Measuring Green IT in Society

Zhiwei Xu
Chinese Academy of Sciences, Beijing

Green IT research must produce scientific, accessible results that are understandable to and usable by consumers, companies, volunteers, and policymakers.

Compared to traditional IT systems, green IT more actively engages users and has greater implications for a sustainable society. Consequently, green IT research must produce scientific, accessible results—that is, results that are understandable to and usable by consumers, companies, volunteers, and policymakers.

A central issue of this research is measurement—the qualitative and quantitative assessment of an IT system's greenness. Real cases in China have demonstrated three types of green IT measurement problems:

- assessing greenness before building a system;
- relating performance to power and energy; and
- effectively evaluating green behaviors.

The first problem is primarily a concern for policymakers, while IT system designers must address the second, and the third mainly applies to household users and appliance manufacturers.

All three problems have real-world implications and currently remain unsolved.

ASSESSING GREENNESS

IT spending in China is currently localized according to user population; for example, more than 70 percent of spending on servers and software is concentrated in northern China and coastal areas—where most users are located. However, many fossil fuel and material resources are located in western China, which has fewer than 20 percent of computing users and accounts for less than 15 percent of IT spending in China.

Since the 1980s, China has relied on transporting materials and energy from the western to the more developed coastal areas, supported by infrastructure investments for roads, railways, pipelines, and electric transmission lines worth trillions of yuan, or hundreds of billions of US dollars.

A debate in China has emerged about potentially shifting computing to western China to save energy and reduce CO_2 emissions. Correlated policies could include recommendations to build new datacenters exclusively in this area.

This is a fundamental question for building up China's national cloud computing infrastructure over the next several decades. Because of recent advances in the telecommunications infrastructure and network computing, moving the servers and datacenters to western China could be advantageous. However, before putting this into large-scale practice, policymakers must scientifically validate that shipping information is indeed more efficient than transporting materials and energy. Some anecdotal evidence is available, but there's no scientific data to support this intuitive alternative. There's not even a scientific answer to the question of how much energy a Web search query requires, including energy used by the client device, network, and datacenter.

RELATING PERFORMANCE TO POWER AND ENERGY

Recently, some of China's top Internet services companies approached the Institute of Computing Technology

0018-9162/12/$31.00 © 2012 IEEE Published by the IEEE Computer Society

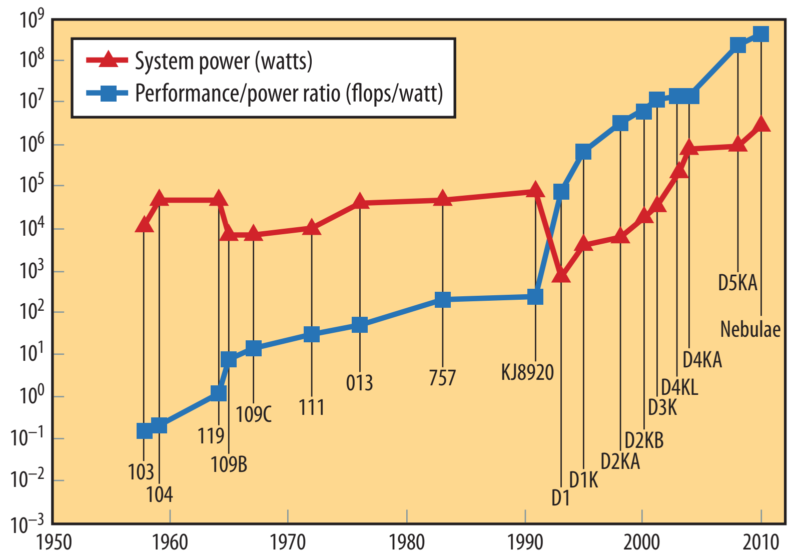

Figure 1. The 50-year trends in the performance-to-power ratio of computers built at the Institute of Computing Technology, Chinese Academy of Sciences.

at the Chinese Academy of Sciences to pose some essential questions. Given that user requests are projected to increase to 50 to 100 billion page views per day in the next 10 years, how much power consumption will a datacenter computer need in 2020? More importantly, what's the best design for a datacenter computer within a realistic power budget?

Researchers need an equation that relates performance to contributing factors, including power and energy. The equation and the involved quantities should be simple and intuitive so they can help researchers choose from millions of design options.

