The Green Index: A Metric for Evaluating System-Wide Energy Efficiency in HPC Systems

Balaji Subramaniam and Wu-chun Feng
Department of Computer Science
Virginia Tech
`{balaji, feng}@cs.vt.edu`

Abstract*— In recent years, the high-performance computing (HPC) community has recognized the need to design energy-efficient HPC systems. The main focus, however, has been on improving the energy efficiency of *computation*, resulting in an oversight on the energy efficiencies of other aspects of the system such as *memory* or *disks*. Furthermore, the energy consumption of the non-computational parts of a HPC system continues to consume an increasing percentage of the overall energy consumption. Therefore, to capture a more accurate picture of the energy efficiency of a HPC system, we seek to create a benchmark suite and associated methodology to stress different components of a HPC system, such as the processor, memory, and disk. Doing so, however, results in a potpourri of benchmark numbers that make it difficult to "rank" the energy efficiency of HPC systems. This leads to the following question: *What metric, if any, can capture the energy efficiency of a HPC system with a single number?

To address the above, we propose *The Green Index (TGI)*, a metric to capture the system-wide energy efficiency of a HPC system as a *single number*. Then, in turn, we present (1) a methodology to compute TGI, (2) an evaluation of system-wide energy efficiency using TGI, and (3) a preliminary comparison of TGI to the traditional performance-to-power metric, i.e., floating-point operations per second (FLOPS) per watt.

I. INTRODUCTION

For decades now, the LINPACK benchmark has been widely used to evaluate the performance of high-performance computing (HPC) systems. The TOP500 list [7], for example, uses the high-performance LINPACK (HPL) benchmark [2] to rank the 500 fastest supercomputers in the world with respect to the *fl*oating-point *o*perations *p*er *s*econd (or FLOPS). While the TOP500 continues to be of significant importance to the HPC community, the use of HPL to rank the systems has its own limitations as HPL primarily stresses only the processor (or CPU component) of the system. To address this limitation and enable the performance evaluation of different *components* of a system, the HPC Challenge (HPCC) benchmark suite [3] was developed in 2003. The HPCC benchmark suite is a collection of benchmarks that provide better coverage and stress-testing of different components of the system. However, the broader acceptance of HPCC as *the* performance benchmark for HPC has been limited, arguably because of its inability to capture performance in a rankable manner, e.g., with a single number. Hence, HPL and its performance metric of *FLOPS* continues to reign.

Today, the HPC community finds itself in a similar situation, but this time, with respect to the greenness of HPC systems. The HPC community has acknowledged this issue as a major problem in designing future exascale systems. Efforts such as the Green500 [9], launched in 2007, seek to raise the awareness of power and energy in supercomputing by reporting on power consumption and energy efficiency, as defined by *fl*oating-point *o*perations *p*er *s*econd per watt or FLOPS/watt. By 2008, the TOP500 also began tracking the power consumption of supercomputers. However, these efforts are inherently limited by the use of LINPACK benchmark for performance measurement because it primarily stresses only the CPU component of an HPC system. Furthermore, as noted in a recent exascale study [1], the energy consumption of a HPC system when executing non-computational tasks, especially data movement, is expected to overtake the energy consumed due to the processing elements. Clearly, there is a necessity to evaluate the energy efficiency of different components of a HPC system.

Like performance benchmarking with HPCC, we propose an approach that evaluates the energy efficiency of different components of an HPC system using a benchmark suite. However, there are seven different benchmark tests in the suite, and each of them reports their own individual performance using their own metrics. We propose an approach that seeks to answer the following questions:

- What metric should be used to evaluate the energy efficiency of different components of the system?

- How should the energy consumed by the different benchmarks be represented as a *single number*?

This paper presents an initial step towards answering the above questions. Specifically, we propose *The Green Index (TGI)* [8] — a metric for evaluating the system-wide energy efficiency of an HPC system via a suite of benchmarks. The chosen benchmarks currently include HPL for computation, STREAM for memory, and IOzone for I/O.

