# Post-Silicon Characterization and On-Line Prediction of Transient Thermal Field in Integrated Circuits Using Thermal System Identification

Minki Cho, *Member, IEEE*, Khondker Zakir Ahmed, *Student Member, IEEE*, William J. Song, Sudhakar Yalamanchili, *Senior Member, IEEE*, and Saibal Mukhopadhyay, *Senior Member, IEEE*

***Abstract*— Thermal system identification (TSI) is presented as a methodology to characterize and estimate the transient thermal field of a packaged IC for various workloads considering chip-to-chip variations in electrical and thermal properties. The time–frequency duality is used to identify the thermal system as a low-pass filter in frequency domain through on-line power/thermal measurements on a packaged IC. The identified characteristic system for an individual IC is used for on-line prediction of the transient thermal field of that specific IC for a power pattern. A test-chip, fabricated in 130-nm CMOS, demonstrates the effectiveness of TSI in post-silicon characterization and prediction of transient thermal field. The application TSI in thermal analysis of multicore processors is presented.**

***Index Terms*— Fourier transform, frequency-domain modeling, post-silicon temperature prediction, process variations.**

## I. Introduction

THE characterization of the spatiotemporal variation of on-chip temperature is crucial for thermal-aware design, assembly, and management of a chip [1], [2]. The temperature pattern is generated by the interaction of time-varying power patterns and the thermal properties (resistivity and heat capacity) of die and package. The thermal properties of the die/package assembly can vary between different instances of same IC or over time [3]. Moreover, imperfections in the manufacturing process lead to die-to-die variations in transistor leakage [1]. The leakage and temperature are positively correlated—a higher temperature results in higher leakage, which further increases the temperature. Hence, for the same dynamic power, chip-to-chip leakage variation leads to variation in the on-chip temperature [4]. Therefore, a new challenge is emerging—post-silicon characterization and prediction of the transient thermal field—where the objective is to characterize the interaction of power and chip-to-chip variations in thermal and electrical properties to predict the transient temperature of a packaged IC for various workloads.

The existing transient thermal simulation methods (finite element/volume or distributed RC), suitable for fine-grain design time transient thermal analysis, require accurate estimation of thermal conductivity and heat capacity of all materials [5]–[7]. Methods for measuring the thermal resistance and capacitance of thermal interface material (TIM), heat sink, convective, and heat spreader were studied [8]–[11]. Poppe *et al*. [9] presented dynamic electrical temperature measurement and Campbell *et al*. [10] presented the flash diffusivity method for measuring thermophysical property. The measurements of thermal resistance and capacitance suffer from repeatability, contamination, pressure, and inaccuracy. Further, the thermal resistances of TIM, heat sinks, and interface change with stacking because of imperfections in attachment and manufacturing. Kurabayashi *et al*. noted that the die-attach resistance differs substantially from the value predicted using the bulk thermal conductivity of the attachment material because of partial voiding and delamination [12]. Therefore, the thermal simulators used during design time are difficult to adopt for post-silicon thermal analysis.

This paper presents a unique approach to address the specific requirements of post-silicon thermal characterization and prediction. The proposed approach, referred to as thermal system identification (TSI), is based on the principles of system identification, frequency-domain signal analysis, and positive feedback systems (Fig. 1). The thermal behavior of a chip considering leakage temperature interaction is represented as a system, where power sources are the inputs and observed temperature values at different locations are the outputs. A unique thermal system $[H(\omega)]$ therefore can be constructed for each chip (TSI). The frequency response of the temperature variation $[T(\omega)]$ over a time interval is computed from the Fourier transform of the power pattern $[P\ (\omega)]$ in that interval and the filter matrix $[T(\omega) = H(\omega) \times P(\omega)]$. The time-domain temperature is obtained from the temperature spectra. This paper presents the mathematical principles of the post-silicon temperature prediction approach and the methodologies for TSI in the frequency domain using sequences of on-chip power and temperature measurements. A test-chip is designed in 130-nm CMOS technology with on-chip poly-resistor based heater and analog and digital temperature sensors [14]–[16] to demonstrate the effectiveness of the proposed approach. The measurement results demonstrate average error of

Manuscript received December 20, 2012; revised June 10, 2013; accepted June 17, 2013. Date of publication August 29, 2013; date of current version December 30, 2013. This work was supported by the Semiconductor Research Corporation under Grant 2084.001. Recommended for publication by Associate Editor B. Courtois upon evaluation of reviewers' comments.

M. Cho was with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA. He is now with the Intel Labs, Pittsburgh, PA 15222 USA (e-mail: mcho8@gatech.edu).

K. Z. Ahmed, W. Song, S. Yalamanchili, and S. Mukhopadhyay are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA (e-mail: khondker@gatech.edu; wjhsong@gatech.edu; sudha@ece.gatech.edu; saibal@ece.gatech.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCPMT.2013.2271504

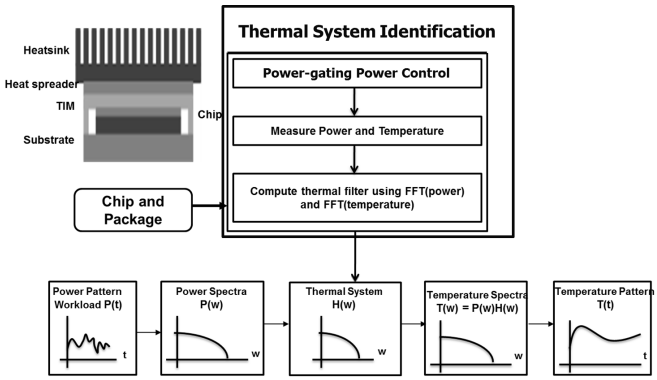


Fig. 1. Methodology of post-silicon prediction of the transient thermal field using TSI.

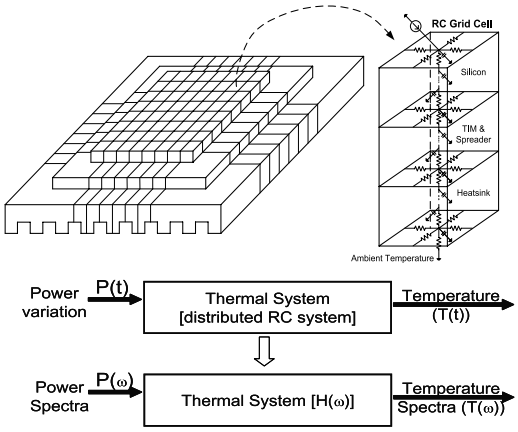


Fig. 2. Distributed RC-based thermal model and the equivalent frequency domain model.

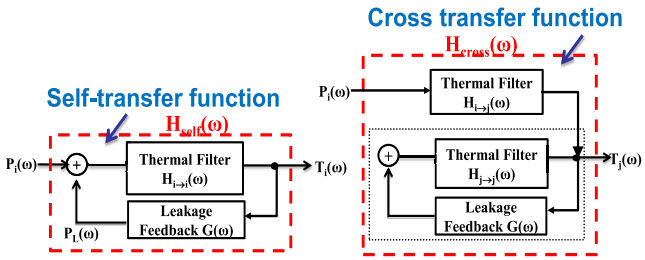


Fig. 3. Mathematical principle of the proposed approach.

