

Trends in Energy-Efficient Computing: A Perspective from the Green500

Balaji Subramaniam[†], Winston Saunders[‡], Tom Scogland[†], and Wu-chun Feng[†]

[†]Department of Computer Science, Virginia Tech

[‡]Data Center and Connected Systems Group, Intel Corporation

{balaji, njstn, feng}@cs.vt.edu, winston.a.saunders@intel.com

Abstract—A recent study shows that computation per kilowatt-hour has doubled every 1.57 years, akin to Moore’s Law. While this trend is encouraging, its implications to high-performance computing (HPC) are not yet clear. For instance, DARPA’s target of a 20-MW exaflop system will require a 56.8-fold performance improvement with only a 2.4-fold increase in power consumption, which seems unachievable in light of the above trend. To provide a more comprehensive perspective, we analyze current trends in energy efficiency from the Green500 and project expectations for the near future.

Specifically, we first provide an analysis of energy efficiency trends in HPC systems from the Green500. We then model and forecast the energy efficiency of future HPC systems. Next, we present *exascalar* – a holistic metric to measure the distance from the exaflop goal. Finally, we discuss our efforts to standardize power measurement methodologies in order to provide the community with reliable and accurate efficiency data.

I. INTRODUCTION

DARPA’s target of a 20-megawatt (MW) exaflop supercomputer has brought energy efficiency to the forefront of the high-performance computing (HPC) community. Historically, the community has pursued performance at any cost. However, the annual infrastructure cost surpassed the acquisition cost of the systems themselves in 2004. By 2008, the annual energy cost to power these systems exceeded the acquisition cost of the systems [5]. Such costs have diminished the efficiency benefits of Moore’s Law¹ related speed-up [7]. Moreover, DARPA’s target of a 20-MW exaflop system [6] would require a 56.8-fold performance improvement with only a 2.4-fold increase in power consumption.

Recently, Koomey et al. [13] surveyed the relationship between computing power and the electricity required to power the devices. The survey concluded that computations per kilowatt-hour has doubled every 1.57 years.² While this trend is encouraging, its implication to high-performance computing (HPC) is not yet clear. The Green500 [10], [3], which ranks the most powerful supercomputers in terms of their energy efficiency, is a data warehouse for analyzing the energy efficiency of supercomputers. Such an analysis will provide insights into the longer-term trends in energy efficiency for HPC.

The Green500 was created in 2007 to raise awareness about energy efficiency to the HPC community and to highlight

its importance. The Green500 encourages the treatment of efficiency metrics such as *operations per watt* to be on par with performance. As the Green500 continues to evolve, we analyze and draw inferences from the data available and inform the HPC community on the trends in energy efficiency. We also strive to standardize the power measurement methodology for large-scale systems to provide the community with accurate and reliable power measurements. Towards these goals, we make the following contributions in this paper:

- A general analysis of the energy efficiency, power consumption, and performance efficiency for different types of HPC systems.
- A model representing the growth in energy efficiency and the commensurate decrease in energy used per computation. We project this trend and discuss the implications of such growth for achieving a 20-MW exaflop system.
- The *exascalar* metric – a holistic measure of *distance* from the goal of an exaflop system.
- A discussion of our efforts to standardize the power measurement methodology for large-scale systems on the Green500. We frame the run rules for the Green500 List by inferring from the instantaneous power profile of high-performance LINPACK (HPL) [2] on different platforms.

Based on our above contributions, we make the following observations and conclusions. First, the energy efficiency of both the Green500 List as a whole as well as the greenest supercomputer continue to improve. Second, heterogeneous (i.e., systems using GPUs or other co-processors) and custom-built (*a la* BlueGene/Q) systems continue to have a better overall energy efficiency than their conventional counterparts. Third, we show that current trends in energy efficiency, projected to the year 2020, and predict that we will be 7.2-fold short of the efficiency required to meet DARPA’s 20-MW exaflop goal. Fourth, we *quantify the distance* from the 20-MW exaflop goal. Fifth, we show that a *unified* standard power measurement methodology can be formulated to provide accurate and reliable power measurements even though the instantaneous power signature of the workload varies drastically on different platforms.

The rest of the paper is organized as follows. Section II offers a high-level analysis of the *greenness* of supercomputers, along with an in-depth analysis of the efficiency characteristics of the different types of machines that achieve high energy

¹Moore’s Law is popularly summarized as *the performance of computing devices doubling every 18 to 24 months*.

²This relationship is popularly referred to as *Koomey’s Law*.

efficiency. Section III studies the trends in energy efficiency and energy-per-operation in HPC and presents the projected implications for achieving an exascale system. It also describes the *exascalar* metric. We discuss our efforts to standardize the power measurement methodology for large-scale systems in Section IV. Related work is discussed in Section V, and finally, Section VI concludes the paper.

