INVITED PAPER

Thermal Modeling, Analysis, and Management in VLSI Circuits: Principles and Methods

Maximum chip performance under peak permissible temperature limits may be achieved with the help of combined electrical and thermal simulation of VLSI circuits.

By Massoud Pedram, *Fellow IEEE*, and Shahin Nazarian

ABSTRACT | The growing packing density and power consumption of very large scale integration (VLSI) circuits have made thermal effects one of the most important concerns of VLSI designers. The increasing variability of key process parameters in nanometer CMOS technologies has resulted in larger impact of the substrate and metal line temperatures on the reliability and performance of the devices and interconnections. Recent data shows that more than 50% of all integrated circuit failures are related to thermal issues. This paper presents a brief discussion of key sources of power dissipation and their temperature relation in CMOS VLSI circuits, and techniques for full-chip temperature calculation with special attention to its implications on the design of high-performance, low-power VLSI circuits. The paper is concluded with an overview of techniques to improve the full-chip thermal integrity by means of off-chip versus on-chip and static versus adaptive methods.

KEYWORDS | Dynamic power; hot spots; leakage power; on-chip temperature; thermal gradient

I. INTRODUCTION

"Smaller and faster" are the chief demands driving today's electronic designs because they generally mean higher performance. However, they also translate into high power densities, higher operating temperatures, and reduced reliability. Furthermore, local hot spots, which have much higher temperatures compared to the average die temperature, are becoming more prevalent in very large scale integration (VLSI) circuits.

Elevated temperatures are a major contributor to lower semiconductor reliability. If heat is not removed at a rate equal to or greater than its rate of generation, junction temperatures will rise. Higher junction temperatures reduce mean time to failure (MTTF) for the devices. Device reliability has a direct impact on the overall system reliability. Removing heat from these devices is thus a major task facing design engineers of modern electronic systems concerned with improving reliability. Understanding the effect of heat on the reliability of electronic products and the integrity of manufacturing processes is critical if problems are to be avoided. This means the need to understand thermal management techniques and the need for comprehensive data has never been greater. With passive cooling methods, the chip temperature is determined by the efficacy of heat transfer out of the device to the ambient. Equilibrium is achieved when the heat generation rate matches the heat transfer rate. The key mechanisms are thermal conduction, thermal convection, and thermal radiation. Of the three methods of heat transfer, radiation is the simplest: it simply needs a large area with good emissivity to transfer a large amount of heat to the surroundings. Conduction and radiation can be implemented with a fully passive heat transfer system, whereas convection is an active method that requires design overhead.

With component packages becoming more compact and having smaller physical profiles, it is no longer sufficient to merely add "a bigger fan" as a downstream fix for thermal problems. Because heat conduction is playing a bigger role while heat convection is playing a lesser role in removing heat, thermal management is best accomplished when it is incorporated starting at the beginning of

Manuscript received August 5, 2005; revised January 29, 2006.
The authors are with the Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089 USA (e-mail: pedram@usc.edu; shahin@usc.edu).
Digital Object Identifier: 10.1109/JPROC.2006.879797

the design cycle. Heat flow must be planned and thermal resistances minimized. In addition, although worst case heating conditions seldom arise in a circuit during its lifetime, when they do arise, they can cause significant problems, ranging from circuit transient timing errors to complete catastrophic burnout. A package designed for the worst case is excessive. To reduce packaging cost without unnecessarily limiting performance, the package should be designed for the worst typical application. Any application that generates more heat than this cheaper package can handle should engage an alternative runtime thermal-management technique. Since typical high-power applications still operate 20% or more below the worst case [1], this can lead to dramatic savings. This is the philosophy behind the thermal design of the Intel Pentium 4. At the same time, the heat flux, or heat load per unit area, for state of the art microprocessors is currently at 10–15 W/cm^2, which is fast exceeding the limit of air cooling.

Temperature difference in the die is especially important for sensitive analog circuits where such differences can easily cause mismatches between signal levels and bias currents, thus degrading the performance of the analog circuit and reducing the noise margins.

Emerging circuit fabrics, such as vertically integrated [three-dimensional (3-D)] integrated circuits (ICs), are significantly impacted by thermal effects. The 3-D architectures, which provide multiple layers of active devices together with high-density local interconnects, offer unique advantages both in terms of density and circuit performance [2], [3]. However, the power density and temperature of these architectures can be quite high. Thermal management is thus of critical importance for the 3-D designs [4], [5].

High on-chip temperatures can give rise to timing failures and reliability concerns. In fact, many of the electronic circuit failures are caused by or related to elevated temperatures, sudden spatial or temporal temperature variations, and presence of hot spots. Temperature variations across a VLSI chip can result in significant timing uncertainty, prompting wider timing margins, and thus, lower circuit performance. Yet another consideration is that the on-chip temperature gradient (the difference in temperatures at different parts of the chip, which are in turn caused by uneven power dissipation, can produce mechanical stress, which may degrade the chip reliability.

Leakage power consumption is known to be highly dependent on the on-chip temperature profile; that is, higher temperature results in larger power dissipation, which in turn increases the on-chip temperatures. This can result in thermal runaway condition. Consequently, power reduction and management interact with thermal effect analysis and control and vice versa.

Dynamic and leakage power are the two main sources of power consumption in VLSI circuits. In many new high-performance designs, the leakage component is comparable to the switching component. Reports indicate that 40% or even a higher percentage of the total power consumption in 90-nm-process technology is due to the leakage of transistors [6]. This percentage is expected to increase with technology scaling. Simulation results in [7] predict that the transistor off-state current per micrometer of transistor width increases by a factor of three to five per generation. As will be shown in Section IV, for a given package, the die temperature can be modeled as a linear function of the total power dissipation of the circuit. At the same time, the leakage power increases exponentially with temperature. These facts clearly motivate the need for leakage power reduction techniques in existing designs.

Fig. 1 illustrates the significant increase in leakage power of a 15-mm die fabricated in a 100-nm technology with a supply voltage of 0.7 V as a function of substrate temperature. If the thermal conductance of the package is not large enough, this exponential dependence will cause thermal runaway where the die temperature increases unbounded and the chip fails [8]. Even when thermal runaway does not occur, the operating temperature of the chip can become larger than the designed value, which will either increase the package cost or degrade the performance as well as the long-term reliability of the chip [9].

This paper focuses on thermal issues and the techniques that deal with them. More precisely, the first part of the paper provides an overview of the major sources of power consumption and their relation to the die temperature. The second part of the paper describes full-chip thermal modeling and electrothermal simulation technique. The third part of the paper focuses on the impact of substrate and interconnect temperatures gate and interconnect delays. The paper is concluded with an overview of dynamic thermal management strategies for microprocessor chips.

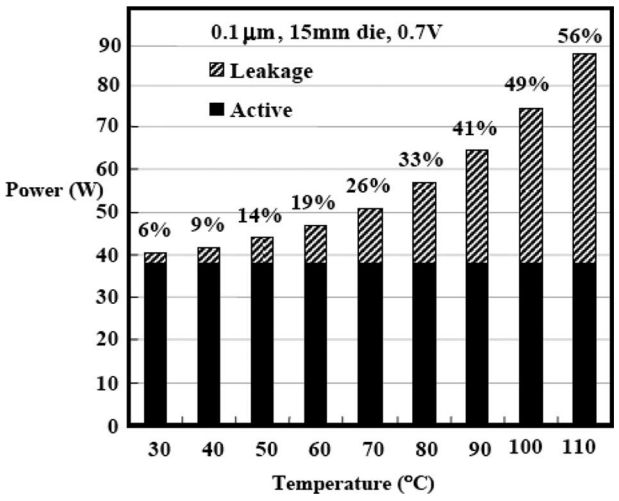

Fig. 1. *Power consumption of a die as a function of temperature. Courtesy of V. De, Intel.*

II. FULL-CHIP TEMPERATURE CALCULATION

Heat is generated in both the substrate and the interconnections. The major source of heat generation is the power dissipation of devices that are embedded in the substrate. Some power dissipation also results from Joule heating (or self-heating) caused by the flow of current in the interconnect network. Although interconnect Joule heating constitutes only a small fraction of the total power dissipation in the chip, the temperature rise in the interconnections due to Joule heating can be significant. This is due to the fact that interconnects are located away from the silicon substrate and the heat sink by several layers of insulating materials which have lower thermal conductivities than that of silicon.

Simply stated, the operating temperature of a VLSI chip can be calculated from the following linear equation:

$$T_{\text{chip}} = T_a + R_\theta \cdot \frac{P_{\text{tot}}}{A} \tag{1}$$

where T_{chip} is the average chip (silicon junction) temperature, T_a is the ambient temperature ($T_a = 25\ ^\circ\text{C}$), P_{tot} (in W) is the total power consumption, A (in cm^2) is the chip area, and R_θ is the equivalent thermal resistance of the substrate (Si) layer plus the package and heat sink ($\text{cm}^2\ ^\circ\text{C/W}$.) As this equation shows, to calculate the chip temperature, one must have calculated power dissipation of the circuit (P_{tot}), constructed the chip thermal model (R_θ), and be given information about the environment (T_a).

The self-heating effect can be analyzed as follows [10]. The metal temperature T_{metal} is given by

$$\begin{aligned} T_{\text{metal}} &= T_{\text{chip}} + \Delta T_{\text{self}} \\ \Delta T_{\text{self}} &= R_E I_{\text{rms}}^2 R_{\theta,\text{self}} \end{aligned} \tag{2}$$

where ΔT_{self} is the temperature rise of the metal interconnect due to the flow of current, R_E is the electrical resistance of interconnect, and $R_{\theta,\text{self}}$ is the thermal impedance of the interconnect line to the substrate.