## II. Mathematical Principles

### A. Frequency-Domain Thermal Modeling

In a RC thermal model, the temperature at a location is the analog of voltage across the thermal capacitance at that location, and the power dissipation causing heat generation is modeled with a current source. We can consider power and temperature as an input and an output of the thermal system, respectively (Fig. 2). The power input at one location impacts the temperature at multiple locations. Hence, the temperature spectra at locations $i$ and $j$ due to a power source at location $i$ is given as follows:

$$T_i(\omega) = P_i(\omega)H_{i\to i}(\omega) \quad \text{and} \quad T_j(\omega) = P_i(\omega)H_{i\to j}(\omega) \tag{1}$$

where $H_{i\to i}$ is the self-transfer function ($H_{\text{self}}$) connecting power and temperature of the same location and $H_{i\to j}$ is the cross transfer function ($H_{\text{cross}}$) connecting power of one location and temperature of another. Note that the self and cross transfer functions need to incorporate both the thermal characteristics and leakage–temperature interactions. To understand this first consider the self-transfer function as follows:

$$\begin{aligned} T_i(\omega) &= \left[P_{D_i}(\omega) + P_{L0_i}(\omega) + F\left(f(T_i(t))\right)\right] H^{\text{TH}}_{i\to i}(\omega) \\ &= \left[\underbrace{P_{D_i}(\omega) + P_{L0_i}(\omega)}_{P_i(\omega)} + \alpha T_i(\omega)\right] H^{\text{TH}}_{i\to i}(\omega) \end{aligned} \tag{2}$$

where $P_{L0}$ is the leakage power profile at room temperature, $f(T)$ is the sensitivity of leakage power to temperature, $P_D$ is the dynamic power profile, and $H^{\text{TH}}_{i\to i}(\omega)$ is the thermal RC properties of the system (excluding leakage feedback). We assume a linear interaction between leakage and temperature to improve analytical tractability. Both the room temperature leakage ($P_{L0}$) and the coefficient ($\alpha$) depends on leakage–temperature interaction. Note that $P_i(\omega) = P_D(\omega) + P_L(\omega)$ is the spectral response of power without leakage–temperature feedback. The thermal system is modeled as (Fig. 3)

$$T_i(\omega) = \underbrace{\left[H^{\text{TH}}_{i\to i}(\omega) \Big/ \left(1 - \alpha H^{\text{TH}}_{i\to i}(\omega)\right)\right]}_{H_{i\to i}(\omega)=H_{\text{self}}(\omega)} P_i(\omega). \tag{3}$$

We now evaluate the temperature of location $i$ due to power source at location $j$. The heat generated in location $i$ propagates to location $j$, which increases the temperature of location $j$. Increase in temperature at location $j$ triggers the leakage feedback loop at location $j$. This results in leakage power at location $j$ and hence, increase temperature of location $i$. The temperature increase in location $j$ due to power of location $i$ is therefore estimated as follows (Fig. 3):

$$\begin{aligned} T_j(\omega) &= P_i(\omega)H^{\text{TH}}_{i\to j}(\omega) + \alpha T_j(\omega)H_{j\to j}(\omega) \\ &= \underbrace{\left[H^{\text{TH}}_{i\to j}(\omega) \Big/ \left(1 - \alpha H_{j\to j}(\omega)\right)\right]}_{H_{i\to i}(\omega)=H_{\text{self}}(\omega)} P_i(\omega) \end{aligned} \tag{4}$$

where $H^{\text{TH}}_{i\to j}$ is the thermal RC properties of the system (excluding leakage feedback) connecting location $i$ to location $j$. The above analysis suggests that if $H_{\text{self}}(\omega)$ and $H_{\text{cross}}(\omega)$ are estimated for a packaged IC it incorporates both the chip/package specific thermal properties and leakage–temperature interactions (i.e., process variations).

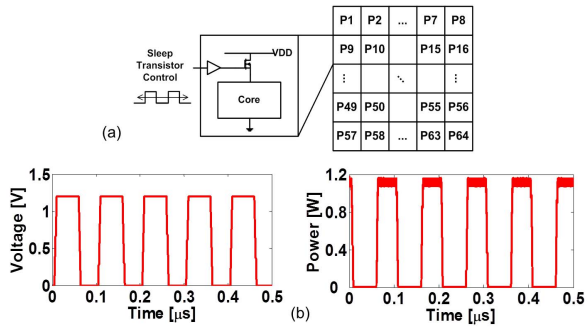


Fig. 4. Power-gating-based approach to power spectra generation. (a) Sleep transistor signal (5 MHz). (b) Power pattern (5 MHz).

### *B. Methods for Thermal System Identification*

The principle discussed above requires frequency responses of the self- and cross-transfer functions for each chip (i.e., TSI). To perform TSI on a multiple output system, the input power source can be excited with various spectra and temperature is measured at all observation points considered. Hence, (2) transforms to the following:

$$T_k^{\text{measured}}(\omega) = P_i^{\text{applied}}(\omega) H_{i\to k}(\omega) \Rightarrow H_{i\to k}(\omega) = \frac{T_k^{\text{measured}}(\omega)}{P_i^{\text{applied}}(\omega)}. \quad (5)$$

The above equation can be used to estimate the thermal filter from an inputs power source at location $i$ to all temperature observation points $k$ (including the observation point at location $i$). As (5) is a division of two complex numbers, both magnitude and phase of the filter response are extracted. The practical challenge in the TSI of digital systems is the generation of power spectra in (5). The accurate approach is to apply sinusoidal power waveforms of different frequencies (small signal analysis). Generating sinusoidal power waveforms in hardware (in a chip) is, however, challenging. We propose to use a pseudosquare wave with varying principle frequency to extract the thermal filer. One way to achieve this goal is to exploit the concept of gating in digital processor cores. We can control the core-level power and clock gating to generate power patterns of desired frequency spectra. To illustrate the approach, we perform SPICE simulation considering power gating as shown in Fig. 4. We use hundreds of 15-stage ring oscillators (ROs) to emulate dynamic power. Each system block is controlled with a periodic sleep control signal of a given frequency that generates periodic power pattern of same frequency. Hence, by controlling the period of the sleep signal we can modulate the spectral behavior of the power patterns.