II. ANALYSIS OF THE GREEN500 LISTS

In this section, we analyze the Green500 List to infer trends on different metrics ranging from efficiency to the architectural composition of the lists. We first start with tracking energy efficiency over time and describe the highlights of the November 2012 release. We seek to uncover the relationship between Moore’s Law and energy efficiency using the feature size of the systems on the list as metrics. Third, we discuss the statistics on the power consumption of the list over the years. Finally, we comment on the architectural composition (i.e., heterogeneous vs. homogeneous systems) of the list and its correlation with efficiency in terms of both energy and performance.

A. Efficiency Trends

The Green500 List has seen steady improvement in the energy efficiency of systems over the years. Figure 1 shows the energy efficiency of all the HPC systems for every Green500 List. The plot clearly shows that the gap between the most energy-efficient systems and the rest of the systems is widening every year. The top ten systems in the last two editions of the Green500 List possess an energy efficiency of more than 2.0 GFLOPS/watt. In the last four lists, we have seen outliers at the top end of the list. BlueGene/Q systems have been largely responsible for this. In previous lists, efficiencies of greater than 2.0 GFLOPS/watt were only achieved by the BlueGene/Q systems. However, in the latest edition of the Green500, several other flavors of systems surpassed the 2.0 GFLOPS/watt mark.




Fig. 1. Energy Efficiency in Green500 Lists

The most energy-efficient system in the latest edition of Green500 List is the Beacon system, which uses Intel Sandy Bridge processors and Intel Xeon Phi co-processors to achieve nearly 2.5 GFLOPS/watt. The SANAM supercomputer occupies the second position; it is another heterogeneous system, which uses Intel Sandy Bridge processor and AMD FirePro GPUs. Titan, the world’s fastest supercomputer, according to the November 2012 release of the Top500 list [4], is the third most energy-efficient system. Each node in Titan consists of AMD Opteron CPUs and NVIDIA Kepler (K20) GPUs. The fourth spot is occupied by Todi, similarly equipped with AMD Opteron CPUs and NVIDIA K20 GPUs, and the remaining top ten spots are occupied by BlueGene/Q systems, as shown in Table I. The rest of the machines that have an efficiency of more than 2.0 GLOPS/watt are all BlueGene/Q systems.

#	Gf/W	Computer
1	2.499	Beacon - Intel Xeon E5-2670 + Intel Xeon Phi
2	2.351	SANAM - Intel Xeon E5-2650 + AMD FirePro S10000
3	2.142	Titan - AMD Opteron 6274 + Nvidia K20
4	2.121	Todi - AMD Opteron 6272 + Nvidia K20
5	2.102	JUQueen - BlueGene/Q
6	2.101	BGQdev - ATI Radeon GPU (Nagasaki U.)
7	2.101	rzuseq - B505, NVIDIA 2090 (BSC-CNS)
8	2.101	BlueGene/Q (IBM T.J. Watson Center)
9	2.101	BlueGene/Q (IBM T.J. Watson Center)
10	2.101	CADMOS BG/Q - BlueGene/Q

TABLE I
THE GREENEST 10 SUPERCOMPUTERS

Moore’s Law promises the increase in transistor density on chip. A by-product of Moore’s law is the decreasing feature size of the processors. We seek to track the correlation between the feature size and increasing trend in energy efficiency on the Green500 List. Figure 2 shows the trend in average feature size for sets of machines on the Green500 and Top500 Lists. In general, the feature size across sets of systems over the years have improved (i.e., decreased in size) on both the lists. Counterintuitively, the systems occupying the top ten ranks on the Green500 Lists have had a larger feature size than the rest of the list in years prior to 2010. However, this trend evened out in 2010, and since then, the greenest ten machines have had smaller feature sizes than the list on an average. This shift highly correlates with the increase in custom or heterogeneous processors, such as BlueGene/Q and GPU-based systems, on the list.

B. Power Consumption Trends

The current fastest system consumes 8.2 MW of power while delivering a performance of 17.59 PFLOPS. With DARPA’s target power envelope for an exascale system at 20 MW, we need a 56.8-fold increase in performance with only a 2.4-fold increase in the power envelope. To understand the trend with respect to the thermal power envelope, we need to track the power consumption of systems even if there is a steady growth in energy efficiency. Figure 3 shows the statistics on the power envelope of systems on the Green500 List. The mean power consumption of the list has been steadily increasing while the maximum power envelope has remained




Fig. 2. Average Feature Size in Nanometers

a constant for the past three lists with K-Computer consuming 12.66 MW, the most power consumed by a system on the list.