A. Power Dissipation Sources

Power dissipation in the substrate of a CMOS VLSI circuit can be calculated as

$$P_{\text{total}} = P_{\text{dynamic}} + P_{\text{short-circuit}} + P_{\text{static}} \tag{3}$$

where P_{dynamic} denotes the dynamic power consumption that occurs when the output signal of a CMOS logic cell makes a transition; $P_{\text{short-circuit}}$ represents the power dissipation of the circuits when both n- and p-transistors transistors are simultaneously conducting, creating a direct path between the supplying power and the ground; and P_{static} is the static power dissipation that is caused by the static current drawn from power supply. This component is mainly due to the direct gate current and the subthreshold conduction current, which are collectively referred to as the leakage current. Each component of the power dissipation in a CMOS circuit is discussed in more detail next.

1) Dynamic Power: Dynamic or switching power is due to the signal switching activity at the output of a CMOS logic cell. The dynamic power component dominates during the active mode of the cell operation. This component is expressed as

$$P_{\text{dynamic}} = 0.5 C_{\text{load}} V_{\text{DD}}^2 f \alpha \tag{4}$$

where f is the clock frequency, and α is the expected number of output transitions in a clock period, and C_{load} is the load capacitance (including gate input and interconnect capacitances). During the output signal transition, the output capacitance is charged to V_{DD} or discharged to ground as follows: while charging, half of the energy supplied by V_{DD} is stored in the output capacitance and the other half is dissipated in the pull-up transistors. While discharging the remaining charge is removed from the output capacitance and dissipated in the pull-down transistors.

2) Short-Circuit Power: Short-circuit power ($I_{\text{short-circuit}} \times V_{\text{DD}}$) is due to direct current flow from the power supply to the ground. The short-circuit current $I_{\text{short-circuit}}$ occurs when pull-up and pull-down networks are conducting simultaneously. Short-circuit current is dependent on the duration of the simultaneous on-times of the pull-up and pull-down networks, the transistor sizes, and the supply voltage level. In general, the short-circuit power of a cell is minimized if the output transition time is larger than its input transition time. The derivation of an exact formula for the short-circuit power is a complicated task; however, simple closed-form expressions have been proposed by making simplifying assumptions and/or considering special cases [11]–[14]. The short-circuit current is found to depend on the carrier mobility and threshold voltage of the transistors both of which vary with temperature.

3) Static Power: Although a source of static power manifests itself in circuits that have constant sources of current between the power supplies, leakage currents are the major sources static power dissipation. Although there are

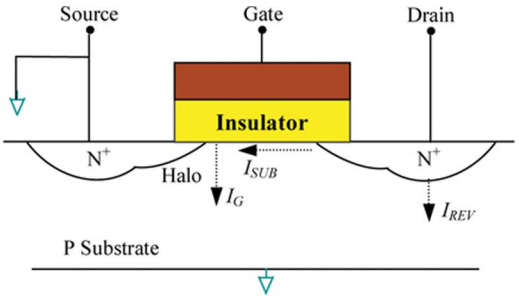

Fig. 2. *Leakage current components in an nMOS transistor.*

sources of leakage current in a CMOS circuit, the three dominant ones are (cf. Fig. 2):

1) reverse-biased junction leakage current (I_{REV});
2) gate direct tunneling leakage (I_G);
3) Subthreshold (weak inversion) leakage (I_{SUB}).

I_{REV} flows from the source or drain to the substrate through the reverse biased diodes when a transistor is off. The magnitude of I_{REV} depends on the area of the drain diffusion and the leakage current density, which is in turn determined by the doping concentration. If both n and p regions are heavily doped, band-to-band tunneling dominates the pn junction leakage [15]. Junction leakage has a rather high temperature dependency as much as 50-100x/100 °C; however, this is generally significant only in circuits designed to operate at high temperatures greater than 150 °C. Junction reverse-bias leakage components from both the source–drain diodes and the well diodes are generally negligible with respect to the other three leakage sources.

Gate direct tunneling leakage I_G flows from the gate through the insulator and to the substrate. In oxide layers thicker than 3–4 nm, I_G is due to the Fowler–Nordheim tunneling of electrons into the conduction band of the oxide layer under a high applied electric field across the oxide layer. In technology node 0.15 μm and lower which have lower oxide thicknesses, direct tunneling through the silicon oxide layer is the leading effect. I_G of a p-transistor is typically one order of magnitude smaller than that of an n-transistor with identical gate oxide thickness T_{ox} when SiO_2 is used as the gate dielectric. The magnitude of I_G increases exponentially with T_{ox} and V_{DD}. For example, for relatively thin oxide thicknesses in the order of 2–3 nm, at $V_{GS} = 1$ V, every 0.2-nm reduction in T_{ox} causes a tenfold increase in I_G [16]. The temperature dependency of I_G is quite weak, i.e., only $\sim 2x/100$ °C.

Subthreshold leakage I_{SUB} is the drain–source current of a transistor operating in the weak inversion region. Unlike the strong inversion region in which the drift current dominates, the subthreshold conduction is due to the diffusion current of the minority carriers in the channel. In current CMOS technologies, I_{SUB} is much larger than the other leakage current components [17]. This is mainly because of the relatively low V_T in modern CMOS devices. According to the BSIM3v3.2 MOSFET transistor model [18], the subthreshold drain current I_{SUB} of a transistor in the normal "off" state $V_{ds} = V_{DD}$ and $V_{gs} = 0$ is expressed by the following equation:

$$I_{sub} = k_{tech}\left(\frac{W}{L}\right)10^{-\frac{V_T}{S}} \tag{5}$$

where k_{tech} is a transistor geometry and CMOS technology dependent parameter, W and L denote the transistor width and length, V_T denotes the threshold voltage of the device and S, which is called the subthreshold swing parameter, is equal to the subthreshold voltage decrease required to increase I_{sub} by a factor of ten. In fact, $S = 2.3nk_BT/q$ where $n \geq 1$ is a device-dependent parameter, k_B is the Boltzmann's constant, T denotes absolute temperature in degrees Kelvin, and q is the electron charge.

It is desirable to have as small an S value as possible, since this is the parameter that determines the amount of voltage swing necessary to switch a MOSFET from off to the on state (Typical values of S for bulk CMOS devices are 70–90 mV/decade). To minimize S, the thinnest possible gate oxide to increase C_{ox} and the lowest possible doping concentration in the channel to decrease C_{dep} should be used. Higher temperatures increase S, which in turn increase the off leakage current.

I_{SUB} is a function of temperature, threshold voltage, device size, and the process parameters out of which the threshold voltage (V_T) is dominant. The subthreshold leakage current increases rapidly with temperature. This is shown in Fig. 3 which illustrates the leakage current versus temperature for several technology nodes. I_{SUB} has a temperature sensitivity of 8-12x/100 °C. The data also confirms that the leakage power increases as the technology moves forward.

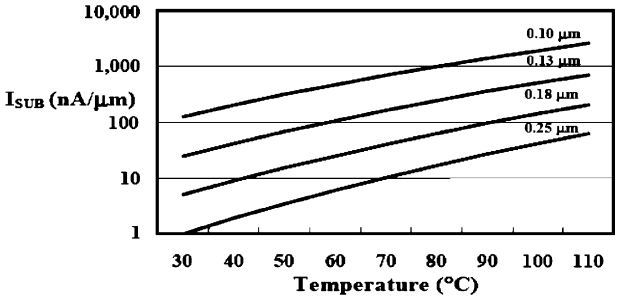

Fig. 3. I_{SUB} ($V_{GS} = 0$) *trend as a function of temperature. Courtesy of Vivek De, Intel.*

It is seen that each component of power consumption is a function of temperature. Reduction of the supply voltage reduces the chip total power consumption, which in turn reduces the chip temperature. As the chip temperature is reduced, the leakage power is reduced dramatically.

B. Full-Chip Thermal Modeling

Key to successful thermal management is the ability to obtain comprehensive and accurate temperature data under as realistic operating conditions. The commonly used method of gathering this temperature data by using point contact methods (thermocouples) is limited by the large number of points to be monitored and the small size of the components. Connecting tens or hundreds of thermocouples is very time consuming. Infrared (IR) thermal imaging is a new technique which addresses these issues by providing comprehensive two-dimensional (2-D) maps of thousands of temperatures in a matter of seconds. This is accomplished without the need to make contact with the components. This approach is, however, expensive and time-consuming and can only be applied postdesign. It is thus important to have full-chip thermal models and simulation tools that can provide the temperature profile of the die.

The heat diffusion equation is in general used to describe the heat conduction in a chip and calculate the temperature profile [19]

$$\rho c_p \frac{\partial T(\vec{r},t)}{\partial t} = \nabla \cdot [k(\vec{r},T)\nabla T(\vec{r},t)] + g(\vec{r},t) \qquad (6)$$

which is subject to the general thermal convection boundary condition

$$k(\vec{r},T)\frac{\partial T(\vec{r},t)}{\partial n_i} = h_i(T_a - T(\vec{r},t)). \qquad (7)$$

In the above equations, T is the temperature (°C), k is the thermal conductivity [W/(m°C)], ρ is the density of material (Kg/m^3), c_p is the specific heat [J/(Kg°C)], g is the power density of the heat sources (W/m^3), and h_i is the heat transfer coefficient in the direction of heat flow $\vec{i}$ on the boundary surface of the chip [W/(m^2°C)]. Note that $h_i = 1/(A_i R_{\theta,i})$ where A_i is the effective area normal to $\vec{i}$ and $R_{\theta,i}$ denotes the equivalent thermal resistance. $\partial/\partial n_i$ is the differentiation operator along the outward direction normal to the boundary surface and T_a is the ambient temperature.