## III. Test-Chip Architecture

A test-chip is designed to verify the proposed TSI approach. The test-chip contains three BJT-based thermal sensors, a digital RO-based temperature sensor, and a polyresistor-based heater to apply a controllable power profile. The heater is used to mimic a digital system block. The test-chip is fabricated in 130-nm CMOS technology and mounted on a printed circuit board for measurement. Fig. 5 shows the die photo and the organization of the test-chip. Extraction

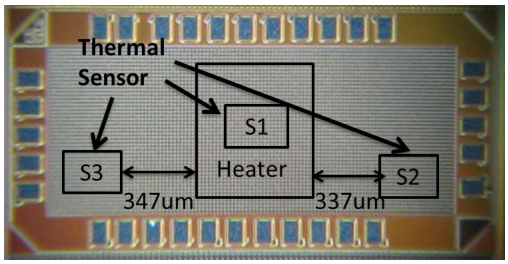


| | |
|---|---|
| Technology | 130nm CMOS |
| Chip size | 2 mm² |
| $V_{CTRL}$ | 1.2 V |
| $V_{DD}$ | 0.55 - 1.2 V |
| $V_{HEATER}$ | 0 - 10 V |
| Max. Power Density in Heater | 130W/cm² |
| Reference Clock Freq. | 10KHz |

Fig. 5. Die micrograph showing the architecture and the measurement conditions.

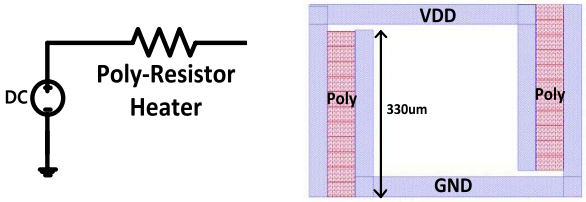


Fig. 6. Schematic view of poly-resistor-based heater.

is performed based on the power/temperature measurement using an oscilloscope. Once the thermal filter is extracted, the temperature prediction is performed using the extracted $H(\omega)$ through MATLAB-based fast Fourier transform and inverse fast Fourier ransform (IFFT) computations.

### *A. Poly-Resistor-Based Heater*

To characterize the thermal system of a chip, power with controllable power spectra needs to be generated. An on-chip heater is designed to mimic the power consumption of large digital functional blocks. To generate a maximum operating power/thermal condition of 130 W/cm$^2$, a resistor of 50 Ω is used as shown in Fig. 6. The power density is controlled by varying the supply voltage across the resistor from 0 to 10 V. The power density is then estimated by the power dissipated by the heater and the given chip area. Current flowing through the heater is limited to meet reliability constrains of the heater's poly and metal components [13].

### *B. Analog Temperature Sensor*

A BJT-based analog temperature sensor is designed (Fig. 7). A supply independent current is generated, mirrored, and pushed through the diode-connected PNP device. The base–emitter voltage, $V_{BE}$ that exhibits negative temperature coefficient, is used as the thermal sensor output. Fig. 7 shows the characterization of the sensor. Measured sensor output decreases by ∼200 mV, which corresponds to a temperature increase of ∼120 °C, for a generated heater power of ∼2 W. Identical sensors are placed in designated locations of the chip

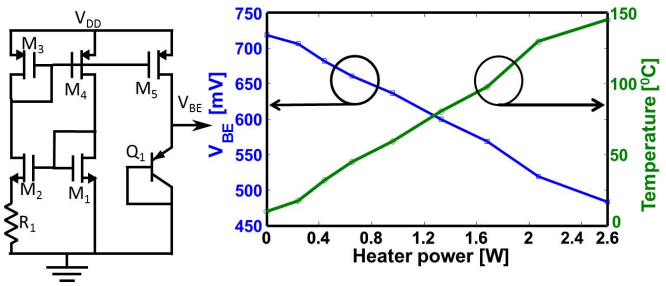


Fig. 7. Schematic view of the analog sensor and measurement results showing the sensor characterization.

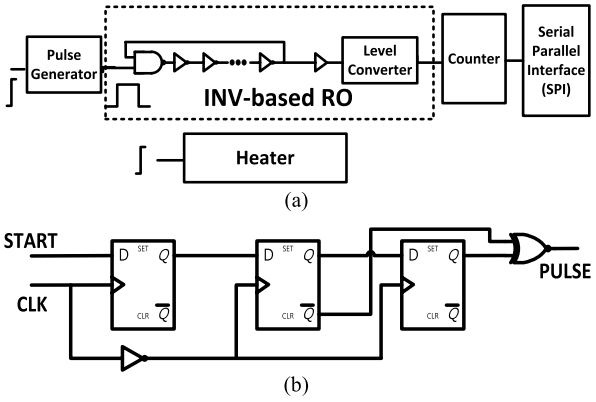


Fig. 8. Digital sensor. (a) Block diagram. (b) Pulse generator.

for studying the spatial effects of thermal events. One sensor (S1) is located in the center of the heater, while the other two sensors (S2, S3) are located ~350 $\mu$m away from the center of the heater as shown in Fig. 5.

### C. Digital Temperature Sensor

A RO-based digital temperature sensor is also designed (Fig. 8). Temperature is calculated by measuring the frequency of a inverter-based ROs. Control circuits include a pulse generator, a heater, level converters, and a counter. To decouple the effect of control circuits, the RO is implemented with 100-stage plus an initial NAND2 gate. The delay of 101-stage RO is dominant in the total delay considering the control circuits. To find the oscillation frequency, we count the transitions of the RO over a given time period using a three flip-flop pulse generator controlled by a reference clock (Fig. 8). A reference clock signal is applied to flip-flop and an inverted clock signal is applied to next two consecutive flip-flops. Once the start signal propagates through continuous flip-flops based on the clock signal, the delayed start signal by a clock signal is generated. Those delayed start signals at the output of flip-flops create a pulse signal, which has a time period of the reference clock through xnor gate. Only test circuits are supplied with variable $V_{DD}$ while a nominal supply is used for the control circuits. A conventional cross-coupled level converter is placed at the output of each RO, which is connected to a multiplexer. A 10-KHz reference clock is used and the SPI block is used to collect the count number to estimate the RO frequency.

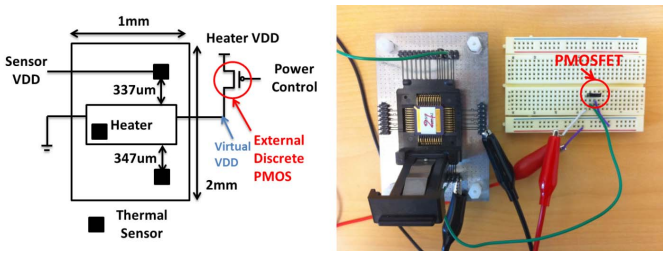


Fig. 9. Schematic view and photo of the experimental hardware.

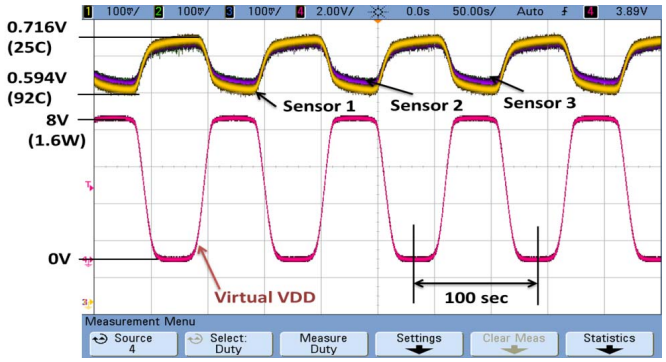


Fig. 10. Oscilloscope waveform of thermal sensor and power.