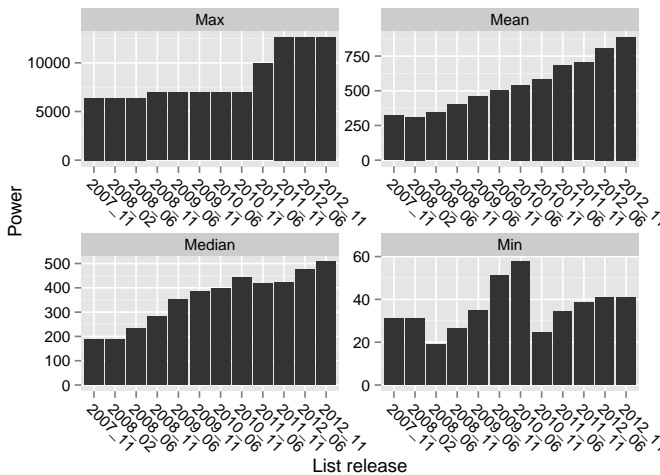


Fig. 3. Total Power Consumption Statistics

C. Heterogeneity and Efficiency Trends

The heterogeneity of the list is increasing every year. We can classify the systems on the Green500 as (1) accelerator or co-processor-based clusters, (2) custom supercomputers, and (3) traditional commodity supercomputers. Even though homogeneous systems still form the majority of the list, the high energy efficiency provided by the accelerators affect the overall energy of the system. At present, the accelerator- or co-processor-based cluster is a commodity cluster enhanced with high-bandwidth accelerators or co-processors, which include AMD and NVIDIA GPUs and most recently the Intel Xeon Phi co-processors. These devices are capable of performing massively parallel simple calculations due to use of large number of simple cores and the high-bandwidth memory such as GDDR or XDR. The custom supercomputers are high-density systems, which include the BlueGene/Qs and the K-Computer. The traditional commodity supercomputers include all the CPU-based machines. An analysis of these classes of computers can provide a better understanding of the efficiency

of these system and insights into how they affect the overall statistics of the Green500 List.

Figure 4 shows the energy and performance efficiency of heterogeneous and homogeneous systems on the list. The energy efficiency is calculated as performance-per-watt and the performance efficiency is calculated as the percentage of theoretical peak performance achieved. The accelerator/co-processor-based systems provide much better energy efficiency compared to rest of list, and as a result, affect the overall energy efficiency of the list, even though they are in the minority. However, these systems do not provide good performance efficiency³ in general as they are not as suitable for running latency-bound tasks.

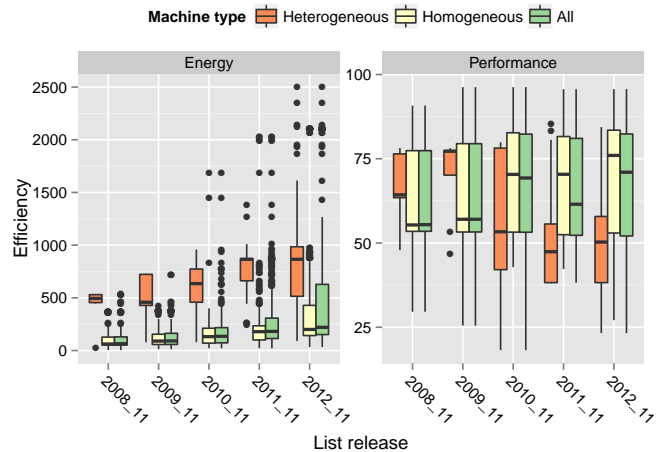


Fig. 4. Heterogeneous vs. Homogeneous Systems - Energy and Performance Efficiency

III. ENERGY-EFFICIENCY PROJECTIONS FOR EXASCALE SYSTEMS

The HPC community has set its sights on an exascale system. However, as suggested by the recent DARPA exascale report, power will be the limiting function in building an exascale system. Achieving a 20-MW exaflop system would require 20-fold improvement in energy efficiency. To keep the community informed about how much closer we are to this goal, we have projected the energy efficiency trends of supercomputers on an annual basis using historical trends. In this section, we model the historical trends in energy efficiency and extrapolate these trends for the future to examine the prospects of achieving the exaflop goal in 2018 and 2020. We also present the exascal metric – a holistic metric to quantify the *distance* between current systems and our exascal goals in terms of both performance and energy efficiency.

³The recently introduced NVIDIA Kepler GPUs and Intel Xeon Phi processors are an exception as they have demonstrated high performance efficiency numbers.




