Real-Time Power Sensors for Intelligent Power Management and Beyond

Srikar Bhagavatula, Byunghoo Jung, and Yung-Hsiang Lu
Purdue University

Editor's notes:
This article presents a low-overhead on-chip power sensor design that enables cross-layer power management policy used in future computing systems.
—Qiang Xu, The Chinese University of Hong Kong

Electricity consumed by data centers alone grew 36% in the United States and 56% worldwide between 2005 and 2010. Accelerated growth in the quality as well as the quantity of mobile platforms like smartphones and tablets has resulted in rapid growth of energy consumption in the mobile sector. In addition, multimedia and social networking have propelled fast expansion of wireless networks as well as cloud storage and services. As a result, the need to reduce power consumption to prolong the uptime of battery-powered devices and to reduce the associated cooling costs in data centers has become a crucial objective of research today with focus on a variety of device, circuit, system level techniques, and algorithms for reducing power consumption.

Power consumption in digital circuits is modeled as

$$P = C \cdot V_{dd}^2 \cdot f + \hat{I}_{sc} \cdot V_{dd} + I_{\text{leak}} \cdot V_{dd}. \quad (1)$$

Improvements in power consumption, therefore, target one of the various parameters appearing in this equation—power supply (V_{dd}), frequency of operation (f), device capacitance (C), and leakage currents (I_{leak}). However, increasing leakage power and the consequent limit of threshold voltage scaling results in shrinking overdrives, imposing limits on voltage scaling. Hence, system-level approaches such as adaptive voltage and frequency scaling (AVFS) are needed to lower power consumption by selective scaling of supply voltage and frequency while reducing the penalties on throughput. Similarly, switching off the power supply to inactive circuit blocks is very effective at reducing leakage power (up to 99% of leakage power). Thus, power gating is a ubiquitous design choice for digital and mixed-signal circuits in scaled technologies.

However, such system-level techniques require real-time feedback with details like power dissipation and on-chip temperature. Although methods for estimation of on-chip temperatures have been proposed, techniques to estimate load currents accurately remain scarce. In addition, among this vast amount of research of low-power design techniques, two questions must be answered: How much power is actually reduced? And at what cost? Hence, there is a need for precise on-chip estimation of load currents in order to maximize power savings and understand the impact of various power saving schemes. It turns out that measuring power or power savings is not as easy due

Digital Object Identifier 10.1109/MDAT.2014.2325534
Date of publication: 19 May 2014; date of current version: 07 August 2014.

Copublished by the IEEE CEDA, IEEE CASS, IEEE SSCS, and TTTC

to many factors and constraints, as explained later in this paper.

We extend our solution for real-time estimation of temperature and current proposed in [1] by first giving a detailed outline of the challenges faced in developing power sensors followed by a brief description of our idea. Results including additional measurements that enable the estimation of $3-\sigma$ error tolerances are then discussed. Finally, we attempt to envision the future of power management.

Challenges in estimating power

Significant efforts have been spent in modeling and simulation of power consumption and subsequent measurement using performance counters which log the activity in a microprocessor and estimate the power consumption. However, results from these models must be validated using experimental data. Moreover, power consumption may vary widely with ambient operating conditions such as the supply voltage or temperature.

Some scoping techniques that sample power consumption (at the rate of a few kilohertz) using a small series resistance have been presented. However, in most microprocessors and digital circuits, the clock rate can vary from hundreds of megahertz to a few gigahertz. A sampling rate of kilohertz cannot ensure that important events such as occasional spikes in power consumption are captured by these sensors. On the other hand, raising the sampling rate to capture such events increases the power overhead significantly. They also employ a separate voltage supply which results in significant area overhead and the inability to obtain block-level power estimates.

Table 1 Challenges making built-in power sensor.

Requirements	Challenges
Low overhead	Power consumption of the measurement circuits
	Area occupied by the measurement circuits
	Performance degradation due to measurement
High accuracy	High sampling rates for fast response
	Tolerance to process variations
	Tolerance to temperature variations

Thermal sensors have widely been deployed in high-performance microprocessors [2]–[6]. However, temperature changes due to power dissipation occur after considerable delay. Coupled with lateral heat spread on chip surface, thermal sensors are characterized by poor temporal and spatial resolutions. Estimation of power from reported temperatures is a resource-intensive process [7] and may also be affected by ambient conditions and cooling; yet, power-based optimization is more widely applicable than temperature management.