## IV. Measurement Results

The chip power is controlled externally by pMOS (FQD11P06 from Fairchild semiconductor) to generate a specific frequency signal as shown in Fig. 9. Fig. 10 shows the measured waveforms of the temperature sensor output and virtual $V_{DD}$ voltage. The power is calculated from the virtual $V_{DD}$ voltage. Virtual $V_{DD}$ voltage, which is the voltage across the poly-resister-based heater, is measured at the drain of the sleep transistor through an external pin. Sinusoidal waves are applied to the gate of sleep transistor during the filter extraction period. Because of the switching characteristics of a MOSFET, the generated power signal has the shape of a square waveform (Fig. 10).

### A. Thermal Filter Extraction

The thermal filter is extracted through the power and temperature measurement. We apply multiple power patterns with different fundamental frequencies (0.01 Hz, 0.1 Hz, 1 Hz, 10 Hz, 100 Hz, 1 KHz, and 10 KHz) to the gate of the sleep transistor. Fourier transforms are performed for the measured power and temperature patterns in time domain. For example, the power spectra converted from 1-Hz power signal shows dominant 1-Hz frequency component (Fig. 11). Each fundamental frequency component of the thermal filter is estimated, and intermediate frequency data points are interpolated. Note that extracting one fundamental frequency at a time can help to eliminate the effect of thermal sensor noise, which can introduce error in the extracted filter.

Fig. 12(a) shows the frequency response of the thermal filters at three different locations. We observe that the thermal systems behave as a low-pass filter. Hence, fast time-varying power input has less impact on the temperature while low frequency power variations are more critical. Sensor 1, which

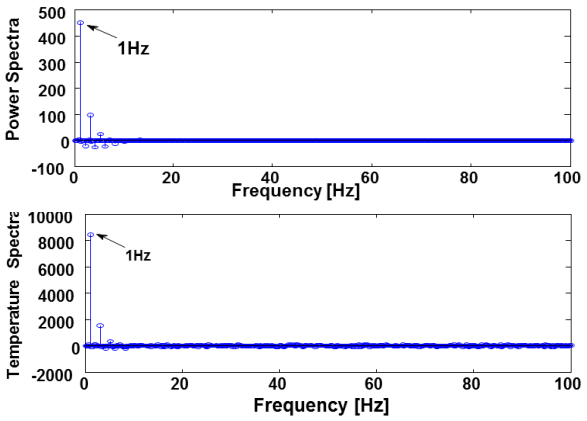


Fig. 11. Measured power and temperature spectral response.

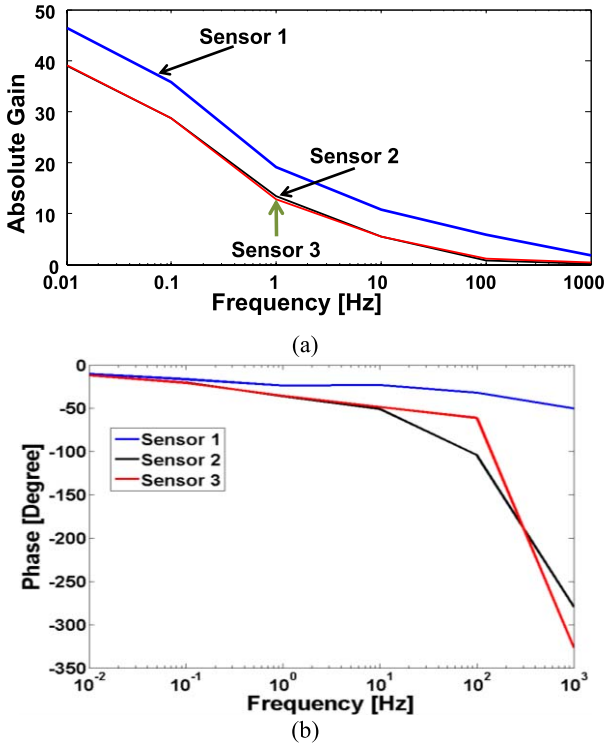


Fig. 12. Extracted thermal filters. (a) Absolute gain. (b) Phase.

is located at the center of the heater, has higher absolute gain than the other filters. Other two filters, which are ∼350 $\mu$m away from the heater, have very similar frequency responses but each of them has a lower gain than the center sensor (Sensor 1). This is because two filters correspond to two thermal sensors placed at a similar distance from the heater. Fig. 12(b) shows the phase responses of the thermal system that is measured from three sensors.

To evaluate the impact of variation, multiple (10) chips are measured while using the same socket, board, and the control PFET device. Fig. 13 shows the frequency response of the thermal filer for sensor 1 across various chips. We observe the absolute gain in low frequency has more variation than in high frequency. This demonstrates the need for chip-package specific extraction of the thermal filter, specifically, at the low frequency regions. Note that the filter response depends on the physical property of the material system and independent of

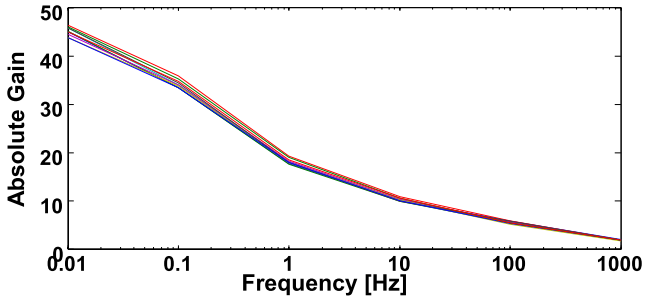


Fig. 13. Extracted thermal filters (for sensor 1) for 10 chips.

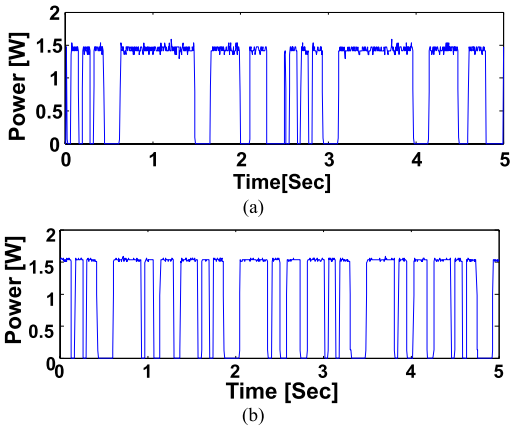


Fig. 14. Applied power patterns. (a) Pattern 1. (b) Pattern 2.

the magnitude of the generated power, floor plan of the chip, and architecture. The latter factors modulate the power profile and hence, temperature profile but not the filter response.

