

The Green500 List: Year Two

Wu-chun Feng and Heshan Lin
{feng, hlin2}@cs.vt.edu

Department of Computer Science
Virginia Tech
Blacksburg, VA 24060

Abstract

The Green500 turned two years old this past November at the ACM/IEEE SC|09 Conference. As part of the grassroots movement of the Green500, this paper takes a look back and reflects on how the Green500 has evolved in its second year as well as since its inception. Specifically, it analyzes trends in the Green500 and reports on the implications of these trends. In addition, based on significant feedback from the high-end computing (HEC) community, the Green500 announced three exploratory sub-lists: the Little Green500, the Open Green500, and the HPCC Green500, which are each discussed in this paper.

1. Introduction

The focus of “performance at any cost” in high-end computing (HEC) has led to supercomputers that consume vast amounts of electrical power and require large cooling facilities to ensure proper operation [4], [5], [15]. To address this inequity, the Green500 List was created [7] to provide a complementary view to the TOP500 [17]. Since its debut in November 2007, the Green500 has garnered increasing interest from industry, government labs, and academia. The most recent release of the Green500 at ACM/IEEE SC|09 marked its two-year anniversary.

With two year’s worth of Green500 data now available, this paper tracks the progress of energy-efficient supercomputing on the Green500, analyzes trends in the Green500, and reports on the implications of these

This work was supported in part by NSF grant CCF-0848670 and by Supermicro.

trends. In addition, reflections from the first year of the Green500 [8] brought with it new feedback from the HEC community, which strongly indicated that the single Green500 List would *not* be sufficient in addressing the needs of the HEC community interested in the pursuit of sustainable green supercomputing. Of particular interest to the HEC community were the following:

- Lowering the “barrier to entry” to the Green500 by being more inclusive with the definition of a supercomputer.
- Rewarding innovation, not only in hardware, but also in software via alternative energy-efficient algorithms, power-aware compilation, or adaptive run-time systems, for example.
- Investigating the possibility of using a different benchmark(s) to evaluate the energy efficiency of supercomputers.

Each of these items will be addressed in greater detail later in this paper.

The rest of the paper is organized as follows. Section 2 provides background information on the need for a Green500 List. Section 3 presents an analysis of trends in the Green500 List and discusses their implications. In addition, the section also discusses how the list is evolving with the announcement of three new exploratory lists. Finally, Section 4 presents concluding remarks and discussion of future work.

2. Background

While both the TOP500 and Green500 focus on *performance*, the former uses *speed* as its performance metric, as measured by *floating-point operations per second* (i.e., FLOPS) while the latter uses *energy effi-*

ciency, as measured by *FLOPS/watt*. Why should we care about energy efficiency? As Belady notes in [5] and as shown in Figure 1, the *annual energy cost* for a 1U server in 2008 surpassed its purchase cost. This, coupled with the annual infrastructure cost for a 1U server exceeding its purchase cost in 2004, contributes to a fast-rising total cost of ownership (TCO).




Figure 1. Annual Amortized Costs in the Data Center.

The “Green Destiny” cluster [6], [9] aggressively addressed the above issues by *turbocharging* low-power x86-compatible Transmeta processors via proactive dynamic compilation at run time and by *consolidating* common hardware infrastructure via blade technology in order to massively improve processor performance, particularly floating-point operations, and significantly reduce per-node power consumption, respectively. The use of low-power blades by Green Destiny enabled compute nodes to be packed more densely, i.e., increased FLOPS per square foot, while *simultaneously reducing power density*, i.e., decreasing watts per square foot. While the HEC community clearly understood the benefits of the former, i.e., increasingly computational density, the latter, i.e., reducing power density, was largely ignored. As a result, though the subsequent proliferation of server blades created more densely-packed compute nodes or servers, it also caused a much faster than expected increase in the power density of compute servers [3], [18], as shown in Figure 2.