These reasons call for the development of built-in power sensors. However, measuring power consumption is challenging. Table 1 summarizes the challenges which can be classified into two categories: overhead and accuracy. To begin with, we discuss the Heisenberg effect: it is impossible to measure anything without perturbing the system being measured. With power sensors, this effect is due to additional power consumed by the sensor, which must be minimized. It should also be able to provide highly accurate estimates with low area overheads, as estimation accuracy determines the confidence with which a corrective action may be initiated. As power management techniques (such as power gating) are widely adopted, the power consumption of a subsystem can change multiple times within a microsecond. Hence, the measurement circuit must have fast response times. Moreover, with device scaling, tolerance to process variations becomes vital. Due to these challenges, few successful studies showing low-overhead, high-accuracy power sensors have been reported.

On-chip power sensor

We have developed one such sensor that provides real-time on-chip estimates of power in [1], utilizing the widespread practice of gating power supply with large MOSFETs known as sleep transistors. As sleep transistors operate in linear region during theon-state, the drain-to-source voltage drop (V_{DS}) varies linearly with current flowing through the transistor which is being consumed by the circuit under test. Therefore, it can be treated as a resistor with anon-resistance of R_{sleep}. A concurrent study reported in [8] also attempts to utilize this voltage drop to proportionally integrate a capacitor and provide a quantized two-bit output. However, this two-bit output of the monitor may increase the number of iterations the power management loop

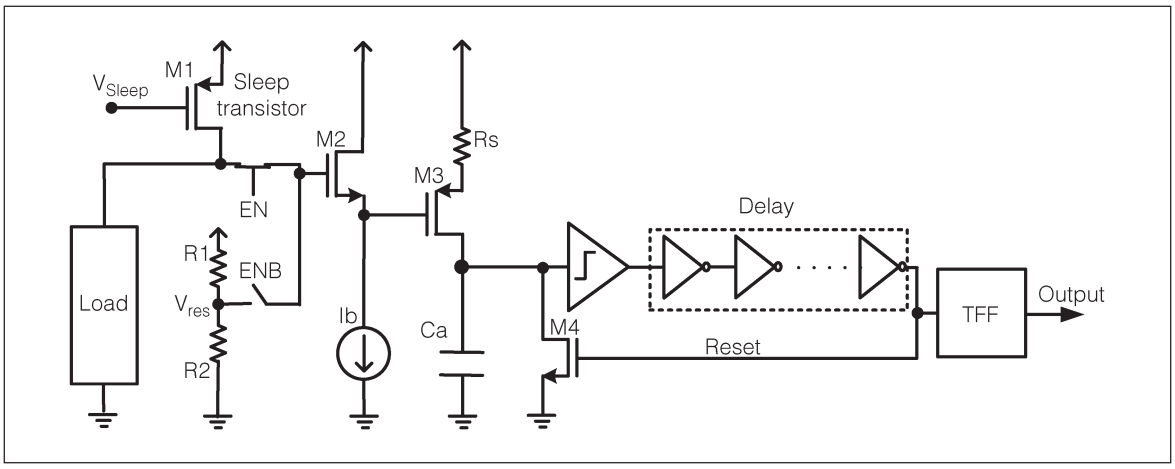

Figure 1. Circuit schematic of the power and temperature sensor [1].

has to go through before converging to the optimal operating point. In addition, its response time also trades off steeply with power resolution.

Figure 1 shows the schematic of our proposed power sensor. Voltage drop V_{DS} is sensed through a source follower (gain of A_{sf}), an amplified, and then converted into a current that is proportional to the load current by a common-source FET with a transconductance of G_m. This current is used to charge a capacitor (with a capacitance C_a). When the voltage at this capacitor reaches the threshold voltage of a comparator ($V_{th;COM}$), the capacitor is reset via a delay chain of inverters. By making the discharge time negligible compared to its charging time, an inverse relationship is ensured between the time period of the voltage waveform and the charging current, and therefore, between the time period at capacitor and the load current. A T-Flip flop at the end of the delay chain converts this waveform into a square pulse waveform to be input to a reciprocal pulse counter. As the rate of charging of the capacitor is proportional to load current, the frequency of output pulses is proportional to load current

$$I_{\text{load}} = I_0 + \frac{2 \cdot C_a \cdot V_{th,\text{COM}} \cdot f_{\text{out}}}{R_{\text{sleep}} \cdot A_{sf} \cdot G_m} = I_0 + K \cdot f_{\text{out}} \quad (2)$$

where I_{load} is the load current and I_0 is the zero error, and f_{out} is the frequency of output pulses. From (2), it can be seen that the output frequency is proportional to the load current, but the proportionality constants are susceptible to process and temperature variations. In order to obtain a power estimate tolerant to such variations, we devised a two-point calibration technique, which requires an estimate of on-chip temperatures and a temperature-tolerant voltage comparator.