Fig. 5. FLOP/Joule and Joules/FLOP Trend in HPC

TABLE II
LINEAR REGRESSION PARAMETERS FOR FLOP/JOULE AND JOULES/FLOP TREND

—	Green5	Green10	Green50	Green100
Slope - FLOP/Joule	287.39	278.68	204.60	148.86
Intercept - FLOP/Joule	34.61	-87.61	-152.80	-87.43
R ² - FLOP/Joule	0.96	0.93	0.85	0.84
Slope - Joules/FLOP	-257.12	-306.57	-453.50	-550.02
Intercept - Joules/FLOP	2235.95	2670.80	4158.9	5220.08
R ² - Joules/FLOP	0.91	0.94	0.97	0.98

A. Tracking Koomey's Law for HPC: Historical and Projected Trends for Energy Efficiency in HPC

Koomey et al. [13] have uncovered the relationship between computing power and the electricity required to deliver the performance. They have concluded that computations per kilowatt-hour over time has doubled every 1.57 years. One of the limitation of this work is that it focuses on peak power and performance. However, this trend may not hold for systems running at less than peak power. In this section, we explore whether such trends hold in the HPC domain.

We analyze eight different Green500 Lists to infer insights into the trends in energy efficiency for HPC. Specifically, we seek to understand the change in the FLOP/joule and joules/FLOP ratios over the years. Figure 5 shows the data from the Green500 Lists. Both graphs in the figure present the average ratio over the 5, 10, 50, and 100 most energy-efficient supercomputers from the lists between June 2009 and November 2012. We analyze the averages in order to reduce the bias towards small systems (i.e., systems that consume relatively less power). We observe an increasing trend for computations per joule and a decreasing trend for the joules per computation in all cases. Figure 5 also presents the linear-regression fits (labeled with suffix *-lin*) for the trends. A typical regression has the form:

$$\text{Metric} = \text{Slope} * \text{Year} + \text{Intercept} \quad (1)$$

where Metric can be either FLOP/joule or joules/FLOP, and Year is a number between one and eight. We would like to note that not all observations are linear. However, linear regression provides a good approximation of the trends. Table II shows

all the linear regression parameters for the data presented in Figure 5. The R² (i.e., the quality of fit) for the regression varies from 0.84 to 0.96 for the FLOP/joule trend and from 0.91 to 0.98 for the joules/FLOP trend. The slope of lines indicates how much each metric increases every six months. For example, the FLOP/joule ratio increases by 287.39 every six months for the Green5 data. The Green5 data is the fastest growing trend for the FLOP/joule ratio, and the Green100 data is the fastest decreasing trend for the joules/FLOP ratio.

The FLOP/Joule metric also represents the energy efficiency of the system:

$$\begin{aligned} \text{FLOPS/Watt} &= \frac{\text{Floating Point Operations/Second}}{\text{Joules/Second}} \\ &= \frac{\text{Floating Point Operations}}{\text{Joules}} \end{aligned} \quad (2)$$

Extrapolation of energy efficiency for an exascale system is one of the valuable insights we gain from these models. For a system to achieve exaflop while consuming 20 MW and 100 MW power, respectively, it should have an energy efficiency of 50 GFLOPS/watt and 10 GLOPS/watt, respectively. For frame of reference, the current greenest supercomputer as of November 2012 has an energy efficiency of 2.49 GFLOPS/watt. Assuming that the energy efficiency will improve at the same rate as predicted by the linear models, Table III shows the energy efficiency trend extrapolated to years 2018 and 2020. When the linear model for energy efficiency for Green5 data is extrapolated to year 2018 and 2020, it predicts that the average energy efficiency of the greenest five supercomputers will be 5.78 and 6.93 GFLOPS/watt, respectively. As expected the averages of the other groups (which is more representative

TABLE III
ENERGY EFFICIENCY TRENDS EXTRAPOLATED TO YEAR 2018 AND 2020
(IN GFLOPS/WATT)

—	In 2018	In 2020
Green5	5.78	6.93
Green10	5.48	6.60
Green50	3.93	4.75
Green100	2.88	3.48

of the list) are predicted to have lesser energy efficiency in the years 2018 and 2020.




Fig. 6. A Comparison of Size (Represented by Total Power) Versus the Energy Efficiency of the System

The above analysis might be biased for an exascale projection as smaller machines tend to occupy the top spots on the Green500 Lists. For example, the greenest supercomputer consumes 44.89 kilowatts whereas the fastest supercomputer consumes 8.20 megawatts. While this trend is changing (e.g., Titan supercomputer occupies the third rank on the November 2012 Green500 List), the majority of the systems in the top 20 ranks of the Green500 List are relatively small systems. In order to look at the Green500 trends in an unbiased manner, we can plot the total power consumed by each machine versus its energy efficiency. The total power has a strong correlation with the size of the machine and its performance. Figure 6 presents the power versus the energy efficiency of all the systems in Green500 Lists. The plot clearly shows an increase in the energy efficiency of the system across all sizes every year. The graph also shows two points to represent a system which achieves exaflops at 20 and 100 MWs (labeled as 2018_11). The trend indicates that we will reach the 100-MW goal in 2018. However, the 20-MW goal will require the growth of energy efficiency at a much higher rate.

