Interconnect Lifetime Prediction for Reliability-Aware Systems

Zhijian Lu, *Student Member, IEEE*, Wei Huang, *Member, IEEE*, Mircea R. Stan, *Senior Member, IEEE*, Kevin Skadron, *Senior Member, IEEE*, and John Lach, *Senior Member, IEEE*

***Abstract*—Thermal effects are becoming a limiting factor in high-performance circuit design due to the strong temperature dependence of leakage power, circuit performance, IC package cost, and reliability. While many interconnect reliability models assume a constant temperature, this paper analyzes the effects of temporal and spatial thermal gradients on interconnect lifetime in terms of electromigration, and presents a physics-based dynamic reliability model which returns reliability equivalent temperature and current density that can be used in traditional reliability analysis tools. The model is verified with numerical simulations and reveals that blindly using the maximum temperature leads to too pessimistic lifetime estimation. Therefore, the proposed model not only increases the accuracy of reliability estimates, but also enables designers to reclaim design margin in reliability-aware design. In addition, the model is useful for improving the performance of temperature-aware runtime management by modeling system lifetime as a resource to be consumed at a stress-dependent rate.**

***Index Terms*—Dynamic stress, dynamic reliability management (DRM), dynamic thermal management (DTM), electromigration (EM), reliability-aware design, temperature gradients.**

I. Introduction

DUE to increasing complexity and clock frequency, temperature has become a major concern in IC design. Higher temperatures not only degrade system performance, raise packaging costs, and increase leakage power, but they also reduce system reliability via temperature enhanced failure mechanisms such as gate oxide breakdown, interconnect fast thermal cycling, stress-migration, and electromigration (EM). The introduction of low-k dielectrics in the future technology nodes will further exacerbate the thermal threats [1]. In this paper, we study temperature-related EM failure on interconnects.

Manuscript received October 26, 2005; revised May 24, 2006. This work is supported in part by the National Science Foundation under Grant CCR-0105626, Grant CCR-0133634, and Grant CCF-0429765, by a Grant from Intel MRL, by the Army Research Office (ARO) under Grant W911NF-04-1-0288, and by MARCO IFC.

Z. Lu was with the Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA 22904 USA. He is now with Marvell Technology Group, Santa Clara, CA 95054 USA (e-mail: zhijian@marvell.com).

W. Huang was with the Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA 22904 USA. He is now with the Department of Computer Science, University of Virginia, Charlottesville, VA 22904 USA (e-mail: wh6p@virginia.edu).

M. R. Stan and J. Lach are with the Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA 22904 USA (e-mail: zl4j@virginia.edu; wh6p@virginia.edu; mircea@virginia.edu; jlach@virginia.edu).

K. Skadron is with the Department of Computer Science, University of Virginia, Charlottesville, VA 22904 USA (e-mail: skadron@cs.virginia.edu).

Digital Object Identifier 10.1109/TVLSI.2007.893578

The field of temperature-aware design has recently emerged to maximize system performance under lifetime constraints. Dynamic thermal management (DTM) techniques [2], [3] are being developed to address this problem. While the dynamic temperature profile of a system is workload dependent [3], [4], several efficient and accurate techniques have been proposed to simulate transient chip-wide temperature distribution [4], [5], providing design-time knowledge of the thermal behavior of different design alternatives. Currently, DTM studies assume a fixed maximum temperature, which is unnecessarily conservative. To better evaluate these techniques and explore the design space, designers need better information about the lifetime impact of temperature.

Failure probability in VLSI interconnects due to EM is commonly modeled with lognormal reliability functions. The variability of lifetime is strongly dependent on the interconnect structure geometries and weakly dependent on environmental stresses such as current and temperature [6], while median time to failure (MTF) is determined by current and temperature in the interconnect. In this paper, we use MTF as the reliability metric and investigate how it is affected by temporal and spatial thermal gradients. Historically, Black [7] proposed a semi-empirical temperature-dependent equation for EM failures

$$T_f = \frac{A}{j^n} \exp\left(\frac{Q}{kT}\right) \tag{1}$$

where T_f is the time to failure, A is a constant based on the interconnect geometry and material, j is the current density, Q is the activation energy (e.g., 0.6 eV for aluminum), and kT is the thermal energy. The current exponent n, has different values according to the actual failure mechanism. It is assumed that $n = 2$ for void nucleation limited failure and $n = 1$ for void growth limited failure [8]. Black's equation is widely used in thermal reliability analysis and design.