Temperature-tolerant low-power comparator

A process- and temperature-tolerant inverter comparator [1] is designed to monitor the capacitor (C_a). As shown in Figure 2, this inverter comparator consists of two inverters with voltage-controlled resistances to the supply nodes. One

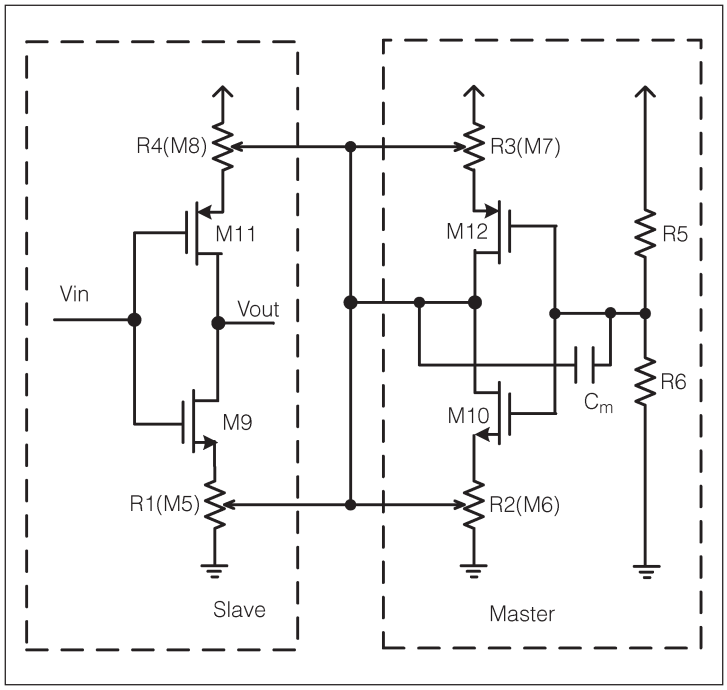

Figure 2. Circuit schematic of the process- and temperature-tolerant comparator [1].

inverter stage acts as the master switch and is fed by a resistor–divider voltage (set to $V_{dd}/2$). Its output controls the resistances of the four MOSFETs (R1–R4) that connect these switches to supply rails and acts as automatic calibration against temperature variations. Input to the slave inverter is connected to C_a and triggers when the node voltage reaches $V_{dd}/2$ tolerant to temperature variations [1].

Temperature estimation

This sensor can also be used to estimate the temperature by disconnecting the source follower from the sleep transistor and connecting it to a resistor divider that provides a temperature-tolerant voltage input to the sensor. As threshold voltages of MOSFETs decrease linearly with increasing temperature, the charging current in this mode of operation can be approximated to a linear function of temperature. Thus, the time period (t_p) is given by

$$t_p \approx \frac{C_a \cdot V_{th,\mathrm{COM}} \cdot (I_c - KT)}{I_c^2} \quad (3)$$

where I_c is the charging current at zero Kelvin, and k is a process-dependent constant.

Calibration

The equation relating temperature (T) to output time-period (t_p) can be obtained by measuring output time periods (t_{p1}, t_{p2}) at two different test temperatures (T_1 and T_2)

$$T = a_1 + b_1 \cdot t_p. \quad (4)$$

At each of these two test temperatures (T_i), output frequency is also measured in the power sensor mode at two different current loads (I_1 and I_2). Therefore, at each given temperature, the slope K (at T_i) and intercept I_0 at (T_i) for the linear equation between I_{load} and f_{out} are obtained. As shown in our previous work [1], the slope and intercept also vary linearly with temperature for small ranges in input voltage. Hence, following equations are obtained:

$$K1 = a2 + b2 \cdot T \quad (5)$$

and

$$I_0 = a_3 + b_3 \cdot T \quad (6)$$

where a_1, b_1, a_2, b_2, a_3, and b_3 are process-dependent constants.

Thus, the process dependency can be measured and is captured in the constants a_1–b_3. The temperature sensor then captures variations due to on-chip temperatures and enables a variation-resilient estimate of the load current (I_{load}) from measured quantities (t_p and f_{out}) as follows:

$$I_{\mathrm{load}} = a_2 + b_2 \cdot (a_3 + b_3 \cdot t_p) + \big(a_1 + b_1 \cdot (a_3 + b_3 \cdot t_p)\big) \cdot f_{\mathrm{out}}. \quad (7)$$

Analysis

This power and temperature sensor was designed in 45-nm silicon-on-insulator (SOI) technology and occupies 0.0196-mm^2 on-chip area, which includes the area of the sleep transistor (0.008 mm^2). Thus, the sensor adds an area overhead of only 0.012 mm^2. In addition, as sleep transistors are always designed to have low current-resistance (IR) drops (tens of millivolts) inon-state, the input range does not change, making for an easy replication of the sensor at multiple locations on chip without redesign.