Continuing to “feed the supercomputing beast” is *not* an environmentally sustainable solution as the demand for energy continues to grow rapidly. As noted




Figure 2. ASHRAE Projection of Heat Density.

earlier, the annual cost to power and to cool a 1U server has already exceeded the actual cost of the 1U server. By 2014, the infrastructure and energy cost (I&E), as shown in Figure 1, will account for 75% of the total cost of a 1U server. This is in stark contrast to the early 1990s when I&E only accounted for a mere 20% of the total cost of a 1U server.

Consequently, the energy efficiency of supercomputing systems needs to be improved in order to make more efficient use of electricity. This is particularly true of many existing datacenters and HEC centers that are already constrained by the physical wattage entering their respective buildings, e.g., the National Security Agency (NSA) [10], [19]. In the meantime, inefficient datacenters continue to be constructed, with the data-center proprietors oftentimes negotiating extraordinarily favorable energy deals with electricity suppliers to build or upgrade power substations near their datacenters. Examples of such datacenters arguably include, and are obviously not limited to, the following: Oak Ridge National Laboratory [20], National Center for Supercomputing Applications [16], and the aforementioned National Security Agency [10], [19].

And when not enough cheap power can be built at or near the datacenters, many institutions simply *move* their datacenters to the power source. Examples arguably include, and again are not limited to, the following: Google [2], [15], Microsoft [21], and NSA [14].

Hence, the mission of the Green500 is to *raise awareness in the energy efficiency of supercomputing and drive energy efficiency as a first-order design constraint* (on par with performance).

3. Analysis and Evolution of the Green500

In this section, we first present an analysis of the Green500, which ranks the energy efficiency of top 500 supercomputers in the world. We then discuss the evolution of the Green500 into three new exploratory lists.

3.1. Analysis

Overall Energy Efficiency. As shown in Figure 3, where the “MFLOPS/W” is plotted over the Green500 rank for machines on the three November lists, the overall energy efficiency of supercomputers has continued to improve over the past two years. While improvements are observed over the full range of machines, the majority comes from the top half of the lists, similar to what we observed in our last year’s report [8]. To obtain further insight, Figure 4 plots the maximum and average energy efficiency of the past five releases. Most notably, *both metrics have more than doubled over the two years*. More specifically, from November 2007 to November 2009, the *maximum energy efficiency* increased by 102% from 357 MFLOPS/W to 723 MFLOPS/W while the *average energy efficiency* increased by 108% from 60 MFLOPS/W to 125 MFLOPS/W.

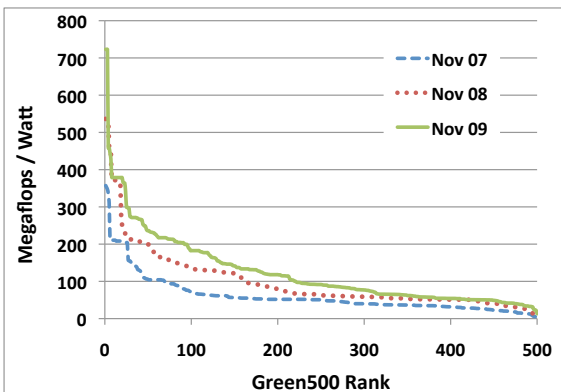


Figure 3. Energy Efficiency by Green500 Rank

After “Year One” of the Green500 [8], we surmised that the growth of overall energy efficiency might follow Moore’s Law, as strongly suggested by the first year’s 75% improvement on the measured average energy efficiency. While the average energy efficiency improved only 37% in “Year Two” of the Green500, it still resulted in a Moore’s Law effect with respect

to average energy efficiency, which doubled over the 24-month existence of the Green500.




Figure 4. Maximum and Average Energy Efficiency per Green500 Release.