However, Black's equation assumes a constant temperature. For interconnects subject to temporal and/or spatial thermal gradients, two questions need to be answered: 1) Is Black's equation still valid for reliability analysis in these cases? and 2) If Black's equation is valid, what temperature should be used? Though in absence of clear answers in the literature, in practice, Black's equation is still widely assumed, and the worst case temperature profile is usually used to provide safeguard, resulting in pessimistic estimations and unnecessarily restricted design spaces. As an example of why worst case temperature assumptions may severely underestimate reliability, we use the *Hotspot* toolset [4], a validated architecture-level compact thermal model, to simulate a processor running the Spec2000 benchmarks. The temperature and the power of the hottest block

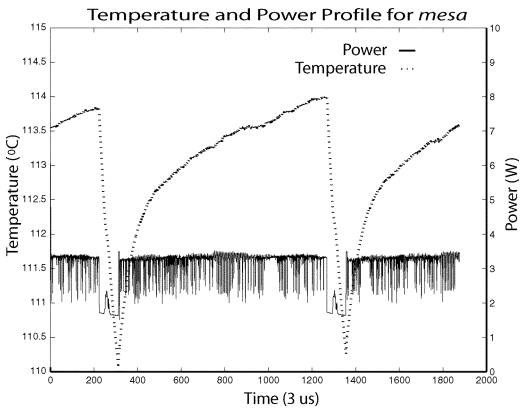

Fig. 1. Simulated temperature/power profile for an integer unit running the *mesa* Spec2000 benchmark. [9].

(i.e., the integer unit) for one benchmark are plotted in Fig. 1. In this case, the substrate temperature varies between 110 and 114 °C, and the maximum power is more than 1.5 times the minimum power. We can see that for only a small portion of time is the program running at the worst case temperature.

In the first part of this paper, we answer the previous two questions. We find that, for EM subject to time-varying stresses, Black's equation is still valid, but only with the reliability equivalent temperature/current density derived from a dynamic reliability model presented in this paper [9]. For EM subject to spatial thermal distribution, Black's equation cannot be applied directly. Instead, we give the bounding temperatures which can be used in Black's equation to bound the actual lifetime subject to nonuniform temperature distribution. Therefore, our results can be seamlessly integrated into current reliability analysis tools based on Black's equation [10]. In addition, while designers are currently constrained by constant, worst case temperature assumptions, the analysis presented in this paper provides more accurate, less pessimistic interconnect lifetime predictions. This results in fewer unnecessary reliability design rule violations, enabling designers to more aggressively explore a larger design space. One limitation in the application of our results is that our analysis is currently based on two-terminal interconnects, such as those seen in global signal interconnects and power/ground distribution networks. Recently, Alam *et al.* [10] proposed lifetime predictions for multiterminal interconnects. Our future work will include extending our current findings to multiterminal interconnects.

Worst case power dissipation and environmental conditions are rare for general purpose microprocessors. Designing the cooling solution for the worst case is wasteful. Instead, the cooling solution should be designed for the worst "expected" case. In the event that environmental or workload conditions exceed the cooling solution's capabilities and temperature rises to a dangerous level, on-chip temperature sensors can engage some form of DTM [4], [11], which sacrifices a certain amount of performance to maintain reliability by reducing circuit speed whenever necessary.

However, existing DTM techniques do not consider the effects of temperature fluctuations on lifetime and may unnecessarily impose performance penalties. In the second part of this paper, we propose runtime dynamic reliability management (DRM) techniques based on our dynamic reliability model [12]. By leveraging this model, one can dynamically track the "consumption" of chip lifetime during operation. In general, when temperature increases, lifetime is being consumed more rapidly, and *vice versa*. Therefore, if temperature is below the traditional DTM engagement threshold for an extended period, it may be acceptable to let the threshold be exceeded for a time, while still maintaining the required expected lifetime. In effect, lifetime is modeled as a resource that is being "banked" during periods of low temperature, allowing for future withdrawals to maintain performance during times of higher operating temperatures. Using EM, as an example, we show the benefits of lifetime banking by avoiding unnecessary DTM engagements while meeting expected lifetime requirements.

