Making a Case for a Green500 List*

Sushant Sharma[1], Chung-Hsing Hsu[1], and Wu-chun Feng[2]

[1]Los Alamos National Laboratory
Advanced Computing Lab.
Los Alamos, NM 87545 USA
{sushant, chunghsu}@lanl.gov

[2]Virginia Polytechnic Institute and State University
Dept. of Computer Science
Blacksburg, VA 24061 USA
feng@cs.vt.edu

Abstract

For decades now, the notion of "performance" has been synonymous with "speed" (as measured in FLOPS, short for floating-point operations per second). Unfortunately, this particular focus has led to the emergence of supercomputers that consume egregious amounts of electrical power and produce so much heat that extravagant cooling facilities must be constructed to ensure proper operation. In addition, the emphasis on speed as the *performance metric has caused other performance metrics to be largely ignored, e.g., reliability, availability, and usability. As a consequence, all of the above has led to an extraordinary increase in the total cost of ownership (TCO) of a supercomputer.*

Despite the importance of the TOP500 List, we argue that the list makes it much more difficult for the high-performance computing (HPC) community to focus on performance metrics other than speed. Therefore, to raise awareness to other performance metrics of interest, e.g., energy efficiency for improved reliability, we propose a Green500 List *and discuss the potential metrics that would be used to rank supercomputing systems on such a list.*

1 Motivation

Would it be correct to say that supercomputers today have reached efficiency levels that no one could have ever imagined decades ago? Depending on the perspective, one could argue that the answer might be yes as well as no. "Yes" if one considers efficiency as *only* the ability to perform a certain number of instructions per second on a given supercomputer. "No" if one starts to consider other factors such as reliability, availability, and total cost of ownership (TCO) [3], just to name a few.

Currently, the focus of the TOP500 List (http://www.top500.org/) is solely on the performance metric of speed, as defined by FLOPS, short for floating-point operations per second. While this focus has led to supercomputers that can complete hundreds of trillions of floating-point operations per second, it has also led to supercomputers that consume egregious amounts of electrical power and produce so much heat that extravagant cooling facilities must be constructed to ensure proper operation.

System	CPUs	MTB(I/F) (Hours)	Power (kW)	Space (Sq Ft)
ASC Q	8,192	6.5	3800	20,000
ASC White	8,192	40 ('03) 5.0 ('01)	1000	10,000
PSC Lemieux	3,016	9.7	N/A	N/A

MTB(I/F): Mean Time Between (Interrupts/Failures)

Table 1. Reliability and Availability of HPC Systems.

For instance, Seager of Lawrence Livermore National Laboratory (LLNL) notes that the large consumption of electricity to power and cool his supercomputers leads to exorbitant energy bills, e.g., $14M/year ($8M to power and $6M to cool) [15]. While at Los Alamos National Laboratory (LANL), the building for the ASC Q supercomputer cost nearly $100M to construct. Even with such extravagant facilities in place, the excessive heat generation impacts the reliability and availability of such systems, as shown in

*This work was supported by the DOE ASC Program through Los Alamos National Laboratory contract W-7405-ENG-36. Available as LANL technical report LA-UR 06-0793.

Table 1 [14].[1] Therefore, not too surprisingly, all of the above results in an astronomical increase in the total cost of ownership (TCO).

With the above considerations in mind, we argue for *the need to maintain a list where the performance metric of interest is not only speed but also energy efficiency as it relates to reliability and availability*. Therefore, we propose a *Green500 List* and discuss the potential metrics that would be used to rank supercomputing systems on such a list.

2 Background

Efforts towards building energy-efficient supercomputers include Green Destiny [3, 17], a 240-processor supercomputer that consumed just 3.2 kilowatts (kW) of power when booted diskless.[2] Although this low-power supercomputer was criticized for its computing ineptitude, Green Destiny *with its customized high-performance code-morphing software* produced a Linpack rating, i.e., 101 Gflops, that was equal to that of a contemporary 256-processor SGI Origin 2000 at the time. Furthermore, the extraordinarily low power consumption of Green Destiny resulted in an extremely reliable supercomputer that had *no unscheduled downtime in its 24-month existence*. It is also important to note here that Green Destiny never required any special cooling or air filtration in order to keep it running.