The contributions of this paper are as follows:

- A methodology for aggregating the energy efficiency of different benchmarks into a single metric—TGI.
- A preliminary analysis of the goodness and scalability of the TGI metric for evaluating system-wide energy efficiency.
- A comparison of the TGI metric to traditional energy-efficient metrics such as the performance-to-power ratio.

The rest of the paper is organized as follows. Section II describes the proposed TGI metric. Section III presents an evaluation of the TGI metric. Specifically, we propose the desired property for an energy-efficient metric and analyze the validity of the TGI metric. We also look into appropriate weights which can be used with TGI. In Section IV, we present experimental results, including a description of the benchmarks that we used to evaluate the TGI metric and a comparison of TGI with the performance metrics used in the benchmarks. Related work in this area is described in Section V. Section VI concludes the paper.

II. THE GREEN INDEX (TGI) FOR HPC

Formulating a canonical *green metric* for a diverse benchmark suite that stresses different components of an HPC system is a challenging task. To the best of our knowledge, there exists no methodology that can be used to combine all the different performance outputs from the different benchmark tests and deliver a *single number* of energy efficiency to look at. Despite arguments that energy efficiency can only be represented by a vector which captures the effect of energy consumed by a benchmark suite, we seek the "holy grail" of a single representative number with which to make comparisons.

The earliest metric for comparing system performance is the *Standard Performance Evaluation Corporation (SPEC) rating* [6]. As shown in Equation (1), the SPEC rating defines the performance of a system under test, relative to a reference system, where time is used as the unit of performance. A SPEC rating of 25 means that the system under test is 25 times faster than the reference system. The rating is relative to a reference system in order to normalize and ease the process of comparison with other systems. The reference system for each benchmark is different. For example, the reference system for SPEC CPU2000 is a Sun Ultra5_10 workstation whereas SPEC CPU2006 uses the Sun Ultra Enterprise 2 workstation as its reference machine.

$$\text{SPEC rating} = \frac{\text{Performance of Reference System}}{\text{Performance of System Under Test}} \tag{1}$$

We propose the Green Index (TGI) metric in an effort to capture the "greenness" of supercomputers by combining all the different performance outputs from the different benchmark tests. We believe that TGI is the first effort towards providing a single number for evaluating the energy efficiency of supercomputers using a benchmark suite.

Following an approach similar to the SPEC rating, the energy efficiency of a supercomputer is measured with respect to a reference system by providing a relative performance-to-watt metric. The TGI of a system can be calculated by using the following algorithm:

1) Calculate the energy efficiency (EE), i.e., performance-to-power ratio, while executing different benchmark tests from a benchmark suite on the supercomputer:

$$\text{EE}_i = \frac{\text{Performance}_i}{\text{Power Consumed}_i} \tag{2}$$

where each i represents a different benchmark test.

2) Obtain the relative energy efficiency (REE) for a specific benchmark by dividing the above results with the corresponding result from a reference system:

$$\text{REE}_i = \frac{\text{EE}_i}{\text{EE}_{Ref_i}} \tag{3}$$

where each i represents a different benchmark test.

3) For each benchmark, assign a TGI component (or weighting factor W) such that the sum of all weighting factor is equal to one.

4) Use the weighting factors and sum across product of all weighting factors and corresponding REEs to arrive at the overall TGI of the system.

$$TGI = \sum_i W_i * \text{REE}_i \tag{4}$$

TGI allows for flexibility in green benchmarking as it can be used and viewed in different ways by its consumers. For example, although we use the performance-per-watt metric for energy efficiency in this paper, the methodology used for computing TGI can be used with any other energy-efficient metric, such as the energy-delay product. The advantages of using TGI for evaluating the energy efficiency of supercomputers are as follows:

1) Each weighting factor can be assigned a value based of the specific needs of the user, e.g., assigning a higher weighting factor for the memory benchmark if we are evaluating a supercomputer to execute a memory-intensive application.
2) TGI can be extended to incorporate power consumed outside the HPC system, e.g., cooling.
3) TGI provides a single number that can be used to gauge the energy efficiency of a supercomputer with respect to different benchmark tests.