The concept of DRM is first introduced by Srinivasan *et al.* [3]. In their work, they proposed a chip level reliability model and showed the potential benefits of trading off reliability with performance for individual applications. They assumed an oracular algorithm for runtime management in their study, and they did not consider the effects of inter-application thermal behaviors on reliability. In this paper, we focus on practical runtime management techniques for the worst case on-chip component (i.e., hottest interconnect) to exploit both intra- and inter-application temperature variations. Ramakrishnan and Pecht [13] proposed to monitor the life consumption of an electronic system and project the system lifetime based on the monitoring results. We take a similar approach to monitoring the stresses on the circuit continuously, and we also intelligently adapt the circuit operation to maximize the circuit performance without reducing reliability.

The rest of this paper is organized as follows. Section II introduces a stress-based analytic model for EM, which serves as the base model in this paper. In Section III, we extend this model to cope with time-varying stresses (i.e., temperature and current) and derive a formula to estimate interconnect lifetime, which we analyze in Section IV. In Section V, we analyze the impact of nonuniform temperature distribution on lifetime prediction due to EM. In Section VI, we exploit our proposed dynamic reliability model in runtime thermal management and propose a banking-based dynamic reliability management technique to improve system performance while maintaining lifetime constraints. Finally, we summarize the paper in Section VII.

II. Analytic Model for EM

Clement [14] provides a review of 1-D analytic EM models. Several more sophisticated EM models are also available [8], [15]. In this paper, we only discuss the EM-induced stress build-up model of Clement and Korhonen [16], [17], which has been widely used in EM analysis.

EM is the process of self-diffusion due to the momentum exchange between electrons and atoms. The dislocation of atoms causes stress build-up according to the following equation [16], [17]:

$$\frac{\partial \sigma}{\partial t} - \frac{\partial}{\partial x}\left(\left[D_a\left(\frac{B\Omega}{kTl^2\varepsilon}\right)\right]\left(\frac{\partial \sigma}{\partial x} - \frac{qlE}{\Omega}\right)\right) = 0 \quad (2)$$

where $\sigma(x,t)$ is the stress function, and an interconnect failure is considered to happen when $\sigma(x,t)$ reaches a threshold (critical) value σ_{th}. D_a is the diffusivity of atoms, a function of

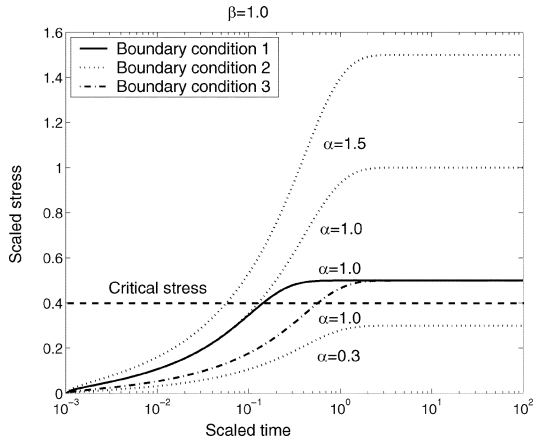

Fig. 2. EM stress build-up for different boundary conditions and α values. All processes have $\beta = 1$ [α and β are defined in (4)] [9].