Figure 4 also illustrates non-linear growth in the maximum energy efficiency across Green500 Lists. After an initial leap, there was an apparent slowdown from June 2008 to June 2009. One possible explanation of the initial rapid growth is the “plucking” of low-hanging fruit in energy efficiency, e.g., using existing low-power microprocessors to build supercomputers. After the initial stage, further improvements would require novel energy-centric architectural designs, and hence, would take longer to achieve. One example of such architectural innovations is the IBM QPACE (quantum chromodynamics parallel computing on the Cell). Built on the energy-efficient PowerXCell 8i processors connected by a network of FPGA (Field Programmable Gate Arrays), the IBM QPACE delivered a 35% boost in the MFLOPS/W rating over the top-ranked green machine on the previous list, i.e., June 2009, and landed at the top of the November 2009 release of the Green500.

Energy Efficiency vs. Speed. Many in the HEC community have inquired about how the energy efficiency of supercomputers correlates to their performance. Intuitively, one would think that lower-ranked machines in the TOP500 list tend to have smaller system sizes, and hence, can be more energy efficient because of the overhead of scaling a system. More specifically, as the system size increases, the FLOPS measured by LINPACK increases sub-linearly, while the power consumption increases perfectly linearly to super-linearly. To verify our intuition, we took a statistical approach to the problem, as outlined below.

The linear correlation between two variables is com-

monly measured by the Pearson Correlation Coefficient [1], which is defined as the covariance of the two variables divided by the product of their standard deviations as follows:

$$\rho_{X,Y} = \frac{cov(X,Y)}{\sigma_x\sigma_y} = \frac{E[(X - \mu_X)(Y - \mu_Y)]}{\sigma_x\sigma_y} \quad (1)$$

where E is the expectation, μ is the mean, and σ is the standard deviation. The Pearson Correlation Coefficient yields a value between -1 and 1. A value greater than 0 indicates a positive correlation between the variables, whereas a value less than 0 indicates a negative correlation. The closer that the absolute value of the coefficient is to 1, the stronger is the correlation.

We calculate the Pearson Correlation Coefficients between the registered MFLOPS/W and the LINPACK performance of machines across the past three November lists. Table 1 shows our results.¹ Surprisingly, all the correlation coefficients are greater than 0, suggesting that there is no direct support to the hypothesis that smaller machines are more energy efficient in the examined lists. However, the relatively low coefficient values (i.e., less than 0.3) also indicate weak linear association between energy efficiency and performance.

	Correlation Coefficient	P Value
Nov. 2007	0.284	9.749e-11
Nov. 2008	0.292	2.816e-11
Nov. 2009	0.252	1.101e-08

Table 1. Correlation between Energy Efficiency (MFLOPS/W) and LINPACK Performance (TFLOPS).

Figure 5 plots the registered MFLOPS/W over the LINPACK performance of machines from the November 2009 list. As expected from the low coefficient value, the linear regression line fits poorly with the observed MFLOPS values. It is worth noting that the slope of the regression line is increasing, which suggests the energy efficiency tends to be better for the faster machines. This counterintuitive result can be attributed to the fact that newly built large-scale machines strongly tend towards being more energy efficient than earlier-built small machines. For example, Roadrunner, the 2nd most powerful supercomputer in the world, also ranks high in Green500 (i.e., 6th).

Figure 5 also shows that the most energy-efficient

¹The P value associated with the correlation analysis of each November list is also presented. The P value is a statistic that measures the significance of the correlation. In general, a P value less than 0.05 means that the analysis results is statistically significant.

machine, IBM QPACE powered by the PowerXCell 8i processor and FPGA, is nearly 3 times more efficient (723 MFLOPS/watt vs. 253 MFLOPS/watt) but more than 40 times slower (42 TFLOPS vs. 1759 TFLOPS) than the most powerful supercomputer, i.e., Jaguar at Oak Ridge National Laboratory.




Figure 5. Energy Efficiency vs. Performance for the November 2009 Release.