### *B. Temperature Prediction From Arbitrary Power Variation*

The application of the extracted filters in predicting the transient temperature variations for arbitrary power profile is validated. The arbitrary power profiles are generated, which include different frequency components [Fig. 14(a) and (b)]. The power profiles are generated by controlling a sleep transistor signal fed by Keithly 2602B power supply. Sleep transistor signal patterns are programmed through ExpressTSP [17], which is an in-built program in the power supply. While the power is applied to the heater, the temperatures from the three analog temperature sensors are measured. Initailly, the frequency spectrum of the measured power is multiplied to the frequency response of the thermal filter to obtain the frequency spectrum of the predicted temperature. Next, the predicted temperature spectrum is converted to time domain using IFFTs to obtain the transient temperature variation. Fig. 15 shows predicted temperature waveforms for three different sensors with two different power patterns. Sensor 1 shows higher temperature than other two sensors for two power patterns because of spatial effects. For all different power patterns the estimation error is close to ∼2 °C. The average temperature tends to be overestimated because of the noise effect in the temperature sensor.

### *C. Application of the Digital Sensor*

The digital sensor is implemented with inverter-based ROs and a counter as explained in Section III-C. Temperature of

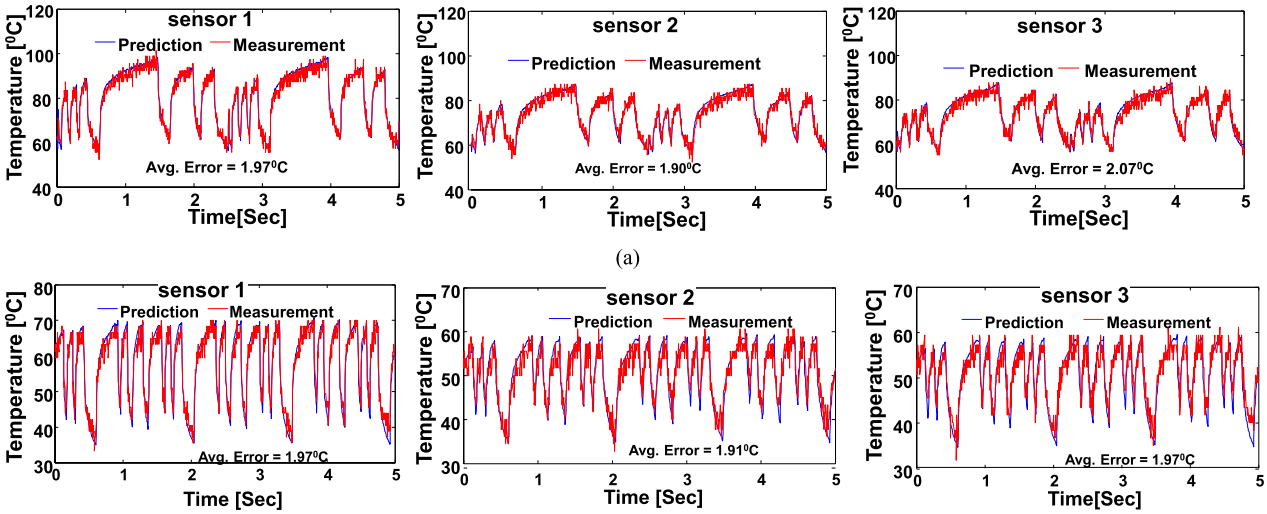


Fig. 15. Accuracy of the temperature prediction at different sensor locations. (a) Power pattern 1. (b) Power pattern 2.

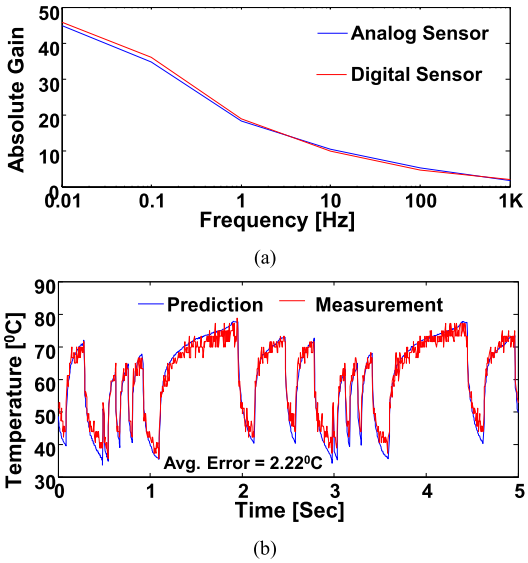


Fig. 16. Application of digital sensor. (a) Comparison of the extracted thermal filters. (b) Estimated temperature.

the digital sensor is estimated based on the frequency of the RO. The thermal filter extracted from the digital sensor closely matches the filter extracted from the analog sensor as shown in Fig. 16(a).Therefore, the temperature can be predicted with very good accuracy using the thermal filter extracted from the digital sensor [Fig. 16(b)]. Therefore, the proposed TSI-based temperature prediction methods can be used with both digital and analog sensors. The delay chain based digital temperature sensors are much simpler to integrate with logic circuits and provide a lower cost and lower power (compared with analog sensors) approach to temperature measurements. The demonstration of the similar accuracy of the TSI with digital and analog temperature sensors shows the applicability of the TSI in fine-grain temperature prediction in microprocessors.

## V. Applications to Many-Core Processor

In this section, the TSI-based approach is applied to the post-silicon thermal prediction of future many-core processor.

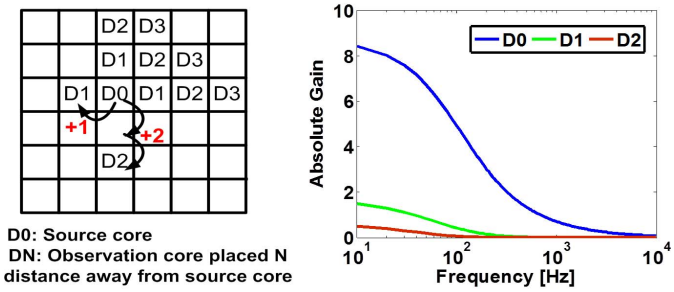


Fig. 17. Low-pass filter behavior of self (D1) and cross (D1, D2) transfer function in multicore system.

