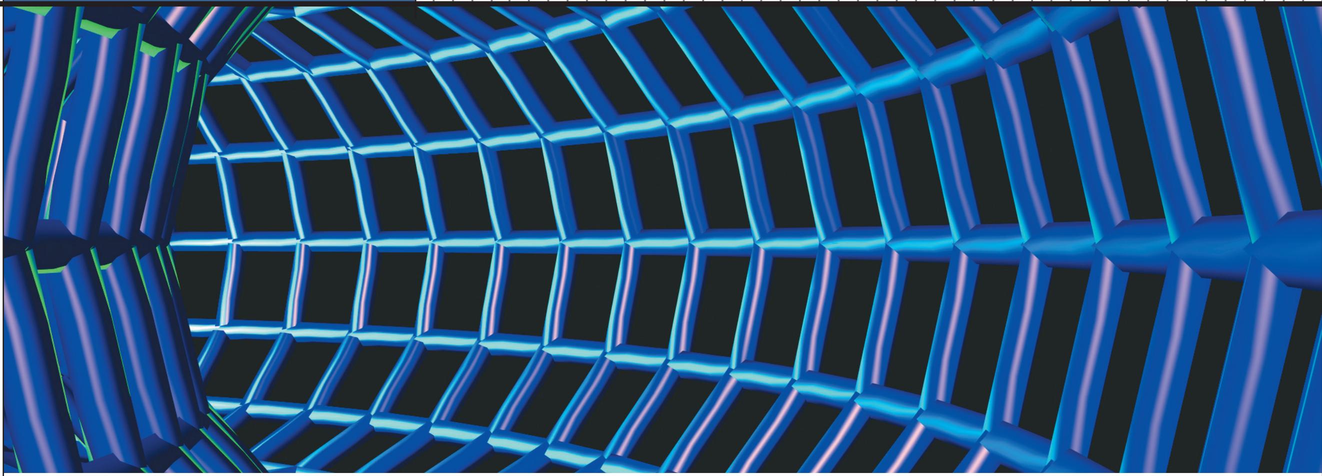

Trends in Server Energy Proportionality

Frederick Ryckbosch, Stijn Polfliet, and Lieven Eeckhout, *Ghent University, Belgium*

Server energy proportionality, as quantified by the proposed EP metric, has improved significantly, from 30-40 percent in 2007 to 50-80 percent today, but much more can be done to move systems closer to ideal.

Energy efficiency has emerged as a major design driver in servers, as it impacts capital costs as well as powering and cooling expenses.

Luiz Barroso and Urs Hölzle[1] demonstrated that server capital costs dominate the total cost of ownership in a classic datacenter, with 69 percent of the monthly TCO related to server purchase and maintenance costs. In contrast, the cost of all infrastructure and power to host a contemporary datacenter with commodity-based lower-cost servers or higher power prices is more than twice the purchase and maintenance costs. The researchers concluded that, with electricity and construction expenses rising, datacenter facility costs—which are proportional to power consumption—will become an increasingly larger part of the TCO. In other words, a datacenter's TCO primarily will be a function of its power consumption, and the purchase and maintenance costs will matter less. Besides cost considerations, improving energy efficiency reduces carbon dioxide emissions, leading to greener IT.

Improving a server's energy efficiency is nontrivial and presents many challenges. Peak power consumption affects capital costs for power distribution, the supply units, and the cooling infrastructure. High power consumption leads to increased power density and temperature, which affects cooling costs and possibly hardware reliability and availability. Total energy consumption impacts the operating expense for powering the servers. Further, servers in large datacenters rarely are completely idle and seldom at or near maximum utilization—they typically operate in the 10-50 percent utilization range.[2] This insight motivated Barroso and Hölzle[3] to make the case for servers that consume energy proportional to their utilization level or load. The key is that servers should be optimized both to reduce peak power consumption and to maximize performance at lower utilization levels.

In December 2007, the Standard Performance Evaluation Corporation released SPECpower_ssj2008, an industry-standard benchmark for server power and performance.[4] SPECpower measures performance and power at different utilization levels, from which it computes an overall metric. The metric thus includes some notion of energy proportionality but does not explicitly quantify it.

We propose a metric to quantify a server's energy proportionality. Using the EP metric on published power and performance data, we evaluated how energy proportionality has evolved over time, examined how the EP metric relates to SPECpower, and quantified how further improving servers' EP can reduce total energy consumption.

