

Chipping Away at Greenhouse Gases

Power-saving processor algorithms have the potential to create significant energy and cost savings.

THE INFORMATION TECHNOLOGY industry is in the vanguard of “going green.” Projects such as a \$100 million hydro-powered high-performance data center planned for Holyoke, MA, and green corporate entities such as Google Energy, the search giant’s new electrical power subsidiary, are high-profile examples of IT’s big moves into reducing the greenhouse gases caused by computers.

However, the true benefits of such projects are likely to be limited; most users in areas supplied by coal, oil, or natural gas-fired power plants would likely find it difficult to change to a fully sustainable supply source.

These market dynamics have not been lost on government research directors. Agencies such as the U.S. National Science Foundation (NSF) have begun encouraging just the sort of research into component-level power management that might bring significant energy savings and reduced climatic impact to end users everywhere without sacrificing computational performance.

In fact, the NSF has held two workshops in the newly emphasized science of power management, one in

The screenshot shows the Granola application interface. On the left, a list of statistics is displayed with corresponding icons: a lightbulb for kWh savings, a dollar sign for USD savings, a leaf for CO2 savings, and a percentage sign for CPU energy savings. At the bottom left, a graphic shows a row of trees representing 151,738 trees offset. On the right, there is a large yellow padlock icon with a blue 'g' on it. Below the padlock is a login form with fields for Username and Password, a 'Sign In' button, and links for 'Forgot your password?' and 'Need an account?'.

granola

- You'll save 234.6 kWh yearly**
Enough to power 31 electric furnaces, a space heater, and 2 refrigerators for an hour
- You'll save 28.15 USD yearly**
Enough for a monkey wrench to throw in the gears, 3 shirts from the thrift store, and a political bumper sticker
- You'll save 319.0 lbs CO2 yearly**
As much as a 500-mile flight, a tree, and 14 miles in a compact car
- You've saved 45.1% CPU energy**
...and you didn't even notice!

151,738 trees
You and the Granola community will offset 151,738 trees worth of CO2 this year.

By logging into your Granola account, you can track your savings across systems through the Granola website. Click the lock above to access your savings information.

Username: *****@acm.org
Password: *****
Logged in! **Sign In**
[Forgot your password?](#)
[Need an account?](#)

An intelligent power-management application, Granola uses predictive algorithms to dynamically manage frequency and voltage scaling in the chips of consumer PCs.

2009 and one in 2010. Krishna Kant, a program director in the Computer Systems Research (CSR) cluster at the NSF, says the power management project is part of the NSF’s larger Science, Engineering, and Education for Sustainability (SEES) investment area.

“There are some fundamental ques-

tions that haven’t been answered, and NSF funding might help answer them,” Kant says. “These have been lingering for quite some time. For instance, when you look at the question of how much energy or power you really need to get some computation done, there has been some research, but it tends

to be at a very, very abstract level to the extent it's not very useful."

Thermal Head Start

However abstract the state of some of the research into power management might be, basic computer science has given the IT industry a head start over other industries in addressing power issues. Whereas an auto manufacturer could continue to make gas-guzzling vehicles as long as a market supported such a strategy, two factors in particular have focused microprocessor designers' efforts on the imperatives of power efficiency.

One of the factors is the thermal limitations of microprocessors as each succeeding generation grew doubly powerful per unit size. The other is the proliferation of laptops and mobile computing devices, which demand advanced power management features to extend battery life. Kirk Cameron, associate professor of computer science at Virginia Polytechnic Institute, says this shift in product emphasis has given engineers working on power management theories more tools with which to work on the central processing unit (CPU); these chips are also installed on desktop machines and servers as chip manufacturers design one family for numerous platforms, based on overall market demand. Examples of these tools include application programming interfaces such as Intel's SpeedStep and AMD's PowerNow, which allow third-party software to dynamically raise or lower the frequency of cycles and the voltage surging through the

processor, depending on the computational load at any given time.