The term $\nabla \cdot [k(\vec{r},T)\nabla T(\vec{r},t)]$ in (6) can be replaced by $k(T)\nabla^2 T(\vec{r},t)$ for homogeneous materials, resulting in a second-order parabolic partial differential equation

$$\rho c_p \frac{\partial T(\vec{r},t)}{\partial t} = k(T)\left(\frac{\partial^2 T(\vec{r},t)}{\partial x^2} + \frac{\partial^2 T(\vec{r},t)}{\partial y^2} + \frac{\partial^2 T(\vec{r},t)}{\partial y^2}\right) + g(\vec{r},t). \qquad (8)$$

The heat flow described by this differential equation has a similar form to that for electrical current, and there is a well-known duality between them. The heat flow (W) passing through a thermal resistor (°C/W) is equivalent to the electrical current (ampere) through an electrical resistance (ohm), and the temperature difference (°C) corresponds to voltage difference (volt). There is also the thermal equivalent capacitance (J/°C) where the heat is absorbed for the electrical capacitance (farad). More precisely, $C_\theta = \rho c_p \Delta H$, which is the thermal capacitance, is modeled by an electrical capacitance whereas $g\Delta H$, which is the heat flow coming from power generated by logic cells in control volume $\Delta H = \Delta x \Delta y \Delta z$, is modeled by an electrical current source i_p. Various thermal resistances, which are defined in x, y, and z directions with values inverse-linearly proportional to k, and distances are replaced by the corresponding electrical resistances R_E. Ambient temperature is expressed using an independent voltage source v_0. Node temperatures will then correspond to node voltages in the electrical network constructed in this way (cf. Fig. 4.)

Now consider the general case of multiple heat sources in the substrate connected via thermal resistances to each other and to the ambient environment as depicted in Fig. 5. Here, node n_i represents a circuit block (logic cell or groups of logic cells depending on the granularity of the thermal model.) The power consumption of each circuit block is represented as a current source i_k associated with the corresponding node. Between neighboring nodes, a thermal resistance r_{ij} is added to model the lateral heat conduction path. Thermal resistances are also added between nodes and the ambient voltage terminal r_{i0} to capture the vertical component of the thermal resistance between the circuit block and the ambience (includes the effects of substrate, package, and heat sink).

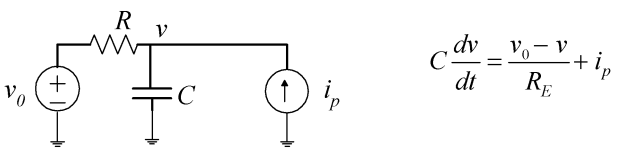

Fig. 4. *Simple electrical model of heat flux and temperature with the corresponding differential equation for a single heat source on the chip.*

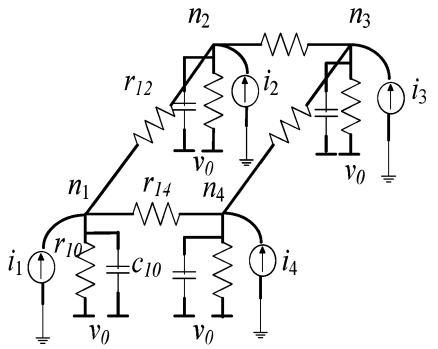

Fig. 5. *Equivalent circuit model to temperature distribution accounting for both lateral and vertical heat flux and heat absorption.*

A thermal capacitor c_{io} at each node is included to model the heat absorption and storage in the substrate and thereby derive the chip temperature evolution over time. The ambient temperature is modeled as an independent fixed voltage source v_0.

One may use dc (or ac) analysis functions available in SPICE-type circuit simulators to calculate chip temperature [20], [21]. Notice that it is critical to consider the effect of metal interconnect on heat distribution, since the metal interconnects tend to provide a low thermal impedance path for heat flux among various parts of the substrate.

The one-dimensional (1-D) heat diffusion equation in metal interconnection under the steady state can be written as

$$\frac{\partial^2 T(x)}{\partial x^2} = -\frac{g(x,T)}{k_m} \tag{9}$$

where $g(x,T)$ is the temperature-dependent power density of heat sources $(\mathrm{W/m^3})$ at x and k_m is the thermal conductivity of the material $[\mathrm{W/(m^\circ C)}]$. For the interconnects, $g(x,T) = J_{\mathrm{rms}}^2 \rho(x,T)$, where J_{rms} is the rms current density $(\mathrm{A/m^2})$ and $\rho(x,T)$ is the temperature-dependent metal resistivity at x $(\Omega\mathrm{m})$.

Consider an interconnect line of length l, width w, thickness t_m, and thermal conductivity k_m that passes over the silicon substrate with an insulator of thickness t_{ins} and an insulator thermal conductivity k_{ins} in between. The interconnect line is connected to the substrate by vias/contacts at its two ends. (cf. Fig. 6.) Power dissipation P_{gen} in a metal segment of length $\Delta(x)$ at coordinate x can be expressed as

$$P_{\mathrm{gen}}(x) = \Delta R_E(x,T) I_{\mathrm{rms}}^2$$

where I_{rms} is the root mean squared current passing through the line, and $\Delta R_E(x)$ is the interconnect electrical resistance of metal segment $\Delta(x)$ located at coordinate x where

$$\begin{aligned}\Delta R_E(x,T) &= R_0(1+\beta T_{\mathrm{metal}}(x))\Delta x \\ &= \frac{\rho_{\mathrm{i}}}{w t_m}(1+\beta T_{\mathrm{metal}}(x))\Delta x\end{aligned}$$

where $R_0 = \rho_{\mathrm{i}}/w t_m$ is the resistance per unit length of metal interconnect at a reference temperature (Ω/m), ρ_{i} is the electrical resistivity of the interconnect at the reference temperature, β is the temperature coefficient of the electrical resistance $(1/^\circ\mathrm{C})$, and T_{metal} is the interconnect temperature $(^\circ\mathrm{C})$.

$P_{\mathrm{loss}}(x)$, the heat flow between the interconnect and the substrate over the length of $\Delta(x)$, is given by

$$P_{\mathrm{loss}}(x) = \frac{T_{\mathrm{metal}}(x) - T_{\mathrm{chip}}(x)}{\Delta R_\theta(x)}$$

where T_{metal} is the interconnect temperature, T_{chip} is the underlying substrate temperature, and ΔR_θ is the insulator thermal resistance of metal segment Δx

$$\Delta R_\theta(x) = \frac{t_{\mathrm{ins}}}{k_{\mathrm{ins}} w \Delta x}.$$

k_{ins}^*, the *effective* insulator thermal conductivity, is a shape-dependent parameter that accounts for the geometrical shape of the heat conducting body on its thermal conductivity [22]. The net heat energy generation per unit volume is

$$g(x,T) = \frac{P_{\mathrm{gen}}(x) - P_{\mathrm{loss}}(x)}{w t_m \Delta x}. \tag{10}$$

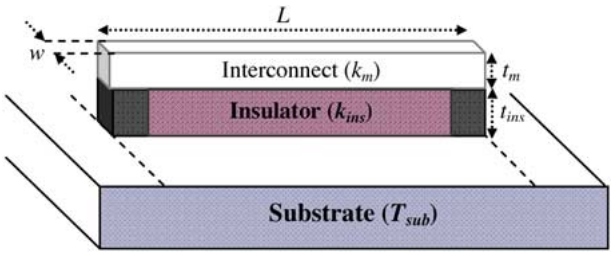

Fig. 6. *An interconnect line passing over the substrate, separated by an insulation layer.*

Since the length of global interconnects can be assumed to be much larger than its thickness and width, the thermal gradients along thickness and width can be ignored, i.e., a 1-D formula can be found to represent the heat diffusion

$$\frac{d^2T_{\text{metal}}(x)}{dx^2} = \frac{1}{k_m}\left(\left(\frac{k^*_{\text{ins}}}{t_m t_{\text{ins}}} - \frac{\rho_i \beta I^2_{\text{rms}}}{w^2 t_m^2}\right) T_{\text{metal}}(x) - \frac{k^*_{\text{ins}}}{t_m t_{\text{ins}}} T_{\text{chip}}(x) - \frac{\rho_i I^2_{\text{rms}}}{w^2 t_m^2}\right). \quad (11)$$

This equation is the basis of the interconnect temperature calculations.

One may assume that $T_{\text{chip}}(x)$ is constant for all positions along the length of the line. Then, by solving (11) subject to boundary condition (i.e., given temperatures at the two ends of the line), the thermal profile of interconnection lines can be found. Alternatively, (11) may be solved and for arbitrary chip temperature profile under the line extending between the two end points. Detailed calculations can be found in [23]. This reference also presents a temperature-aware interconnect delay model based on distributed RC-π stages and provides a closed-form expression for the temperature-aware Elmore delay of interconnect. Assuming a fixed driver resistance, which simplifies the delay expression, the authors show that about 5% delay degradation is achieved for each 20 °C temperature increase in the constant temperature along the interconnect line. Furthermore, the line resistances scales linearly with temperature, which implies that the IR drop in interconnect lines increases. The actual nonuniform temperature profile along the metal line can greatly impact the results. For example, the application of exponential temperature distributions in the substrate underneath the interconnect line shows that the delay degradation becomes highly dependent on the specific thermal distribution. These results also point to the inaccuracy of uniform temperature assumption for the purpose of interconnect delay calculation.