With a V_{dd} of 1.2 V, this sensor was tested at various temperatures from 25 °C to 85 °C for load currents ranging from 0 to 5 mA (Figure 3a). For the given sleep transistor design, a current load of 3 mA corresponded to an IR = current(I)*Resistance(R) voltage drop across a given resistor drop across sleep transistor of 15 mV. Hence, the power sensor has sufficient dynamic range to monitor average power for most circuits. The sensor output was monitored by a reciprocal pulse counter implemented on an FPGA running at 500 MHz. The output which is a digital count is to be sent to a power management unit (PMU) which runs power saving algorithms to implement techniques such as AVFS or power gating.

The sensors output pulsewidth was lower than 54 ns under test conditions. So, theoretically, the highest achievable conversion speed would be as high as 18 MHz. However, in order to reduce the effect of supply noise and sampling rates, the output is averaged over a window of 0.5 s (2 MHz) which provides a resolution of 0.1 ns. With on-chip frequency counters with higher resolutions (order of tens of picoseconds), response times better than 0.5 μs can be achieved. The current overhead of this sensor is 100 μA at 1.2-V V_{dd}.

At all the test temperatures, the sensor output showed a linear response with load current

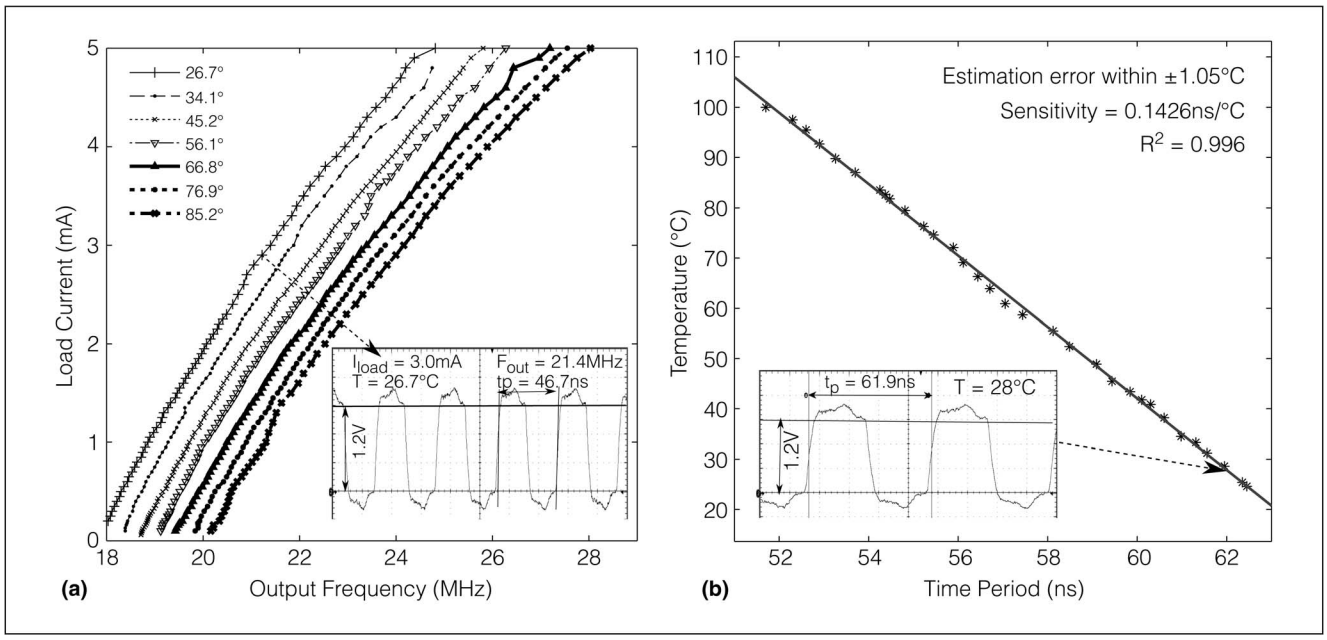

Figure 3. Sensor measurement results in (a) the power sensor mode at various temperatures and (b) the temperature sensor mode [1].