0018-9162/11/$26.00 © 2011 IEEE

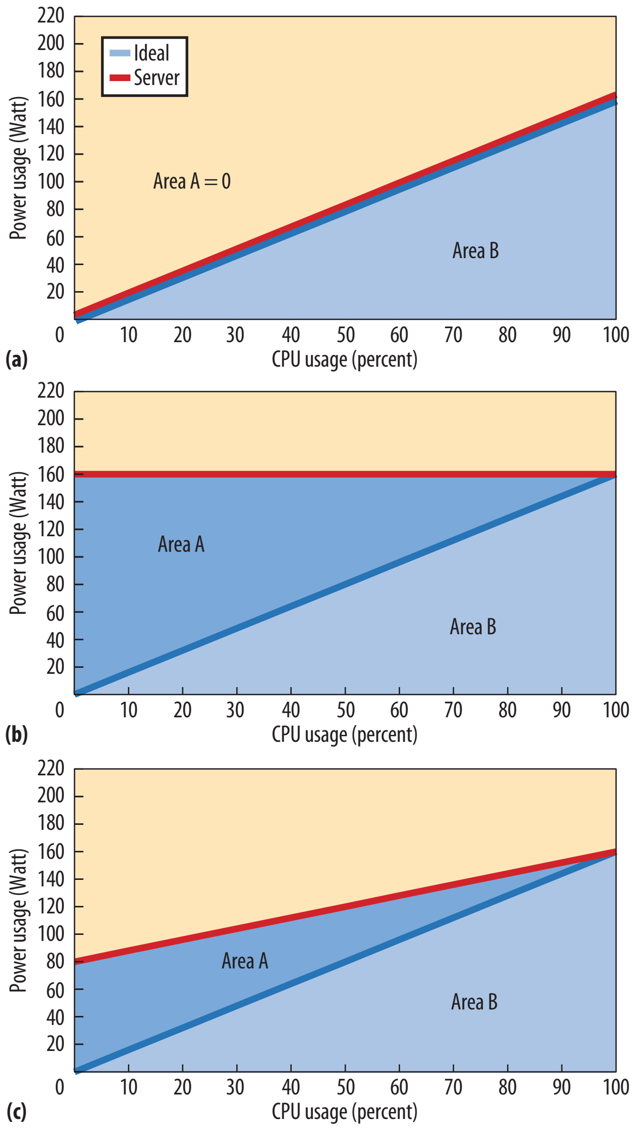

Figure 1. Energy proportionality (EP) is computed as 1 minus area A divided by area B; area A is defined as the area between the server and ideal curves, and area B is defined as the area below the ideal curve. (a) Perfect energy-proportional system (EP = 1). (b) Non-energy-proportional system (EP = 0). (c) 50 percent energy-proportional system (EP = 0.5).

EP METRIC

An analogy can easily explain the EP concept. A car's engine consumes fuel when the car is being driven, and even more fuel when the car accelerates. However, the engine also consumes fuel when it is running but idle—for example, when the car is stopped at a traffic light or stop sign. In these situations, the energy consumption is not proportional: the engine consumes energy (fuel) even though the car does not make physical progress. Similarly, a computer system consumes energy when it is idle—the system is powered on but does not do any useful work.

Ideally, an energy-proportional system would consume zero power when completely idle, and it would consume power proportional to its utilization level when doing useful work—for example, when operating at 30 percent of its peak performance, the system should consume 30 percent of its peak power. In practice, however, power consumption is higher than what the ideal scenario would suggest.

Informed by the seminal work on energy proportionality by Barroso and Hölzle,[3] we define a server's EP as 1 minus the integral of the relative delta in power consumption for the ideal energy-proportional server, across a range of utilization levels. As Figure 1 shows, EP is computed as 1 minus the area between the server's power consumption and the ideal power consumption curve (area A) divided by the area under the ideal curve (area B). An EP of 1 means that the server consumes power proportional to its load (Figure 1a). An EP of 0 indicates that the server consumes a constant amount of power irrespective of its load (Figure 1b). Figure 1c represents a server with an EP of 0.5; this server consumes 50 percent of its peak power at zero load. Intuitively, EP can serve as a quantitative metric for how closely a server's energy proportionality approaches perfect scaling across different server utilization levels.

Figure 2 quantifies system energy proportionality of 213 systems under test from 20 vendors between the fourth quarter of 2007 and early January 2011 (www.spec.org/power). It is encouraging to observe that server manufacturers are designing increasingly energy-proportional systems. Whereas older systems (circa 2007) had an EP in the 30-40 percent range, current systems have an EP in the 50-80 percent range; some servers even have an EP close to 90 percent.