Assuming a constant current density in all metal layers of a signal net, it is found that the heat diffusion length is larger for the higher level metal layers due to their higher underlying insulator thickness. For interconnects whose lengths are comparable to the heat diffusion length, the line temperature does not reach the estimated maximum peak value. Therefore, techniques have been suggested to lower the peak temperature by adding extra dummy vias separated by a distance less than the thermal diffusion length [24]. Due to the variation of switching activities and hence power consumption of chip blocks, it is, however, more realistic to consider nonuniformity for the temperature profile of the underlying substrate.

In [25], it is shown that temperature-dependent electrothermal simulation is necessary for accurate estimates of multilevel interconnect temperatures under aggressive thermal conditions. In particular, it is shown that for 90- and 65-nm nodes, the temperature rises within the interconnect stacks are less than 13 °C and 33 °C, respectively. However, metal temperatures increase significantly (i.e., by hundreds of Celsius degrees) beyond the 45-nm node owing to the combined effects of increasing metal resistivity, increasing current density, increasing number of global metal levels, and decreasing ILD thermal conductivity. The maximum metal temperature occurs at the topmost global wires. It is thus concluded that even after considering densely embedded vias, the interconnect temperature is expected to increase significantly with scaling, due to increasing current density, increasing surface and grain boundary contributions to metal resistivity, and decreasing ILD thermal conductivity.

C. Thermal and Electrothermal Simulation

Many research works have concentrated on thermal modeling and simulation at circuit or gate level. Due to the often huge number of thermal components and power sources at full circuit or gate level, different methods have been proposed to increase the efficiency of thermal circuit simulation while maintaining the accuracy. The thermal simulation methods while may be roughly classified into two categories.

The first category of methods is based on the discretization of differential operators or the field quality. The corresponding thermal simulators solve the heat conduction problem numerically by using various techniques such as the finite difference [26]–[28], the finite-element [29], [30], or the boundary-element techniques [31]. The advantage of these methods is their high accuracy and ability to handle different heat sources and different types of boundary conditions. The main drawback of these methods is the enormous sizes of the resulting thermal circuits due to volume meshing. Different techniques have been devised to tackle this shortcoming. Examples include 3-D thermal ADI in [32] and model order reduction in [21] and [33]. The second category of methods is based on Green function formulation [34]–[37], which provides a fast, yet less accurate thermal simulation, due to the simplified 2-D modeling of the thermal problem.

There are two major approaches to doing electrothermal simulation of a given circuit. In the first approach, the thermal problem is mapped into an equivalent electrical problem and an electrical solver performs cosimulation of both electrical and thermal subsystems that coexist in an IC chip. In the other approach, two independent simulators, one thermal simulator and one electrical circuit simulator, iterate interactively to deliver the solution of the given electrothermal problem. Both approaches have advantages and disadvantages of their own; however, the latter approach, which is known as the *relaxation method*, is more desirable. With the relaxation method, existing software packages can be used for basic simulations and the electrical and thermal model of a specific chip can be

constructed separately. The most important representatives of relaxation methods are [26], [30], and [38]–[40]. The advantage is the relative simplicity of the implementation; the drawback is that very fast changes cannot be considered, and in case of strongly coupled thermal problems, the simulator coupling frequently cannot achieve convergence.

ILLIADS-T [26] is a well-known chip-level electrothermal timing simulator for CMOS VLSI circuits. Given the chip layout, the packaging specification, and the periodic input signal pattern, ILLIADS-T finds the on-chip steady-state temperature profile and identifies the resulting temperature-dependent circuit performance and reliability. Each gate in ILLIADS-T is viewed as a heat source. Given the input patterns, the average power dissipation of each gate is calculated by a timing/power simulator. To increase the computational efficiency, the power and temperature calculations are decoupled. An iterative approach is utilized to calculate the steady-state temperature profile from the power dissipation of the gates. During the thermal simulation, the chip temperature profile is found by solving the 3-D heat diffusion equation for the chip substrate. However, a 1-D *effective* heat transfer macromodel is applied at the package and heat sink boundaries for computational efficiency. As another example, given a circuit netlist and its layout information, the relaxation-based electrothermal simulator of [40] calculates the average power of each circuit component at the initial temperature using HSPICE as an embedded circuit simulator. Next, the simulated results of power information are fed to the thermal simulator to calculate the temperature profile, which is then used to update component temperatures in the circuit. Following this first iteration, the second circuit simulation is executed. This process continues until the temperature converges to a steady value or the iteration time reaches its predetermined maximum allowable limit.

In the *direct method*, the thermal system is represented by an electrical model network that has common thermal nodes with the electrical-only network. The iterative solution takes place simultaneously for the thermal and electrical subnetworks. Representatives of this method are discussed in [29] and [41]–[43]. The advantage is the capability of considering very fast changes; the drawback is that the direct method requires a more complex implementation than the relaxation method.

In the *direct method*, the thermal system is represented by an electrical model network that has common thermal nodes with the electrical-only network.

Direct electrothermal simulation of the ICs means that electrical equivalent circuit must also model the thermal behavior of the circuit components and structures, including the chip, the package, and the heat sink. A good example is reference [43], where the electrical-only and the thermal equivalent circuits are connected in one network via the thermal nodes, whose voltages represent the temperature values. The two subnetworks form a large network for which the network equations have to be solved, providing simultaneously the temperature and voltage values of the thermal and electrical subnetworks respectively. Notice that for this method to be effective, the device models have to be electrothermal, i.e., each device needs to have a thermal node besides the electrical nodes. The power dissipation has to be input to the thermal node as a current while the device function will have to depend on the temperature (voltage) data of this node. Also note that because the time required for the on-chip temperature to reach its steady state is orders of magnitude larger than the clock signal period in digital circuits, a moving window average power dissipation value (rather than the instantaneous power) can be used for finding the temperature of thermal nodes.

III. TEMPERATURE EFFECTS ON CIRCUIT PERFORMANCE METRICS

Temperature has a direct and often substantial impact on nearly all of the key figures of merit (performance parameters) of a VLSI circuit, including circuit speed, lifetime, power dissipation, and power plane integrity. The chip temperature is in turn set by the power dissipation in the substrate and interconnects as well as the physical layout, routing resources, and power distribution network in the chip. These interactions are discussed next.

A. Effect on Circuit Reliability

Electromigration (EM) occurs when the current density in metal interconnects becomes too high. In the case of power and ground tracks, the "electron wind" induced by the current flowing through a track causes metal ions in the track to migrate. This migration creates "voids" in the "upwind" direction, while metal ions can accumulate "downwind" to form features called "hillocks" and "whiskers" [44], [45]. The mean-time-to-failure (MTTF) due to EM is calculated by the well-known Black's equation [46]

$$\mathrm{MTTF} = AJ^{-n}e^{Q/kT}$$

where A is a process- and geometry-dependent constant, J is the dc (average) current, exponent n is two under

normal use conditions, Q is the activation energy for grain-boundary diffusion and is equal to ~ 0.7 eV for Al–Cu, k is Boltzmann's constant, and T denotes the metal temperature.

To analyze the effect of temperature on circuit reliability, the authors of [47] define

$$\gamma_j = J_{\max}(T_{\text{junc}})/J_{\max}(T_{\text{spec}})$$

where $J_{\max}(T_{\text{spec}})$ is the maximum current density defined at the specification temperature ($T_{\text{spec}} = 120\ ^\circ\text{C}$), and $J_{\max}(T_{\text{junc}})$ is a "new" current limit calculated using the actual maximum temperature T_{junc} so that the same MTTF is achieved from Black's equation. The higher the chip temperature becomes, the smaller the actual $J_{\max}$ target will be, which is represented by $\gamma_j < 1$. When the maximum temperature equals T_{spec}, γ_j becomes 1. The authors report that in high-performance SoC, local temperature exceeds specification temperature, which enforces exponentially lower current density limits, $J_{\max}(T_{\text{junc}})$ getting as low as 30% of $J_{\max}(T_{\text{spec}})$. Typically, a chip has some maximum junction temperature tolerance, which is considered carefully and is typically met by the chip designers. If, however, a local area on the chip heats up, the chance of EM-induced failures increases exponentially, greatly reducing the lifetime of the chip.

Thermal effects impact interconnects design and reliability in the following ways. Firstly, they limit the maximum allowable RMS current density in the interconnect lines in order to limit the temperature increase due to self-heating phenomenon. Secondly, interconnect lifetime (reliability), which is limited by EM, has an exponential dependence on the inverse metal temperature. Hence, the temperature rise of metal interconnects due to self-heating phenomenon also imposes a limit on the maximum allowed average current density.

B. Effect on Propagation Delays and Signal Integrity

According to [47], the logic gate delay change is about 4% with 40 °C temperature difference in a 130-nm industrial process. The wire resistance increased by about 12% for 40 °C around the nominal temperature. Delay change for the wire resulted in about 5% for 40 °C using the same process. Clock skew can be increased by as much as 10% of the clock cycle time when the junction temperature changes in the substrate by as much as 40 °C (i.e., we have a nonuniform thermal distribution in the chip).

The existence of high thermal gradients on the substrate creates nonuniform temperature profiles along the length of the global interconnect lines, which are located above the substrate. This inherently leads to nonuniform interconnect resistance profiles. The nonuniform resistance profile of the global interconnects will in turn strongly impact many aspects of interconnect performance modeling and optimization. In addition, as feature size scales to sub-90-nm dimensions, in spite of an increase in the number of metal layers that will be available in advanced technology nodes, the top metal layers may get closer to the substrate and find stronger coupling with the substrate, and thereby impact interconnect performance analysis. Clearly, the dependence of interconnect performance on nonuniform temperature distributions along the length of global wires will have a big impact on the solutions to many physical design and layout-optimization problems, including clock skew control, wire sizing, layer assignment, crosstalk effects, and buffer insertion as shown in [48]. On the other hand, the interconnect load capacitance is the predominant part of the load capacitance in dynamic power component; therefore, interconnect thermal effects play a crucial role in the reliability and performance of the whole system.