With efforts such as Green Destiny from 2001-2002, microprocessor vendors have been slowly giving up on the power-hungry, clock-speed race and focusing more on efficient processor design. For example, in October 2004, Intel announced that after years of promoting clock speed as the most important indicator of processor performance, it now believes that introducing multicore products and new silicon features are the best ways to improve processor performance [11]. A month later in November 2004, the energy-efficient IBM BlueGene/L debuted at #1 on the TOP500 Supercomputer List using slowly-clocked 700-MHz PowerPC processors in spite of the availability of PowerPC processors with much higher clock speeds, and hence, more power-hungry appetites. More recently, PA Semi announced its PWRficient™Processor Family, which is based on the Power Architecture™(licensed from IBM). As noted by the company's renowned CEO, Dan Dobberpuhl, PA Semi is aiming to "really drive a breakthrough in performance per watt." [16]. Thus, the above evidence indicates that the commercial industry is moving more towards lower-power and more energy-efficient (but still high-performing) microprocessors.

An alternative approach towards energy-efficient HPC is to use existing power-hungry microprocessors but to leverage an interface to the microprocessor that allows for the dynamic scaling of a microprocessor's clock frequency and supply voltage, as the power consumption of a microprocessor is directly proportional to the clock frequency and the square of the supply voltage. Such research has gained significant traction in the HPC community [2, 4, 5, 6, 7, 9].

Irrespective of the approach towards energy-efficient supercomputing, we believe that there exists a need to develop an alternative to the TOP500 Supercomputer List: the *Green500 Supercomputer List*. But creating such a list means determining what metric(s) to use to rank the supercomputers. *The purpose of this paper is to decide on such a metric and to use that metric to rank supercomputers relative to energy efficiency.*

2.1 Which Metric?

Supercomputers on the TOP500 List use *FLOPS* — short for floating-point operations per second — as the evaluation metric for performance relative to speed. However, the HPC community now understands that supercomputers should not be evaluated solely on the basis of speed but should also consider metrics related to usability, availability, and energy efficiency. With respect to the latter, researchers have borrowed the ED^n metric[3] from the the circuit-design domain in order to quantify the energy-performance efficiency of different systems [1, 8, 12, 13].

In [7], Cameron et al. propose a variant to the ED^n metric. Specifically, they introduce a weighting variable called ∂ that could be used to put more emphasis on energy E or on performance D, depending on what is of interest to the end user. In short, the end user is allowed to choose the value for ∂. What this means is that the end user can ultimately choose from an infinite number of variants of the ED^n metric, but it still leaves the problem of what value of ∂ should the end user choose and what value, if any, should be used to order the Green500 Supercomputer List. On the other hand, Hsu and Feng demonstrate how various ED^n metrics are arguably biased towards massively parallel supercomputing systems [10]. Rather than use an ED^n-based metric, they ultimately "fall back" to using

[1] Arrhenius' equation, as applied to microelectronics, projects that the failure rate of a compute node in a supercomputer doubles with every 10^{o}C (18^{o}F) rise in temperature.

[2] 3.2 kW is roughly equivalent to the power draw of two hairdryers.

[3] E is the energy being used by a system while running a benchmark, and D is the time taken to complete that same benchmark.

the $FLOPS/watt$ metric for energy efficiency. All this suggests that there is still no consensus amongst HPC researchers on which metric to choose for calculating the energy efficiency of a supercomputer.

Rather than simply adopt an energy-efficiency metric and apply it to our tested systems (and even the systems on the TOP500 List), we first present the results of *various* energy-efficiency metrics across a multitude of parallel-computing systems. Next, we provide some analysis and insight into what factors should be considered when comparing the energy efficiency of different supercomputers, using the currently available metrics. Based on our analysis, we then make a case for a Green500 Supercomputer List, an energy-efficient list that will implicitly capture the performance metrics of speed and energy usage. In addition, we will also discuss (1) *how* the results from a particular efficiency metric vary when only CPU power consumption is used instead of total system power and (2) *when* CPU power consumption should be used in calculating energy efficiency instead of total system power.