The thermal property of a many-core chip is modeled as a multiple input, multiple output (MIMO) system, where the temperature of an observation point is affected by multiple input power sources. Since distributed-RC network is a linear system, superposition principle can be applied to the thermal system, i.e., the temperature at one location is the additive response of all power sources in the system. Assume that there are $M$ power sources organized into $m \times m$ 2-D grid. We further assume that there are $L$ numbers of observation points organized in $l \times l$ grid. The temperature at the observation point $(i, j)$ in frequency domain is estimated as follows:

$$\begin{aligned} T_{ij}(k) &= P_{11}(k)H_{11\to ij}(k) + P_{12}(k)H_{12\to ij}(k) + \cdots \\ &\quad + P_{mm}(k)H_{mm\to ij}(k) \\ &= \sum_{p,q=1}^{m} P_{pq}(k)H_{pq\to ij}(k). \end{aligned} \tag{6}$$

Note that $H_{ij\to ij}(\omega)$ is the self-transfer function and $H_{pq\to ij}(\omega)$ $(\forall p, q \neq i, j)$ is the cross-transfer function defined in Section II. The above formulation leads to the 2-D filter matrix for the MIMO system, as follows:

$$\begin{pmatrix} T_{11}(\omega) \\ \vdots \\ T_{ll}(\omega) \end{pmatrix} = \begin{bmatrix} H_{11\to 11}(\omega) & \cdots & H_{mm\to 11}(\omega) \\ \vdots & \vdots & \vdots \\ H_{11\to ll}(\omega) & \cdots & H_{mm\to ll}(\omega) \end{bmatrix} \begin{pmatrix} P_{11}(\omega) \\ \vdots \\ P_{mm}(\omega) \end{pmatrix}. \tag{7}$$

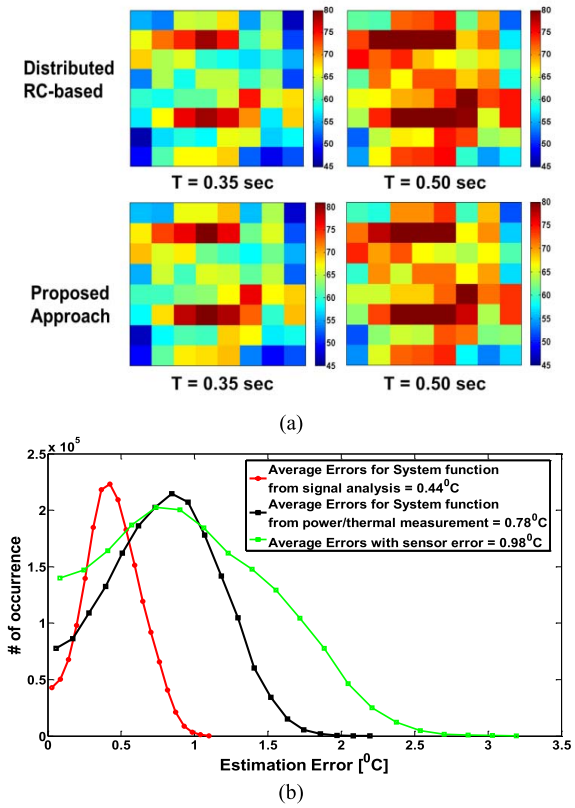


Fig. 18. Accuracy of the proposed approach. (a) Estimation error in instances of the spatial thermal field at different time points. (b) Core-level error statistics considering 64 cores and 60 random workloads (500 ms of real-time simulation).

For the MIMO system, each filter function [and $H_{pq\rightarrow ij}(\omega)$ ($\forall p, q$ and $\forall i, j$)] can be extracted following the methods proposed in (5) by turning off all but one core at a time and controlling the core gating signal for that core.

The application of the proposed TSI to future multicore system is verified using distributed-RC-based thermal simulation platform [19]. We consider distributed RC-based 3-D model of the thermal system including chip, TIM, heat spreader, and heat sink (Fig. 2). The circuit simulator, HSPICE, is used for solving the distributed-RC grid in the time domain. The power profiles are applied as current sources. The chip is modeled as a homogenous 64-core processor with private caches designed in predictive 22-nm technology (total chip area is 400 mm$^2$, each core and private cache is $\sim$6.25 mm$^2$). Each core is modeled as close to Intel Nehalem architecture [18] running at 3.0 GHz. Fig. 17 shows the thermal filter for different locations of interest when a power source is applied at core D0. The gain at the observation point continues to decrease in all frequency as it moves away from the source.

To verify the application of the MIMO-based TSI, power traces of SPEC 2006 benchmark suites are generated using cycle-level architecture timing simulation (Zesto [19]) and power (McPAT [20]) considering x86 architecture [21], [22]. Each benchmark is run or repeated for 0.5 s of real time. We first create several (60) workloads by randomly assigning the power traces of different applications (0.5 s of real time data) to different cores and use them for thermal analysis. The same patterns are also run through the baseline distributed-RC-based thermal simulator. Fig. 18(a) shows examples of spatial thermal maps at different time instants generated from distributed-RC-based simulation and the TSI approach, both transient and spatial temperature variation are accurately predicted using TSI. Fig. 18(b) shows the estimation error considering temperatures of all cores, all time points, and all random patterns. We observe that average error is $<$1 °C between detailed RC-based thermal simulation (SPICE) and the proposed TSI-based model. We study the effect of sensor errors (assuming uniform random error in the range of $-3$ °C to $+3$ °C) during TSI on temperature prediction. As shown in the Fig. 18(b) even with the random sensor errors, the average prediction error remains within $\pm 1$ °C but the spread in the error distribution increases marginally.

The ability of TSI in predicting the effect of variations in process corners and thermal conductivity is studied. Normal random die-to-die variation of $V_{\text{th}}$ is considered for analysis. Each $V_{\text{th}}$ point generated from this normal distribution represents a unique die, for each of such die we consider three different thermal conductivities. TSI is used to extract the unique thermal system for all of these die/package conditions. The low-$V_t$ die and lower conductivity material (TIM/spreader) increase the gain in the low-frequency range of the filter transfer function [Fig. 19(a) and (b)]. Next the same power pattern is applied to all such unique thermal systems. Fig. 19(c) shows the die-to-die distribution of the peak temperature (considering the entire workload and all cores). The leakage variation induces a significant spread in operating temperature of different dies for the same workload and the effect is much stronger when thermal conductivity is low. Note that the presilicon methods would have predicted the same maximum temperature for all such instances.

## VI. Discussion

Finite difference, finite element, and finite volume methods are used for accurate, fine-grain, steady-state, and transient thermal analysis [5]. Less accurate simulators use distributed-RC grid for thermal modeling using circuit solvers like SPICE or specialized tools such as Hotspot [29]. Several time-domain methods were proposed in recent years for fast transient temperature simulations [25], [26], and fast spatiotemporal analysis considering multilayers of power and materials (e.g., ThermalScope [27]). Frequency-domain methods are also proposed for fast computation of the steady-state spatial thermal map [7], [23], [24].