Circuit simulation plays an important role during early design steps and final verification phases. The accuracy of circuit simulation tools depend on the implemented device models. Hence, for accurate performance (timing) analysis precise gate as well as interconnect delay models must be considered. Temperature dependence, noise effects, and process variations are among the factors that need to be taken into account in a complete model. The circuit delay has a specifically strong dependence on temperature; therefore, it is important to explore how the thermal effects can affect the circuit timing characteristics, such as the path delay, and the criticality of path.

C. Effect on Power Dissipation

As stated previously, temperature has a strong effect on subthreshold leakage, which tends to be a major component of the circuit power dissipation in sub-90-nm CMOS designs. (At 90-nm-process nodes, leakage accounts for 25 to 40% of total power. At 65-nm processes, leakage accounts for 50 to 70% of total power.) Since leakage is critically dependent on operating temperature and power supply, the authors of [49] present a full chip leakage estimation technique which accurately accounts for power supply and temperature variations. State of the art techniques are used to compute the thermal and power supply profile of the entire chip. Closed-form models are presented which relate leakage to temperature and V_{DD} variations. These models coupled with the thermal and V_{DD} profile are used to generate an accurate full chip leakage estimation technique considering environmental variations. The results of this approach are demonstrated on large-scale industrial designs.

The authors in [50] present a framework for full-chip estimation of subthreshold leakage power distribution considering both within-die and die-to-die variations in process (P), temperature (T), and supply voltage (V). Using this framework, a quantitative analysis of the relative sensitivities of subthreshold leakage to P-T-V

variations has been presented. It was shown that for accurate estimation of subthreshold leakage, it is important to consider die-to-die temperature variations which can significantly increase the leakage power due to electro-thermal couplings between power and temperature. Furthermore, the full-chip leakage power distribution arising due to both within-die and die-to-die P-T-V is calculated, which is subsequently used to estimate the leakage constrained yield under the impact of these variations. The calculations show that the yield is significantly lowered under the impact of within-die and die-to-die process and temperature variations.

D. Effect on P/G Integrity

Uneven power consumption in various logic blocks in an IC chip results in two effects: 1) appearance of hot spots and temperature gradients and 2) supply voltage variations due to current demands of logic blocks and design of a power/ground distribution network. To accommodate variations in local temperatures and supply voltage levels, designers have traditionally been forced to pad logic cell characteristics and design margins. However, creating the power distribution network using excessively conservative design practices can result in loss of valuable silicon real estate, increasing congestion, and thereby, resulting in performance loss.

Deep submicrometer (DSM) IC chips are prone to resistive voltage drop effects, which are caused by the resistance associated with the network of wires used to distribute power and ground from the external pins to the internal circuitry [51]. Every power and ground track segment has a small amount of resistance associated with it, i.e., the logic block that is closest to the external power or ground pins receives nearly perfect supply voltage levels. However, logic blocks that are further down the power-ground distribution are presented with a degraded supply voltage levels. In the remainder of this section we focus on voltage drops on the V_{DD} rail. Similar statements apply to the Gnd rail.

In general, the voltage drops on the power rail can be in the form of a self-induced IR drop from the external power pin to the power terminal of a logic block due to the current that is drawn by the logic block itself. In addition, they can in the form of a transient IR drop caused by switching currents that flow into the other logic blocks further up the power distribution network. Indeed, the transitory power surge in a logic block (due to switching activity in that block) can momentarily reduce the voltage supply level to blocks down the power supply chain.

It is worthwhile to point out that accepted techniques for power management in VLSI circuits (including clock and power gating) tend to also increase the on-chip thermal gradients and power surges (with ensuing transitory supply voltage drops.) As an example, when a large logic block that was previously clock gated comes online, it will suddenly draw a large amount of current from it power terminal, which will cause voltage drop not only on its own power terminals but also on power terminals of every other logic block down the power distribution network. With reduced supply voltage level comes slower logic switching times, which means some block in the circuit can suffer a setup time failure. There is also an increase in the interconnect delays associated with wires that are driven by underpowered logic cells. Furthermore, when the supply is reduced, the logic cells become more susceptible to noise because of reduced dc noise margin.

The issue is in fact even more complicated that what is described above, and that is because the inductive voltage drop (Ldi/dt noise) can add to this dc drop and cause even more performance degradation and thus further increase the chances for circuit failure. Of course, insertion of the right amount of decoupling capacitors at the right places in the supply distribution network can help alleviate the inductive voltage drop effects [52].

Voltage drop effects are becoming increasingly significant, because the resistivity of the power and ground tracks rises as a function of decreasing feature sizes (track widths) and increasing chip temperatures. These effects can be minimized by increasing the width of power tracks (which reduces the power track resistances) and/or by increasing the spacing between logic blocks (which reduces power density and hence reduces chip temperature). However, the first approach tends to use up valuable real estate on the silicon, which typically causes routing congestion problems whereas the second approach increases interconnect delays (and power consumption due to increased physical capacitance of longer logic signal tracks.) Thus, implementing an optimal power network requires the balancing of many diverse factors [53], [54].

EM in power and ground tracks can cause significant timing problems because the increased track resistance associated with void results in a voltage drop. This will consecutively cause longer propagation delays and higher noise susceptibility in affected logic blocks. Power and ground EM can also cause major catastrophic errors to occur, because the voids may eventually lead to open circuits while the hillocks and whiskers may cause short circuits to neighboring wires [55], [56].

E. Putting It Together

Fig. 7 illustrates the interaction among thermal analysis, timing and power analyses, and power-ground (P/G) network analysis. Chip thermal analysis needs power consumption values from the circuit power analyzer and the chip thermal model. In return, it calculates the temperature profile of the chip, which will in turn impact the leakage power dissipation values, the gate drive strength, and the interconnect resistances. Hence, the thermal analysis results directly feed into the power analysis, timing analysis, and P/G analysis modules. The P/G network analysis also needs the current densities in the various sections of the power distribution network, which comes

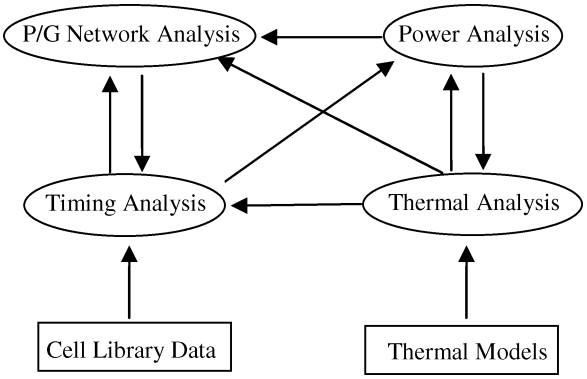

Fig. 7. *Thermal analysis affects timing (with signal integrity), power and P/G network analyses.*

from the results of the power analysis. The timing analysis module provides information about the signal transition times and arrival times of various gates to the power analysis and P/G network analysis modules. Finally, the P/G network analysis module provides voltage drop and ground bounce information to the timing analysis module.

The real picture is indeed more complicated than that described above. This is due to the rising impact of global and random sources of variations on performance characteristics of VLSI circuits [57], which tend to increase the criticality of temperature dependencies in the circuit. This is in turn because variations in V_T or L_{eff} of transistors and voltage drops on power supply lines reduce the noise margins of logic cells, leaving less room to accommodate temperature effects on parasitic *RC* values and circuit performance.

Up until recently, corner-based timing and signal integrity analysis techniques were used as relatively fast techniques to address the concerns related to various sources of variation in VLSI circuits. In general, corner-based techniques tend to overestimate circuit delay and noise effects. These techniques can also result in underestimation of circuit delay and noise because these metrics are nonmonotone functions of some circuit parameters. Exacerbating the situation, it is nontrivial to find the worst case value for each circuit parameter that would result in the worst case delay or noise. Statistical analysis [58], [59] has been used to address the above-mentioned shortcomings of corner-based approaches. Each analysis tool in Fig. 7 must therefore have a statistical analysis component to effectively address the shortcoming of traditional corner-based analyzers.

The lack of integration between the various analysis tools can result in a huge number of "false errors." Engineers often over-design to avoid false errors. For example, a perceived power plane integrity problem may be addressed by excessively wide power rails. This in turn can cause design to fail to meet its area or timing constraints and to become congested. Compensating for these side effects can cause ripples throughout the whole design resulting in excessive margins in the timing goals, excessive margins on signal integrity, etc. It is easily seen that this kind of situation is undesirable, since it makes it nearly impossible to achieve design closure.

IV. TEMPERATURE CONTROL

The chip temperature can be controlled by off-chip (system-level or board-level) means, or on-chip thermal management techniques. The on-chip techniques an in turn be divided into static (design time) or dynamic (runtime) strategies and structures. These techniques are described next.

A. Package/System-Based Thermal Management

One method involves the use of thermal feedback (from a thermistor, for instance), to change the fan speed and keep the maximum temperature of the chips on a printed circuit board below a threshold temperature. Another method of local temperature control involves the use of a solid-state heat pump (or thermoelectric cooler), which employ the Peltier effect. Such a device can electronically pump heat from an IC chip to a heat sink [60]. The drawback of this system is the added space, weight, and expense of the thermoelectric cooler and heat sink. Yet another method of local temperature control involves the use of phase-change material [61]. Such a method holds the chip temperature at a nearly fixed temperature by absorbing heat into a material that changes its physical state. The drawback of such a system is that it only works for a limited time. In addition, there is the added system complexity and the containment problem of the phase change material.