The remainder of the paper is organized as follows. Section 3 presents our experimental setup, experimental results, and discussion of the results. In particular, we present and discuss the results of applying various efficiency metrics across systems with different architectures. Section 4 shows how a subset of supercomputers from the TOP500 list would be ranked in a Green500 Supercomputer List. Finally, we present our conclusions in Section 5.

3 Experiments

In this section, we first describe the experimental set-up that we use to take our performance measurements. Then, we present and discuss the results of applying different energy-efficient metrics across a variety of parallel-computing platforms.

3.1 Experimental Set-Up

Figure 1 shows the hardware set-up for our experiments. At the upper left is the profiling computer, which records the data that is measured on our Yokogawa digital power meter, shown in the lower left of the figure. To measure the power consumption of the parallel-computing system that appears on the right-hand side of the figure, we connected our power meter to the same power strip as the parallel-computing system. The power meter then continuously samples the instantaneous wattage at a rate of 50 kHz (i.e., every 20 μs) and delivers the readings to the profiling computer.

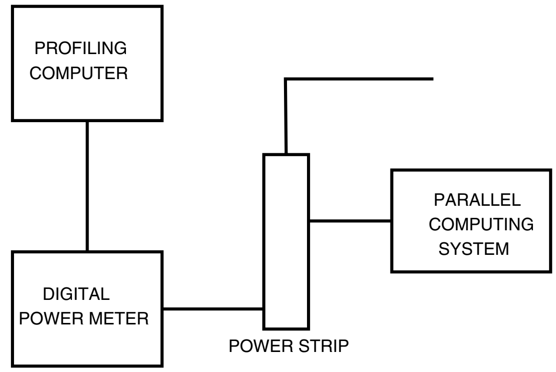

Figure 1. Experimental Set-Up for Benchmark Tests.

For the purposes of comparing results, we kept the software configuration across all the parallel-computing systems as similar as possible. Relative to the application software set-up, we chose the Linpack benchmark to evaluate the FLOPS performance of our parallel computing systems (for legacy reasons relative to the TOP500 Supercomputer List), compiled Linpack using the latest GotoBLAS library, and ran Linpack using LAM-MPI 7.1.1. The problem size for Linpack was kept the same for all the test runs on all the systems. Relative to the system software, we ran SUSELinux OS (9.3/10.0) with the Linux 2.6.x kernel.

Each parallel-computing system under evaluation consisted of exactly four processors. However, the hardware configuration of each system varied in terms of topology and processor architecture. With respect to topology, the systems varied from having four single-processor systems interconnected via Gigabit Ethernet to two dual-processor systems interconnected via Gigabit Ethernet to a single quad-processor system.

With respect to processor architecture, we drew from three 64-bit families: Intel Pentium4, AMD Athlon64, and AMD Opteron, as shown in Table 2. Cluster *C1* consists of four uniprocessor nodes, each with the latest 3.6-GHz Pentium4 processor, interconnected via Gigabit Ethernet. Cluster *C2* is a single SMP node consisting of four 2.0-GHz Opteron processors. Cluster *C3* consists of four uniprocessor nodes, each with 2.4-GHz Athlon64 processors, interconnected via Gigabit Ethernet. Cluster *C6* contains dual 2.0-GHz Opteron processors in each node with Gigabit Ethernet interconnecting the nodes. Finally, Clusters *C4*, *C5*, and *C7* are static power-aware variants of *C3*. That is, we ran Linpack on cluster *C3* four times but at different frequencies each time, i.e., 2.4, 2.2, 2.0, and 1.8 GHz.