The primary advantage of TSI over the above approaches for post-silicon thermal analysis is that the TSI performs temperature prediction using the transfer function extracted from the full thermal system (i.e., stacks of heat sink, spreader, TIM, and chip), instead of computing thermal resistance and capacitances of individual materials in isolation. Therefore, the effects of any nonuniformity and/or uncertainty in the thermal properties of the materials are captured in the extracted transfer function. As the leakage temperature interaction is considered as a part of the extracted thermal system, the effect of process variation of individual chips is also incorporated. Therefore, TSI can accurately predict the temperature variations for

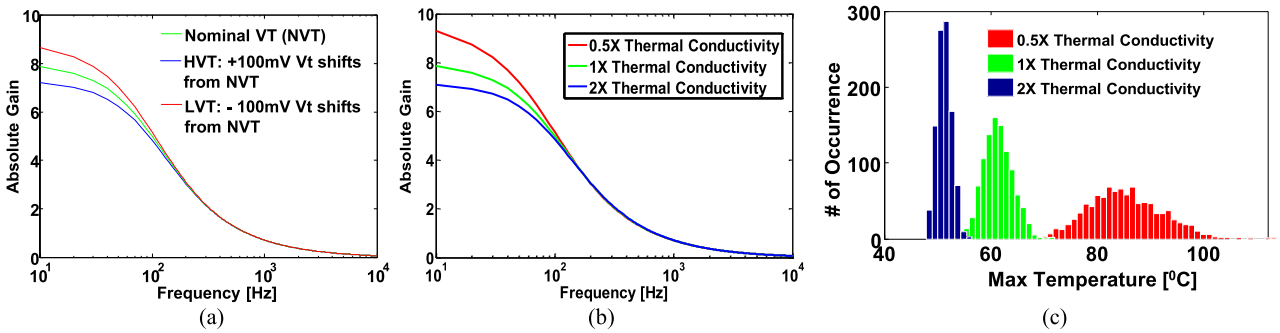


Fig. 19. Effect of post-silicon thermal characterization with chip-to-chip variation in leakage and thermal conductivity. (a) Leakage-temperature interactions. (b) Material properties on the extracted thermal filters. (c) Statistical variation in peak temperature of the 64-core processor for same dynamic power pattern considering process and material variations.

different dies/packages while running the same workload and dissipating same dynamic power.

In the existing methods, leakage power is updated in each time-step based on the current thermal map to model the leakage–temperature interactions [27]–[30]; hence, higher estimation accuracy requires finer time-step and longer simulation time. TSI incorporates the leakage–temperature interaction in the thermal filter, which needs to be extracted only once for a chip and package. Once the filter is extracted, the temperature for varying power pattern is estimated in the frequency domain. Hence, estimation accuracy considering leakage–temperature interaction is less sensitive to time-step and the simulation time can be lower than the distributed-RC-based simulations. In the time-domain methods the temperature of one node depends on the temperature of neighboring nodes. In the TSI approach, this interdependence is precharacterized into the filter responses. Hence, the temperature of each node can be computed independently. Exploiting this property, the TSI can be parallelized to reduce simulation time. The temperature of one core can be computed without computing temperature of other cores but still accounting for the power of all cores. This provides opportunity for localized thermal analysis.

The TSI approach is limited by the accuracy of the temperature and power sensors as shown in Fig. 19(b). Interestingly, by recognizing the thermal system is a low-pass filter, TSI can also be used to eliminate the effect of high-frequency noise in the sensors as shown in Fig. 15. The granularity of the extracted transfer function in the frequency domain also impacts the accuracy. A finer granularity of the frequency while extracting the transfer function improves estimation accuracy. There are also practical limitations on fine-grain control of the frequency components of the generated power spectra. Fine-grain analysis also increases the memory requirement. Storing a limited number of coefficients can be used to reduce the memory requirement.

Theoretically TSI can be scaled to arbitrarily fine spatial grids; however, as the number of filters is proportional to the number of grids, TSI may not be a suitable approach for microscale design time thermal analysis. For regular grids and spatially homogeneous material properties, the required transfer functions can be reduced. On the other hand, run-time thermal management is performed based on only a limited number of on-chip sensors. TSI creates a filter matrix connecting the power control points to the sensor locations. Any arbitrary number of sensors and nonuniform sensor placement can be handled using TSI. Therefore, TSI is most useful for run-time macroscale transient thermal analysis.

## VII. Conclusion

This paper, presented TSI—a methodology for post-silicon thermal prediction. The proposed method identified the frequency-domain response of the thermal system of a packaged die and used the extracted filter for chip-specific analysis of transient thermal field. A test-chip successfully demonstrated the effectiveness of TSI. The post-silicon characterization can benefit thermal design and management at chip/system level. Future high core count processors will be asymmetric and heterogeneous and will host time-varying workloads. Thermal limits will drive power management and proactive/predictive thermal management is becoming increasingly desirable. The post-silicon thermal analysis using TSI opens up innovative avenues for creating application spaces that were matched to the thermal properties of the die. For example, the post-silicon characterization of a multicore chip can be used by operating systems to schedule workloads because the identification of the chip thermal system enables schedulers to reason about the thermal consequences of scheduling a specific workload on a target chip. The integration of TSI with innovative thermal management is an important future work.

## References

[1] S. Borkar, "Thousand core chips: A technology perspective," in *Proc. IEEE/ACM Design Autom. Conf.*, Jun. 2007, pp. 746–749.

[2] D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," in *Proc. IEEE Int. Symp. High Perform. Comput. Archit.*, Jan. 2001, pp. 171–182.

[3] T. R. Conrad and R. L. Shook, "Impact of moisture/reflow induced delaminations on integrated circuit thermal performance," in *Proc. Electron. Compon. Technol. Conf.*, 1994, pp. 527–531.

[4] M. Cho, N. Sathe, A. Raychowdhury, and S. Mukhopadhyay, "Optimization of burn-in test for many-core processors through adaptive spatiotemporal power migration," in *Proc. IEEE ITC*, Nov. 2010, pp. 1–9.

[5] Y. Cheng, C. Tsai, C. Teng, and S. Kang, *Electrothermal Analysis of VLSI Systems*. New York, NY, USA: Springer-Verlag, 2000.

[6] Y. Zhan and S. Sapatnekar, "High-efficiency Green function-based thermal simulation algorithms," *IEEE Trans. Comput. Aided Design Integr. Circuits Syst.*, vol. 26, no. 9, pp. 1661–1675, Sep. 2007.

[7] R. Cochran and S. Reda, "Spectral techniques for high-resolution thermal characterization with limited sensor data," in *Proc. IEEE 46th ACM Design Autom. Conf.*, Jul. 2009, pp. 478–483.

[8] *ASTM: Standard Test Method for Thermal Transmission Properties of Thermally Conductive Electrical Insulation Materials*, International Standard ASTM D5470-06, 2006.

[9] A. Poppe and V. Szekely, "Dynamic temperature measurement: Tools providing a look into packaging and mount structures," *Electron. Cooling*, vol. 18, no. 1, pp. 1–3, May 2002.

[10] S. Smith and R. C. Campbell, "Flash diffusivity method: A survey of capabilities," *Electron. Coolings*, vol. 8, no. 2, pp. 1–32, May 2002.

[11] S. Kim and R. L. Webb, "Analysis of convective thermal resistance in ducted fan-heat sinks," *IEEE Trans. Compon. Packag. Technol.*, vol. 29, no. 3, pp. 439–448, Sep. 2006.

[12] K. Kurabayashi and E. Goodson, "Precision measurement and mapping of die-attach thermal resistance," *IEEE Trans. Compon. Packag. Manuf. Technol. Part A*, vol. 21, no. 3, pp. 506–514, Sep. 1998.

