Exclusive license pending

LOS ALAMOS NATIONAL LABORATORY

ENERGYFIT

COOL FAST RELIABLE COMPUTING

Reduces a processor's energy consumption by up to 70%

Diminishes greenhouse gas emissions

Improves the reliability of computing systems

Los Alamos
NATIONAL LABORATORY
EST. 1943

LA-UR-06-1360

EnergyFit
Cool Fast Reliable Computing

Technology Transfer Contacts:
Laura Barber
tmt-2@lanl.gov

Technical Contact:
Chung-hsing Hsu

ABOUT THE COVER

Developer Wu Feng illustrates the cool, fast approach to reliable computing made possible by the EnergyFit software, which reduces a processor's energy consumption, with minimal impact on its performance. EnergyFit can be applied to any computing system—laptops, embedded systems, desktop computers, search-engine farms, Internet service provider servers, and so on—as represented by the systems shown.

Los Alamos NATIONAL LABORATORY EST.1943

Executive Summary

EnergyFit
Cool Fast Reliable Computing

Features

EnergyFit is self-aware, self-adapting software that *automatically* reduces the energy consumption of a processor in a computing system by as much as 70%, with minimal impact on performance. In addition to reducing energy bills, EnergyFit ultimately lowers temperatures in computer systems, which in turn leads to better reliability. That is, as per Arrhenius' equation, for every 10°C decrease in temperature, the long-term reliability of the electronics in the system doubles.

Applications

- Supercomputing centers
- Financial data centers
- Laptops
- Desktop computers
- Data-center servers
- Internet service providers
- Search-engine farms
- Sensor networks
- Embedded processing systems (e.g., cell phones, automobiles, airline reservation systems, multimedia portable players, network routers on the Internet, video-gaming consoles such as Sony PlayStation Portable, or anything that has an embedded processor in it)

Benefits

- Delivers substantial energy savings and, hence, lowers energy bills
- Enhances reliability by reducing energy consumption and temperature
- Allows the footprint of a data center to be shrunk because of the reduced thermal envelope of processors
- Lowers the total cost of ownership
- Provides an environmentally friendly solution that can generate sustainable revenue

Summary

At a high level, our invention is like an autopilot program for a car, where the goal is to minimize fuel usage while getting to the destination in time. Depending on traffic and road conditions, the autopilot adjusts the car speed for the best fuel efficiency and timeliness. Under certain traffic and road conditions, speeding up the car will *not* ultimately help in getting to the destination faster. Furthermore, the repeated acceleration and deceleration wastes fuel and puts unnecessary wear and tear on the car. The best solution is to try to set the car to a speed that is appropriate (to get there on time) yet slowly varying (for fuel efficiency).

Because traffic and road conditions vary over time, the autopilot program has to predict upcoming conditions and their impact on getting to the destination (based on past observations). EnergyFit's novelty is in its ability to predict the traffic and road conditions and their impact at a sufficiently precise level but to do so in a computing system.[1] Previous research in this area, virtually none of which has been deployed, has specifically targeted embedded systems rather than computing systems in general. Furthermore, this previous research frequently makes mispredictions, thus resulting in late arrival—sometimes very late.

EnergyFit is based on a fundamental theorem that we developed to characterize the "best driving pattern in real time" for a processor. EnergyFit not only predicts future traffic and road conditions well and gets to the destination on time but also minimizes the total fuel usage for the trip—but of course, in the context of a processor rather than a car. More concretely, many of today's general-purpose processors from vendors such as AMD, IBM, and Intel provide a mechanism called dynamic voltage and frequency scaling (DVFS) that allows the processor's voltage supply and speed to be adjusted on the fly, but only with respect to a certain set of frequency-power tradeoffs called *settings*. The faster a processor runs, the higher the required supply voltage and the more power it consumes; the slower it runs, the lower the required supply voltage and the less power it consumes.[2] Our autopilot program—EnergyFit—predicts the traffic and road conditions in a computing system to intelligently decide *when* to change a setting and *what* to change the setting to in order to minimize the overall energy consumption while meeting a deadline. Since the performance impact on each program running at the same processor frequency is different and is not known a priori, EnergyFit *learns it on the fly*. EnergyFit then repeatedly applies the fundamental theorem to get the best schedule of settings (given the performance model it has learned and constructed thus far).[3]