B. Design-Time Thermal Management

Designers can employ a number of techniques to correct thermal problems on IC chips. One technique is to do temperature-aware physical design planning by equalizing the temperature across the chip during the initial macrocell placement. Designers can also make designs more robust by limiting the maximum power draw that is sustained over a period of say tens of microseconds. They can make the VLSI interconnects more robust by widening traces and by doing buffer insertion and sizing. Dummy vias in the higher metal layers may be used to reduce temperatures on interconnect without impacting their electrical resistance and capacitance. This is because the vias will reduce the thermal resistance between the metal line and the substrate. By including a static thermal calculator, design teams can consider various packaging and on-chip electrothermal parameters. Invoking such a tool from the early stages in the physical-synthesis flow allows for thermally conscious physical design planning and global placement to minimize the magnitude of nonuniform

temperature profiles. A detailed discussion of design optimization for improving the full-chip thermal integrity is provided elsewhere in this special issue.

What is needed is an RTL-to-GDSII system containing an embedded thermal analysis engine that is integrated with various components of the optimization flow [62]. The target system is capable of addressing the impact of temperature on on-chip parasitic parameters and incorporating the thermal effects into the design optimization algorithms. The end goal is to produce an EDA methodology, design flow, and tool suite that include thermally aware analysis and optimization techniques designed to reduce the impact of on-chip temperature and hot spots on various circuit performance parameters.

C. Adaptive Thermal Management

Current thermal solutions are designed to limit the peak processor power dissipation to ensure its reliable operation under worst case scenarios. However, the peak chip power and ensuing peak temperature are hardly ever observed. Dynamic thermal management (DTM) has been proposed as a class of microarchitectural solutions and software strategies to achieve the highest chip performance under a peak temperature limit. Furthermore, as stated earlier, the power density across the chip is nonuniform, resulting in localized hot spots. DTM solutions must address this phenomenon as much as they tackle system-wide temperature violations.

When the chip approaches the thermal limit, a DTM controller initiates hardware reconfiguration, slowdown, or shutdown to lower the chip temperature. Possible response mechanisms include microarchitectural adaptations (e.g., clock throttling, register file resizing, limiting the issue width of a processor, and computation migration to auxiliary hardware), and/or on-the-fly performance adjustment via dynamic power management (DPM), dynamic voltage scaling (DVS), clock/power gating.

Dynamic power management (DPM) is a feature of the runtime environment of a *power-managed circuit* (PMC) that adaptively reconfigures itself to provide the requested services and performance levels with a minimum number of active components or a minimum activity level on such components. DPM encompasses a set of techniques that achieve energy-efficient computation by selectively turning off (or reducing the performance of) circuit components when they are idle (or partially unexploited.) The fundamental premise for the applicability of DPM is that circuit (and its functional blocks) experience nonuniform workloads during operation time. Such an assumption is valid for many circuits. A second assumption of DPM is that it is possible to predict, with a certain degree of confidence, the fluctuations of workload.

The early works on DPM focused on predictive shutdown approaches [63], [64] which make use of "time-out" based policies. A power management approach based on discrete time Markovian decision processes was proposed in [65]. The discrete-time model requires policy evaluation at periodic time instances and may thereby consume a large amount of power even when no change in the system state has occurred. To surmount this shortcoming, a model based on continuous-time Markovian decision processes (CTMDP) was proposed in [66]. The policy change under this model is asynchronous and is thus more suitable for implementation as part of a real-time operating system environment. Other approaches such as adaptive learning based strategies [67], session clustering and prediction strategies [68], online strategies [69], Petri-based DPM techniques [70], hierarchical system decomposition and modeling [71], [72], and combination of CTMDP-based approaches and timeout based power management techniques [73] have also been utilized to find a DPM policy of power-managed computer system.

Dynamic power management (DPM) is a feature of the runtime environment of a *power-managed circuit* (PMC) that adaptively reconfigures itself.

The key differences between DPM and DTM can be stated as follows. Localized heating occurs much faster than chip-wide heating. Additionally, power dissipation is spatially nonuniform across the chip, resulting in emergence of hot spots and spatial temperature gradients that can cause timing errors or even physical damage. These effects evolve over time scales of hundreds of microseconds, which implies that power-management techniques, in order to be used for thermal management, must directly target the *spatiotemporal behavior* of the chip temperature. In fact, many DPM techniques have little or no effect on substrate temperature, because they do not reduce the power density in hot spots, reduce power dissipation with too short a timing granularity, or do not attempt to reduce power dissipation when no positive timing slack is present. Furthermore, DPM techniques attempt to lower the sum total of energy consumed over the entire application run, while DTM techniques must ensure that a thermal limit is not exceeded. The power control algorithm tracks the power consumption of the entire chip as a whole (which is a *minsum optimization* problem), while the temperature control algorithm concentrates on the power consumption of specific localized structures on chip (which is a *minmax optimization* problem). Finally, DPM algorithms seek to minimize energy while meeting a task completion deadline whereas with DTM algorithms, there is no minimal performance target other than not exceeding a temperature threshold. In other

words DPM is a *constrained* optimization problem whereas DTM is often formulated as an *unconstrained* optimization problem.

Despite the long-standing concern about thermal effects, only a few studies have been published on DTM, and they are mostly in the computer architecture field. Gunther *et al.* [74] describe the thermal design approach for the Pentium 4, where thermal management is accomplished via global clock gating. Lim *et al.* [75] propose a heterogeneous dual-pipeline processor for mobile devices in which the standard execution core is augmented by a low-power, single-issue, in-order pipeline that shares the fetch engine, register files, and execution units but deactivates out-of-order components like the renamer and issue queues. Skadron *et al.* [76], [77] study the architecture-level thermal management and modeling based on an equivalent circuit of thermal resistances and capacitances. Architecture-level thermal management that uses runtime knowledge of application behavior and the current thermal status is incorporated with dynamic voltage and frequency scaling techniques to ensure safe operation. Various architectural mechanisms for DTM are evaluated to regulate the on-chip temperatures. An architectural thermal model for microprocessor, HotSpot, which constructs a multilayer lumped thermal *RC* network, is described in [77] to model the heat dissipation path from the silicon die through the cooling package to the ambient.

There have been a couple of research results on enforcing the temperature limits of the processors. The trigger mechanism used to cool the microprocessor's temperature with DTM has been described in [78] by using Wattch [79]. A temperature threshold can be set to a higher level and used as a gauge of how successful the trigger mechanism is in dealing with the increase in temperature. In [80], an integrated dynamic thermal management circuit for system-on-chip is proposed that incorporates on-chip power/speed modulation and integrated multistage fan controllers. When the chip approaches the thermal limit, DTM controls each component's operation speed with the feedback from present temperature information. The work presented in [81] tackles the performance optimization problem for disk drives by studying the relationship between capacity, performance, and the thermal characteristics of disk drives. This reference provides solution for temperature-aware disk drive design while presenting the roadmap based on models in technology scaling of the fundamental parameters from the thermal perspective.

Reactive DTM techniques, which have limited time to respond to a thermal emergency, preclude the use of mechanisms such as dynamic voltage scaling (DVS) and register file resizing that have high invocation overhead. In contrast, predictive DTM techniques are able to employ DVS and register file resizing. A predictive DTM approach that exploits certain properties of multimedia applications is presented in [82]. Two adaptation techniques, which can impact hot spot problems, i.e., instruction window resizing and deactivating appropriate register file ports, are studied. This reference shows that a combination of DVS and architectural adaptation with predictive algorithm is highly effective in controlling the chip temperature.

V. CONCLUSION

Tight interactions between average and peak power dissipation, delay and signal integrity, P/G network design, and chip temperatures strongly argues for the need to have a tightly coupled set of power, timing, and signal integrity analysis and optimization tools and flows that are built on top of a unified design environment, which in turn combines disparate physical, electrical and thermal data models into one database. To fully account for the impact of voltage drop effects, for example, it is important to have an environment that can derate for timing—on a cell-by-cell basis—based on actual voltage drops and local temperatures. The timing analysis engine should then make use of this derated timing data to identify potential changes to the critical paths. In turn, the optimization engine should make appropriate modifications to address potential setup or hold problems that appear as a result of the timing changes. This requires a design environment in which the power analysis, voltage drop analysis, derating calculations, timing analysis, and optimization engines all work seamlessly together.

Temperature-dependent design issues are becoming dominant factors in resolving signal- and power-integrity issues in future generations of VLSI chips. There is thus a clear need for effective temperature-aware analysis and optimization tools and design flows that would in turn enable the design of high-performance and reliable IC chips. ■

REFERENCES

[1] E. Rohou and M. Smith, "Dynamically managing processor temperature and power," in *Proc. FDDO-2*, 1999.

[2] S. I. Souri, K. Banerjee, A. Mehrotra, and K. C. Saraswat, "Multiple Si layer ICs: Motivation, performance analysis, and design implications," in *Proc. Design Automation Conf.*, 2000, pp. 873–880.

[3] S. Das and A. Chadrakasan, "Three-dimensional integrated circuits: Performance, design methodology, and CAD tools," in *Proc. IEEE Annu. Symp. VLSI*, 2003, pp. 13–18.

[4] M. B. Kleiner, S. A. Kuhn, P. Ramm, and W. Weber, "Thermal analysis of vertically integrated circuits," in *Tech. Dig. Int. Electron Devices Meeting*, 1995, pp. 487–490.

[5] A. Rahman and R. Reif, "Thermal analysis of three-dimensional (3-D) integrated circuits (ICs)," in *Proc. Int. Interconnect Technology Conf.*, 2001, pp. 157–159.