(Figure 3). However, the accuracy of the sensor is limited by linearity of slope (K_1) and the intercept (I_0) with temperature variations. Hence, the target accuracy of the estimates being within ±10% of the load is limited to current values less than 3.3 mA (Figure 4). In temperature sensor mode of operation, output time period is measured at various temperatures from 22 °C to 100 °C where the sensor shows linear response with $R^2 > 0.99$. The estimation accuracy in this mode was within ±1.05 °C of the on-chip temperature, with a $3-\sigma$ error within 4.05 °C over a range of temperatures from 20 °C to 100 °C. This accuracy is also sufficient for thermal management in microprocessors [6].

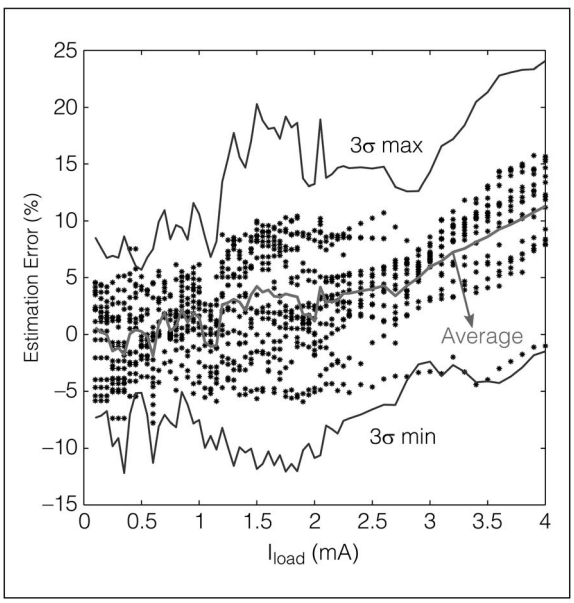

Figure 4. Measured errors in estimating I_{load} at various temperatures from 20 °C to 85 °C.

Compared to the best reported thermal sensors [2]–[6] (Table 2), that occupy on-chip areas in excess of 0.08 mm^2, have power overheads in the order of a few milliwatts and conversion times in excess of 10 μs, this sensor offers a low-cost (0.3 mW), real-time (0.5 μs) alternative to on-chip power management. Energy per conversion (conversion time × power overhead) is used as a figure of merit (FOM) to compare thermal sensors. However, it does not account for the time lag between power consumption and the associated rise in temperature that also contributes to obtaining a power estimate. Hence, we define FOM_2 as

$$FOM_2 = \text{inaccuracy} \times \text{response time} \times \text{power overhead} \quad (8)$$

where response time is accounted as the time elapsed between the occurrence of a power event and a corresponding change in the sensor output so that the PMU may initiate a response. Wang et al. [7] suggest that in order to estimate power within 4%

Table 2 Performance summary and comparison with the state of the art.

Metric	[2]	[3]	[4]	[5]	[6]	[8]	This work
Technology	32nm	0.16μm	0.18μm	65nm	0.7μm	0.35μm	45nm
Area overhead	0.02mm^2	0.08mm^2	0.18mm^2	0.008mm^2	0.8mm^2	0.05mm^2	0.012mm^2
Power overhead	3.78mW	5.1μW	30μW	500μW	159μW	244μW	120μW (310μW)[1]
Conversion time	10-100μs	5.3ms	12.5μs	2.3μs	2.2ms	4.32μs	0.5μs
Temperature inaccuracy	1.5°C -10-110°C	±0.15°C -55-125°C	±0.5°C 0-100°C	±0.15°C -45-130°C	±1.5°C 0-110°C	-	±1.05°C 22-100°C
Current inccuracy[2]	-	-	-	-	-	-	±10%
Conversion Energy	37.8nJ	26.5nJ	0.375nJ	1nJ	350.9nJ	1.05nJ	0.15nJ
FOM$_2$	228nJ	1.36nJ	1.82nJ	30.04nJ	23.68nJ	-	0.015nJ

error margin with preverified on-chip models it can take up to 1.5 ms.

Future of power management

Such power sensors have profound implications beyond just understanding the power consumption of a chip. A typical power management control loop is shown in Figure 5. The SoC is run by a clock with time period < 5 ns which also runs the power management algorithm. Such algorithms can find opportunities for voltage and frequency scaling opportunities within tens of cycles. Phase-locked loops (PLLs) with locking times in microseconds and voltage regulators with settling times in tens of nanoseconds have been reported. However, sensors in use today still have response times in the order of milliseconds, which are slower than the rest of the control loop by three orders of magnitude. A few algorithms to predictively manage power have been explored that appear to sidestep the need for power sensors. However, these algorithms either rely on past power consumption values to predict future power consumption and are, therefore, dependent on sensors, or they require extensive training [10]. The performance of such learning-based algorithms might also degrade if the characteristics of the workload change over time.