When *actively running* a computer program (e.g., Google), EnergyFit can reduce energy consumption by as much as 70%, with a less than 5% impact on performance (e.g., a Google search takes 1.05 second instead of 1 second). When tuned for better *performance*, EnergyFit can reduce energy consumption by 54%, with a less than 2% impact on performance. Tuned for better *energy efficiency*, it can reduce energy consumption by nearly 80%, with a less than 15% impact on performance. *On average*, during program execution, the default settings in EnergyFit save about 20% of processor energy usage (with only a 3% impact on performance) across a set of scientific programs from physics, chemistry, and other compute-intensive disciplines.

Even more notable are our internal projections of the benefit of EnergyFit in data centers that must operate 24/7 but are not computationally active 24/7. In such an environment, EnergyFit could reduce a processor's energy consumption (without performance degradation) by as much as 90% when data-center resources are computationally idle.

Technical details about EnergyFit, can be found in the Appendix.

---

[1] If a processor's clock frequency is viewed as the rotational speed of a car's wheels, current computing systems *always* have their wheels spinning, irrespective of whether the "car" is moving (i.e., spinning its wheels while staying in place). Clearly, this is energy inefficient and wasteful of energy.

[2] More precisely, the faster that a processor runs (i.e., higher frequency), the higher the supply voltage that is needed to drive the higher frequency, and hence, the higher the power consumption because the power consumption of a processor is proportional to its frequency and the square of its supply voltage.

[3] The two biggest problems with current solutions are (1) they do not formulate a general performance model to predict future traffic and road conditions for the processor, and (2) they assume that changing processor voltage and frequency settings is nearly instantaneous. The latter problem results in ill-advised changes in settings, e.g., oscillatory changes in settings that end up consuming more energy because transitioning between settings is not negligible, ~20,000,000 processor cycles. This problem is particularly pronounced in Transmeta LongRun (http://www.transmeta.com/crusoe/longrun.html).

Competition

Manufacturer	Brand Name	Model Number
Open-Source Community	SpeedswitchXP	1.4
Transmeta	LongRun	2.0

Comparison matrix

All of the technologies in the matrix leverage a nearly ubiquitous mechanism called *dynamic voltage and frequency scaling (DVFS)*, which is branded as PowerNow! by AMD, Enhanced SpeedStep by Intel, and PowerTune by IBM. In each category, the “winner” is highlighted.

Key:
- CPU = central processing unit (or more generally, processor)
- DVFS = dynamic voltage and frequency scaling
- HPC = high-performance computing
- OS = operating system
- x86 = Intel-compatible architecture

Category	**EnergyFit 1.0**	**LongRun 2.0**	**SpeedswitchXP 1.4**
Platform Dependence	Any processor with DVFS	Transmeta CPU with DVFS	Any x86 CPU with DVFS
OS Dependence	Linux (32- or 64-bit)	Windows XP (32-bit only)	Windows XP (32-bit only)
Energy Adaptation	Can occur at any time	Generally only when CPU is idle	Only when CPU is idle
Performance Loss on HPC Codes	Tunable*	~2%	~2%
Energy Savings on HPC Codes	20% average (70% maximum)**	–5% to 5%	1%
Applicability	Any computing system	Laptops and embedded systems	Laptops only
Cost	Proprietary ($)	Proprietary ($$)	Free

* In some cases, EnergyFit actually exhibits both performance improvement and energy savings simultaneously.

** These numbers result from EnergyFit running in its default configuration. If EnergyFit is tuned for better energy efficiency at the expense of performance, it can reduce energy consumption by as much as 80% when a program is actively running on the processor.

