Seeing Green: Autonomic Software for Energy-Efficient Computing

Abstract

This paper presents an innovative *software* solution that transparently self-adapts to produce as much as 40% energy savings while maintaining performance on commodity computing systems. What makes this solution so unique is that it saves this energy *while* applications are actively running on the computing system. Current state-of-the-art solutions only save energy when the computing system is idle. This software solution, which we call EnergyFit™, can run as a user-level library or a user-level service (or daemon) and does *not* require any modification to applications or the computing system in order to realize energy-efficient performance, e.g., nearly 25% energy savings on SPECjbb2005 (Java Business Benchmarks).

Introduction

Today, being green in computing is all the rage. Everywhere you turn, information technology (IT) companies tout how green their solutions are. Why is being green so important? The reasons include, but are not limited to, (1) lower electrical bills, (2) enhanced IT reliability, (3) improved battery life in mobile IT devices, (4) delayed investment in infrastructure expansion for data centers, and (5) positive PR for being green.

To address the above, we present a novel software approach to green computing called EnergyFit™. *EnergyFit™ runs on commodity computing systems to save energy and maintain high performance while programs are actively running*, as illustrated conceptually in Figure 1. By way of analogy, our software-only approach to green computing is akin to automatically transforming today's existing computing systems, which oftentimes have the ecological footprint of a Hummer, into more efficient computing systems with an ecological footprint of a Toyota Prius while delivering the performance of a Maserati GranTurismo S.

Background

[FADE IN.] The year is 2002. Media from around the world descend upon the small town of Los Alamos, NM, where a population of 18,000 people and Los Alamos National Laboratory consume a whopping 350,000 MWh of energy per year.[1] They are there to witness the antithesis of energy-consuming ways: *Green Destiny*, a 240-node supercomputer in 5 square feet and consuming a mere 3.2 kW of power, i.e., equivalent to two hairdryers. This event resulted in a tidal wave of media coverage, including *The New York Times, CNN, BBC News,* and *PCWorld*, because of how disruptive this technology was to conventional wisdom. *[FADE OUT.]* Back then, however, the notion of being green in computing, particularly supercomputing, was heresy; all anyone cared about was speed … the need for speed. Furthermore, the IT community was uncomfortable with the notion of adopting unconventional and only "pseudo-commodity" technology. As a result, we also sought to complement the

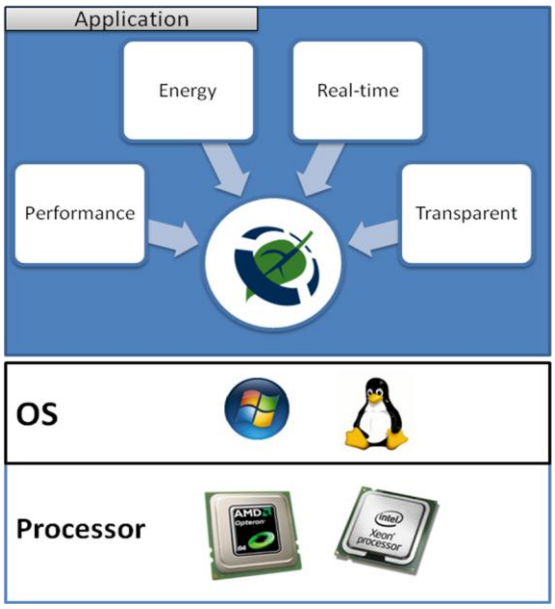

Figure 1. Cross Platform Solution

[1] Source: Los Alamos Sustainable Energy Network, 2005.

above architectural approach by inventing ***EnergyFit™***, software that automatically improves "the gas mileage of your high-octane computer."

EnergyFit: Earning Green Dollars with Green Energy in IT

Because of the economies of scale in IT, exotic architectural solutions for green computing, such as Green Destiny, arguably do *not* provide the best return on investment due to their higher price and performance-price ratio. However, a software-based solution that runs on commodity computing systems can deliver a significantly higher return on investment due to lower electrical power draw, reduced cooling requirements, enhanced reliability, and deferred infrastructure expansion in the case of a datacenter. Sounds awesome. But how? It turns out that this is a particularly difficult problem to solve.

First, we observe that the power consumption of processor is proportional to its clock frequency and the square of its voltage supply.[2] Thus, to reduce power consumption, we can reduce the clock frequency and voltage. However, this in turn reduces the performance of the processor because performance is also proportional to the clock frequency. As a result, not only can this result in a loss of performance, it can also consume *more* energy because energy is the integral of power over time, as shown in Figure 2, where reducing the power caused the program to take longer to execute and consume *more* energy, i.e., the boxed area represented with diagonal lines.

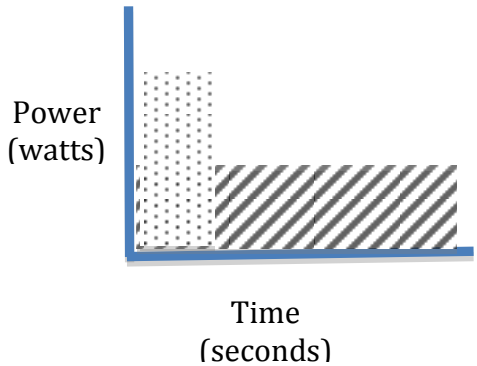

Figure 2. Power Consumption vs. Execution Time

Second, the time to scale a frequency (and its corresponding voltage) can run as high as 10 *milliseconds*, which translates into 20,000,000 clock cycles on a 2.0-GHz processor. Thus, scaling the frequency and voltage is very expensive operation.