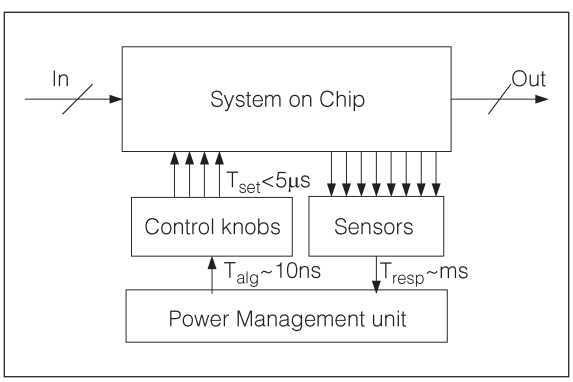

Figure 5. Block diagram of the power management control loop.

Hence, sensors with fast response times and fine spatial resolution are the crucial next step in enabling the shift away from the coarse-grain power management of the present to a fine-grain power

management in the future (Figure 6) [9]. With the advent of many-core processors into mainstream computing, scheduling algorithms that implement core hopping, thread stopping, global and local DVFS techniques to reduce the incidence of hot spots, and improve reliability as well as throughput will become more popular. Such algorithms can achieve higher savings and throughputs if we can measure power consumption and predict these thermal events before they occur and prevent critical temperature from being reached by switching to a lower activity level. As the number of cores increases, local chip temperatures depend on chip-level heat dissipation and cooling effects rather than just the local (core-level) heat dissipation. As a result, reliability of thermal sensors for power management decreases. At the same time, the number of computations required to back annotate temperature values to local power dissipation increases with the number of sources of power dissipation, which leads to higher power management overheads. Therefore, sensors that rely on true power estimation such as the one presented here are essential in next-generation computing.

Furthermore, there has been a growing recognition of the need to define efficiency of algorithms not just in terms of the orders of computational complexity, but also in terms of energy efficiency [11]. Power sensors with quick response times become essential not just to rate the efficiency of algorithms, but also to validate the

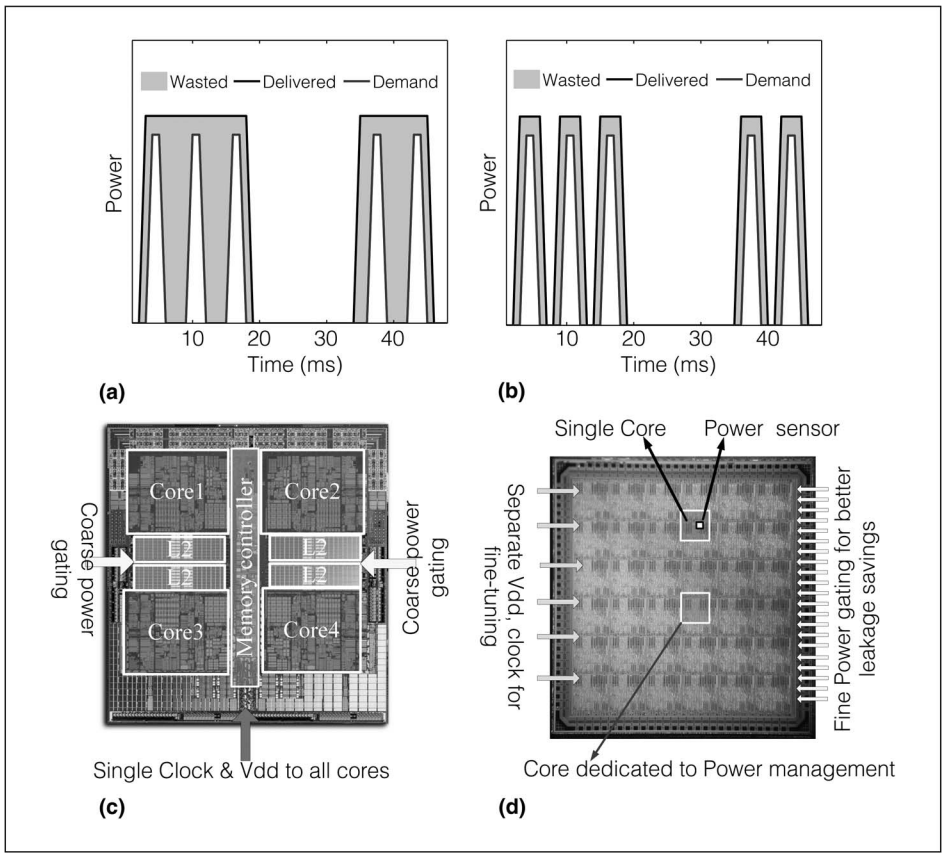

Figure 6. (a) and (c) Coarse-grained and (b) and (d) fine-grained power management in time and spatial domains.