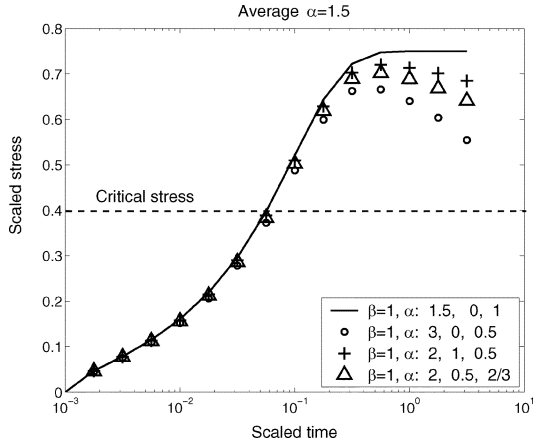

Fig. 3. EM stress build-up under time-dependent current stress. In each EM process, α [defined in (4)] oscillates between two values with different duty cycles. The time dependence of α is given in the legend. All curves have the same average value of α. The solid line is the stress build-up with a constant value of α. [9].

temperature. B is the appropriate elastic modulus, depending on the properties of the metal and the surrounding material and the line aspect ratio. Ω is the atom volume. ε is the ratio of the line cross-sectional area to the area of the diffusion path. l is the characteristic length of the metal line (i.e., the length of the effective diffusion path of atoms). q is the effective charge. E is the applied electric field, which is equal to ρj, the product of resistivity and current density. The term qlE/Ω corresponds to the atom flux due to the electric field, while $\partial\sigma/\partial x$ corresponds to a back-flow flux created by the stress gradient to counter balance the EM flux. The total atomic flux at a specific location in the interconnect is proportional to the sum of these two components

$$J = \left[D_a\left(\frac{B\Omega}{kTl^2\varepsilon}\right)\right]\left(\frac{\partial\sigma}{\partial x} - \frac{qlE}{\Omega}\right). \tag{3}$$

Equation (2) states that the mechanical stress build-up at any location is caused by the divergence of atomic flux at that point, or $\partial\sigma/\partial t = \nabla J$. If we assume a uniform temperature across the interconnect characteristic length and let $\beta(T) = D_a\left(B\Omega/kTl^2\varepsilon\right)$ (which we refer to as the temperature factor throughout the paper) and $\alpha(j) = qlE/\Omega$, we obtain the following simplified version, the solution of which depends on both temperature and current density:

$$\frac{\partial\sigma}{\partial t} - \beta(T)\frac{\partial}{\partial x}\left(\frac{\partial\sigma}{\partial x} - \alpha(j)\right) = 0. \tag{4}$$

Clement [16] investigated the effect of current density on stress build-up using (4), assuming that temperature is unchanged (i.e., $\beta(T) =$ constant), for several different boundary conditions. He found that the time to failure, derived from this analytic model, had exactly the same form as Black's (1). The exponential component in Black's equation is due to the atom diffusivity's (D_a's) dependency on temperature by the well-known Arrhenius equation: $D_a = D_{ao}\exp\left(-Q/kT\right)$.

At any time t, the maximum stress along a metal line can be found at the boundaries of the interconnect line. Fig. 2 shows the numerical solutions for (4) at one end of the line (i.e., $x = 0$) for different boundary conditions and α values, all with $\beta = 1$. The three boundary conditions shown here are similar to those discussed in [16] for finite length interconnect lines. It indicates that both boundary conditions and current density (α) affect the stress build-up rate (i.e., the larger the current, the faster the stress builds up). Also, seen from the figure is that the stress build-up saturates at a certain point. This is because, in saturation, the atom flux caused by EM is completely counter balanced by the stress gradient along the metal line. It is believed that the interconnect EM failure occurs whenever the stress build-up reaches a critical value, σ_{th} (as shown in Fig. 2). If the saturating stress is below the critical stress, no failure happens. In the following discussion, we assume that the saturating stress in an EM process is always above the critical stress.

III. EM Under Dynamic Stress

A. Time-Dependent Current Stress

Clement [16] used a concentration build-up model similar to the one discussed here to verify that in the case in which temperature is kept constant, the average current density can be used in Black's equation for pulsed dc current. As for ac current, an EM effective current is used by the average current recovery (ACR) model [18], [19]. In this paper, we do not distinguish between these two cases. We only consider the change of EM effective current due to various causes (e.g., phased behaviors in many workloads). This is because the time scale of the current variation studied in this paper is usually much longer than that of the actual dc/ac current changes in the interconnects.