III. EVALUATION OF "THE GREEN INDEX" (TGI)

Before analyzing the space of weights, it is necessary that we understand the desired property of the metric. As mentioned earlier, we consider performance-to-power ratio as the energy efficiency metric. Usually performance in HPC community is measured in rates. Performance (measured in rate) to power ratio indirectly measures the operation for each joule of energy consumed as shown in Equation (5) in case of the $FLOPS/watt$ metric.

As a result, we consider the energy consumed for a given amount of work as the main indicator of the energy efficiency. Therefore, the metric should be chosen in a such a way that it is inversely proportional to energy consumed.

$$\begin{aligned}\text{FLOPS/watt} &= \frac{\text{Floating-Point Operations Per Second}}{\text{Joules/Second}} \\ &= \text{Floating-Point Operations Per Joule}\end{aligned} \tag{5}$$

$$AM = \frac{\sum \mathrm{X}_i}{n} \tag{6}$$

A. Arithmetic Mean

The simplest way to assign the TGI component (weighting factor) is to use the arithmetic mean i.e., assign an equal weighting factor to all the benchmarks. The arithmetic mean (AM) for a data set X_i where $i \geq 1$ in general is defined by Equation (6)

$$\text{TGI using Arithmetic Mean} = \frac{\sum_{i=0}^{n} \text{REE}_i}{n} \tag{7}$$

$$\begin{aligned}\text{TGI using Arithmetic Mean} &= \frac{\sum_{i=0}^{n} \frac{\text{EE}_i}{\text{EE}_{Ref_i}}}{n} \\ &= \frac{\sum_{i=0}^{n} \frac{\text{Performance Metric}_i}{\text{Power Consumed}_i}}{\text{EE}_{Ref_i} * n} \\ &\propto \sum_{i=0}^{n} \frac{1}{\text{Power Consumed}_i}\end{aligned} \tag{8}$$

In our case, AM can be applied to TGI by assigning equal weights to all the benchmarks in Equation (4) as shown in Equation (7). Equation (7) can be rewritten as Equation (8). Equation (8) satisfies the desired property as it is inversely proportional to the energy consumed by each benchmark given the performance. However, the benchmarks have different properties such as execution times, energy consumption and power consumption and arithmetic mean might not be the best candidate to capture the *true* energy efficiency of the system. We look at different ways of assigning weights in order to take the application properties into account. The rest of section gives a preliminary view on the effects of using different weights.

$$WAM = \sum_i W_i * \mathrm{X}_i \tag{9}$$

$$W_{t_i} = \frac{t_i}{\sum_i t_i} \tag{10}$$

$$W_{e_i} = \frac{e_i}{\sum_i e_i} \tag{11}$$

$$W_{p_i} = \frac{p_i}{\sum_i p_i} \tag{12}$$

B. Weighted Means

The weighted means are used for data where all the data points do not contribute equally to the final average. Each data point is assigned a different weight to account for its contribution to the average. The sum of weights assigned to each of the data point must be equal to one. Equation (9) shows the weighted arithmetic mean (WAM), where $\sum_i W_i = 1$.

We discuss the use of different weights such as time, energy and power in this section. Weights corresponding to time, energy and power are given by Equations (10), (11) and (12) where t_i, e_i and p_i refer to the execution time, energy and power consumed by each benchmarks respectively. Equations (13), (14) and (15) show the derivations of TGI using W_{t_i}, W_{e_i} and W_{p_i} respectively for a given performance where M_i is the metric used by the benchmark. TGI using time as weight given the performance has the desired property. However, using energy and power as weights cancels the effect of the the energy component of the benchmarked systems. We compare TGI using all these weights with the traditional performance to power metric in Section IV.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

The reference system used in all of the experiments is called SystemG. The cluster consists of 324 Mac Pros, each with two 2.8-GHz quad-core Intel Xeon 5462 processors and 8GB of RAM. QDR InfiniBand is used as the interconnect. We use 128 nodes from the cluster for a total of 1024 cores.