B. The Exascalar: A Holistic Measure of Distance from Exascale Systems

The ultimate goal of supercomputing is to deliver the highest performance. It is also now well established that this goal cannot be achieved economically without accompanying high energy efficiency. Supercomputing leadership requires both highest efficiency and performance. In 2011, an analysis of the top rated supercomputers was proposed which looked

at the performance and efficiency of the Green500 List simultaneously [16]. Since both efficiency and performance leadership are required to achieve the exascale goal of 10^{18} flops in a power envelope of 20 MW, it is useful to synthesize the performance and efficiency information in the Green500 into a single metric. We present the *exascalar* metric – a holistic measure of distance from the Exaflop goal. More specifically, the *exascalar* metric is the scalar distance between a system and the 20-MW exaflop goal in a performance versus energy efficiency plot as shown in Equation 3, where $Perf_i$ is performance of the system and EE_i is the energy efficiency of the system. This holistic *standard ruler* will serve to measure the (logarithmic) distance of any specific supercomputer to the exascale goal in terms of both energy efficiency and performance. An inherent advantage of this ranking is that for two systems with equivalent performance, the one with higher efficiency would have the better exascalar (lower absolute value is better). Such ranking encourages simultaneous performance and energy efficiency growth.




Fig. 7. The Exascalar Plot for November 2012 Green500 List

Figure 7 shows the most recent exascalar plot [14] as derived from the November 2012 Green500 List. The graph plots each system based on both its efficiency and performance. The form of the graph was chosen with efficiency as the independent and performance as the dependent variable, since for a fixed power envelope performance scales proportionally with efficiency. This dependence is emphasized by the lines of constant power shown. The top exascalar has an absolute value of about 2.2 (i.e., about 2.2 orders of magnitude as a linear measure from the goal of 10^{18} flops in 20 MW). The green trend-line shows the historical top exascalar from November 2007 to present. The overall performance trend shows the expected strong (though not completely monotonic) dependency on efficiency. The column of points on the right hand side are the IBM BlueGene/Q machines, all of which have almost equivalent energy efficiency but different scales (represented by the performance of the system). The most

$$Exascalor = \sqrt{(\log_{10}(Perf_{System}) - \log_{10}(Perf_{@20MWExaftop}))^2 + (\log_{10}(EE_{System}) - \log_{10}(EE_{@20MWExaftop}))^2} \quad (3)$$

performant system according to the Top500 list has just slightly higher efficiency and performance, thus continuing the expected exascalor trend.

The overall evolution of the *exascalor* for large-scale systems is shown in Figure 8. It shows steady increase at the rate of about a factor of ten every 2.7 years, or doubling every 0.8 years. The median *exascalor* trend line increases at a lower rate as might be expected from similarly observed slower trends in median performance and efficiency [14]. Eight of the top ten systems according to *exascalor* distance changed in the last one year.




Fig. 8. The Exascalor Trend

The locus of the Green500 data points, when viewed as an *exascalor* plot, have the approximate shape of a triangle. We refer to this triangle as the *exascalor* triangle. This triangle can be understood by a simple taxonomy [15] shown in Figure 9. The boundaries of the triangle can be ascribed to simply obtained parameters of the Green500 List. The Base of the triangle is the lower performance cut-off which comes from pressure on the population exerted by the Top500 list. The right hand face of the triangle, which tends to push right over time, represents the increasing highest efficiency systems, and the hypotenuse is constrained by power (and scale) limitations of practical systems.

Similarly each corner is of specific interest. The top corner can be thought of as the point of supercomputing leadership capability. Historically there are not always systems in near proximity to this point, but the gaps shows a strong trend of shrinking over time [15] indicating the lag in reaching scale with new technology is shrinking. The lower right corner is where new innovation in higher efficiency (the top of the Green500 List) has usually first appeared, and hence serves as a kind of *technology doorway* provided by the Green500 List to the Top500. It is also interesting to note that systems in the lower left hand corner of the graph consume almost a factor of 100x more energy than the systems in the lower

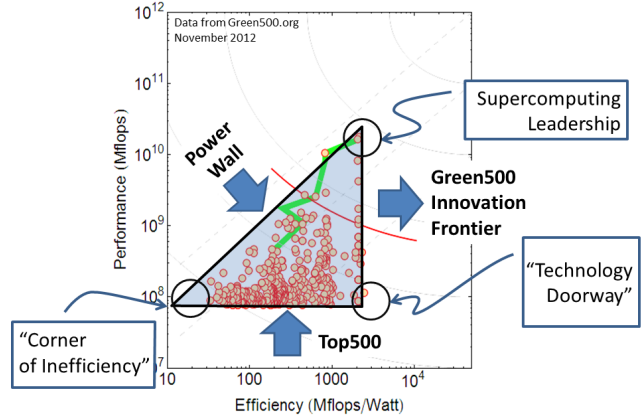


Fig. 9. The Exascalor Triangle

right, although both deliver comparable performance.