Advantages

EnergyFit is more energy efficient.

The web sites for Transmeta LongRun (http://www.transmeta.com/crusoe/longrun.html) and SpeedswitchXP (http://www.diefer.de/speedswitchxp/) note that their technologies *depend on processor idle time (or low processor utilization) to save energy*. For example, the Transmeta web site notes the following:

The LongRun algorithm in the Code Morphing Software monitors the Crusoe processor and dynamically switches between the following points as runtime conditions change.

- If no idle time is detected during a workload, the frequency/voltage point is increased (if possible). Translation: Higher speed and higher power consumption.
- If idle time is detected, LongRun may decide that energy is being wasted and will decrease the frequency/voltage level. Translation: Lower speed and lower power consumption.

Similarly, SpeedswitchXP "switches between the lower and maximum speed according to current CPU utilization."

In stark contrast, EnergyFit relies on a fundamental theorem that enables it to predict the traffic and road conditions in a computing system in order to intelligently decide *when* to change a setting and *what* to change the setting to in order to minimize the overall power and energy consumption while meeting a deadline. As a result, when executing programs, *EnergyFit saves 10 to 15 times more energy than LongRun and 10 to 70 times more energy than SpeedswitchXP.* (As a side note, our research has shown that LongRun's approach to DVFS is suboptimal and can result in a processor that is actually less energy efficient than if the processor were to run without LongRun. For instance, note that in the comparison matrix, the energy savings are sometimes negative, i.e., *more* energy is used when using LongRun on HPC codes.

EnergyFit is more versatile.

SpeedswitchXP applies only to notebook computers (i.e., laptops) running on an Intel-compatible processor. LongRun applies only to laptops and embedded systems running on a Transmeta processor. In contrast, EnergyFit can be applied to any computing system in general—laptops, embedded systems, desktop computers, search-engine farms, Internet service provider (ISP) servers, and so on.

EnergyFit is more platform independent.

EnergyFit, LongRun, and SpeedswitchXP all rely on DVFS, which has become a nearly ubiquitous mechanism in commodity processors. However, while EnergyFit can run on any processor with DVFS, LongRun runs only on niche Transmeta processors that use the Windows operating system, and SpeedswitchXP runs only on Intel-compatible processors that use the Windows XP operating system specifically.

EnergyFit is more tunable and controllable.

LongRun does not provide any override mechanisms to tune or control the tradeoff between performance and energy consumption. As a result, there are times when LongRun actually worsens performance and consumes more energy.

SpeedswitchXP provides four coarse settings to control the tradeoff between performance and energy consumption:
(1) maximum performance, which keeps the CPU at maximum speed; (2) battery optimized, which keeps the CPU at a predetermined lower speed; (3) maximum battery, which keeps the CPU at a predetermined lower speed and allows further throttling, depending on remaining battery power; and (4) dynamic switching, which switches between the lower and maximum speed, according to current CPU utilization.

Using a fundamental theorem developed by its inventors, EnergyFit provides tight control over performance (either by default or override) and allows both voltage and frequency to vary over time. In contrast to SpeedswitchXP, which tends only to toggle between a lower frequency and the maximum frequency, EnergyFit will change to whatever voltage and frequency settings are appropriate to maintain control over performance while reducing energy consumption as much as possible. (Recall that a processor's power consumption, i.e., rate of energy consumption, is proportional to its frequency and the square of its supply voltage.)

Applications

The principal application of EnergyFit is energy-efficient computing. Using less energy means that power plants need to generate less energy, which in turn, reduces greenhouse gas emissions and improves the quality of our air. Being energy efficient in computing also has several additional beneficial consequences. First, as noted by the government-backed program EnergyStar, energy efficiency helps the economy by saving consumers and businesses billions of dollars in energy costs.[4] Second, reduced energy consumption ultimately leads to lower temperatures in computer systems, which in turn leads to better reliability. That is, as per Arrhenius' equation, the long-term reliability of the electronics in the system doubles for every 10°C decrease in temperature.