We evaluate the TGI metric on our Fire cluster. It is an eight-node cluster, where each node consists of two AMD Opteron 6134 operating at 2.3 GHz and has 32 GB of memory. The core count for the entire cluster is 128. The cluster is capable of delivering 910 GFLOPS on the LINPACK benchmark.

The energy consumption is measured using a "Watts Up? PRO ES" power meter. The power meter is connected between the power outlet and the system as shown in Figure 1.

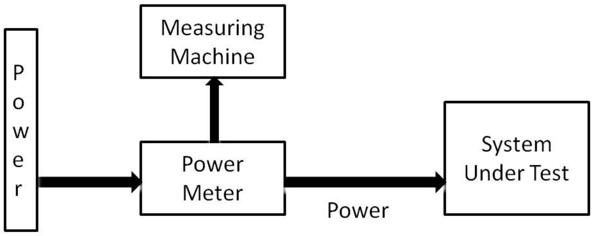

Fig. 1. Power Meter Setup

A. Benchmarks Used

To test the effectiveness of TGI, we need benchmarks which stress on different components of the system. We chose to use HPL [2], STREAM [5] and IOzone [4] benchmarks as they stress on CPU, memory and I/O subsystem respectively for a preliminary analysis of TGI. Note that TGI is neither limited by the metrics used in each benchmarks nor by the number of benchmarks.

HPL is a popular benchmark in HPC community as exemplified by the Top500 list. It solves a dense linear system of equations of the form $Ax = b$ of the order N. It uses LU factorization with row partial pivoting of matrix A and the solution x is obtained by solving the resultant upper triangular system. The data is distributed on a two-dimensional grid using a cyclic scheme for better load balance and scalability. The HPL benchmark reports its performance as gigaflops (or Gflops). The scalability of the energy efficiency of the benchmark on Fire cluster is shown in Figure 2.

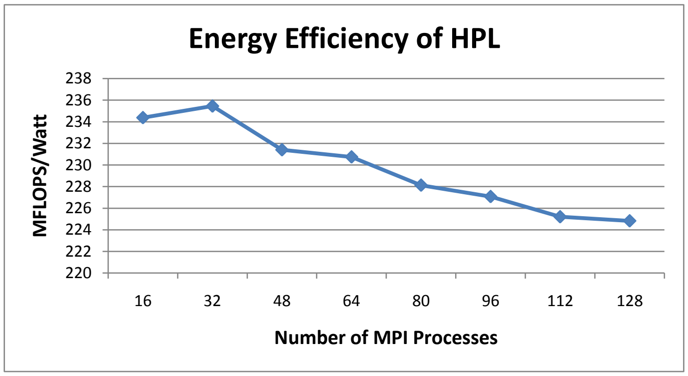

Fig. 2. Energy Efficiency of HPL

STREAM is a synthetic benchmark, which is used to measure the sustainable memory bandwidth of a system and the performance for the computations. There are four different computations performed by the benchmark: *Copy*, *Scale*, *Add*, and *Triad*. We are mainly interested in *Triad* as it is widely accepted that "multiply and accumulate" is the most commonly used computation in scientific computing. *Triad* scales a vector A and adds it to another vector B and writes the result to a third vector C, as shown in Equation (16).

$$Triad : C = \alpha A + B \tag{16}$$

The STREAM benchmark reports its memory performance as megabytes per second (MBPS). The energy efficiency of the benchmark for different number of processes on the Fire cluster is shown in Figure 3.