IV. POWER-MEASUREMENT METHODOLOGY FOR THE GREEN500 LIST

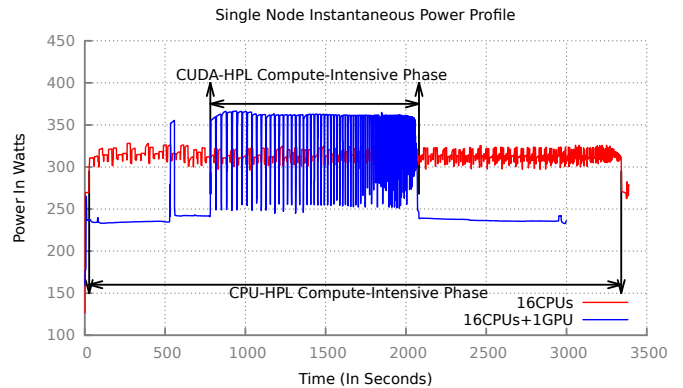


Fig. 10. Instantaneous Power of HPL Benchmark on Single Node Setup

The number of heterogeneous systems listed on the Green500 List is increasing every year. The instantaneous power signature of the same workload can be drastically different on various platforms such GPU-, MIC- and CPU-based systems. There is a necessity to define a *unified* power measurement methodology which will provides reliable and accurate power measurement. The definition of a *unified* power measurement methodology for all classes of large-scale systems is one of the main challenges faced by the Green500 List. The Green500 team strives to define power measurement methodology for large-scale systems and in the process we update our run rules to keep pace with the heterogeneity of the list. In this section, we discuss our efforts to define the

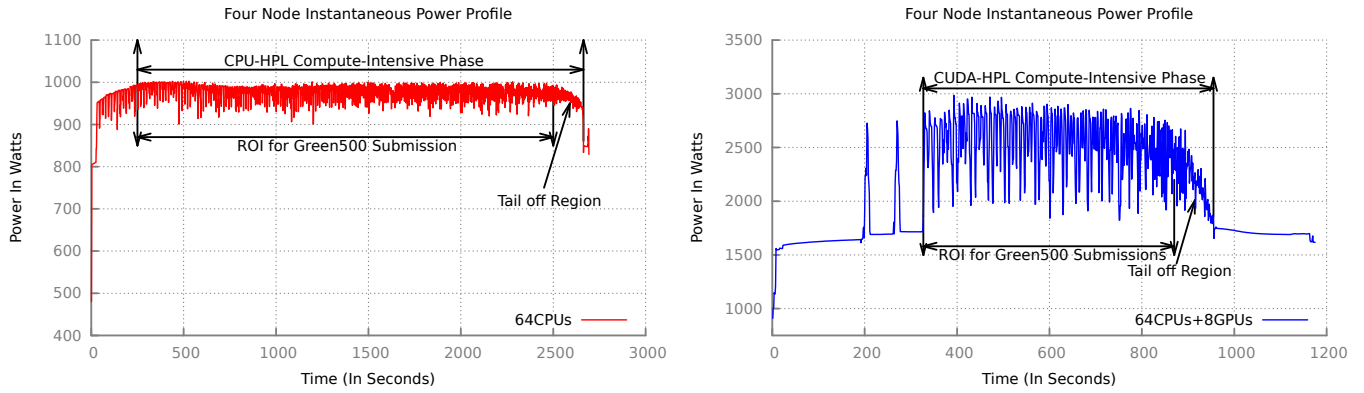


Fig. 11. Instantaneous Power of HPL Benchmark on Multi-Node Setup. Left: CPU-HPL, Right: CUDA-HPL. ROI = Region of Interest. The instantaneous power profile tails off at the end of the compute-intensive phase

run rules for the Green500 List by drawing inference from instantaneous power profile analysis.