We numerically solve (4) with different time-dependent α functions, and the results are plotted in Fig. 3.[1] The stress build-ups for all EM processes in Fig. 3 overlap before saturations (or before reaching the critical stress), since they have the same average current. Thus, the EM process under time-varying current stress can be well approximated by average current. Note that the curves in Fig. 3 diverge after they reach their maximum stress. This is because the time-varying current could not create a stable counterbalancing stress gradient for

[1]For example, the numbers after the circle represent the case in which α is a square-wave function and varies between 3 and 0 with a duty cycle of 0.5. This representation of the time-dependent square-wave function is used in other figures throughout this paper.

EM. However, we are only interested in the EM process before reaching the critical stress when EM failure occurs.

B. Time-Dependent Thermal Stress

If the temperature (or $\beta(T)$) of the interconnect is time dependent, EM stress build-up can be derived as (detailed derivation can be found in [9], [20])

$$\sigma_0(t) = \sigma_1\left(\left(\frac{1}{\beta_1}\right)\int_0^t \beta(T(t))dt\right). \tag{5}$$

If we assume that the stress build-up reaches a certain threshold ($\sigma_{\rm th}$) at which an EM failure occurs, we have

$$\int_0^{t_{\rm failure}} \beta(T(t))dt = \varphi_{\rm th} \tag{6}$$

where $\varphi_{\rm th}$ is a constant determined by the critical stress (i.e., $\varphi_{\rm th} = \sigma_1^{-1}(\sigma_{\rm th})\,\beta_1$). If an average value of $\beta(t)$ exists, we obtain a closed form for the time to failure

$$t_{\rm failure} = \frac{\varphi_{\rm th}}{E(\beta(T(t)))} \tag{7}$$

where $E(\beta(t))$ is the expected value for $\beta(t)$, and $\beta(t)$ is the temperature factor, as defined in (4), having the form $\beta(T(t)) = A'\left(\exp\left(-Q/kT(t)\right)/kT(t)\right)$, where A' is a constant. In comparison with Black's equation, (7) indicates that the average of temperature factor β should be used.

One way to interpret (6) is to consider interconnect time to failure (i.e., interconnect lifetime) as an available resource, which is consumed by the system over time. Then the $\beta(t)$ function can be regarded as the consumption rate.

Let MTF(T) be the time to failure with a constant temperature T. We have $\beta(T) = \varphi_{\rm th}/{\rm MTF}(T)$ by (7). Substitute this relation in (7) again and consider the time-varying temperature, and we obtain an alternative form for (7)

$$t_{\rm failure} = \frac{1}{E\left(\frac{1}{{\rm MTF}(T)}\right)}. \tag{8}$$

Equation (8) can be used to derive the absolute time to failure provided that we know the time to failure for different constant temperatures (e.g., data from experiments). Notice that $\beta(T)$ is a convex function within the operational temperatures. A constant temperature is always better in terms of EM reliability than oscillating around that temperature.

Similar to the methods for verifying the "average current model," we obtain numerical solutions for the stress build-up equation using different square waveforms for β. Fig. 4 compares these results and shows that the time to failure will be the same as long as the EM processes exhibit the same *average* value of β.

C. Combined Dynamic Stress

In reality, both temperature and current change simultaneously. In most cases, the variation of temperature on chip reflects changes in power consumption, thus, directly relating to current flow in the interconnects. It can be shown that the lifetime of an interconnect subject to simultaneous thermal and current variations can be estimated by (detailed derivation can be found in [9], [20])

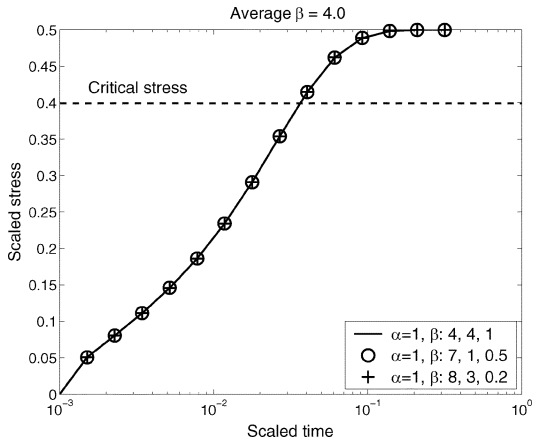

Fig. 4. EM stress build-up at one end of the interconnect with different time-dependent β functions (square waveform). The solid line is the case with a constant value of β equal to the average value of β in other curve. [9].