The Green500 list (www.green500.org) identifies the world's top 500 most energy-efficient supercomputers according to their megaflops per watt. Such benchmarking can generate concrete numbers and reveal historical trends in performance-to-power ratios. For example, Figure 1 shows a trend for computers developed at the Institute of Computing Technology over the past 50 years: the performance-to-power ratio increased from 0.15 flops per watt for the Model 103 system in 1958 to 423 million flops per watt for the Dawning Nebulae system in 2010 (ranked fourth in the Green500 list).

Research to maintain this momentum is ongoing, aiming to reach 30 gigaflops per watt for large production systems by 2020. The drawback is that such a benchmarking approach can measure a system only after it's built.

Such measurements should be complemented by equations that can relate performance to power, energy, and the underlying contributing factors. A good analogy is John Hennessey and David Patterson's cycles per instruction (CPI) equation (C. O'Hanlon, "A Conversation with John Hennessy and David Patterson," *ACM Queue*, vol. 4, no. 10, 2007, pp. 14-22):

> execution time/program = instructions/program × clocks/instruction × time/clock.

This equation relates a processor's performance to three contributing factors: the program's instruction count, the clock CPI, and the clock cycle time. Microprocessor system design often uses this formula.

Little's law (J.D.C. Little, "Little's Law as Viewed on Its 50th Anniversary," *Operational Research*, vol. 59, no. 3, 2011, pp. 536-549) provides a starting point for relating a datacenter computer's performance to its contributing factors.

Consider a datacenter computer over the time period $[0, T]$ when it executes the set of threads $\{\tau_1, \tau_2, \ldots, \tau_N\}$. Let L be the number of active threads and λ be the thread arrival rate, averaged over $[0, T]$. Here, W represents the waiting time of a thread, and E represents the energy consumed by a thread, averaged over $\{\tau_1, \tau_2, \ldots, \tau_N\}$. Little's law shows that $\lambda = L/W = L \times (E/W) \times (1/E)$. In other words, a datacenter's throughput performance λ is related to three contributing factors:

> throughput = parallelism × watts per thread × threads per joule.

This formula suggests that in designing an energy-efficient high-throughput datacenter under a power-consumption constraint, it's beneficial to increase parallelism and threads per joule while controlling watts per thread. A potential approach is to design the macroarchitecture to enable parallelism with the microarchitecture design targeting threads per joule.

MEASURING GREEN BEHAVIORS

Designing green IT systems offers an opportunity to induce green behavior in both non-IT systems and society. Because China's urbanization will add hundreds of millions of households to cities in the coming decades, measuring the electrical consumption of home appliances is an essential research area.

Can we acquire massive and accurate field behavior data for every lamp, refrigerator, air conditioner, and other appliance from 100 million households by implementing one sensor per home? This could lead to at least two benefits. First, obtaining such data from billions of appliances could help induce positive behavior from users. Household users could *grid search*—that is, automatically rank and search good electricity consumption behaviors on billions of appliances, similar to how we do Web searches of billions of webpages today. Second, such data could give the appliance industry the capability of large-scale beta testing, helping manufacturers to evaluate the performance of their products and improve their R&D and leading to future design improvements.

The Institute of Computing Technology recently launched wuTV, a long-term research project. Here, wu means the physical world. In contrast to Internet TVs or cloud TVs that connect to the Internet, wuTV connects to the physical world (such as household appliances) and human society (such as users of household appliance). Having wuTV in a home not only measures field behavior data of appliances and users but also serves as a gathering point to enable sharing and using such data through Internet cloud services.

WuTV's aim is to have a single device computationally differentiate and identify each individual appliance's accurate behavior. It utilizes the physical laws of electricity as well as resources from the physical world, society, and cyberspace.

Preliminary research shows that such a ternary computing approach is promising in that researchers can use it to design appliances with continuously variable electric currents (J. Liu, L. Nie, and Z. Xu, "The Input-Sensing Problem in Ternary Computing and Its Application in Household Energy-Saving," *Proc. IEEE/ACM Int'l Conf. Green Computing and Communications*, IEEE CS, 2011, pp. 131-138).

In addition to concerns related to engineering and technology, green IT research also must address relevant societal challenges. The Chinese Academy of Sciences advocates a ternary computing concept (Z. Xu and G. Li, "Computing for the Masses." *Comm. ACM*, vol. 54, no. 10, 2011, pp. 129-137) that focuses on solving these problems by utilizing resources from the physical word and society as well as cyberspace. ■

Zhiwei Xu *is a professor and CTO of the Institute of Computing Technology at the Chinese Academy of Sciences, Beijing. Contact him at zxu@ict.ac.cn.*