Cluster Name	Processors	Topology	Total Memory (GB)	Linpack (GFlops)	Avg. Power Used (Watts)	Time taken (secs)	ED ($*10^6$)	ED^2 ($*10^9$)	ED^3 ($*10^{12}$)	Flops/ Watt	$V_{\partial=0.5}$	$V_{\partial=-0.5}$
C1	3.6 GHz Pentium4	4 * 1P	4.0	19.550	713.20	315.84	71.14	22.47	7.09	27.41	14.92	33.92
C2	2.0 GHz Opterons	1 * 4P	4.0	12.370	415.90	499.36	103.70	51.79	25.86	29.74	56.74	47.20
C3	2.4 GHz Athlon64	4 * 1P	4.0	14.310	668.50	431.56	124.50	53.73	23.19	21.41	43.17	66.87
C4	2.2 GHz Athlon64	4 * 1P	4.0	13.400	608.50	460.89	129.26	59.57	27.46	22.02	51.84	68.45
C5	2.0 GHz Athlon64	4 * 1P	4.0	12.350	560.50	499.79	140.00	69.97	34.97	22.03	66.07	74.10
C6	2.0 GHz Opterons	2 * 2P	4.0	12.840	615.30	481.01	142.36	64.48	32.94	20.87	60.54	77.44
C7	1.8 GHz Athlon64	4 * 1P	4.0	11.230	520.90	549.87	157.49	86.60	47.62	21.56	88.97	84.29

$$V_{\partial}=E^{(1-\partial)}\ D^{2(1+\partial)}$$

Table 2. Efficiency of different clusters according to various metrics.

3.2 Results and Analysis

Table 2 shows the measured performance numbers (i.e., Linpack rating, average power used, and time taken) and derived efficiency numbers (i.e., ED^n and Flops/Watt) of our clusters, relative to different metrics that exist today. Table 3 provides a Green500 summary of how our clusters ranked according to different efficiency metrics.

Looking at last column of Table 3, Cluster *C2* — the single quad-processor node — clearly consumes the least amount of power while Cluster *C1* — a cluster of uniprocessor nodes — consumes the most. Yet in spite of these disparate power-consumption numbers, both *C1* and *C2* consistently finish in the top two relative to efficiency.

When using any of the standard ED^n metrics, Cluster *C1* always outranks Cluster *C2* because these metrics place greater emphasis on performance than on power consumption. This emphasis becomes much more evident as the value of n increases.

Perhaps the most interesting observation from Table 3 revolves around Cluster *C3*. Though the raw power consumption of *C3* is second only to *C1* (see the last column of Table 3), Cluster *C3* ranks in the *top three* for every ED^n metric but ranks amongst the *bottom two*, relative to the Flops/Watt metric. This extremely large change in rankings is due to the fact that the ED^n metrics place greater emphasis on performance (as n increases) while the Flops/Watt metric effectively "penalizes" Cluster *C3* for consuming too much power, and hence, generating more heat and reducing reliability.

Now, let us examine the efficiency metrics at the extremes. Relative to the Flops/Watt metric, we see that Cluster *C2* is the most efficient and is 39% better than Cluster *C7*, which comes out to be the least efficient, relative to the Flops/Watt metric. On the other hand, relative to the ED^n metrics, the difference in efficiency between the most efficient cluster and the least efficient cluster is 121.3% for the ED^1 metric and increases up to 571.6% for ED^3. Does this huge variation fairly capture the difference in efficiency of the clusters at the extremes?

In our experiments, we also generated data to rank the systems according to more sophisticated variants of ED^n, i.e., $V_{\partial=0.5}$ and $V_{\partial=-0.5}$. For details of how these variants work, please refer to [7]. The results are presented in last two columns of Table 2. Looking at the resultant rankings from these metrics in Table 3, the negative value of ∂ produces the same ranking as the other ED^n metrics. For positive values of ∂, the metric places greater emphasis on performance, i.e., D, than on energy consumption, i.e., E. Therefore, as ∂ increases towards one, the metric approaches the