$$
\begin{aligned}
\text{TGI using } W_{t_i} &= \frac{\sum_{i=0}^{n} W_{t_i} * EE_i}{EE_{Ref_i}} \\
&= \frac{1}{\sum_i t_i * EE_{Ref_i}} * \sum_{i=0}^{n} \frac{t_i * M_i}{t_i * \mathrm{p}_i} \\
&\propto \sum_{i=0}^{n} \frac{1}{p_i} \qquad (13)
\end{aligned}
$$

$$
\begin{aligned}
\text{TGI using } W_{e_i} &= \frac{\sum_{i=0}^{n} W_{e_i} * EE_i}{EE_{Ref_i}} \\
&= \frac{1}{\sum_i e_i * EE_{Ref_i}} * \sum_{i=0}^{n} \frac{e_i * M_i * t_i}{t_i * \mathrm{e}_i} \\
&\propto \frac{1}{\sum_i e_i} \qquad (14)
\end{aligned}
$$

$$
\begin{aligned}
\text{TGI using } W_{p_i} &= \frac{\sum_{i=0}^{n} W_{p_i} * EE_i}{EE_{Ref_i}} \\
&= \frac{1}{\sum_i p_i * EE_{Ref_i}} * \sum_{i=0}^{n} \frac{p_i * M_i}{t_i * \mathrm{p}_i} \\
&\propto \frac{1}{\sum_i p_i} \qquad (15)
\end{aligned}
$$

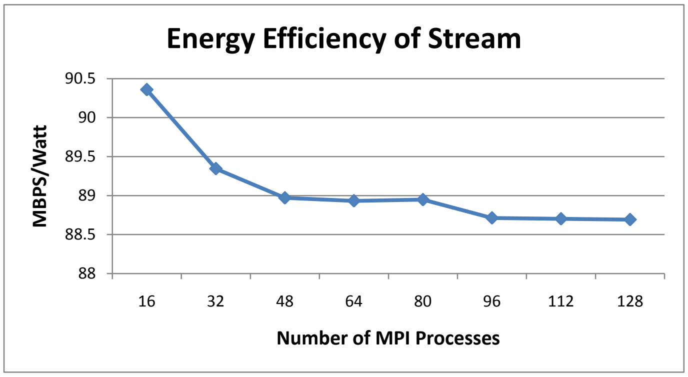

Fig. 3. Energy Efficiency of Stream

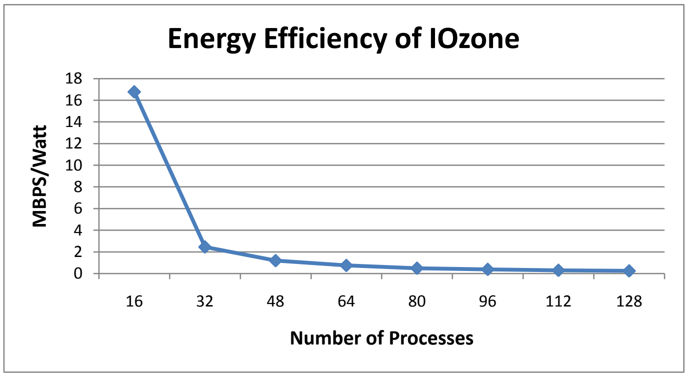

Fig. 4. Energy Efficiency of IOzone

IOzone benchmark stresses on the IO subsystem by performing a variety of file operations. The tool allows us to test the IO performance with various file sizes using typical file system operations such as reads and writes. We perform only the write test using IOzone for simplicity of evaluation as we can extend TGI methodology for any number of tests. The benchmark reports the performance results in MBPS like the STREAM benchmark. The energy efficiency of the benchmark on different number of nodes on the Fire cluster is shown in Figure 4.

B. *Discussion on TGI Metric*

In this section, we evaluate the TGI metric on the Fire cluster. We further make an initial attempt to understand the effects of different weights as discussed in Section III. Table I shows the performance achieved and power consumed by the individual benchmarks on SystemG.