Commodity vs. Custom. Figure 6 compares the energy efficiency achieved by the top-ranked commodity and custom-built machines over the past two years. The custom-built machines in our study refer to specialized architectures including IBM Blue Gene/L and Blue Gene/P, IBM BladeCenter QS22, and IBM QPACE. As can be seen, both curves show similar growing trends. However, the top-ranked custom-built machines maintain a MFLOPS/W rating that is *two times higher* on average than the most energy-efficient commodity machines.

As we observed last year [8], leading-edge commodity machines started to catch-up to previous generation custom-built machines from two years ago. Specifically, a supercomputer based on Intel’s 45nm low-power quad-core Xeon surpassed the overall energy efficiency of IBM Blue Gene/L. In addition, a commodity machine based on the Intel Xeon E5540/E5550 CPU and the AMD Radeon HD4870 GPU accelerator achieved 379.2 MFLOPS/W, making it more efficient than IBM Blue Gene/P (370 MFLOPS/W).

Power Consumption by Country. Figure 7 depicts the aggregate power consumption of the top 10 power-consuming countries in the Green500. As can be seen in the figure, the aggregate power consumption of the top 10 countries remained approximately the same from November 2008 to November 2009. This is surprising as the aggregate power consumed by all countries




Figure 6. Energy Efficiency of the Top-Ranked Commodity and Custom-Built Machines.

increased 35% from 200.4 MW to 272.2 MW in the same period (not shown in the figure). The “freeze” in the aggregate power consumption of top 10 countries can be attributed to the following:

- Newly built supercomputers are more energy efficient, a trend suggested by the continuing improvement of the average energy efficiency, as shown in Figure 4.
- The TOP500 machines are more spread out across different countries. For instance, the top 10 countries accounted for 90.6% of the total installations in November 2008 but only 88.4% in November 2009.




Figure 7. Top Ten Countries Ranked by Aggregate Power Consumption

The United States remains the most power-hungry country. However, the aggregate power consumption of the United States experienced a slowdown in the second

year of the Green500. In fact, the power consumption of the United States decreased slightly from 125.6 MW to 122.9 MW from November 2008 to November 2009. This can be attributed to similar reasons that we discussed above. Nonetheless, the United States continued to dominate the number of installations of top 500 machines, accounting for 56.6%, 58.2% and 55.4% of all systems in the November 2007, November 2008 and November 2009 lists, respectively.

As of the November 2009 list, two countries made their first appearances in the top 10 most power-consuming countries: Canada and New Zealand. These two countries are also the second and the third fastest growing countries in power consumption, respectively. That is, Canada increased 400% from 0.7 MW to 3.5 MW, and New Zealand increased 120% from 2.0 MW to 4.4 MW. The country with the greatest relative increase in total power consumption is South Korea, whose aggregate power increased 17-fold from 0.1 MW to 1.7 MW.

Overall Observations. The total power consumption of the supercomputers on the Green500 continued to grow — from 161.2 MW in November 2007 to 200.4 MW in November 2008 to 272.2 MW in November 2009. The total power consumption grew even faster in the second year, with a 35.8% increase compared to a 24.3% increase in the first year. This increasing trend suggests that satisfying the energy requirement of supercomputers will remain challenging in the near future. Thus, it is imperative to consider energy efficiency as a first-order design constraint.

In the first year of the Green500 [8], we noticed skewness in the improvement of energy efficiency in supercomputers, i.e., the top-ranked machine improved much faster than the bottom-ranked machine (50.1% vs. 11.2%). In the second year, the skewness was dramatically reversed; the energy efficiency of the bottom-ranked machine achieved an astounding 3.2-fold improvement, far surpassing the improvement of the top-ranked machine (35%).

Finally, 15% of the most power-hungry supercomputers on the current Green500 list consume more than half the total power of all the supercomputers on the same list. This percentage is the same as that of the November 2008 list.

