

PUE for end users - Are you interested in more than bread toasting?

Kay Grosskop

Software Improvement Group, Amsterdam, The Netherlands

Email: k.grosskop@sig.eu

Abstract—The Power Usage Effectiveness (PUE) indicator for efficiency of data center infrastructure has been very successful. But focusing solely on PUE tends to restrict action to data center infrastructure management and in some situations even gives a perverse stimulus against optimization at the IT equipment and software levels. Despite the high relevance, no accepted metric has emerged to support optimization and allow the rating of the energy efficiency of the whole stack.

This paper presents a metric, the *Consumption Near Sweet-spot* (CNS), that for a part can fill this gap. It captures how well the system-relative energy efficiency optimum and its utilization are aligned. A strong point is, that it allows a comparison of functionally very different services. The metric is compared to the *Fixed to Variable Energy Ratio* (FVER) metric for data centers recently proposed by the British Computer Society.

I. INTRODUCTION

The PUE is probably the single best known energy efficiency metric for computing infrastructure in use today. The PUE metric makes inefficiencies at the data center infrastructure level visible and allows decision makers to express requirement and achieved improvements with a conceptually simple, widely understood indicator.

But as the popularity of the PUE metric as a steering instrument has grown also its limitations have become more important. First, since it has not been designed to support end-to-end optimization of the whole computing stack, it cannot serve as an instrument to optimize the IT hardware and software that constitute the other main layers in a data center computing infrastructure (Figure 1). Even worse, infrastructure efficiency is typically higher if the data center runs at full capacity and hence optimizing solely for PUE values may result in a perverse stimulus to keep consumption of the IT as high as possible. Second, it does not relate the consumed energy to any useful work done in the data center. As far as concerned to the PUE, you could also operate bread toasters instead of doing any useful computation and still obtain a good efficiency rating. But the energy efficiency of an IT service should be expressed in terms of how much energy is used for a certain task that has some value for an end user like for example streaming a video or completing a monetary transaction.

Since IT equipment and Software are both interesting vectors of optimization it would be very useful to provide decision makers with an equally clear and widely applicable metric as PUE that targets the efficiency of the whole stack. However, two challenges to do so have proven to be hard to overcome:




Fig. 1. The PUE indicator as part of an end-to-end service efficiency

- 1) Finding a unifying, yet meaningful unit of 'tasks accomplished' for functionally different applications that would allow comparison of different systems as opposed to tracking a system's energy efficiency relative to itself.
- 2) Defining some absolute (theoretical) optimum for the energy efficiency of a system in a similar way as it exists for the PUE, where 1 is the best possible value.

The approaches discussed in this article try to avoid these challenges. They use a system-relative definition of *unit of work* and reject the idea that a generic, normalizing unit is necessary in order to establish a useful end-to-end metric for efficiency. Moreover, instead of relating system efficiency to some absolute optimum they attempt to measure known sources of inefficiency. A well known source of inefficiency is the inability of systems to scale with load. Barroso and Hölzle have coined the term *energy proportionality* [1] to point out that efficient systems and components should be able to scale energy consumption according to the amount of work done and when idle (unutilized) they should not consume any energy at all. This is an important property to have, since many real systems are actually underutilized most of the time, are unable to scale down their power needs and hence operate generally in a very inefficient mode. (Figure 2)

Yet, many of today's most efficient systems are far from energy-proportional but instead are optimized by shaping the workload in a way that the system utilization is constantly high. Our metric is designed to acknowledge this fact and to support *two* different strategies for optimization:

- 1) Strive for energy proportionality.
- 2) Raise system utilization in order to let the system operate in an efficient load region.

The second strategy would for example use workload placement and performance tuning. The first will often boil down to reducing energy consumption for an idle or underutilized system.




