the Green500 data.

Acknowledgements

- Jeremy Archuleta and Thomas Scogland
 - The folks that get the *real work* done behind the scenes
- Kirk W. Cameron
 - Co-Lead on The Green500 List
- Visits and feedback from the tens of thousands people who make *The Green500 List* a part of their day ...

Conclusion

- The average overall power efficiency has doubled every 18 months.
 - Sustainable? Unlikely. “Low-hanging fruit” gains have been made.
- The power efficiency of commodity supercomputers on the Green500 now surpasses the power efficiency of the custom-built IBM Blue Gene/L.
- The U.S., U.K., and Japan are by far the largest power consumers. In addition, China shows astronomical growth in its power consumption.
- *Expect substantial changes in **The Green500 List: Year Two**.*
 - Adapt and adopt a number of the suggestions made at the Green500 BoF at SC|08.