[13] M. Parameswaran, A. M. Robinson, D. L. Blackburn, M. Gaitan, and J. Geist, "Micromachined thermal radiation emitter from a commercial CMOS process," *IEEE Electron Device Lett.*, vol. 12, no. 2, pp. 57–59, Feb. 1991.

[14] B. Sinharoy, R. N. Kalla, J. M. Tendler, R. J. Eickemeyer, and J. B. Joyner, "POWER5 system microarchitecture," *IBM J. Res. Develop.*, vol. 49, nos. 4–5, pp. 505–521, Jul. 2005.

[15] M. Ware, K. Rajamani, M. Floyd, B. Brock, J. Rubio, F. Rawson, and J. Carter, "Architecting for power management: The IBM POWER7 approach," in *Proc. IEEE Int. Symp. HPCA*, Jan. 2010, pp. 1–11.

[16] R. McGowen, C. A. Poirier, C. Bostak, J. Ignowski, M. Millican, W. H. Parks, and S. Naffziger, "Power and temperature control on a 90-nm Itanium family processor," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 229–237, Jan. 2006.

[17] (2012). *Keithley Test Script Processor Manual* [Online]. Available: http://www.keithley.com/products /dcac/currentvoltage/gpmp-mn=2602A

[18] *Intel 64 and IA-32 Architecture Optimization Reference Manual*, Intel Corp., Santa Clara, CA, USA, Nov. 2009, pp. 49–61.

[19] G. Loh, S. Subramaniam, and Y. Xie, "Zesto: A cycle-level simulator for highly detailed microarchitecture exploration," in *Proc. IEEE Int. Symp. Perform. Anal. Syst. Softw.*, Apr. 2009, pp. 53–64.

[20] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. IEEE/ACM Int. Symp. Microarchit.*, Dec. 2009, pp. 469–480.

[21] W. Song, M. Cho, S. Yalamanchili, S. Mukhopadhyay, and A. Rodrigues, "Simulation infrastructure for power and temperature modeling in many-core processors and linear system modeling," in *Proc. SRC TECHCON*, Sep. 2011, pp. 1–4.

[22] M. Cho, W. Song, S. Yalamanchili, and S. Mukhopadhyay, "Thermal system identification (TSI): A methodology for post-silicon characterization and prediction of the transient thermal field in multicore chips," in *Proc. IEEE SEMI-THERM*, Mar. 2012, pp. 118–124.

[23] T. Kemper, Y. Zhang, Z. Bian, and A. Shakouri, "Ultrafast temperature profile calculation in IC chips," in *Proc. THERMINIC*, Sep. 2006, pp. 133–137.

[24] A. Nowroz, R. Cochran, and S. Reda, "Thermal monitoring of real processors: Techniques for sensor allocation and full characterization," in *Proc. Design Autom. Conf.*, 2010, pp. 56–61.

[25] Y. K. Cheng and S. M. Kang, "An efficient method for hotspot identification in ULSI circuits," in *Proc. IEEE Int. Conf. Comput. Aided Design*, Nov. 1999, pp. 124–127.

[26] D. Schweitzer, "A fast algorithm for thermal transient multisource simulation using interpolated Zth functions," *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 32, no. 2, pp. 478–483, Jun. 2009.

[27] N. Allec, Z. Hassan, S. Li, R. P. Dick, and R. Yang, "ThermalScope: Multi-scale thermal analysis for nanometer-scale integrated circuits," in *Proc. IEEE Int. Conf. Comput. Aided Design*, Nov. 2008, pp. 603–610.

[28] H. Wang, D. Li, S. Tan, and A. Gupta, "Composable thermal modeling and characterization for fast temperature estimation," in *Proc. Electr. Perform. EPEPS*, Oct. 2010, pp. 185–188.

[29] W. Huang, K. Sankaranarayanan, R. J. Ribando, M. R. Stan, and K. Skadron, "Accurate, Pre-RTL temperature-aware processor design using a parameterized, geometric thermal model," *IEEE Trans. Comput.*, vol. 57, no. 9, pp. 1277–1288, Sep. 2008.

[30] P. Zhou, Y. Ma, Q. Zhou, and X. Hong, "Thermal effects with leakage power considered in 2D/3D floorplanning," in *Proc. IEEE Int. Conf. Comput. Aided Design Comput. Graph.*, Oct. 2007, pp. 338–343.

**Minki Cho** (S'08–M'13) received the B.E. degree in electronics engineering from Sogang University, Seoul, Korea, in 2006, and the M.S. and Ph.D. degrees from the Georgia Institute of Technology, Atlanta, GA, USA, in 2009 and 2012, respectively, both in electrical and computer engineering.

He is currently with Intel Circuit Research Laboratory, as a Research Scientist. In 2011, he interned with the Intel Circuits Research Laboratory. His current research interests include low-power digital circuit design and reliability of digital circuit in nanometer nodes.

**Khondker Zakir Ahmed** (S'12) received the B.Sc. degree in electrical and electronic engineering from the Bangladesh University of Engineering and Technology, Dhaka, Bangladesh, in 2004. He is currently pursuing the Ph.D. degree in electrical and computer engineering with the Georgia Institute of Technology, Atlanta, GA, USA.

He was a Summer Intern with Circuit Research Laboratory, Intel, Hillsboro, OR, USA, in 2013, where he worked on integrated voltage regulators. His current research interests include power management circuits for low voltage energy harvesting systems, and integrated voltage regulators for digital loads.

**William J. Song** received the B.S. degree from the School of Electrical Engineering, Yonsei University, Seoul, Korea. He is currently pursuing the Ph.D. degree with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA.

He is currently an IBM/SRC Graduate Fellow. He had held internships with Sandia National Laboratories, Albuquerque, NM, USA, and AMD, Bellevue, WA, USA.

**Sudhakar Yalamanchili** (S'79–M'82–SM'91) received the Ph.D. degree in electrical and computer engineering from the University of Texas at Austin, Austin, TX, USA.

He became a Faculty Member of electrical and computer engineering with the Georgia Institute of Technology (Georgia Tech), Atlanta, GA, USA, in 1989, where he is currently a Joseph M. Pettit Professor of computer engineering. Since 2003, he has been a Co-Director with the NSF Industry University Cooperative Research Center on Experimental Computer Systems, Georgia Tech.

Dr. Yalamanchili was the General Co-Chair of the 2010 IEEE/ACM International Symposium on Microarchitecture and the Program Committees for the 2011 International Symposium on Networks on Chip, the IEEE/ACM International Symposium on Microarchitecture, and the Micro Top Picks from Computer Architecture Conferences in 2011.

**Saibal Mukhopadhyay** (S'99–M'07–SM'11) received the Ph.D. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, USA, in 2006.

He is currently an Associate Professor with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA. His current research interests include analysis and design of low-power and robust circuits in nanometer technologies and 3-D circuits and systems.

Dr. Mukhopadhyay was a recipient of the IBM Faculty Partnership Award in 2009 and 2010, the National Science Foundation CAREER Award in 2011, and the Office of Naval Research Young Investigator Award in 2012.