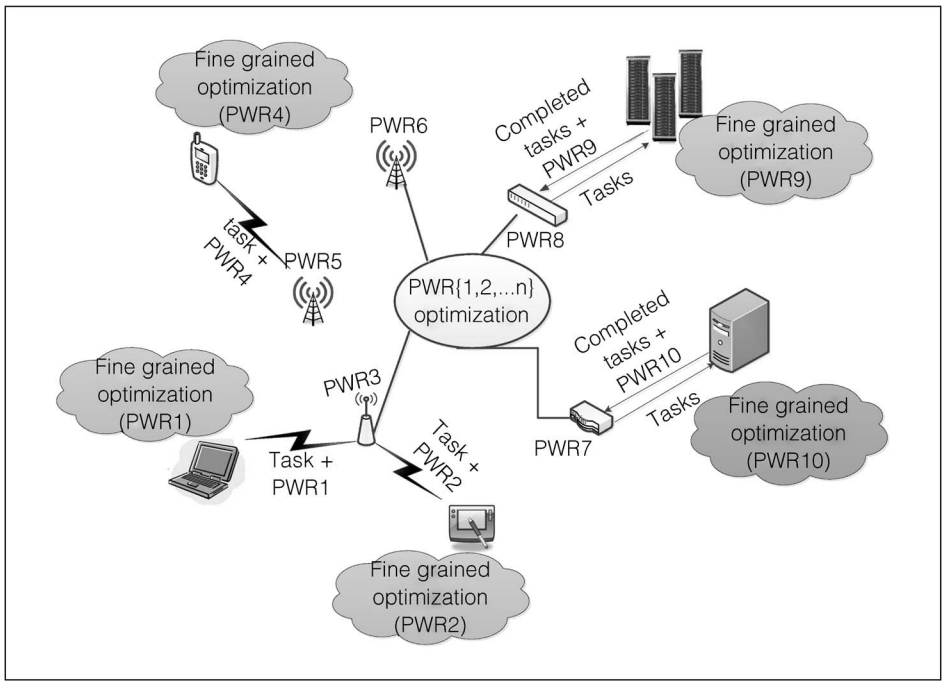

Figure 7. Envisioned future of power management in a connected world.

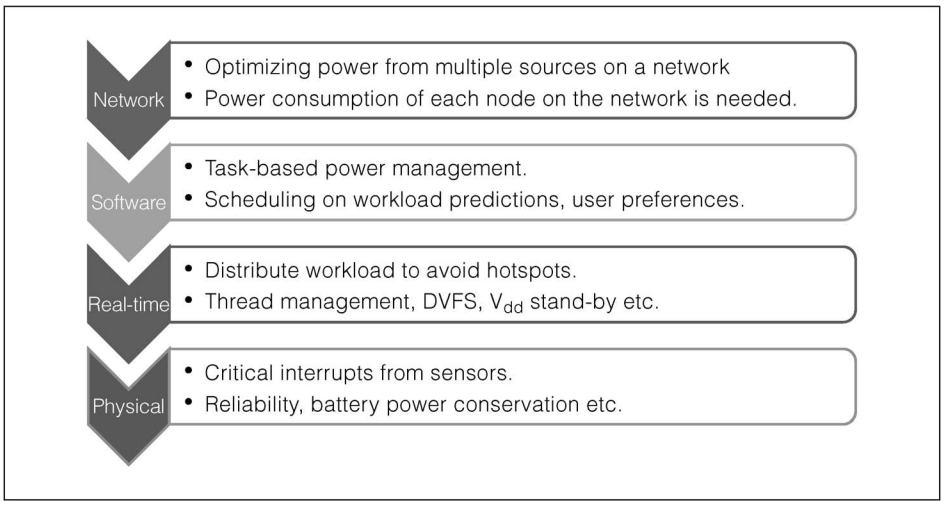

Figure 8. Cross-layer coordination for smarter power management.

principles based on which we define such metrics. In addition, as testing increases in complexity, sensors needed for online testing (after the chip has been shipped) also become relevant. These power sensors enable the implementation of such online test and debug schemes due to the ready availability of output in the form of digital readouts.