[4] http://www.energystar.gov/index.cfm?c=about.ab_index: "In 2004 alone, Americans saved enough energy to power 24 million homes and avoid greenhouse gas emissions equivalent to those from 20 million cars—all while saving $10 billion."

Today, EnergyFit *transparently* and *automatically* delivers energy efficiency to laptops, desktop computers, data-center servers, and supercomputers. For laptops, EnergyFit reduces energy consumption primarily to extend battery life and to help the laptop run cooler and last longer.

Desktop computers now number nearly one billion worldwide, and EnergyFit helps them run cooler and more reliably, which potentially translates into gigantic energy savings. Why? Because most desktop computers are left turned on 24 hours a day (16 hours of which are largely idle) and because there are nearly one billion of them in the world,[5] EnergyFit would reduce processor power consumption by 20% on average during the day and by 80% on average during the evening and overnight, thus reducing the aggregate power consumption of desktop processors by 60% overall.

Assuming a *conservative* power consumption of 60 watts per processor, the aggregate power consumption of *just* the world's desktop processors is 60 billion watts, or 60 gigawatts (GW). Using EnergyFit to reduce power consumption by 60% overall would save 36 GW of power (or approximately the power output of 16 Hoover Dams); at an average cost of $0.10 per kilowatt-hour, this translates into a savings of $3,600,000 per hour or nearly *$32 billion per year*. Furthermore, by automatically saving energy, EnergyFit would also substantially reduce greenhouse gas emissions, which are arguably fueling global climate warming, thus providing an easy way to reduce air pollution.

For large-scale data centers like those found in financial districts, scientific institutions, search-engine farms, ISPs, and supercomputing centers, *EnergyFit offers all the aforementioned benefits—better energy efficiency, improved reliability, lower operational costs and total cost of ownership—and from an environmentally friendly viewpoint, reduced greenhouse gas emissions and reduced air pollution.* In addition, EnergyFit offers further reductions in operational costs due to lower cooling requirements and substantially improved space efficiency. Relative to the former, for every watt of power that such a large-scale data center consumes, nearly an additional watt of power is required to *cool* it. EnergyFit not only decreases the amount of electricity needed to power a data center, but also eliminates nearly an equal amount of power for cooling it. As to space efficiency, data centers that run cooler can be more densely packed, resulting in greater compute power per square foot.

Other applications

In general, EnergyFit can be run on any processor, not just general-purpose processors .

[5] Source: Avi Mendelson, Intel.

EnergyFit's impact could be much more substantial in the embedded systems market, which delivers orders of magnitude more processors than does the market for traditional general-purpose processors. For example, many processors embedded in the following systems could greatly benefit from EnergyFit: cell phones, PDAs, iPods, iNanos, MP3 players, Space Shuttle control systems, automobiles, GPS systems, cash registers, and so on. For those processors, EnergyFit would extend battery life and enhance reliability.

Summary

As noted in a recent LinuxHPC headline, "EnergyFit Raises the Gas Mileage of Your High-Octane Computer." That is, EnergyFit intelligently delivers nearly maximum performance, with minimal energy usage for any given computing system. Using less energy means that power plants need to generate less energy, which in turn reduces greenhouse gas emissions and prevents air pollution.

EnergyFit is self-aware, self-adapting software that *transparently* and *automatically* reduces a processor's power and energy consumption by as much as 70%, with minimal impact on performance—benefits that serve the environmentally friendly ideals of reducing greenhouse gas emissions and preventing air pollution as well as the following pragmatic customer needs:

1. Lower temperatures and, hence, better reliability (For every 10°C decrease in temperature, the long-term reliability of the electronics in the system doubles.)
2. Lower operational costs due to lower electrical usage, for both powering and cooling a computing system, and due to better reliability and longevity of computer components
3. Improved space efficiency as processors' lower temperatures and smaller thermal envelopes allow them to be more densely packed. (Space is money. In Silicon Valley, space costs $200/square foot.)

