# The Green Index (TGI): A Metric for Evaluating Energy Efficiency in HPC Systems

Wu Feng and Balaji Subramaniam

VirginiaTech
Invent the Future

SyNeRG
synergy.cs.vt.edu

# Metrics for Energy Efficiency

- Energy-Delay Product (EDP)
  - Used primarily in circuit design
  - Variants can be used to emphasize more on energy or performance
- FLOPS/Watt
  - Used in high-performance computing
  - Metric for ranking in Green500 list

C. Hsu, W. Feng, J. Archuleta, "Towards Efficient Supercomputing: A Quest for the Right Metric," *Workshop on High-Performance, Power-Aware Computing* (in conjunction with the *IEEE International Parallel & Distributed Processing*), April 2005.

# Issue with Traditional Metrics?

- "In fact, in many cases it has be become clear to the panel that the non-computational aspects of the energy problem, especially the energy in data transport, will dwarf the traditional computational component in future." – DARPA-IPTO Exascale Computing Study.

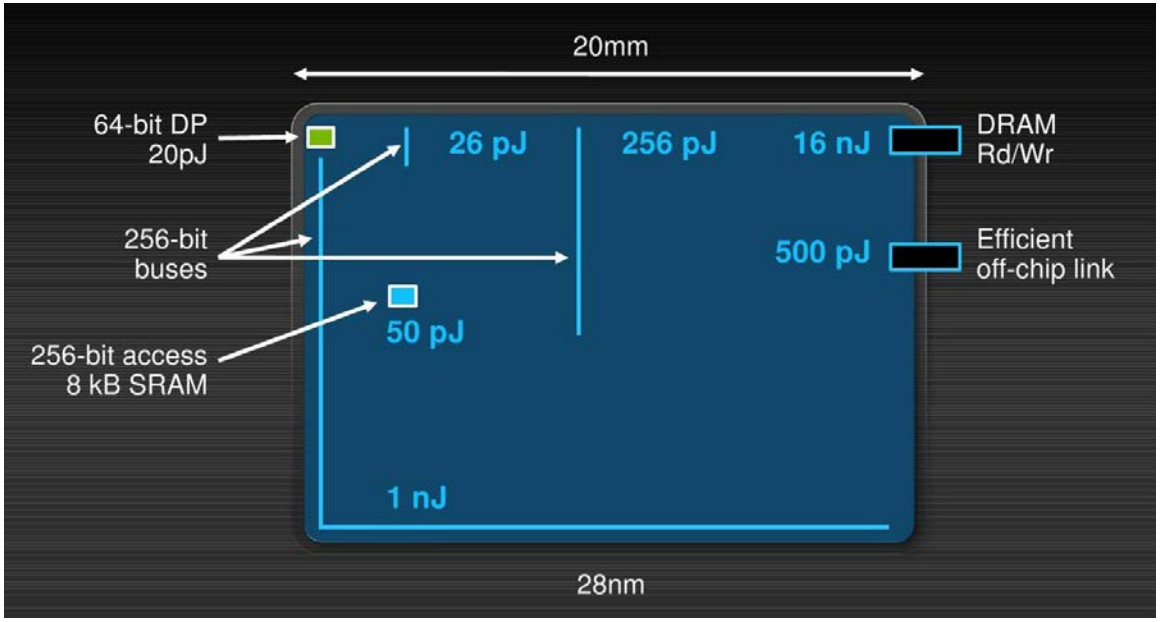


Source: William J. Dally, GPU Computing to Exascale and Beyond, Keynote Talk at SC 2010

# Issue with Workload?

- LINPACK
  - Focus on CPU
    - … but what about the other subsystems?

EE HPC WG – TOP500 – Green Grid – Green500 Goals

- Identify ***workloads*** for exercising other sub-systems; e.g., memory, storage, I/O
- Refine ***methodology*** for measuring the power of supercomputers
- Identify appropriate ***metrics*** for energy efficiency

# Our Approach: New Workload and Metric

- Use multiple benchmarks
  - Stress different components
- Arrive at a system-wide energy efficiency metric
  - Represent all components of the system

# Challenges

- Different benchmarks use different metrics
  - HPL uses FLOPS
  - STREAM uses mega bytes per second (MBPS)
- How to combine these metrics? Should they be combined?
  - Which metric should be used?
  - How to arrive at a single number?