Figure 5 shows TGI for different number of cores on the Fire cluster. Each point in Figure 5 represents TGI calculated while executing HPL, STREAM and

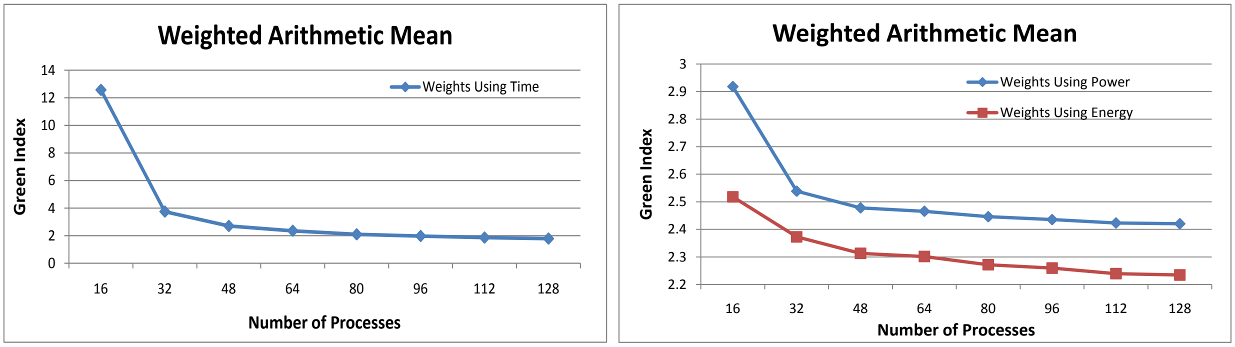

Fig. 6. TGI using Weighted Arithmetic Mean

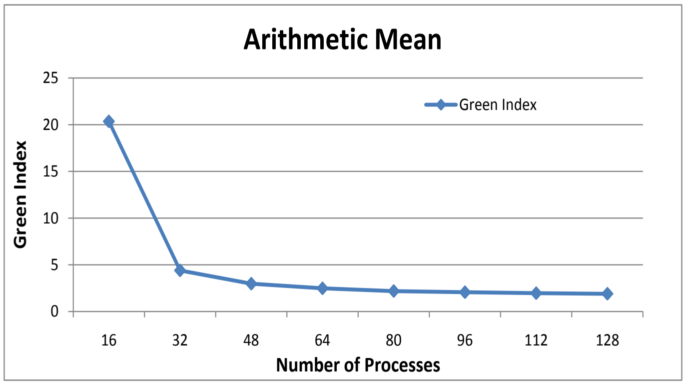

Fig. 5. TGI using Arithmetic Mean

Benchmark	**Performance**	**Power**
HPL	8.11 TFLOPS	80.83 KW
STREAM	1146268.657 MBPS	34.27 KW
IOzone	456.2336 MBPS	1.52 KW

TABLE I
PERFORMANCE ON SYSTEMG

IOzone using a particular number of cores in the cluster. We expect the TGI metric to be bound by benchmark with least REE (refer Section II) to provide a view of system-wide energy efficiency. The fact that TGI follows a similar trend to the energy efficiency of IOzone (Figure 4) is a indicator of the goodness of the metric. To further understand the correlation between the energy efficiency of the benchmarks and TGI, we used the Pearson correlation coefficient (PCC). Equation (17) shows the PCC where X_i and Y_i are the data, $\bar{X}$ and $\bar{Y}$ are the respective means of data sets, S_X & S_Y are the respective standard deviations and n is the number of samples. The value of PCC can lie between -1 and $+1$.

$$\text{PCC} = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y_i - \bar{Y})}{(n-1)S_X S_X} \qquad (17)$$

The correlation between TGI and energy efficiency of IOzone, Stream and HPL are .99, .96 and .58 respectively. The trend followed by TGI is different from the trend followed by energy of HPL (Figure 3), as shown by the correlation, which throws light on the deviation of system-wide energy efficiency from traditional MFLOPS/Watt. The TGI metric using arithmetic mean is useful in capturing the system-wide energy efficiency.