To quantitatively understand the impact that EnergyFit could have, note that Intel reported that the total power consumption of *only* the world's CPUs (i.e., processors) was 160 megawatts in 1992 and 9000 megawatts in 2001, as shown in the figure (p. 13). Today, the power consumption of all the world's processors—laptops, desktop computers, and high-end servers, for example—is conservatively 84 gigawatts, as noted in the table (p. 13). Given that most of the world's 1.3 billion processors[6] are turned on 24 hours a day (but often idle for at least 16 of those hours), EnergyFit would reduce processor power and energy consumption an average of 20% during the day and an average of 80% in the evening and overnight, thus reducing the aggregate power consumption of all the processors

[6] Source: Avi Mendelson, Intel.

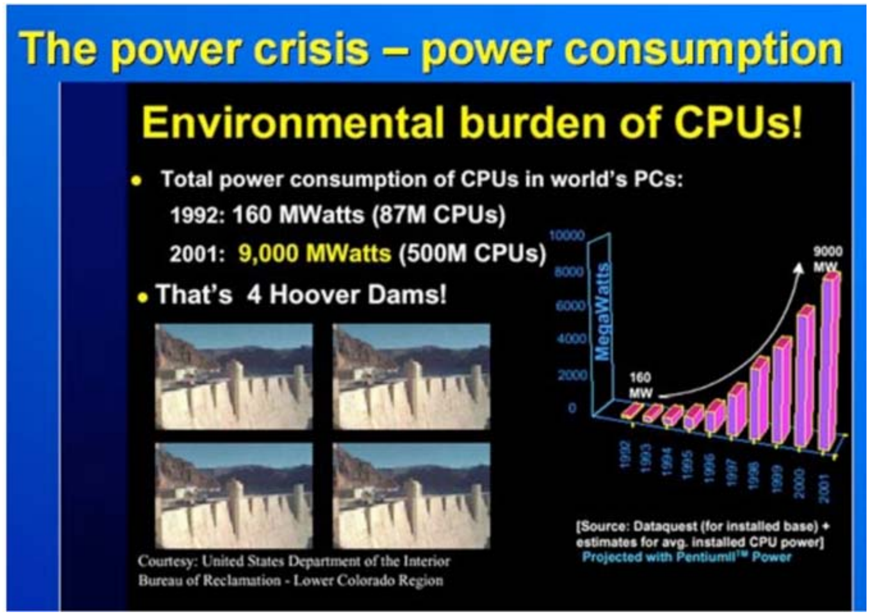

Power Consumption of the World's Processors through 2006

Year	Power (in megawatts)	# CPUs (in millions)
1992	160	87
1993	250	106
1994	392	128
1995	613	156
1996	959	189
1997	1,501	230
1998	2,349	279
1999	3,676	339
2000	5,752	412
2001	9,000	500
2002	14,083	607
2003	22,038	737
2004	34,485	896
2005	53,962	1,088
2006	84,439	1,321

in the world by 60% overall. Therefore, *EnergyFit could potentially reduce the world's aggregate power consumption on processors by 50 gigawatts, the equivalent of enough electricity to power 50,000,000 average U.S. homes or the power of 22 Hoover Dams!* Could EnergyFit have saved California from the rolling blackouts of 2000, 2001, and 2003, when the power deficits each summer reached "only" as high as 7 gigawatts?[7]

Finally, we close with a quote from a letter of reference from AES Corporation. "EnergyFit represents not only an enabling technology to improve efficiency and the total cost of ownership of high processing computer systems, but it also represents an environmentally conscientious technology that can be deployed profitably to address the issue of global warming."

7 http://www.tomorrowsworld.org/cgi-bin/tw/tw-mag.cgi?category=Magazine13&item=1104084018