Fig. 2. Energy consumption per load level for a typical and an ideal energy-proportional system (left) and the resulting efficiency curves (right)

II. CNS

The basic idea behind our new metric can be summarized as follows. The efficiency of a system or service can be expressed by the amount of energy it takes to deliver some unit of work. This efficiency will vary over time depending on the *load level*, that is, the amount of work the system has to do at a given moment. But there will be a maximum efficiency for some system among all different load levels. This is the system-relative optimum: its *sweet spot*. It may be possible that this is still far from what is theoretically possible, but we know at least that the system can reach this efficiency in practice. A system is overall efficient if it operates most of the time close to its sweet spot.

The *consumption near sweet spot* (CNS) is computed as the ratio between the system's average consumption and its optimal consumption per unit of work. This is expressed in the following equation:

$$CNS = \frac{EU_S}{EU_{avg}} \quad (1)$$

Where EU is the energy consumed to deliver a single unit of work, i.e. the system's efficiency during a certain period. EU_{avg} is the average efficiency over an extended, representative period (e.g. a week, a month, or a year) and EU_S is the efficiency at the sweet spot, measured in a small time window when the system performs at its highest efficiency.

Since optimal consumption is always smaller than or equal to average consumption, the CNS can theoretically range between 100% (average consumption coincides with the optimum) and 0% (optimal consumption is negligible with respect to average consumption).

The metric has several strengths: It covers both of the aforementioned optimization strategies. It actually stimulates moving top efficiency and typical load regions towards each other.

Moreover it is an indicator that allows direct comparison between services. Both the type and volume of the unit of work have been fully factored out.

III. COMPARISON WITH FVER

The Data Centre Specialist Group of the British Computer Society has proposed FVER the *Fixed to Variable Energy Ratio* [2], a very similar metric to CNS although there are some important differences.

The reasoning behind FVER is that an optimal energy efficient system should behave energy-proportional. Hence the higher the variation in energy consumption at different load levels (typically between idle and max load) the more efficient it is. Energy consumption that does not vary with load (the fixed part) is suspect for being wasted.

The efficiency is expressed as a ratio between the fixed and the variable part of consumption.

$$FVER = \frac{EU_{fixed}}{EU_{variable}} \quad (2)$$

Where EU_{fixed} is the consumption when idle and $EU_{variable}$ is the difference between this idle baseline and the maximum energy consumption per unit of work. (The original formula is slightly simplified here for the sake of the discussion. Note also, that in contrast to the CNS, FVER is formulated in a way that a smaller value represents a more efficient system)

Like CNS, FVER allows for comparison of different systems by abstracting over the specific workload of a given system. It captures the concept of end-to-end energy-efficiency of a system in terms of useful work delivered to the end user.

But it has a serious shortcoming because it only takes into account the scaling behavior of the system and not the usage profile. The FVER value will not change whether I operate a system most of the time at high utilization or whether it is mostly in low utilization (and probably inefficient) mode. As such it rewards only the first of the optimization strategies mentioned in the introduction and many highly efficient systems will score low.

In contrast, CNS quantifies the extent to which the energy scaling behavior of a system matches the variability in its workload. For systems with very constant, high workload, limited scalability can already result in good CNS values. But systems that have strong fluctuations in workloads a high CNS can only be obtained with flexible scaling behavior.

IV. FUTURE WORK

We already applied the CNS metric in assessments of two industry systems. By collecting a larger number of energy profiles and CNS metrics for services across functional domains and with a wide range of workload profiles and technology footprints, we want not only get more experience in application, but also build up a register of multiple systems that will list the CNS together with two other energy efficiency indicators: the average energy consumption per unit of work and the total energy footprint of a service or system.

The CNS metric was developed together with Dirk Harryvan (Mansystems) and Jeroen Arnoldus and Joost Visser (SIG).

REFERENCES

- [1] L. A. Barroso and U. Hözl, "The case for energy-proportional computing," *IEEE Computer*, vol. 40, no. 12, pp. 33–37, 2007.
- [2] L. Newcombe, Z. Limbuwala, P. Latham, and V. Smith, "Data centre fixed to variable energy ratio metric (dc-fver)," 2012. [Online]. Available: <http://dcsbg.bcs.org/groundbreaking-white-paper-new-dc-metric-available-review>