Benchmark	**Time**	**Energy**	**Power**
IOzone	.99	.92	.99
Stream	.97	.98	.98
HPL	.61	.83	.66

TABLE II
PCC BETWEEN ENERGY EFFICIENCY OF INDIVIDUAL BENCHMARKS AND TGI METRIC USING DIFFERENT WEIGHTS

Figure 6 shows TGI for using different weights on different number of cores on the Fire cluster. We use PCC to study the correlation between the energy efficiency of individual benchmarks and TGI metric computed using different weights (shown in Table II). TGI using time as weights shows similar correlation to individual benchmarks when compared to TGI using arithmetic mean. TGI using energy and power as weights show higher correlation with the energy efficiency of the HPL benchmark which is not a desired property as HPL benchmark does not have the least REE.

V. RELATED WORK

Existing benchmark metrics for evaluating energy efficiency have been studied in earlier literature. In Hsu et al. [11], for example, metrics like the *e*nergy-*d*elay *p*roduct (EDP) and the FLOPS/watt ratio were analyzed and validated on several platforms. The authors concluded that FLOPS/watt metric was best for evaluating the energy efficiency.

The problem of capturing the overall performance in a single number has been well studied. The study of different means to arrive at a single number for performance results of a benchmark suite was presented by Smith [13]. The use of weighted means to arrive at a aggregate metric was investigated by John [12]. The author analyzed different weights for different central tendency measures and concluded that both arithmetic and harmonic means can be used to summarize performance if appropriate weights are applied. Further, the author analyzed the suitability of a central tendency measure to provide a single number with a set of speedups and provide appropriate weights corresponding to the mean.

The MFLOPS metric was examined by Giladi [10]. The author used statistical analysis to validate the MFLOPS metric over different problem sizes and applications. The author further analyzed the correlation between MIPS and MFLOPS and found that there is either a strong or a lack of correlation.

In this paper, we proposed a metric to capture the energy efficiency of all the components of system in a *single number* even while using benchmarks which produces different metric as output.

VI. CONCLUSION AND FUTURE WORK

We proposed TGI, a metric to aggregate the energy efficiency of a benchmark suite in a *single number*, and discussed its advantages. We further presented preliminary evaluation of the goodness of the metric and initial investigation into the use of weights to accommodate the different application properties such as execution time, power and energy consumed by the benchmark. Preliminary results describing the goodness of the metric and the correlation of TGI with conventional metrics were presented.

As a future work, we want to establish the general applicability of TGI by benchmarking more systems. The suitability of TGI to various kind of platforms, such as GPU based system, is of particular interest. We want thoroughly investigate the suitability of different weights for TGI. Further, we would like to extend TGI metric to give a center-wide view of the energy efficiency by including components such as cooling infrastructure.

ACKNOWLEDGEMENT

This project was supported in part by the National Science Foundation (NSF) under the grant CCF-0848670.

REFERENCES

[1] DARPA IPTO Exascale Computing Study: Technology Challenges in Acheiving Exascale Systems.

[2] High Performance LINPACK (HPL). Available at http://www.netlib.org/benchmark/hpl.

[3] HPC Challenge Benchmarks. Available at http://icl.cs.utk.edu/hpcc.

[4] IOZone Filesystem Benchmark. http://www.iozone.org.

[5] STREAM Benchmark. http://www.cs.virginia.edu/stream.

[6] The Standard Performance Evaluation Corporation (SPEC). http://www.spec.org/.

[7] The Top500 list. Available at http://top500.org.

[8] W. Feng. Personal Communication with Sicortex.

[9] W. Feng and K. Cameron. The Green500 List: Encouraging Sustainable Supercomputing. *Computer*, 40(12):50–55, 2007.

[10] R. Giladi. Evaluating the Mflops Measure. *Micro, IEEE*, 16(4):69–75, 1996.

[11] C. Hsu, W. Feng, and J. S. Archuleta. Towards Efficient Supercomputing: A Quest for the Right Metric. In *Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'05) - Workshop 11 - Volume 12*, page 230.1. IEEE Computer Society, 2005.

[12] L. K. John. More on Finding a Single Number to Indicate Overall Performance of a Benchmark Suite. *ACM Computer Architecture News*, 2004.

[13] J. E. Smith. Characterizing Computer Performance with a Single Number. *Communications of the ACM*, 31:1202–1206, Oct. 1988. ACM ID: 63043.