EP VERSUS SPECPOWER

The SPECpower_ssj2008 benchmark defines a server's power efficiency as the sum of the performance measured at each utilization level divided by the sum of the average power at each utilization level, including active idle. Performance is measured as the number of transactions completed per second over a fixed period of time. The benchmark starts its execution with a calibration phase to determine the system's maximum throughput; it then measures performance (throughput) and power at each utilization level starting at maximum load and decreasing in 10 percent increments.

Because SPECpower includes power and performance numbers at different utilization levels, it arguably already includes some notion of energy proportionality. Why, then, is there a need for the EP metric? Figure 3 plots EP versus SPECpower score for the same 213 systems under test. Although SPECpower correlates well with EP, the correlation is not perfect: a system with a high EP does not necessarily have a high SPECpower score, and vice versa. The key difference is that EP focuses on energy proportionality only and quantifies how a server compares to the ideal energy-proportional system, whereas SPECpower quantifies average power and performance across different server load levels.

ROOM FOR IMPROVEMENT

Although energy proportionality has improved dramatically in recent years, a gap remains between contemporary servers and the ideal energy-proportional system. How much more energy (and cost) can be saved by making servers even more energy proportional?

To address this question, we compared power consumption at different utilization levels for the server with the highest EP and the server with the highest SPECpower score, as Figure 4 shows. Assuming that servers operate in the 10-50 percent utilization range most of the time[3]—in fact, we assumed that server operation is uniformly distributed between 10 and 50 percent—we derived how much total energy can be saved by making the server more energy proportional. For the server with the highest EP score, ideal energy proportionality can potentially reduce total energy consumption by 34 percent; for the server with the highest SPECpower score, total energy consumption can be reduced by around 50 percent. Improving energy proportionality further should significantly reduce energy and cost.

Although we likely will never achieve 100 percent energy proportionality—there will always be some overhead for keeping a server running—improving a system's energy proportionality remains an important goal. Barroso and Hölzle[3] reported that the CPU accounts for 50 percent of total system power at peak performance for a Google server. At lower utilization levels, however, it accounts for less than 30 percent, with the remaining 70 consumed by DRAM, hard drives, power supplies, and so on.

Looking at a large-scale datacenter, the building's mechanical and electrical infrastructure consumes considerable power, including chillers, computer room air conditioning, uninterruptible power systems, power distribution units, humidifiers, and so on. This suggests that achieving energy proportionality at the system level will require improvements across the entire system. Researchers are accordingly pursuing ways to increase energy proportionality in the CPU,[5] the network,[6] and storage devices;[7] others advocate rapidly transitioning an entire server between a high-performance active state and a near-zero-power idle state in response to instantaneous changes in load.[8]

Energy proportionality should be a key target in future server design. Although energy proportionality, as quantified by the proposed EP metric, has improved significantly in recent years, much more can be done to move systems closer to ideal energy proportionality. Toward this end, we are pursuing a workload

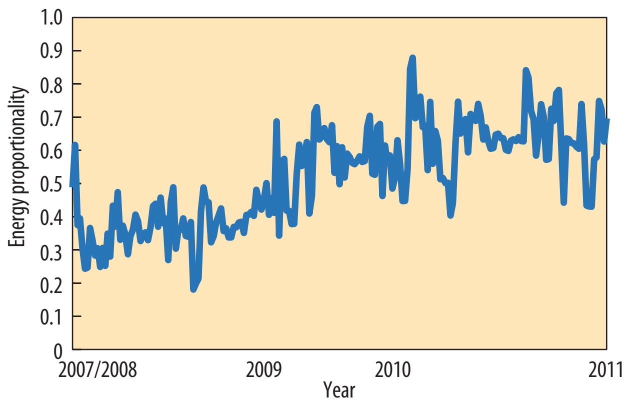

Figure 2. EP over time for 213 systems under test from 20 vendors. Servers are becoming increasingly energy-proportional, with some having an EP close to 90 percent.

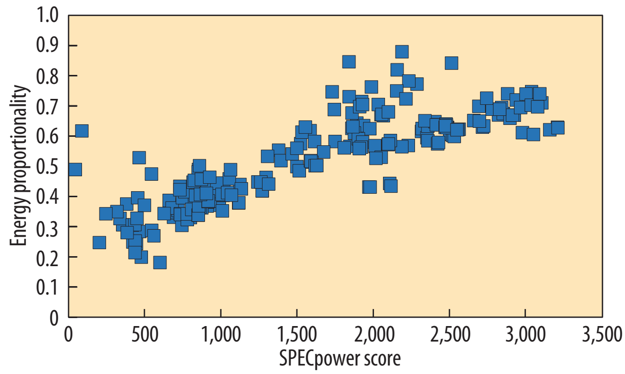

Figure 3. EP versus SPECpower score for the systems under test. EP focuses on energy proportionality only and quantifies how a server compares to the ideal energy-proportional system, whereas SPECpower quantifies average power and performance across different server load levels.