	Green500 Ranking						TOP500 Ranking	Power500 Ranking
Rank	ED	ED^2	ED^3	$V_{\partial=-0.5}$	$V_{\partial=0.5}$	Flops/ Watt	Flops	Watts
1	C1	C1	C1	C1	C1	C2	C1	C2
2	C2	C2	C2	C2	C3	C1	C3	C7
3	C3	C3	C3	C3	C4	C5	C4	C5
4	C4	C4	C4	C4	C2	C4	C6	C4
5	C5	C5	C5	C5	C6	C7	C2	C6
6	C6	C6	C6	C6	C5	C3	C5	C3
7	C7	C7	C7	C7	C7	C6	C7	C1

Table 3. Ranking of different clusters according to various metrics.

limit E^0D^4 and behaves more like the standard Flops metric, which is used for TOP500 List. This analysis is supported by the numbers in the Flops column and the $V_{\partial=0.5}$ column in Table 3. That is, for positive values of ∂, the ranking order is nearly identical to the TOP500 ranking, which only takes performance into consideration. Based on the above discussion, we believe that Flops/Watt is a more balanced metric for comparing the efficiency of cluster supercomputers.

Finally, despite using small four-processor clusters, it is important to note that the ED^n and Flops/Watt metrics already produce noticeably different results in comparing the efficiency of systems. If this is the case now when there is not as much difference between the cluster systems, then the difference in rankings will be even more pronounced when these metrics are used to rank, say the TOP500 supercomputers — a topic which will be discussed in Section 4.

Metric Usage

Efficiency metrics are not only used for comparing the efficiency of supercomputers, but they also play a major role in comparing or evaluating different power-aware techniques that are used to reduce the power utilization of systems while minimizing the impact on overall execution performance. As additional research on power-aware techniques is completed, the need for metrics to evaluate these techniques will be felt more than ever.

When comparing two different power-aware algorithms, it is very important to understand what the metric that is used to compare the results is reporting and what that metric should actually report. We all know that the CPU consumes the major amount of power in a system. However, the actual percentage power that a CPU consumes varies from system to system. Figure 2 shows the percentage of power that the CPU consumes in three different systems in our lab. In Cluster *C3*, processors consume 43.50% of the total cluster power; while in Cluster $C2$, processors consume a whopping 64.57% of the total cluster power. And in the case of our four-processor laptop cluster ($L - Cluster$), the percentage of power drawn by the processor shoots all the way up to 83.37%.

Most power-aware algorithms try to save power in a system by slowing down the processors and reducing the power consumption of processors only. Now, if a power-aware algorithm 'x' reports 50% *total* power savings on Cluster *C3* cluster, and another algorithm 'y' reports 50% *total* power savings on Cluster *C2*, what does it mean? It means that algorithm 'x' is more effective in saving processor power than algorithm 'y'. It can also be inferred that while comparing two algorithms using a particular metric, it is not always fair to use total system power to calculate power (or energy) efficiency.

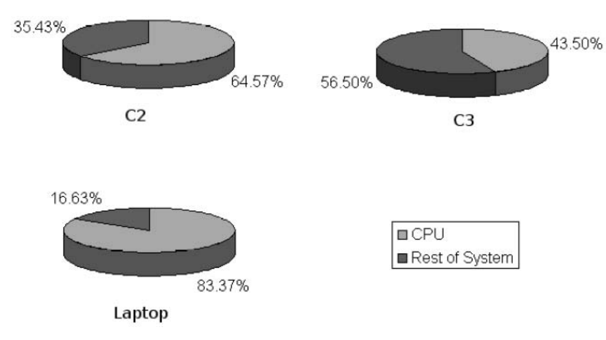

Figure 2. Per-Node Power usage of Clusters.

The decision to use "total system power" or "only the CPU power" depends on the type of systems that are being compared. One should also keep in mind that the processor power usage is not measured directly, it is oftentimes inferred from the total power usage of the system, as shown in [9]. This will result in some information loss. We can say that, if the processors in the systems under comparison consume a similar percentage of total system power, the total system power should be used in the metric for calculating the efficiency. On the other hand, if the percentage of power consumption for processors differ in systems under comparison, it would be much fairer to use only processor power consumption for comparison purposes.