3.2. Evolution

As introduced in Section 1, the HEC community raised the following suggestions for the Green500: (1) lowering the “barrier to entry” to the Green500, (2)

rewarding innovation, not only in hardware, but also in software, and (3) investigating the possibility of using a different benchmark(s) to evaluate the energy efficiency of supercomputers. We address each of these in greater detail below.

Exploratory Lists: Little, Open, HPCC. First, with energy efficiency being the focus for the Green500, the HEC community questioned why the Green500 only evaluated the energy efficiency of the top 500 fastest supercomputers in the world. Why not look at the most energy-efficient supercomputers in the world instead? However, this question in turn raised the following issue: What defines a “supercomputer”? Or more specifically, how far should the Green500 “lower the bar” of eligibility in terms of what would constitute a supercomputer? Lowering the bar of eligibility ultimately favors smaller supercomputers, as we argue it should, since the FLOPS/watt metric scales “sub-linearly” as machine size increases [12]. That is, FLOPS increases *less than* perfectly linearly as machine size increases whereas watts increases perfectly linearly to super-linearly due to added aggregating infrastructure. To address the above, we introduced the *Little Green500*. To be eligible for the Little Green500, a supercomputer must be as “fast” as the 500th-ranked supercomputer on the TOP500 that was listed 18 months prior to the release of the Little Green500.

Second, the HEC community felt that the Green500 should promote and reward energy-efficient innovation, particularly with respect to algorithms and software. For instance, automated run-time systems that optimize for energy efficiency while still delivering high performance can improve the energy efficiency of a supercomputer, as measured in FLOPS/watt, by 20% on average [11] by significantly reducing the denominator of the performance metric. Analogously, alternative algorithms that deliver the same solution faster, e.g., mixed-precision arithmetic [13], can increase the energy efficiency of a supercomputer by significantly increasing the numerator of the FLOPS/watt metric. Consequently, we introduced the *Open Green500* as an exploratory list that allows for mixed-precision algorithms and innovative software to measurably improve the FLOPS/watt metric.

Third, while the LINPACK benchmark has been the de facto standard in measuring the performance of supercomputers, its utility beyond scientific computing arguably remains questionable. As a result, the Green500 intends to leverage alternative benchmark suites, namely the *HPCC benchmark suite*, to create another exploratory list called HPCC Green500. The

biggest challenge here is converging on a “single number” metric for HPCC performance, and hence, for HPCC energy efficiency.

Little Green500 vs. Current Green500? Given that the Little Green500 is the only exploratory list that is currently populated (albeit with a subset of eligible machines), we present a comparative analysis between the Little Green500 List and the current Green500 List. For the November 2009 release, the minimum performance requirement for an eligible machine in the Little Green500 was 9 TFLOPS (i.e., the LINPACK performance of the 500th-ranked machine in the June 2008 TOP500 list), thus resulting in a total of 878 eligible machines from the past lists.

First and foremost, we are interested in knowing how much the Little Green500 deviates from the current Green500. To this end, Figure 8 plots MFLOPS/W distributions for both. As expected, the machines in the Little Green500 have better energy efficiency than those in the current Green500, especially in the bottom half of the lists.

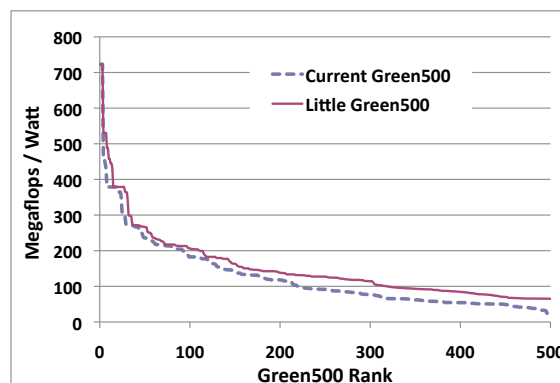


Figure 8. Energy Efficiency Distribution of Current Green500 and Little Green500.