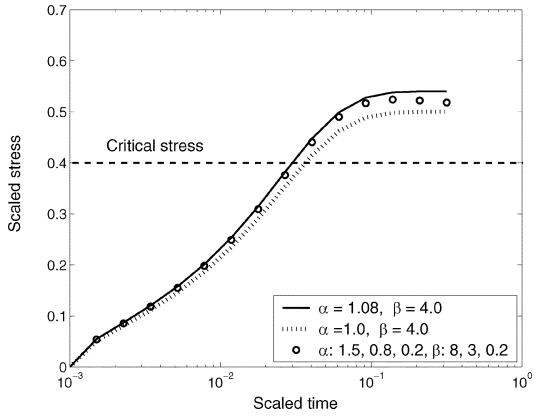

Fig. 5. EM stress build-up at one end of the interconnect with time-varying α (current) and β (temperature) functions (i.e., square waveforms). The circles represent the numerical solution for time-varying α and β. The solid line is with constant values of α and β calculated according to (10). The EM process (dotted line) simply using the average current is also shown. [9].

$$t_{\rm failure} = \frac{C}{j_{\rm equivalent}^n \beta(T_{\rm equivalent})} \tag{9}$$

where C is a constant, and $j_{\rm equivalent}$ and $T_{\rm equivalent}$ are reliability equivalent current density/temperature defined as

$$j_{\rm equivalent} = \frac{E\left[j(t)\beta(t)\right]}{E\left[\beta(t)\right]},\ \beta(T_{\rm equivalent}) = E(\beta(T(t))). \tag{10}$$

Fig. 5 compares the stress build-ups for different dynamic current and temperature combinations. These results illustrate that the EM process under dynamic stresses can be well approximated by an EM process with a constant temperature (i.e., $T_{\rm equivalent}$) and a constant current [i.e., $I_{\rm equivalent}$ as defined in (10)]. Therefore, for an interconnect with concurrent time-dependent temperature and current stresses, time to failure has the same form as Black's equation, except that the reliability-equivalent current [the actual current modulated by the temperature

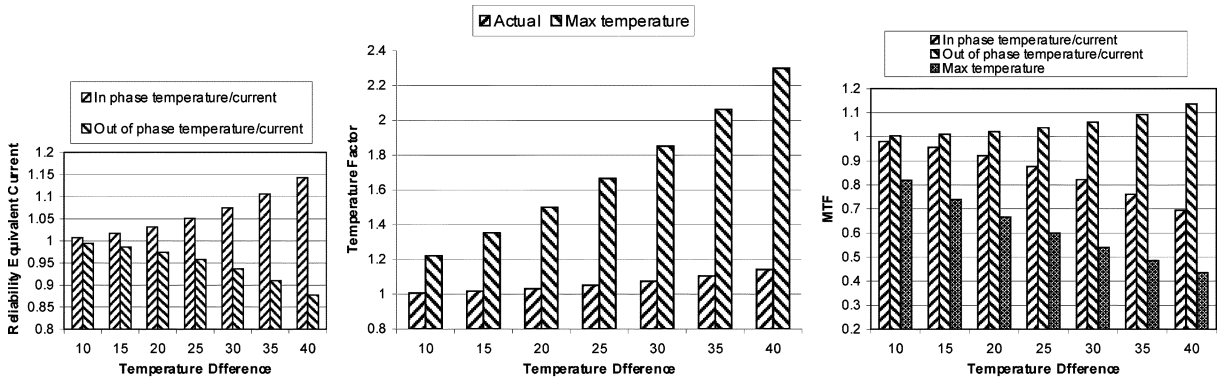

Fig. 6. Comparison of electric current, temperature factor (β), and MTF for different peak-to-peak temperature cycles. All results are normalized to the average current and/or temperature case. (a) Ratio of reliability equivalent current (our model) to average current. Both cases of current variation (in and out of phase with temperature) are included. (b) Ratios of temperature factor (β) using average temperature, max temperature, and our model. (c) Comparison of MTF for four different calculations: average temperature/average current, maximum temperature/average current, our model for current in phase with temperature, and our model for current out of phase with temperature. [9].

factor β (i.e., weighted averaging by β)]) and the mean value of the temperature factor are used.