Again looking at Figure 2, it is clear from this figure that it would not be fair to use total system power usage to compare two algorithms, if one of them is running on *L*-cluster whereas other is running on *C2* and *C3* clusters.

4 TOP500 as Green500

We all know that TOP500 list of supercomputers have systems that are quite diverse in both their hardware and software architecture. The first question that everyone should ask in order to rank the TOP500 systems in a Green500 list is *"Which metric should be used?"*. What characteristics should a metric have so that it is not biased towards a particular type of system? Based on the discussion in Section 3, we believe that *FLOPS/Watt* is a better metric to rank the TOP500 supercomputers as part of a Green500 Supercomputer List.

Table 4 presents some of the TOP500 supercomputers and their peak power usage. It also shows the results of using the FLOPS/Watt metric on them. The Linpack performance for these supercomputers was taken from TOP500 list that was released during SC|05 in November 2005. The sources for power usage are various presentations, articles in magazines, newspapers, and the web sites of several of these supercomputers.

Here we elaborate on how we derived the peak power consumption of some of the systems. The web site of ASC Purple[4] reports that the power usage of Purple as 7.5 MW, which is the total power that is required for powering and cooling the supercomputer. The Eurekalert web site[5] reports the power usage of Purple as 8 MW. Looking at the comparisons given in a presentation at one of the BlueGene/L workshops,[6] the power consumption of Purple just to run the machine (but not cool it) is 4.5 MW. Given the numbers reported by Lawrence Livermore National Laboratory for their supercomputers, i.e., for every watt of power consumed by the system, 0.7 watt of power is required to cool it, the number 4.5 MW is consistent with the 7.5 and 8 MW reported from other sources. We used 7.6 MW (i.e., $4.5 * 1.7$) as the total power required to run and cool Purple.

Supercomputer Name	Peak Linpack Performance (GFlops)	Peak Total Power Usage (kW)	MFlops/ Watt	TOP-500 Rank
BlueGene/L	367000	2500	146.80	1
ASC Purple	77824	7600	10.24	3
Columbia	60960	3400	17.93	4
Earth Simulator	40960	11900	3.44	7
MareNostrum	42144	1071	39.35	8
Jaguar-Cray XT3	24960	1331	18.75	10
ASC Q	20480	10200	2.01	18
ASC White	12288	2040	6.02	47

Table 4. TOP500 power usage.

For BlueGene/L, Eurekalert reports a total of 2.5 MW of power required to run and cool the system, and a BlueGene/L presentation reports 1.2 MW to run only the system.

For the Columbia supercomputer at NASA Ames Research Center, Jack Dongarra in one of his presentations[7] reports the total power usage just to run the system is 2 MW. Given that the thermal design power (TDP) of Itanium-2 processors, 10240 of which are used in Columbia, is 130 watts, the power to run just these 10240 processors comes out to be 1.33 MW. So, 2 MW seems reasonable if the other components in Columbia use only 700 kW of power. Now, using the best-case assumption that for every 1 watt of computer power, 0.7 watt is required to cool it, the total power required to run and cool Columbia comes out to be 3.4 MW.

For the Japanese Earth Simulator, the total power usage reported in one of the presentations[8] is 7 MW

[4] http://www.llnl.gov/asc/platforms/purple/

[5] http://www.eurekalert.org/features/doe/2005-06/dlnl-lsb062405.php

[6] http://www.lofar.org/BlueGene/

[7] www.netlib.org/utk/people/JackDongarra/SLIDES/HK-2004.pdf

[8] www.sc.doe.gov/ascr/dongarra.pdf

Relative Rank	TOP500 Order	Green500 Order
1	BlueGene/L	BlueGene/L
2	ASC Purple	MareNostrum
3	Columbia	Jaguar-Cray XT3
4	Earth Simulator	Columbia
5	MareNostrum	ASC Purple
6	Jaguar-Cray XT3	ASC White
7	ASC Q	Earth Simulator
8	ASC White	ASC Q

Table 5. TOP500 Vs Green500.

just to run the system. The total power usage, including power needed for cooling, comes out to be 11.9 MW. The same source also quotes power usage of ASC White as 1.2 MW and for ASC Q as 6 MW, for just powering the machines. Adding in the power to cool the machines increases the total power for White to 2.04 MW and for Q to 10.2 MW.