However, the default power management schemes supported by current operating systems, which allow users to specify either a high-performance or battery-maximizing mode on laptops, for instance, have numerous handicaps, including their static nature. The fact they need to be manually configured hampers their popularity.

Some power-management products, incubated by university researchers, are already available to dynamically manage power within a computer's CPU. Cameron is also the CEO of Miserware, a startup funded in part by an NSF Small Business Innovation Research Grant. Miserware produces intelligent power-management applications—called Granola for consumer PCs and Miserware ES for servers—that use predictive algorithms to dynamically manage frequency and voltage scaling. Company benchmarks claim that users can reduce power usage by 2%–18%, depending on the application in use; best savings are generated by scaling down power during low-intensity activities.

Granola was launched on Earth Day last year, and has 100,000 downloads. Cameron says the dynamic voltage and frequency scaling (DVFS) technology is very stable, available on most systems, and "kind of the low-hanging fruit" in power management.

Susanne Albers, professor of computer science at Humboldt University of Berlin, believes speed scaling will be a standard approach to power manage-

ment for some time. "I am confident that dynamic speed scaling is an approach with a long-term perspective," she says. "In standard office environments the technique is maybe not so important. However, data and computing centers, having high energy consumption, can greatly benefit from it."

Multicore Architectures

Ironically, although the DVFS technology is currently the most ubiquitous power management solution for processors, Cameron and other researchers say new fundamentals of computing architecture will mandate wholly different solutions sooner rather than later.

The onset of mass production of multicore processors, for example, is mandating that researchers begin practically anew in exploring speed scaling approaches.

"Generally speaking, there exists a good understanding of speed scaling in single processor systems, but there are still many challenging open questions in the area of multicore architectures," Albers notes.

"The new technologies bring new algorithmic issues," says Kirk Pruhs, professor of computer science at the University of Pittsburgh, and an organizer of both NSF workshops. For instance, if a heterogeneous-cored processor is programmed correctly, the utility of using frequency and voltage scaling at all might be moot—applications needing lower power can be sent to a slower core.

However, Pruhs says programming these will be "much more algorithmi-

ACM Awards News

2011 ACM Fellows Nominations

The ACM Fellow program was established by the ACM Council in June 1993 to recognize outstanding ACM members for technical, professional, and leadership contributions that advance the arts, sciences, and practices of information processing; promote the free interchange of ideas and information in the field; develop and maintain the integrity and competence of individuals in the field;

and advance the objectives of ACM.

Each candidate is evaluated as a whole individual and is expected to bring honor to the ACM. A candidate's accomplishments are expected to place him or her among the top 1% of ACM members. In general, two categories of accomplishments are considered: achievements related to information technology and outstanding

service to ACM or the larger computing community. A person selected as an ACM Fellow should be a role model and an inspiration to other members.

Nominations and endorsements must be submitted online no later than Sept. 1, 2011. For Fellows Guidelines, go to http://awards.acm.org/html/fellow_nom_guide.cfm/.

Nomination information organized by a principal

nominator should include excerpts from the candidate's current curriculum vitae, listing selected publications, patents, technical achievements, honors, and other awards; a description of the work of the nominee, drawing attention to the contributions which merit designation as Fellow; and supporting endorsements from five ACM members. For the list of 2010's ACM Fellows, see p. 25.

cally difficult for the operating system to manage, and the same thing happens in memories. The fact everything is changing means you have to go back and reexamine all the algorithmic issues that arise.”

In the case of power management in a parallel environment, Cameron says his research has shown that one cannot take the principles of Amdahl's Law for parallelization—which states that any parallelized program can only speed up at the percentage of a given task within that program not run serially—and get a correct assumption about power savings by simply taking into account the processors running a given application.