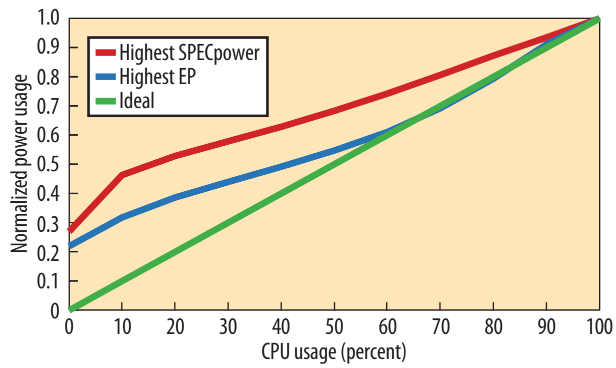

Figure 4. Power consumption, as a function of CPU load, for the system under test with the highest EP and the system with the highest SPECpower score. For the server with the highest EP score, ideal energy proportionality can potentially reduce total energy consumption by 34 percent; for the server with the highest SPECpower score, total energy consumption can be reduced by around 50 percent.

characterization and large-scale system simulation methodology for advising datacenter system integrators to move toward cost- and power-efficient infrastructures. ■

Acknowledgments

The authors thank the anonymous reviewers for their thoughtful comments and suggestions. Frederick Ryckbosch is supported through a doctoral fellowship by the Research Foundation—Flanders (FWO). Stijn Polfliet is supported through a doctoral fellowship by the Agency for Innovation by Science and Technology (IWT). Additional support is provided by the FWO projects G.0232.06, G.0255.08, and G.0179.10, the UGent-BOF projects 01J14407 and 01Z04109, and the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) ERC Grant agreement no. 259295.

References

1. L. Barroso and U. Hölzle, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*, Morgan & Claypool, 2009.
2. X. Fan, W.-D. Weber, and L.A. Barroso, "Power Provisioning for a Warehouse-Sized Computer," *Proc. 34th Ann. Int'l Symp. Computer Architecture* (ISCA 07), ACM Press, 2007, pp. 13-23.
3. L.A. Barroso and U. Hölzle, "The Case for Energy-Proportional Systems," *Computer*, Dec. 2007, pp. 33-37.
4. K.-D. Lange, "Identifying Shades of Green: The SPECpower Benchmarks," *Computer*, Mar. 2009, pp. 95-97.
5. Y. Watanabe, J.D. Davis, and D.A. Wood, "WiDGET: Wisconsin Decoupled Grid Execution Tiles," *Proc. 37th Ann. Int'l Symp. Computer Architecture* (ISCA 10), ACM Press, 2010, pp. 2-13.
6. D. Abts et al., "Energy Proportional Datacenter Networks," *Proc. 37th Ann. Int'l Symp. Computer Architecture* (ISCA 10), ACM Press, 2010, pp. 338-347.
7. J. Guerra et al., "Energy Proportionality for Storage: Impact and Feasibility," *ACM SIGOPS Operating Systems Rev.*, Jan. 2010, pp. 35-39.
8. D. Meisner, B.T. Gold, and T.F. Wenisch, "PowerNap: Eliminating Server Idle Power," *Proc. 14th Int'l Conf. Architectural Support for Programming Languages and Operating Systems* (ASPLOS 09), ACM Press, 2009, pp. 205-216.

***Frederick Ryckbosch** is a PhD student in the Electronics and Information Systems Department at Ghent University, Belgium. His research interests include computer architecture in general, and simulation, modeling, and optimization of high-end computer systems in particular. Ryckbosch received an MS in computer science and engineering from Ghent University. Contact him at frederick.ryckbosch@elis.ugent.be.*

***Stijn Polfliet** is a PhD student in the Electronics and Information Systems Department at Ghent University. His research interests include computer architecture in general, and simulation, modeling, and optimization of high-end computer systems in particular. Polfliet received an MS in computer science and engineering from Ghent University. Contact him at stijn.polfliet@elis.ugent.be.*

***Lieven Eeckhout** is an associate professor in the Electronics and Information Systems Department at Ghent University. His research interests include computer architecture and hardware/software in general, with a focus on performance analysis, evaluation and modeling, and workload characterization. Eeckhout received a PhD in computer science and engineering from Ghent University. He is a member of IEEE and the ACM. Contact him at leeckhou@elis.ugent.be.*