The Register[9] reports the power usage of MareNostrum to be 630 kW just to run the system. This makes the total power usage of MareNostrum including power required for cooling to be 1.07 MW.

The data sheet of Cray XT3[10] reports that each cabinet of XT3 consumes 14.5 kW of power and houses 96 Opteron processors of 2.4 GHz each. TOP500 list reports 5200 AMD Opteron processors of 2.4 GHz each in Jaguar Cray XT3 system. After calculation, the Jaguar system listed in TOP500 list will use about 54 cabinets and consume about 783 kW of power just to run the system. The total power that will include cooling the system comes out to be 1331 kW.

From Table 5, BlueGene/L, because of its arguably ideal mix of performance and extremely low power consumption, is ranked #1 on both the TOP500 list and the Green500 list. The TOP500 list ranks the Japanese Earth Simulator at #7 (or #4 relative to the systems being considered in Table 5), but it is penalized in the Green500 list for consuming an exorbitant amount of power and is ranked second-to-last among the systems we are presenting. ASC White, on the other hand, is ranked higher in the Green500 list because of its lower power usage.

Another somewhat related question that one can ask here is that while using a metric to list TOP500 supercomputers as Green500, should we use CPU power usage or total system-power usage? The answer is that we should use the total system-power usage. The usage of CPU power should be restricted only for comparing various power-aware algorithms as discussed in the previous section. While listing systems on the Green500 list, we are concerned about the total power usage as a result of using the system. We are ranking the systems as a whole and not just the CPUs. However, for the sake of posterity, Figure 3[11], shows the results of GFlops/Watt but relative to only the CPU power consumption, not total power consumption. Most notably, the PowerPC-based architectures, i.e., BlueGene and MareNostrum, achieve a Flops/watt rating that is about an order of magnitude higher than all other CPU architectures.

5 Conclusion

In this paper, we made a case for a Green500 Supercomputer List — a list that would not only take performance (relative to speed) into consideration but would also take energy efficiency into account when ranking supercomputers. In the long run, this list would help the HPC research community and various vendors by focusing their attention towards factors other than just performance (relative to speed).

References

[1] D. Brooks and M. Martonosi. Dynamically exploiting narrow width operands to improve processor power and performance. In *HPCA-5*, Jan. 1999.

[2] W. Feng. The importance of being low power in high-performance computing. *Cyberinfrastructure Technology Watch (CTWatch)*, (3), August 2005.

[3] W. Feng, M. Warren, and E. Weigle. The bladed beowulf: A cost-effective alternative to traditional beowulfs. In *Proceedings of IEEE Cluster*, Sept. 2002.

[4] X. Feng, R. Ge, and K. Cameron. Power and energy profiling of scientific applications on distributed sys-