[6] F. Fallah and M. Pedram, "Standby and active leakage current control and minimization in CMOS VLSI circuits," *IEICE Trans. Electron. (Special Section on Low-Power LSI and Low-Power IP)*, vol. E88-C, no. 4, pp. 509–519, 2005.

[7] B. Chatterjee, M. Sachdev, S. Hsu, R. Krishnamurthy, and S. Borkar, "Effectiveness and scaling trends of leakage

control techniques for sub-100 nm CMOS technologies," in *Proc. Int. Symp. Low-Power Electronics,* 2003, pp. 122–127.

[8] M. L. Mui, K. Banerjee, and A. Mehrota, "Power supply optimization in sub-130 nm leakage dominant technologies," in *Proc. Int. Symp. Quality Electronic Design,* 2004, pp. 409–414.

[9] Int. Technology Roadmap for Semiconductors (ITRS), 2004.

[10] K. Banerjee, A. Mehrotra, A. Sangiovanni-Vincentelli, and C. Hu, "On thermal effects in deep sub-micron VLSL interconnects," in *Proc. Design Automation Conf.*, 1999, pp. 885–891.

[11] H. J. M. Veendrick, "Short-circuit dissipation of static CMOS circuitry and its impact on the design of buffer circuits," *IEEE J. Solid-State Circuits,* vol. SC-19, no. 4, pp. 468–473, Aug. 1984.

[12] A. Alvandpour, P. L. Edefors, and C. Svensson, "Separation and extraction of short-circuit power consumption in digital CMOS VLSI circuits," in *Proc. Low Power Electronics and Design,* 1998, pp. 245–249.

[13] S. Turgis, N. Azemard, and D. Auvergne, "Explicit evaluation of short-circuit power dissipation and its influence on propagating delay for static CMOS gates," in *Proc. IEEE Int. Symp. Circuits and Systems,* 1996, pp. 751–754.

[14] S. H. Jung, J. H. Baek, and S. Y. Kim, "Short circuit power estimation of static CMOS circuits," in *Proc. Asia-Pacific Design Automation Conf.,* 2001, pp. 545–549.

[15] Y. Taur and T. Ning, *Fundamentals of Modern VLSI Devices*. Cambridge, U.K.: Cambridge Univ. Press, 1998.

[16] K. Cao, W.-C. Lee, W. Liu, X. Jin, P. Su, S. Fung, J. An, B. Yu, and C. Hu, "BSIM4 gate leakage model including source drain partition," in *Tech. Dig. Int. Electron Devices Meeting,* 2000, pp. 815–818.

[17] Y. Taur, D. A. Buchanan, W. Chen, D. J. Frank, K. E. Ismail, S. H. Lo, G. Sai-Halasz *et al.* "CMOS scaling into nanometer regime," *Proc. IEEE,* vol. 85, no. 4, pp. 486–504, Apr. 1997.

[18] P. Ko, J. Huang, Z. Liu, and C. Hu, "BSIM3 for analog and digital circuit simulation," in *Proc. IEEE Symp VLSI Technology CAD,* 1993, pp. 400–429.

[19] M. Ozisik, *Boundary Value Problems of Heat Conduction*. Oxford, U.K.: Oxford Univ. Press, 1968.

[20] T. Y. Chiang, K. Banerjee, and K. C. Saraswat, "Compact modeling and SPICE-based simulation for electrothermal analysis of multilevel ULSI interconnects," in *Proc. Int. Conf. Computer-Aided Design,* 2001, pp. 165–172.

[21] T.-Y. Wang and C.-C. Chen, "Spice-compatible thermal simulation with lumped circuit modeling for thermal reliability analysis based on model reduction," in *Proc. Int. Symp. Quality Electronic Design,* 2004, pp. 357–362.

[22] R. V. Andrews, "Solving conductive heat transfer problems with electrical-analogue shape factors," *Chem. Eng. Prog.,* vol. 51, no. 2, pp. 67–71, 1955.

[23] A. H. Ajami, K. Banerjee, and M. Pedram, "Modeling and analysis of non-uniform substrate temperature effects on global ULSI interconnects," *Trans. Comput.-Aided Design Integr. Circuits Syst.,* vol. 24, no. 6, pp. 849–861, Jun. 2005.

[24] T.-Y. Chiang, K. Banerjee, and K. C. Saraswat, "Effect of via separation and low-k dielectric materials on the thermal characteristics of Cu interconnects," in *Tech. Dig. IEEE Int. Electron Devices Meeting,* 2000, pp. 261–264.

[25] S. Im, N. Srivastava, K. Banerjee, and K. E. Goodson, "Scaling analysis of multilevel interconnect temperatures for high-performance ICs," *IEEE Trans. Electron Devices,* vol. 52, no. 12, pp. 2710–2719, Dec. 2005.

[26] Y. K. Cheng, P. Raha, C. C. Teng, E. Rosenbaum, and S. M. Kang, "ILLIADS-T: An electrothermal timing simulator for temperature-sensitive reliability diagnosis of CMOS VLSI chips," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.,* vol. 17, no. 8, pp. 668–681, Aug. 1998.

[27] G. Digele, S. Lindenkreuz, and E. Kasper, "Fully coupled dynamic electrothermal simulation," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.,* vol. 5, no. 3, pp. 250–257, Sep. 1997.

[28] D. Chen, E. Li, E. Rosenbaum, and S.-M. Kang, "Interconnect thermal modeling for accurate simulation of circuit timing and reliability," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.,* vol. 19, no. 2, pp. 197–205, Feb. 2000.

[29] V. A. M. N. Sabry, A. Bontemps, and R. Vahrmann, "Realistic and efficient simulation of electrothermal effects in VLSI circuits," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.,* vol. 5, no. 3, pp. 283–289, Sep. 1997.

[30] C. C. S. Wunsche, C. Clauss, P. Schwarz, and F. Winkler, "Electrothermal circuit simulation using simulator coupling," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.,* vol. 5, no. 3, pp. 277–282, Sep. 1997.

[31] J. P. Fradin and B. Desaunettes, "Automatic computation of conductive conductances intervening in the thermal chain," in *Proc. Int. Conf. Environmental Systems*, 1995.

[32] T. Y. Wang and C. C. P. Chen, "3-D thermal-ADI: A linear-time chip level transient thermal simulator," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.,* vol. 21, no. 12, pp. 1434–1445, Dec. 2002.

[33] L. Codecasa, D. D'Amore, and P. Maffezzoni, "An Arnoldi based thermal network reduction method for electro-thermal analysis," *IEEE Trans. Compon. Packag. Technol.,* vol. 26, no. 1, pp. 186–192, Mar. 2003.

[34] A. Haji-Sheikh, "Peak temperature in high-power chips," *IEEE Trans. Electron Devices,* vol. 37, no. 4, pp. 902–907, Apr. 1990.

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

[36] B. Wang and P. Mazumder, "Fast thermal analysis for VLSI circuits via semi-analytical green's function in multilayer materials," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS),* 2004, pp. II-409–II-412.

[37] Y. Zhan and S. Sapatnekar, "Fast computation of the temperature distribution in VLSI chips using the discrete cosine transform and table look-up," in *Proc. Asia-Pacific Design Automation Conf.*, 2005, pp. 87–92.

[38] Y.-K. Cheng *et al.* "ETS-A: A new electrothermal simulator for CMOS VLSI circuits," in *Proc. ED&TC'96*, pp. 566–570.

[39] T. Li, C. H. Tsai, and S. M. Kang, "Efficient transient electrothermal simulation of CMOS VLSI circuits under electrical overstress," in *Proc. Int. Conf. Computer-Aided Design,* 1998, pp. 6–10.

[40] L. Miao, Z. Runde, and G. Yuanqing. (2000). "A new electrothermal simulator based on relaxation method for integrated circuits with distributed temperatures," in Proc. ICDA. [Online]. Available: http://www.ifip.org/con2000/icda2000/icda-11-3.pdf

[41] C. H. Diaz, S. M. Kang, and C. Duvvury, "Circuit-level electrothermal simulation of electrical overstress failures in advanced MOS I/O protection devices," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.,* vol. 13, no. 4, pp. 482–493, Apr. 1994.

[42] V. Székely, A. Poppe, M. Rencz, A. Csendes, and A. Phi, "Self-consistent electro-thermal simulation: Fundamentals and practice," *Microelectron. J.,* vol. 28, pp. 247–262, 1997.

[43] V. Székely, A. Poppe, A. Phi, A. Csendes, G. Hajas, and M. Renczel, "Electrothermal and logi-thermal simulation of VLSI designs," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.,* vol. 5, no. 3, pp. 258–269, Sep. 1997.

[44] P. C. Li and T. K. Young, "Electromigration: The time bomb in deep-submicron ICs," *IEEE Spectr.,* vol. 33, no. 9, pp. 75–78, Sep. 1996.

[45] B. R. Livesay, N. E. Donlin, A. K. Garrison, H. M. Harris, and J. L. Hubbard, "Dislocation based mechanisms in electromigration," in *Proc. Int. Reliability Physics Symp.,* 1992, pp. 217–227.

[46] J. Black, "Electromigration—A brief survey and some recent results," *IEEE Trans. Electron Devices,* vol. E-16, no. 4, pp. 338–347, Apr. 1969.

[47] T. Sato, J. Ichimiya, N. Ono, K. Hachiya, and M. Hashimoto, "On-chip thermal gradient analysis and temperature flattening for SoC design," *IEICE Trans. Fundam.,* vol. E88-A, no. 12, pp. 3382–3389, Dec. 2005.