The Green500 List uses the high-performance LINPACK (HPL) benchmark [2] as the workload. Henceforth, our discussion will focus on the definition of power measurement methodology for HPL benchmark. The HPL benchmark is a linear algebraic package which solves a dense system of linear equations. It runs in four phases: (1) random square matrix generation, (2) LU factorization of the matrix, (3) backward substitution to solve for the solution, and (4) checking for correctness. We would like to emphasize that only the second and third steps are used for calculating the LINPACK score (in GFLOPS). We need to understand the power signature of the HPL benchmark on different classes of systems in order to define a power measurement methodology for it. To this end, we ran single node and multi-node instances of CPU-HPL [2] and CUDA-HPL [9]. Our single node setup consists of two Intel Xeon E5-2665 processors with 128 GB of CPU memory and one NVIDIA C2075 GPU. The multi-node setup has four nodes with two AMD Opteron 6134 processors, 32 GB of CPU memory and two NVIDIA C2050 GPUs. The four nodes communicate with each other over InfiniBand. We would like to determine *when* to measure power for different classes of systems when it is executing the HPL benchmark.

Figure 10 shows the instantaneous power profile of both the CPU-HPL and CUDA-HPL in a single node setup. We would require the power measurements to occur in the second and third phases of the HPL benchmark as only these phases contribute to the GFLOPS score. Henceforth, we refer to the second and third phases of HPL benchmark as the compute-intensive phases.⁴ The instantaneous power difference between compute-intensive phases and the rest of phases is pronounced for the CUDA-HPL runs. The first three phases of the CPU-HPL have similar power signatures and the fourth phase has different power profile. Due to the third phase (solve phase) of HPL benchmark, we expect the GFLOPS to tail off at the end of the compute-intensive phase [8]. And as instantaneous

⁴Recently, the developers of HPL benchmark have added time stamps in the benchmark to identify this compute-intensive phase.

power has strong correlation with performance, we expect to see a tail off in Figure 10. However, we do not observe such a trend in the instantaneous power profile for the single node setup. Figure 11 shows the power profile for the multi-node setup. Unlike the runs for the single node setup, the power signatures of different phases is more pronounced for both the CPU- and CUDA-HPL benchmarks. We observe that there is a period of gradual increase in power at the start of the CPU-HPL which corresponds to the phase one of the benchmark and there is a tail off at the end of the compute-intensive phase which corresponds to third phase of the benchmark. Similar trends can be seen for CUDA-HPL also.

We have based our power measurement methodology on the observations made from instantaneous power signature analysis of CPU-HPL and CUDA-HPL benchmarks. We expect the power measurements to cover only 20% of compute-intensive phase of the HPL benchmark. As observed, HPL has a tail off power trend within the compute-intensive phase. Therefore we require that the tail off region to not be a part of the measurement. We refer to this region as the region of interest (ROI). The Figures 10 and 11 clearly mark the ROI for the Green500 List. We define the power measurement methodology for the Green500 List as follows:

- 1) The measurement granularity should be at least one sample per second.
- 2) The power measurement should be within the compute-intensive phase of the benchmark.
- 3) The power measurement should cover at least 20% of the compute-intensive phase or one minute (whichever is longer).
 - If the power measurement covers less than 90% of the compute-intensive phase, it should not include the tail off region. We assume that the last 10% of the compute-intensive phase is the tail off region. Therefore, we require the measurement to occur within the first 90% of the compute-intensive phase of the benchmark (i.e., ROI for Green500 submissions).
 - Reporting measurements which includes the tail off

region is allowed only if the power measurement covers more than 90% of the compute-intensive phase.

- 4) The component(s) measured should consume a minimum of 1 kilowatt to reduce any random fluctuations in power measurement.

The Green500 team along with the EEHPC working group [1] will continue to update our power measurement methodology to provide a reliable and accurate overview of the power envelopes of large-scale systems. These reliable power measurements will help us project our energy efficiency trends and view where we stand with respect to 20-MW exaflop goal.

V. RELATED WORK

Bergman et al. [11] discuss the major challenges to achieve an exascale system in their exascale computing study [6]. They classify these challenges as (i) energy and power challenge, (ii) memory and storage challenge, (iii) concurrency and locality challenge and (iv) resiliency challenge. They also provide several discussions on taking the current technology and projecting it to exascale to provide a perspective of what it would yield. In [11], the authors update the exascale study to cover more data and metrics. They also add heterogeneous systems to the types of system analyzed. In this paper, we analyze the general energy efficiency trends in HPC using Green500 data and provide a projection of energy efficiency for systems in 2018 and 2020.

Koomey et al. have uncovered that the number of operations performed for a kilowatt-hour has doubled every 1.57 years [13]. Koomey also provides a methodology to estimate the total power consumption by servers in the US [12]. This study determines the total electricity by using measured data and estimates of power with publicly available and concluded that the total power consumed by servers (including the cooling and the infrastructure) is 2% of the total US power consumption. We model the growth rate for energy efficiency and track Koomey's law for HPC systems in this paper. We also use measured data from the Green500 database.