# The Green Index (TGI)

- Methodology to represent system-wide energy efficiency
  - Use multiple benchmarks
- Single number
  - Combine metrics from each of the benchmark
- Energy efficiency relative to a reference system
  - Similar to SPEC rating

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

# Algorithm To Calculate TGI

- Calculate performance-to-power ratio* for each benchmark

$$EE_i = \frac{\text{Performance}_i}{\text{Power Consumed}_i}$$

- Obtain the relative energy efficiency (REE) w.r.t. the reference machine. REE is dimensionless.

$$REE_i = \frac{EE_i}{EE_{Ref_i}}$$

*i represents the set of all benchmarks and Ref represents the reference machine*

* Note: We use performance-to-power ratio as the metric of choice. However, any other efficiency metric can be used with TGI.

# Algorithm To Calculate TGI

- For each benchmark, assign a weighting factor ($W_i$)
- Sum across the product of REEs and $W_i$ to arrive at TGI for the system

$$TGI = \Sigma W_i * REE_i$$

# Assigning Weights ($W_i$)

- Arithmetic Mean
  - Assign equal weights to all the benchmarks
- Geometric Mean
- Weighted Arithmetic Mean (WAM)
  - Time ($W_{ti}$)
  - Energy ($W_{ei}$)
  - Power ($W_{pi}$)

# Assigning Weights ($W_i$)

- Time as weight

$$W_{t_i} = \frac{t_i}{\sum_i t_i} \qquad \text{TGI} = \frac{1}{\sum_i t_i * EE_{Ref_i}} * \sum_{i=0}^{n} \frac{t_i * M_i}{t_i * \mathrm{p}_i}$$

- Energy as weight

$$W_{e_i} = \frac{e_i}{\sum_i e_i} \qquad \text{TGI} = \frac{1}{\sum_i e_i * EE_{Ref_i}} * \sum_{i=0}^{n} \frac{e_i * M_i * t_i}{t_i * \mathrm{e}_i}$$

- Power as weight

$$W_{p_i} = \frac{p_i}{\sum_i p_i} \qquad \text{TGI} = \frac{1}{\sum_i p_i * EE_{Ref_i}} * \sum_{i=0}^{n} \frac{p_i * M_i}{t_i * \mathrm{p}_i}$$

# Systems Used

- Reference System (System G)
  - Two (2) Intel Xeon 5462 quad-core processors at 2.8 GHz
    - 128 nodes → 1,024 cores
  - 8GB memory
  - QDR InfiniBand

- System Under Test (Fire)
  - Two (2) AMD Opteron 6134 (Magny Cours) oct-core processors at 2.3 GHz
    - 8 nodes → 128 cores
  - 32 GB of memory
  - QDR InfiniBand

# Benchmarks Used in TGI

- High-Performance LINPACK (HPL)
  - Stresses the CPU component of the system
  - Output metric: Floating-point operations per second (FLOPS)
- STREAM Benchmark
  - Stresses the memory sub-system
  - Output metric: Megabytes per second (MPBPS)
- IOZone
  - Stresses the I/O sub-system
  - Output metric: Megabytes per second (MPBPS)