In fact, if the current and the temperature are statistically independent, we have $E[j(t)\beta(t)]/E[\beta(t)] = E[j(t)]$ in (10). In this case, the reliability equivalent current will be reduced to the average current and we get back to the "average current model." On the other hand, if the current is constant, (9) and (10) will lead us to (7). Finally, if both temperature and current are time invariant, Black's equation [(1)] is obtained.

IV. Analysis of the Proposed Model

Equations (9) and (10) form the basis of our proposed EM model under concurrent time-varying temperature and current stress. In this section, we use these equations to evaluate EM reliability. Specifically, we compare the reliability of constant temperature with that of fluctuating temperature, and we show the difference of lifetime projection between our model and the traditional worst case model.

For any two temporal temperature and current profiles, we can easily compare the EM reliability, using our model, by

$$\frac{\mathrm{MTF}_1}{\mathrm{MTF}_2} = \frac{j^2_{\mathrm{equivalent2}} E(\beta(T_2(t)))}{j^2_{\mathrm{equivalent1}} E(\beta(T_1(t)))}$$

where MTF_1 is the time to failure under time-varying temperature profile $T_1(t)$ and electric current profile $j_1(t)$.

In real workload execution, temperature changes along with the changes in power consumption (i.e., current). The possible interactions between temperature and current form a spectrum. We study their impacts on reliability by using temperature/current profiles representing the two extremes of this spectrum: 1) Temperature is in phase with current (i.e., temperature increases as current in the interconnect increases) and 2) Temperature is out of phase with current (i.e., temperature increases while current is reduced). The detailed waveforms used in our evaluation are explained in [9] and [20]. With these temperature/current profiles, we compare the reliability equivalent current with the average current, compare the temperature factor using our model with those of average and maximum temperatures, and finally compare the MTFs calculated in different ways (e.g., using average current/average temperature, reliability equivalent current/average temperature factor (β), or average current/maximum temperature). Our results are reported in Fig. 6, and we summarize our observations as follows.

- As the peak-to-peak temperature difference is small, both the reliability equivalent current and the temperature factor predicted by our dynamic stress model are very close to those calculated from using average current and average temperature. Thus, the MTF predicted by using average temperature/current provides a simple method for reliability evaluation with high accuracy.
- As the temperature difference increases, we can no longer simply use average temperature/current for MTF prediction. Both the reliability equivalent current and the temperature factor increase (degrading reliability) quickly as the temperature difference increases.
- As one expects, using maximum temperature always underestimates the lifetime, resulting in excessive design margins.
- One interesting phenomenon occurs in the case where the current is out of phase with temperature. Recall that the reliability equivalent current is actually a temperature factor weighted average current, and high temperature increases the weights for the accompanied current. Thus, the reliability equivalent current is reduced compared to the case in which temperature/current are synchronized. This brings a nonintuitive effect on the reliability projection—MTF even slightly increases as the temperature cycling magnitude increases.

In the previous discussion, the duty cycle of the current waveform is fixed (i.e., 0.5). We also investigated the effects of different duty cycles, but the data is not shown here due to space limitations. In general, when the temperature change is small (e.g., within 10 °C), using the average temperature to predict lifetime is still a good approximation (less than 5% error) regardless of the duty cycle. While the temperature variation increases, the difference between our model and using average temperature is largest at a duty cycle of about 0.4. On the other hand, the smaller the duty cycle, the larger the difference between our model and using maximum temperature. Thus, using