Next, we compare the normalized MFLOPS of the maximum, average, median and minimum values of both lists, as depicted in Figure 9. Interestingly, the top-ranked machine is the same on both lists. The average and median values deviate moderately by 19% and 27%, respectively. The most striking difference comes from the minimum MFLOPS/W, which differs by 80% between the two lists.

The minimum MFLOPS/W in the Little Green500 is 65.14 MFLOPS/W, which equals that of the 333rd machine on the current Green500. This implies that 167 out of 500 or 33.4% machines on the current Green500 List could NOT make the Little Green500.




Figure 9. Normalized Energy Efficiency of the Current Green500 and the Little Green500.

4. Conclusion and Future Work

With the Green500 List turning two years-old in November 2009, this paper presented an analysis of the Green500 data since the launch of the first list in November 2007. In addition, we reported on constructive feedback from the HEC community and our efforts to address that feedback via three new exploratory lists.

The findings from our analysis can be summarized as follows:

- The overall power efficiency continued to improve at a rate similar to Moore’s Law, i.e., doubling every 24 months.
- The energy efficiency of today’s commodity machines is equivalent to the energy efficiency of custom-built machines from two years ago. (For example, a commodity machine built with the Intel Xeon E5540/E5550 CPU and the AMD Radeon HD4870 GPU accelerator surpassed IBM Blue Gene/P in energy efficiency.
- The exploratory Little Green500, which lowered the “barrier to entry” of supercomputers, differed substantially from the current Green500.
- While the total power consumption of top 500 supercomputers increased considerably, the aggregate power consumption of the top 10 countries remained the same.

In addition to the above analysis, The high-end computing (HEC) community indicated interest in pursuing the following agenda, which in turn, resulted in the launching of three new exploratory lists:

- *The Little Green500*, which lowers the “barrier to entry” to the Green500 by being more inclusive with the definition of a supercomputer.

- *The Open Green500*, which rewards innovation, not only in hardware, but also in software via alternative energy-efficient algorithms, power-aware compilation, or adaptive run-time systems, for example.
- *The HPCC Green500*, which investigates the possibility of using a different benchmark(s) to evaluate the energy efficiency of supercomputers.

However, these new lists may have opened up a Pandora’s box. For instance, while the Little Green500 demonstrates potential in offering complementary information to the current Green500, early indications are that it will “only” be useful in guiding the purchase decisions on smaller, but commodity, supercomputers. (It is important to keep in mind that these are only early indications on a small set of data.) More rigorous data analysis on a larger data set will be needed not only to justify the criteria of choosing candidate machines, i.e., by “lowering the bar”, but also to justify the existence of the exploratory Little Green500. Similar arguments hold for the exploratory Open Green500, but in this case, it is much too early to tell.

While the HPCC benchmark suite offers a more comprehensive view to the performance of a supercomputer, how to use the multi-dimensional information that HPCC produces in order to rank supercomputers remains a particularly challenging and difficult task. One possible solution, which would loosely parallel what HPCC does with performance, is to offer different rankings along individual dimensions, such as sustainable memory bandwidth, Fast Fourier Transform (FFT) performance, and so on. Alternatively, a comprehensive model needs to be developed to take into account all performance metrics.

Irrespective of all the aforementioned lists, the Green500 seeks to raise awareness in the energy efficiency of supercomputing by driving energy efficiency as a first-order design constraint (on par with performance) and encouraging fair use of the list rankings to promote energy efficiency in high-end computing (HEC) systems.

Acknowledgements

The authors thank Jeremy Archuleta, Yang Jiao, Thomas Scogland, and Kirk Cameron for their support in sustaining the Green500 List and the Green500 community for their continued support and feedback, particularly the participants from the Green500 BoFs of 2009.