VI. CONCLUSION

In the past, the HPC community has been particularly guilty of seeking performance above all else. Recent expectations such as the DARPA's 20-MW exascale system goal, the ever increasing cost of electricity, power provisioning infrastructure and cooling and the commensurate decrease in efficiency benefits due to Moore's law related performance gains have brought metrics such as energy efficiency and total cost of ownership to the forefront in HPC community. While recent postulation by Koomey et al. on the growth in computations per kilowatt-hour are encouraging, the implications of such trends to HPC is yet to be understood.

In this paper, we use the Green500 database to track, analyze and project energy efficiency trends in HPC. In general, our study shows that energy efficiency has increased and the energy efficiency of GPU- and MIC-based and custom built system dominate the CPU-based systems. We build a linear

model which tracks the energy efficiency of different sets of machines on the Green500 list. Our projection results indicate that we are 7.2x smaller than the energy efficiency goals of 20-MW exaflop goal efficiency goals in the year 2020. We also present the *exascalar* metric – a holistic metric to measure the *distance* from the 20-MW exaflop goal and show that the top exascalar system is about 2.2 orders of magnitude as a linear measure from the goal. Finally, we propose a *unified* power measurement methodology for accurate and reliable power measurement on large-scale systems. Clearly, the energy efficiency of large-scale systems should grow at a higher rate than current trends to achieve the 20-MW exaflop goal in 2020. The Green500 list will continue to track innovations in energy efficiency of large-scale systems and inform the HPC community about the exciting times that lie ahead.

VII. ACKNOWLEDGEMENTS

This work was supported in part by a gift and equipment donation by Supermicro. We also thank Jason Lockhart, Mark Gardner, Heshan Lin, the Green500 community, and the participants at the Green500 birds-of-a-feather sessions for their continued support and feedback.

REFERENCES

- [1] Energy-Efficient High-Performance Computing Working Group (EEHPC WG). <http://eehpcwg.lbl.gov>.
- [2] High Performance LINPACK (HPL). Available at <http://www.netlib.org/benchmark/hpl>.
- [3] The Green500 List. Available at <http://www.green500.org>.
- [4] The Top500 list. Available at <http://top500.org>.
- [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] K. Bergman, S. Borkar, D. Campbell, W. Carlson, W. Dally, M. Denneau, P. Franzon, W. Harrod, K. Hill, J. Hiller, S. Karp, S. Keckler, D. Klein, R. Lucas, M. Richards, A. Scarpelli, S. Scott, A. Snively, T. Sterling, R. S. Williams, K. Yelick, and P. Kogge. Exascale Computing Study: Technology Challenges in Achieving Exascale Systems.
- [7] M. Bohr. A 30 Year Retrospective on Dennard's MOSFET Scaling Paper. *IEEE Solid-State Circuits Newsletter*, 12(1):11–13, 2007.
- [8] J. Dongarra. LINPACK Benchmark with Time Limits on Multicore and GPU Based Accelerators, June 2010. Available at <http://www.netlib.org/utk/people/JackDongarra/SLIDES/isc-talk-06102.pdf>.
- [9] M. Fatica. Accelerating Linpack with CUDA on Heterogenous Clusters. In *Proceedings of 2nd Workshop on General Purpose Processing on Graphics Processing Units, GPGPU*, 2009.
- [10] W. Feng and K. Cameron. The Green500 List: Encouraging Sustainable Supercomputing. *Computer*, 40(12):50–55, 2007.
- [11] P. Kogge and T. Dysart. Using the TOP500 to Trace and Project Technology and Architecture Trends. In *Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis*, SC, 2011.
- [12] J. Koomey. Growth In Data Center Electricity Use 2005 To 2010. Available at <http://www.mediafire.com/file/zzqna34282fr2f/koomeydatacenterelectuse2011finalversion.pdf>.
- [13] J. Koomey, S. Berard, M. Sanchez, and H. Wong. Implications of Historical Trends in the Electrical Efficiency of Computing. *IEEE Annals of the History of Computing*, 2011.
- [14] W. Saunders. Data Center Knowledge. Available at <http://www.datacenterknowledge.com/archives/2013/01/02/exascalar-results-from-november-2012-part-1>.
- [15] W. Saunders. Data Center Knowledge. Available at <http://www.datacenterknowledge.com/archives/2013/01/28/the-taxonomy-of-exascalar>.
- [16] W. Saunders. Intel Data Stack. Available at <http://communities.intel.com/community/datastack/blog/2011/10/20/rethinking-supercomputer-performance-and-efficiency-for-exascale>.