# Energy Efficiency of Benchmarks (Fire Cluster)

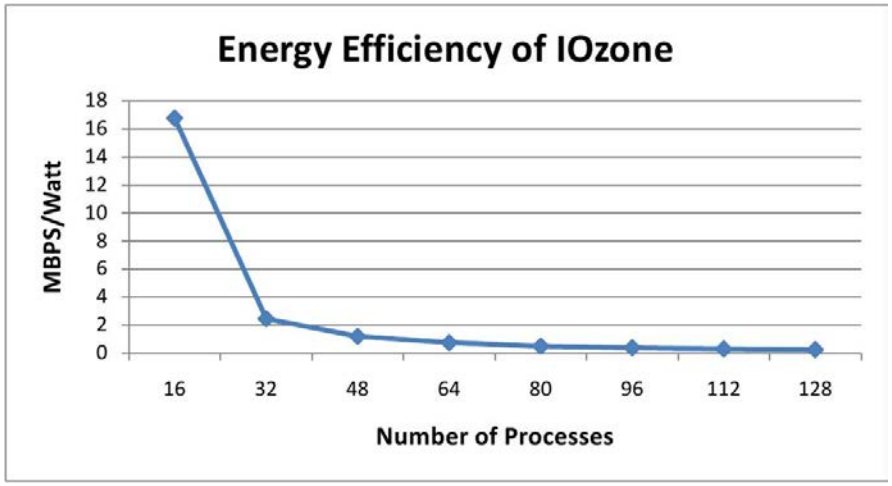


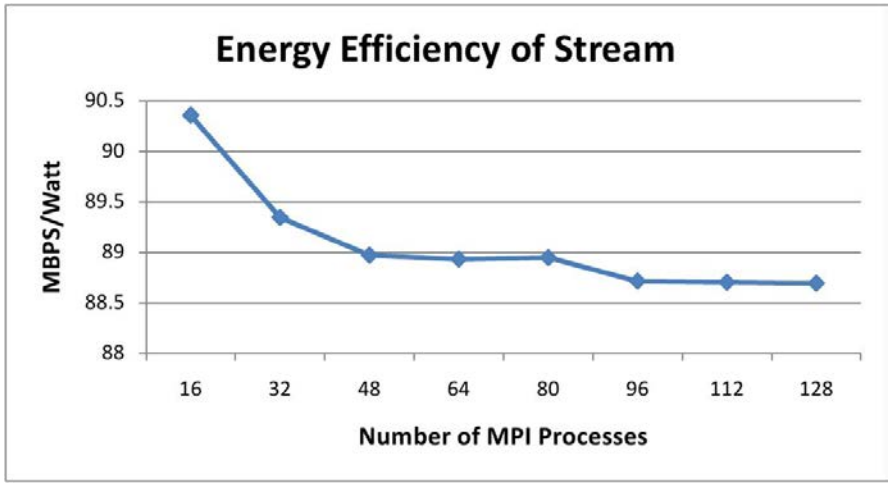


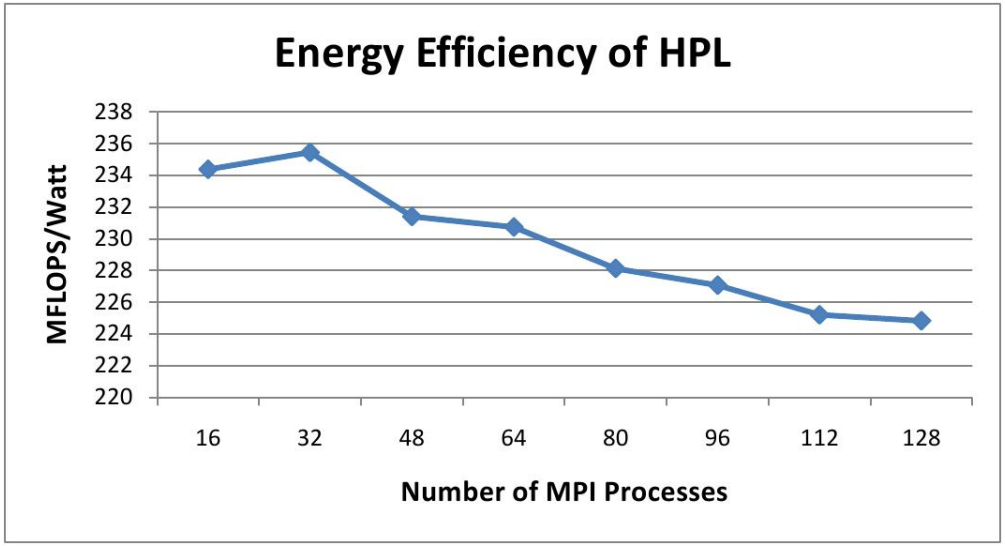

# The Green Index (Scalability)

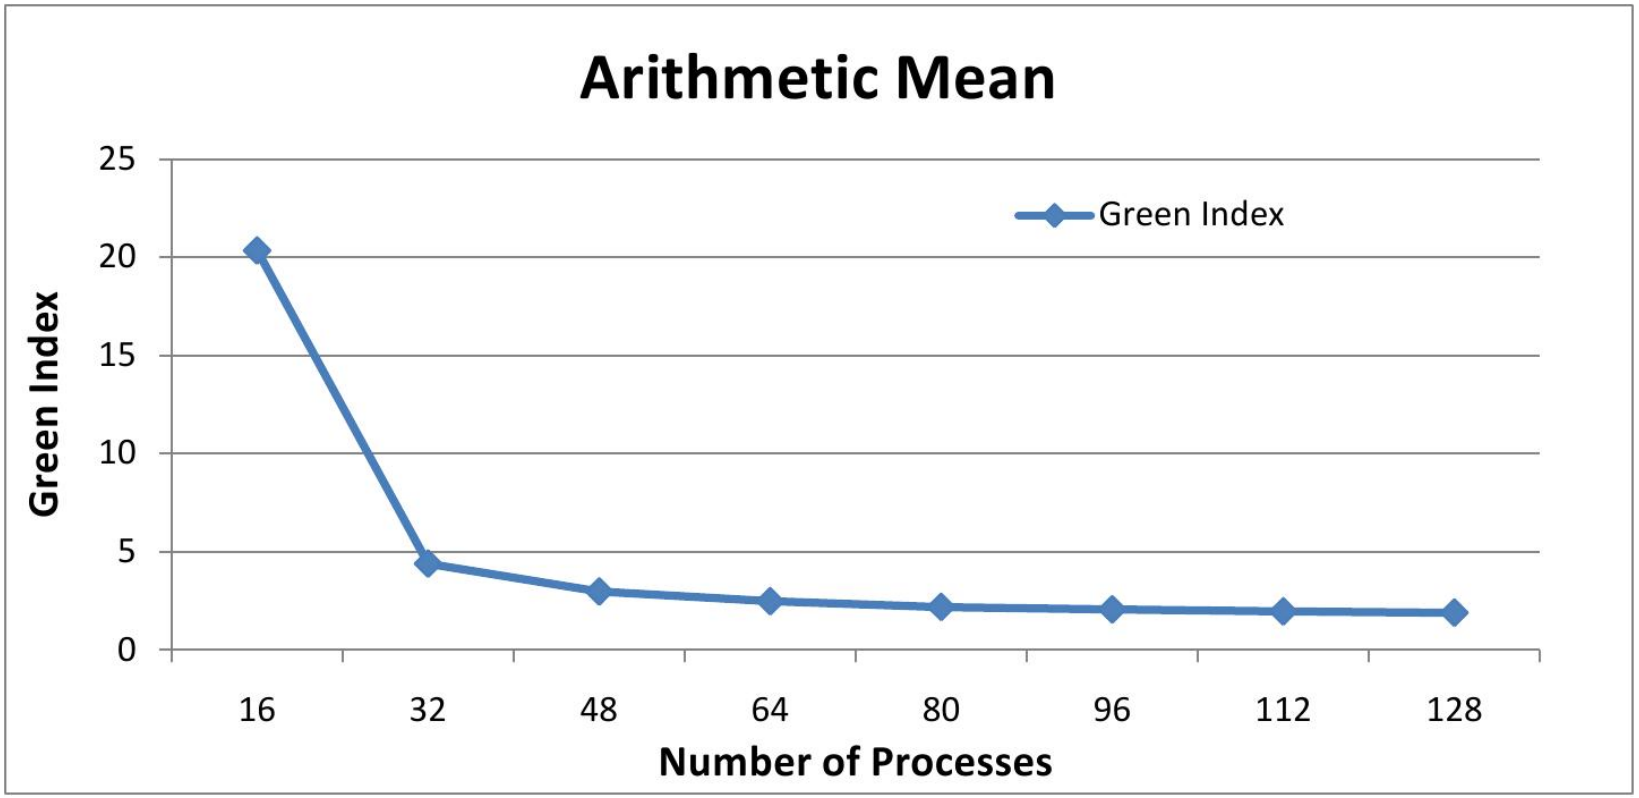

# The Green Index (Weighting Factors)

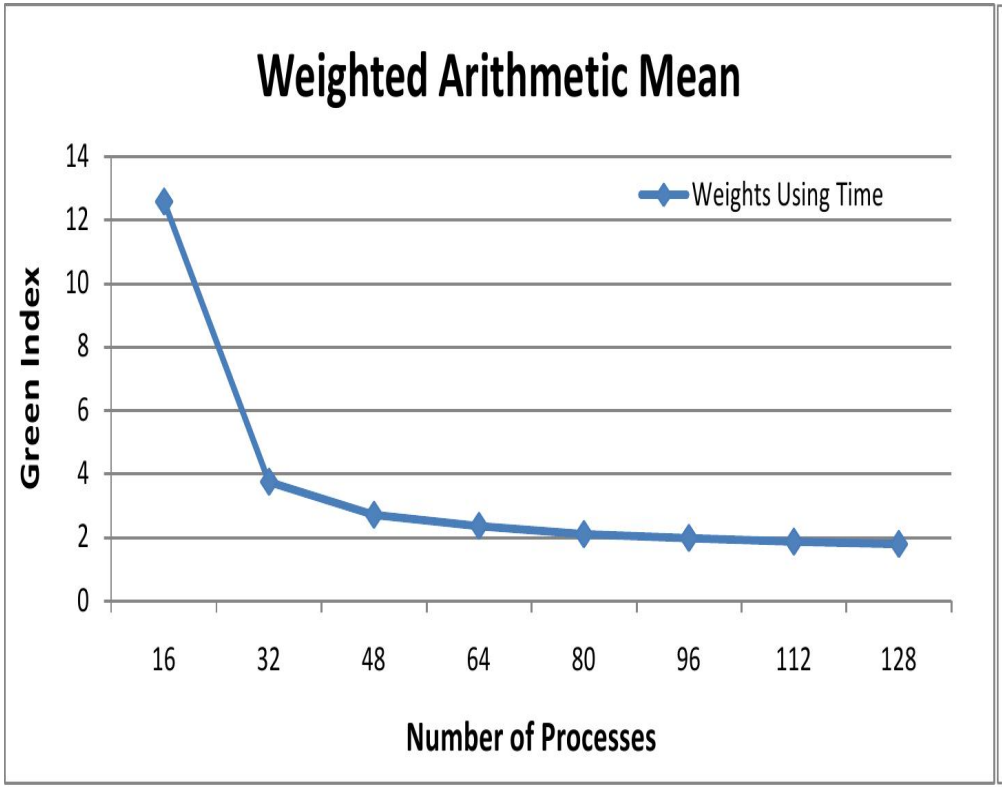


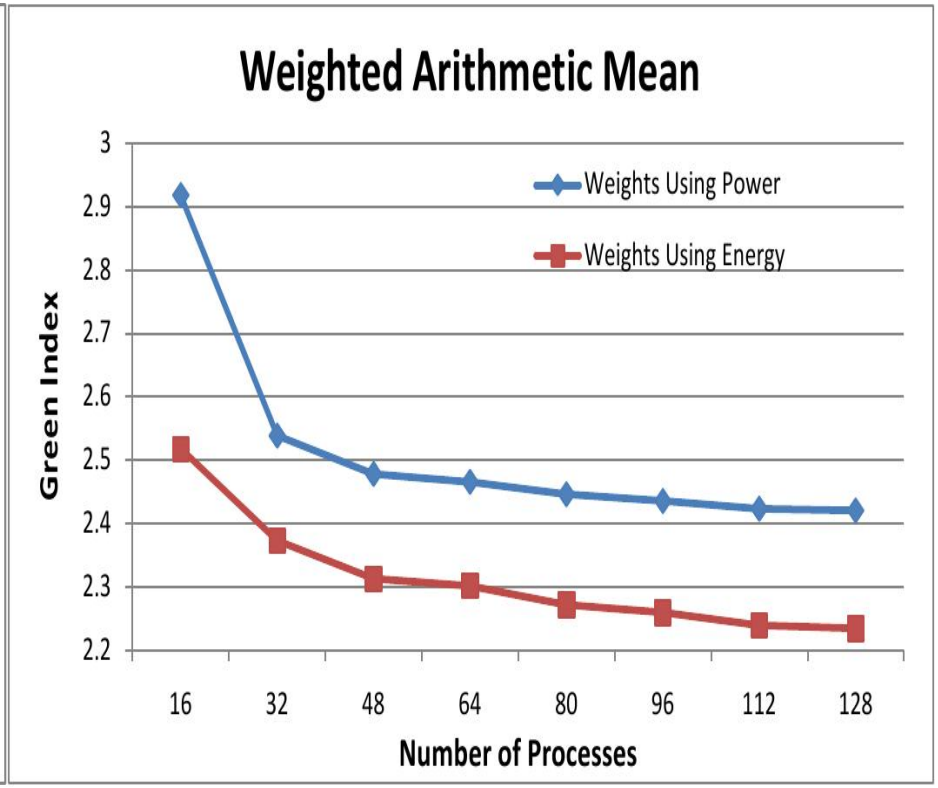

# Future Work

- Benchmark more systems
- Methodology for heterogeneous systems
- Rigorous study for weight factor assignment
- Adapt TGI based on application
  - Assign weights based on component utilization by application?
  - Analyze whether the system is energy efficient for a particular application
- System-wide → Center-wide
  - Include cooling infrastructure …

# Summary

- The Green Index (TGI)
  - A methodology to provide insights into system-wide energy efficiency
  - Comparison with traditional metrics (performance-to-power ratio)
  - Preliminary analysis

For more information,

B. Subramaniam and W. Feng, "The Green Index: A Metric for Evaluating System-Wide Energy Efficiency in HPC Systems," *Workshop on High-Performance, Power-Aware Computing* (in conjunction with the *IEEE International Parallel & Distributed Processing*), May 2012.