The authors also wish to thank the reviewers for their

insightful comments and constructive criticism, which helped to improve the paper and provide future avenues of exploration.

This work was supported in part by NSF grant CCF-0848670 and by Supermicro.

References

- [1] Pearson Product-Moment Correlation Coefficient. http://en.wikipedia.org/wiki/Pearson_product-moment_correlation_coefficient.
- [2] M. Ahmed. Google Search Finds Seafaring Solution. In *The Times*, September 2008. http://technology.timesonline.co.uk/tol/news/tech_and_web/the_web/article4753389.ece.
- [3] American Society of Heating, Refrigerating and Air-Conditioning Engineers. *Datacom Equipment Power Trends and Cooling Applications*. ASHRAE, USA, January 2005.
- [4] Don Atwood and John G. Miner. Reducing Data Center Cost with an Air Economizer. White Paper: Intel Corporation, August 2008.
- [5] C. Belady. In the Data Center, Power and Cooling Cost More Than the IT Equipment It Supports. *Electronics Cooling Magazine*, 13(1), May 2007.
- [6] W. Feng. Making a Case for Efficient Supercomputing. *ACM Queue*, 1(7):54–64, October 2003.
- [7] W. Feng and K. Cameron. The Green500 List: Encouraging Sustainable Computing. *IEEE Computer*, December 2007.
- [8] W. Feng and T. Scogland. The Green500 List: Year One. In *5th IEEE Workshop on High-Performance, Power-Aware Computing (in conjunction with the 23rd International Parallel & Distributed Processing Symposium)*, Rome, Italy, May 2009.
- [9] W. Feng, M. Warren, and E. Weigle. The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs. In *Proc. of the IEEE Int'l Conf. on Cluster Computing*, September 2002.
- [10] S. Gorman. NSA Risking Electrical Overload. In *The Baltimore Sun*, August 2006.
- [11] C. Hsu and W. Feng. A Power-Aware Run-Time System for High-Performance Computing. In *Proc. of SC2005: The IEEE/ACM International Conference on High-Performance Computing, Networking, Storage, and Analytics*, November 2005.
- [12] C. Hsu, W. Feng, and J. Archuleta. Towards Efficient Supercomputing: A Quest for the Right Metric. In *IPDPS '05: Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'05) - Workshop 11*, page 230.1, Washington, DC, USA, 2005. IEEE Computer Society.
- [13] J. Kurzak and J. Dongarra. Implementation of Mixed Precision in Solving Systems of Linear Equations on the Cell Processor: Research Articles. *Concurr. Comput. : Pract. Exper.*, 19(10):1371–1385, 2007.
- [14] M. LaPlante. Spies Like Us: NSA to Build Huge Facility in Utah. *Salt Lake Tribune*, 2009.
- [15] J. Markoff and S. Hansell. Hiding in Plain Sight, Google Seeks More Power. In *The New York Times*, June 2006. <http://www.nytimes.com/2006/06/14/technology/14search.html>.
- [16] J. Meichi. Building the Data Center of the Future. In *NCSA Newsletters*, November 2008. <http://www.ncsa.uiuc.edu/News/Stories/PCF/>.
- [17] H. Meuer. The TOP500 Project: Looking Back over 15 Years of Supercomputing Experience. www.top500.org, 2008.
- [18] R. Schmidt and D. Beaty. ASHRAE Committee Formed to Establish Thermal Guidelines for Datacom Facilities. *Electronics Cooling Magazine*, Feb 2005.
- [19] NSA Electrical Power Upgrade, January 2008. <http://cryptome.org/nsa010208.htm>.
- [20] ASCR Computing News Roundup, October 2007. <http://www.sc.doe.gov/ascr/News/MonthlyNewsRoundup10-07.html>.
- [21] A. Vance. Microsoft's Data Center Offensive Sounds Offensive. In *The Register*, June 2006. http://www.theregister.co.uk/2008/03/03/microsoft_data_center_blowout.