“In Amdahl's Law, you have one thing that changes, the number of processors,” Cameron says. “In our generalization, we ask what if you have two observable changes? You might think you could apply Amdahl's Law in two dimensions, but there are interactive effects between the two. In isolation, you could measure both of those using Amdahl's Law, but it turns out there is a third term, of the combined effects working in conjunction, and that gets missed if you apply them one at a time.”

Doing Nothing Well

In the long term, power management may borrow from sensor networks and embedded systems, which have extensively dealt with power constraints. Both David Culler, professor of computer science at the University of California, Berkeley, and Bernard Meyerson, vice president of innovation at IBM, cite the disproportionately large power demands of processors doing little or no work as an area where great savings may be realized.

Culler says processor design might take a lesson from network sensor design in principle. Measuring performance during active processing “talk” time is misplaced, he says. Instead, efficiency must be introduced while awaiting instruction—“talk is cheap, listening is hard.”

Culler says theories behind effectively shutting down idle processors (“doing nothing well”) essentially fall into two basic camps that “hearken back to dark ages”—the principles following Token Ring or other time

In the long term, processor power management may borrow from sensor networks and embedded systems, which have extensively dealt with power constraints.

division multiplex technologies, or a Carrier Sense Multiple Access approach akin to Ethernet topology, in which nodes about to transmit can first “sense” whether or not a network is idle before proceeding.

He says this principle can apply to any scenario, be it a Wi-Fi network or a bus protocol on a motherboard. “Doing nothing well and being able to respond to asynchronous events anyway is the key to power proportionality, and can apply across the board,” says Culler.

Management From a Chip

Market demand for dynamically provisioned processors is still an unknown. Albers says processor-level power management is not particularly viewed as a critical issue among European users.

“Energy and environmental issues have always received considerable attention in Europe. However, the typical person is probably more concerned about energy consumption in his household and private car than about the consumption of his PC or laptop,” Albers observes.

IBM has placed a bet on combining chip-level energy allotment with the network architectures of homes and offices. The company has introduced fabricating technology for dedicated power management chips that control power usage while they communicate wirelessly in real time with systems used to monitor smart buildings, energy grids, and transportation systems. The main function of power-management chips is to optimize power usage

and serve as bridges so electricity can flow uninterrupted among systems and electronics that require varying levels of current.

Meyerson says that, while reducing battery usage on end user devices may be sexy, “that's not the win for society. The win for society is when there's an area of a building and the sensors over a period of time crawl through all the data of the occupancy of all the offices, and they autonomically adjust for the fact this is Paris in August—and in Paris in August people just aren't showing up.”

IBM estimates the new technology can cut manufacturing costs by about 20% while allowing for the integration of numerous functions, resulting in one chip where previously three or four were needed. Meyerson says the technology can work for any appropriate algorithm researchers can come up with.

“Discovery algorithms that can look ahead and be predictive instead of reactive can be incredibly important,” he says. “What we are doing is ensuring that if they come up with a solution, there's a way to execute it in a single chip, in a very efficient, synergistic way. It is a real footrace to stay ahead of the energy demands of society and IT.” ■

Further Reading

Albers, S.

Energy-efficient algorithms, *Communications of the ACM* 53, 5, May 2010.

Bansal, N., Kimbrel, T., and Pruhs, K.

Speed scaling to manage energy and temperature, *Journal of the ACM* 54, 1, March 2007.

Ge, R. and Cameron, K.W.

Power-aware speedup. *IEEE International Parallel and Distributed Processing Symposium*, Long Beach, CA, March 26–March 30, 2007.

Gupta, R., Irani, S., and Shukla, S.

Formal methods for dynamic power management. *Proceedings of the International Conference on Computer Aided Design*, San Jose, CA, Nov. 11–13, 2003.

Yao, F., Demers, A., and Shenker, S.

A scheduling model for reduced CPU energy. *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, Milwaukee, WI, Oct. 23–25, 1995.

Gregory Goth is an Oakville, CT-based writer who specializes in science and technology.

© 2011 ACM 0001-0782/11/0200 \$10.00