However, despite the above challenges, both AMD and Intel find "dynamic voltage and frequency scaling" (DVFS) important enough to directly support such a mechanism via PowerNow! and Speedstep, respectively. Intel does so because it can use its DVFS mechanism, i.e., Speedstep, to create a software policy, i.e., `cpuspeed`, that scales frequency and voltage when a computing system is *idle*, thus leaving plenty of time to scale the frequency and voltage and reducing power consumption.

Alas, whenever a computing system is *actively* running applications, such as the SPEC Java Business Benchmarks, current policies for power management, such as `cpuspeed` and `ondemand`, only reduce power consumption by 3% and 4%, respectively, as shown in Table 1. *In stark contrast, EnergyFit™ reduces power consumption by a whopping 25%.*

In addition, EnergyFit™ ensures that the performance of a computing system is maintained at a desired level *while* saving power and energy consumption. *The end result is a 23.8% improvement in the performance-per-watt metric*, or more specifically, a 23.8% improvement in "business operations per second (bops) per watt" or bops/watt, as shown in Table 2.

Power Management	Watts	% Power Reduction
None	264	0%
cpuspeed	257	3%
ondemand	253	4%
EnergyFit™	196	25%

Table 1. Power Consumption When Running the Java Business Benchmarks (SPECjbb2005)

Of more recent note is a high-profile pilot project with a major insurance company. In deploying EnergyFit™

[2] Due to the electrical properties of a processor, a higher clock frequency requires a higher voltage to drive that frequency. Thus, both the frequency and voltage must be changed in tandem, i.e., both go up or both go down.

in the field, EnergyFit™ reduced the energy consumption of an Oracle/RedHat-based datacenter platform by 14% with no perceivable impact on performance. This is of particular importance to the insurance company because (1) they are already energy-constrainted and would need to spend $12M to add a 1-MW power substation, and (2) they are expecting a 40% increase in energy cost in 2010 when rate caps are removed.

EnergyFit™: How Does It Work?

The key to EnergyFit™ is in its clever design of an *accurate performance model in real time* that characterizes the profile of the application that is running, the tradeoff between power consumption and performance when scaling frequency and voltage, and the overhead associated with exercising DVFS. This design enables EnergyFit™ to substantially outperform current commercial solutions as well as academic research prototypes based on the traditional proportional-integral-derivative (PID) controller.

Power Management	Bops/Watt
None	100.00%
EnergyFit	123.70%
ondemand	104.37%
powersave	95.20%
cpuspeed	102.56%

Table 2. Normalized Performance-Per-Watt When Running Java Business Benchmarks

In practice, EnergyFit™ runs as a user-level library or a user-level service (i.e., daemon), and thus, it does *not* require any modification to the operating system and does *not* require any modification to applications. Thus, EnergyFit™ is a "drop-in" solution that delivers energy efficiency transparently to end users and systems.

Conclusions

With EnergyFit™ already demonstrating its efficacy in the field, it provides substantial validation for our approach. For the SPEC Java Business Benchmarks (SPECjbb2005), EnergyFit™ delivers a 25% savings in energy consumption while also improving the performance-per-watt metric by nearly 24%. We are already working on a next-generation version of EnergyFit™ that will be able to make more substantial energy savings, particularly in the presence of virtual machines, while still delivering high performance. Preliminary research that applies our techniques to graphics processing units (GPUs) promises to improve the performance-per-watt by as much as *two orders of magnitude.*

Takeaways

- EnergyFit™, a software-based solution that runs as a user-level service, improves the *performance-per-watt* metric by 24% when running the SPEC Java Business Benchmarks.
- Reduced power consumption translates to lower temperatures, which in turn, translates to better reliability.

About the Author

Dr. Wu Feng is the Chief Technology Officer at EnergyWare. He is world-renown for his expertise in high-performance computing and networking, green computing, and bioinformatics and is the author of over 150 publications in these fields. He is also a faculty member with the Department of Computer Science and Department of Electrical & Computer Engineering at Virginia Tech.

His invention of an energy-efficient supercomputer dubbed *Green Destiny* garnered international acclaim with media coverage in hundreds of news outlets, including *The New York Times, CNN,* and *BBC News.* This invention also led to an R&D 100 Award[3] in 2003 due to it having a footprint of only five square feet and consuming only 3.2 kilowatts of power (i.e., two hairdryers). In 2004, he netted *two* R&D 100 Awards for his work on efficient networking and computing with *10-Gigabit Ethernet Adapter: Speed Really Changes Everything* and *mpiBLAST: A High-Speed Software Catalyst for Genetic Research*, respectively. He was also named to *HPCwire's Top People to Watch List* in 2004. More recently, *EcoDaemon*, his automated systems software that simultaneously saves energy and delivers high performance while programs are actively running, won the Southeastern Universities Research Association (SURA) Intellectual Property-to-Market (IP2M) Award amongst more than 60 competing institutions.

Previous professional stints include the University of Illinois at Urbana-Champaign, Purdue University, The Ohio State University, Los Alamos National Laboratory, NASA Ames Research Center, Orion Multisystems, Vosaic, and IBM T.J. Watson Research Center.

About the Company

ENERGYWARE

EnergyWare provides efficient computing solutions that reduce cost and increase value to enterprise IT infrastructure. Their green solutions deliver proven supercomputing technology to commercial customers who demand high performance and efficiency. In addition, EnergyWare products reduce carbon emissions, increase reliability, and significantly lower operational cost while maintaining the performance of enterprise data centers.

[3] An R&D 100 Award provides recognition for an invention that is one of the most innovative with respect to research & development. It is often referred to as the "Oscars for Invention."