As more and more electronic systems are connected through networks, saving the power in just one system regardless of its interactions with the other connected systems is insufficient. As an example, most mobile systems (smartphones, tablets, and laptops) are connected to the Internet through wireless networks. When a mobile user watches a streaming video, power is consumed on the mobile system, as well as wireless access points, network routers, servers, and storage. It is inadequate to separate these connected systems and reduce their power consumption independently. A recent paper [12] proposes the concept of end-to-end energy management, suggesting the need to consider multiple connected systems as a whole for power reduction (Figure 7).

Real-time power sensors are essential components for realizing end-to-end energy management because we are able to monitor the power dissipation of multiple systems as they communicate through networks. Moreover, the premise of cloud computing is the ability to autonomously migrate computing to meet performance requirements and resource constraints. Access to real-time power consumption allows researchers and engineers to dynamically adjust power management strategies across systems to ensure better efficiency.

Thus, tomorrow's power management strategies require coordination across layers and optimization involves a combination of algorithms at the network, software, kernel, and hardware levels (Figure 8). On-chip power sensors that provide real-time power readings with minimal overheads are, therefore, crucial in realizing this in the future. ■

References

[1] S. Bhagavatula and B. Jung, "A low power real-time on-chip power sensor in 45 nm SOI," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 59, no. 7, pp. 1577–1588, Jul. 2012.

[2] J. Shor, K. Luria, and D. Zilberman, "Ratioametric BJT-based thermal sensor in 32 nm and 22 nm technologies," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2012, pp. 210–212.

[3] K. Souri, Y. Chae, and K. Makinwa, "A CMOS temperature sensor with a voltage calibrated inaccuracy of $\pm 0.15\ ^{\circ}C\ (3\sigma)$ from -55 to $125\ ^{\circ}C$," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2012, pp. 208–210.

[4] C. Wu, W. Chan, and T. Lin, "A 80 kS/s 36 μW resistor-based temperature sensor using BGR-free SAR ADC with a unevenly-weighted resistor string in 0.18 μm CMOS," in *Proc. Symp. VLSI Circuits*, Jun. 2011, pp. 222–223.

[5] S. Hwang, J. Koo, K. Kim, H. Lee, and C. Kim, "A 0.008 mm^2 500 μW 469 kS/s frequency-to-digital converter based cmos temper-ature sensor with process variation compensation," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 60, no. 9, pp. 2241–2248, Sep. 2013.

[6] A. Heidary, G. Wang, K. Makinwa, and G. Meijer, "A BJT-based CMOS temperature sensor with a 3.6 $pJ.K^2$ resolution FoM," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2014, pp. 224–225.

[7] H. Wang, S. X.-D. Tan, X.-X. Liu, and A. Gupta, "Runtime power estimator calibration for high-performance microprocessors," in *Proc. Design Autom. Test Eur.*, Mar. 2012, pp. 352–357.

[8] N. Mehta, G. Naik, and B. Amrutur, "*In-situ* power monitoring scheme and its application to dynamic voltage and threshold scaling for digital CMOS

integrated circuits,'' in *Proc. Int. Symp. Low Power Electron. Design*, Aug. 2010, pp. 259–264.
[9] V. De, ''Fine-grained power management,'' in *Proc. Int. Solid-State Circuits Conf. Tech. Forum*, Feb. 2013.
[10] H. Jung and M. Pedram, ''Supervised learning based power management for multicore processors,'' *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 29, no. 9, pp. 1395–1408, Sep. 2010.
[11] K. Kant, ''Toward a science of power management,'' *Computer*, vol. 42, no. 9, pp. 99–101, Sep. 2009.
[12] Y. H. Lu, Q. Qiu, A. R. Butt, and K. W. Cameron, ''End-to-end energy management,'' *Computer* vol. 44, no. 11, pp. 75–77, Nov. 2011.

Srikar Bhagavatula is currently working toward a PhD at the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. His research interests include on-chip power estimation and power management techniques. Bhagavatula has a BTech in electrical engineering from the Indian Institute of Technology Bombay (IIT Bombay), Mumbai, India (2006).

Byunghoo Jung is an Associate Professor in the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. His research interests include analog and radio-frequency (RF) circuit design for communication systems. Jung has a PhD in electrical and computer engineering from the University of Minnesota, Minneapolis, MN, USA.

Yung-Hsiang Lu is an Associate Professor in the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. His research focuses on energy efficiency of computing systems. Lu has a PhD in electrical engineering from Stanford University, Stanford, CA, USA. He is a senior member of the IEEE.

■ Direct questions and comments about this article to Srikar Bhagavatula, Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA; sbhagava@purdue.edu.