[9] http://www.theregister.co.uk/2005/04/13/barcelona_supercomputer/

[10] http://www.cray.com/products/xt3/

[11] www.netlib.org/utk/people/JackDongarra/SLIDES/hpcasia-1105.pdf

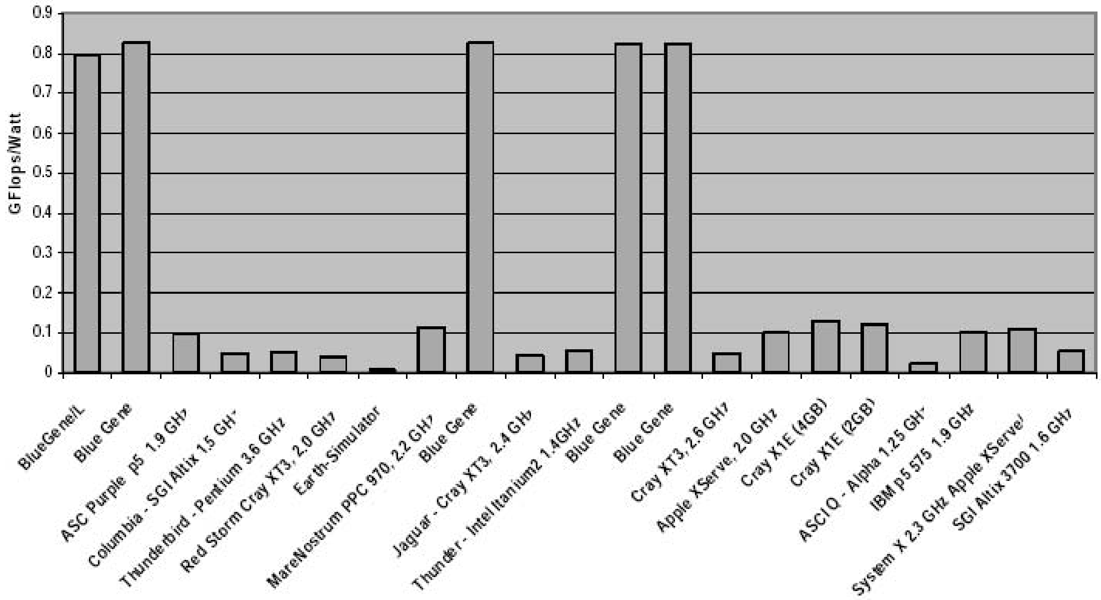

Figure 3. Green500 results (Relative to CPU only).

tems. In *International Parallel and Distributed Processing Symposium (IPDPS)*, Apr. 2005.

[5] V. Freeh, D. Lowenthal, F. Pan, and N. Kappiah. Using multiple energy gears in MPI programs on a power-scalable cluster. In *Principles and Practices of Parallel Programming (PPoPP)*, June 2005.

[6] V. Freeh, D. Lowenthal, R. Springer, F. Pan, and N. Kappiah. Exploring the energy-time tradeoff in MPI programs on a power-scalable cluster. In *International Parallel and Distributed Processing Symposium (IPDPS)*, Apr. 2005.

[7] R. Ge, X. Feng, and K. Cameron. Improvement of power-performance efficiency for high-end computing. In *The First Workshop on High-Performance, Power-Aware Computing (HP-PAC)*, Apr. 2005.

[8] R. Gonzalez and M. Horowitz. Energy dissipation in general purpose microprocessors. *IEEE Journal of Solid-State Circuits*, 31(9):1277–1284, Sept. 1996.

[9] C. Hsu and W. Feng. A power-aware run-time system for high-perfromance computing. In *Supercomputing Conference 2005 (SC05)*, Nov. 2005.

[10] C. Hsu, W. Feng, and J. Archuleta. Towards efficient supercomputing: Choosing the right efficiency metric. In *The First Workshop on High-Performance, Power-Aware Computing (HP-PAC)*, Apr. 2005.

[11] T. Krazit. Intel shelves plans for 4-GHz P4. http://www.pcworld.com/news/article/0,aid,118165,00.asp, October 2004.

[12] A. Martin. Towards an energy complexity of computation. *Information Processing Letters*, 77:181–187, Feb. 2001.

[13] A. Martin, M. Nystrm̈, and P. Pénzes. ET^2: a metric for time and energy efficiency of computation. Kluwer Academic Publishers, 2002.

[14] D. Reed. High-end computing: The challenge of scale. Director's Colloquium, Los Alamos National Laboratory, May 2004.

[15] M. Seager. What are the future trends in high-performance interconnects for parallel computers? In *IEEE Symposium on High-Performance Interconnects (HotI)*, Aug. 2004.

[16] A. Vance. DEC veterans prepare chip challenge for Intel, Amd, Ibm and Sun. http://www.theregister.co.uk/2005/10/24/pasemi_power/, October 2005.

[17] M. Warren, E. Weigle, and W. Feng. High-density computing: A 240-node beowulf in one cubic meter. In *Proceedings of SC2002*, Nov. 2002.