[48] A. H. Ajami, M. Pedram, K. Banerjee, and L. P. P. P. van Ginneken, "Analysis of non-uniform temperature-dependent interconnect performance in high performance ICs," in *Proc. Design Automation Conf.,* 2001, pp. 567–572.

[49] H. Su, F. Liu, A. Devgan, E. Acar, and S. Nassif, "Full chip leakage estimation considering power supply and temperature variations," in *Proc. Symp. Low Power Electronics and Design*, Aug. 2003, pp. 78–83.

[50] S. Zhang, V. Wason, and K. Banerjee, "A probabilistic framework to estimate full-chip subthreshold leakage power distribution considering within-die and die-to-die P-T-V variations," in *Proc. Symp. Low Power Electronics and Design,* 2004, pp. 156–161.

[51] G. Bai, S. Bobba, and I. N. Hajj, "Simulation and optimization of the power distribution network in VLSI circuits," in *Proc. IEEE Int. Conf. Computer-Aided Design,* 2000, pp. 481–486.

[52] P. Heydari and M. Pedram, "Ground bounce in digital VLSI circuits," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.,* vol. 11, no. 2, pp. 180–193, Apr. 2003.

[53] N. E. Evmorfopoulos, D. P. Karampatzakis, and G. I. Stamoulis, "Voltage-drop-constrained optimization of power distribution network based on reliable maximum current estimates," in *Proc. Int. Conf. Computer Aided Design,* 2004, pp. 479–484.

[54] X.-D. Tan and C.-J. Shi, "Efficient very large scale integration power/ground network sizing based on equivalent circuit modeling," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.,* vol. 22, no. 3, pp. 277–284, Mar. 2003.

[55] S. M. Alam, C. L. Gan, F. L. Wei, C. V. Thompson, and D. E. Troxel,

"Circuit-level reliability requirements for Cu metallization," *IEEE Trans. Device Mater. Reliab.*, vol. 5, no. 3, pp. 522–531, Sep. 2005.

[56] T. Y. Wang, J. L. Tsai, and C. C. P. Chen, "Thermal and power integrity based power/ground networks optimization," in *Proc. Design, Automation and Test in Europe*, 2004, pp. 830–835.

[57] S. Nassif, "Modeling and analysis of manufacturing variations," in *Proc. CICC*, 2001, pp. 223–228.

[58] C. Visweswariah, K. Ravindran, K. Kalafala, S. G. Walker, and S. Narayan, "First-order incremental block-based statistical timing analysis," in *Proc. Design Automation Conf.* 2004, pp. 331–336.

[59] S. Abbaspour, H. Fatemi, and M. Pedram, "VGTA: Variation-aware gate timing analysis," in *Proc. Int. Conf. Computer Design: VLSI in Computers and Processors*, 2005, pp. 351–356.

[60] W. Stubstad, "The application of thermoelectric spot cooling to electronic equipment," *IRE Trans. Product Eng. Prod.*, pp. 22–29, Dec. 1961.

[61] D. Pal and Y. Joshi, "Application of phase change materials for passive thermal control of plastic quad flat packages (PQFP): A computational study," in *Proc. 11th Annu. IEEE Semiconductor Thermal Measurement and Management Symp.*, 1995, pp. 65–71.

[62] A. Ajami, "Thermal management takes center stage in IC design," *Electron. Des.*, Oct. 2005.

[63] M. Srivastava, A. Chandrakasan, and R. Brodersen, "Predictive system shutdown and other architectural techniques for energy efficient programmable computation," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 4, no. 1, pp. 42–55, Mar. 1996.

[64] C.-H. Hwang and A. Wu, "A predictive system shutdown method for energy saving of event-driven computation," in *Proc. Int. Conf. Computer-Aided Design*, 1997, pp. 28–32.

[65] L. Benini, G. Paleologo, A. Bogliolo, and G. De Micheli, "Policy optimization for dynamic power management," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 18, no. 6, pp. 813–833, Jun. 1999.

[66] Q. Qiu, Q. Wu, and M. Pedram, "Stochastic modeling of a power-managed system-construction and optimization," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 20, no. 10, pp. 1200–1217, Oct. 2001.

[67] E.-Y. Chung, L. Benini, and G. D. Micheli, "Dynamic power management using adaptive learning trees," in *Dig. Tech. Papers IEEE/ACM Int. Conf. Computer-Aided Design*, 1999, pp. 274–279.

[68] Y. Lu and G. DeMicheli, "Adaptive hard disk power management on personal computers," in *Proc. Great Lakes Symp. VLSI*, 1999.

[69] S. Irani, S. Shukla, and R. Gupta, "Online strategies for dynamic power management in systems with multiple power saving states," *IEEE Trans. on Embedded Computing Systems*, 2003, pp. 325–346.

[70] Q. Wu, Q. Qiu, and M. Pedram, "Dynamic power management of complex systems using generalized stochastic Petri nets," in *Proc. 37th Design Automation Conf.*, 2000, pp. 352–356.

[71] Z. Ren, B. H. Krogh, and R. Marculescu, "Hierarchical adaptive dynamic power management," in *Proc. Conf. Design, Automation and Test in Europe*, 2004, pp. 136–141.

[72] P. Rong and M. Pedram, "Hierarchical dynamic power management with application scheduling," in *Proc. Int. Symp. Low Power Electronics and Design*, 2005, pp. 269–274.

[73] ——, "Determining the optimal timeout values for a power-managed system based on the theory of Markovian processes: Offline and online algorithms," in *Proc. Design Automation and Test in Europe*, 2006, pp. 1128–1133.

[74] S. Gunther, F. Binns, D. M. Carmean, and J. C. Hall. (2001, Feb.). "Managing the impact of increasing microprocessor power consumption," *Intel Tech. J., Q1.* [Online]. *5(1)*. Available: ftp://download.intel.com/technology/itj/q12001/pdf/art_4.pdf

[75] C.-H. Lim, W. Daasch, and G. Cai, "A thermal-aware superscalar microprocessor," in *Proc. Int. Symp. Quality Electronic Design*, 2002, pp. 517–522.

[76] K. Skadron *et al.* "Temperature-aware computer systems: Opportunities and challenges," *IEEE Micro*, vol. 23, no. 6, pp. 52–61, Nov.–Dec. 2003.

[77] K. Skadron, M. R. Stan *et al.* "Temperature-aware microarchitecure," in *Proc. Int. Symp. Computer Architecture*, 2003, pp. 2–13.

[78] D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," in *Proc. Int. Symp. High-Performance Computer Architecture*, 2001, pp. 171–182.

[79] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: A framework for architectural-level power analysis and optimization," in *Proc. Int. Symp. Computer Architecture*, Jun. 2000, pp. 83–94.

[80] H. Chiueh, J. Draper, and J. Choma, Jr., "A dynamic thermal management circuit for system-on-chip design," *J. Analog Integr. Circuits Signal Process.*, vol. 36, no. 1–2, pp. 175–181, Jul.–Aug. 2003.

[81] S. Gurumurthi, A. Sivasubramaniam, and V. K. Natarajan, "Disk drive roadmap from the thermal perspective: A case for dynamic thermal management," *Comput. Archit. News*, vol. 33, no. 2, pp. 38–49, May 2005.

[82] J. Srinivasan and S. V. Adve, "Predictive dynamic thermal management for multimedia applications," in *Proc. 17th Annu. Int. Conf. Supercomputing*, 2003, pp. 109–120.

ABOUT THE AUTHORS

Massoud Pedram (Fellow, IEEE) received the B.S. degree in electrical engineering from the California Institute of Technology, Pasadena, in 1986 and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989 and 1991, respectively.

He joined the Department of Electrical Engineering—Systems at the University of Southern California, Los Angeles, where he is currently a Professor. He has published four books and more than 300 journal and conference papers. His current work focuses on developing design methodologies and techniques for low-power electronics and timing, power, and temperature analysis of CMOS very large scale integration (VLSI) circuits.

Dr. Pedram has served on the technical program committee of a number of conferences, including the Design automation Conference (DAC), Design and Test in Europe Conference (DATE), Asia-Pacific Design Automation Conference (ASP-DAC), and International Conference on Computer Aided Design (ICCAD). He cofounded and served as the Technical Co-chair and General Co-chair of the International Symposium on Low Power Electronics and Design (SLPED) in 1996 and 1997, respectively. He was the Technical Program Chair and the General Chair of the 2002 and 2003 International Symposium on Physical Design. His research has received a number of awards including two DAC Best Paper Awards, a Distinguished Paper Citation from ICCAD, two ICCD Best Paper Awards, and an IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS Best Paper Award. He is a recipient of the NSF's Young Investigator Award (1994) and the Presidential Faculty Fellows Award (a.k.a. PECASE Award) (1996). He is the IEEE Circuits and Systems Society VP of Publications. He is also an Advisory Board Member of the ACM Interest Group on Design Automation.

Shahin Nazarian received the B.Sc. degree from Amirkabir University of Technology (Tehran Polytechnic), Tehran, Iran, in 1996 and the M.Sc. degree from the University of Tehran in 1999, both in electrical engineering and electronics. He also has the M.S. degree in Computer Engineering from the University of Southern California, Los Angeles. He is currently working toward the Ph.D. degree in the Electrical Engineering Department at the University of Southern California.

During the summer of 2003, he was with Lucent Technologies, Bell Labs Innovations, where he engaged in systems-on-chip (SoC) testing based on the IEEE P1500 SECT standard. He worked on crosstalk-aware logic cell timing analysis at Magma Design Automation in the summer of 2004. His research interests include signal integrity analysis, timing, and power analysis considering process variations, as well as testing and validation of integrated circuits.