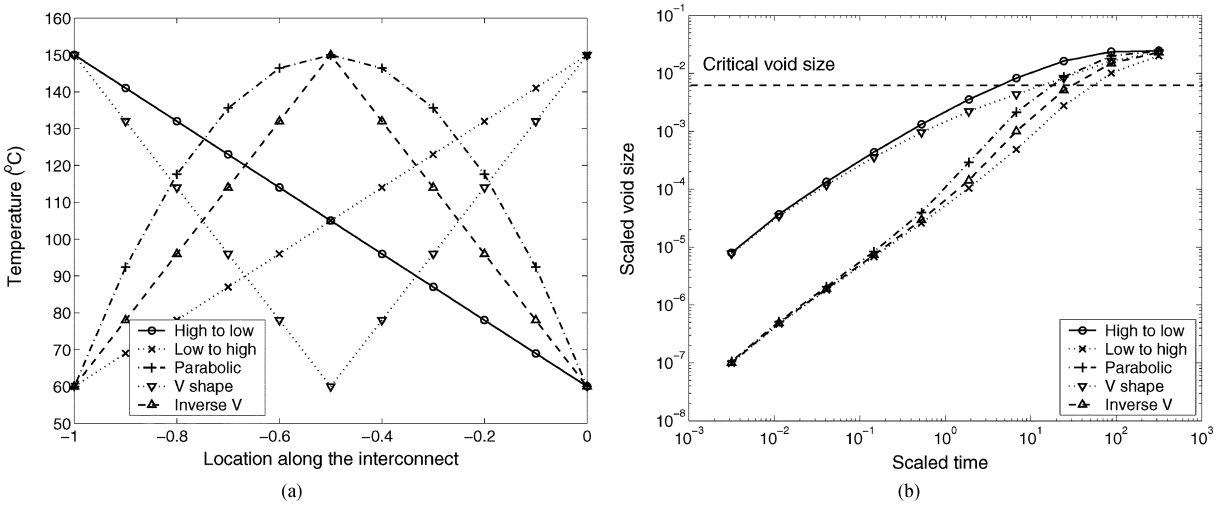

Fig. 7. Effects of nonuniform spatial temperature distribution on EM induced void growth. (a) Various temperature profiles along a 100 μm Cu interconnect (left end is the cathode). (b) Void growth with different spatial temperature profiles.

maximum temperature is reasonable only when the duty cycle is large (i.e., higher temperature dominates almost the entire cycle).

V. EM Under Spatial Temperature Gradients

In addition to temporal temperature variations, large temperature differences across the chip are commonly seen in modern VLSI design. Ajami *et al.* [21] showed that nonuniform temperature has great impacts on interconnect performance. In this section, we will illustrate the importance of considering spatial temperature gradients for interconnect reliability.

A. EM Model With Spatial Thermal Gradients

Due to the exponential dependence of diffusivity on temperature, EM in interconnects with spatial temperature gradients has quite different characteristics than those with constant temperature. Guo *et al.* [22] reported that EM in aluminum interconnect is strongly affected by the relative direction of electron wind and thermal gradients, while Nguyen *et al.* [23] found that temperature gradients greatly enhance EM in aluminum interconnect. Following the stress build-up model introduced in Section II, the atomic flux due to EM can be modeled by $J = \beta(T)\left(\partial\sigma/\partial x - \alpha(j)\right)$, and the stress build-up at a specific location is caused by the divergence of atomic flux at that location, i.e., $\partial\sigma/\partial t = \nabla J$. When the temperature is uniform across the interconnect, i.e., $\beta(T)$ is independent of location, (4) is obtained. When the temperature is not uniform, the following equation is derived to describe the stress build-up under thermal gradients

$$\frac{\partial\sigma}{\partial t} - \beta(T(x))\frac{\partial}{\partial x}\left(\frac{\partial\sigma}{\partial x} - \alpha(j)\right) - \frac{\partial\beta(T(x))}{\partial x}\left[\frac{\partial\sigma}{\partial x} - \alpha(j)\right] = 0 \tag{11}$$

where σ, β, and α are defined in Section II. When compared with (4), (11) introduces a third term $(\partial\beta(T(x))/\partial x)\left[\partial\sigma/\partial x - \alpha(j)\right]$, which captures the atomic flux divergence induced by spatial thermal gradients along the interconnect. Although temperature gradient itself will cause migrations of atoms from high temperature to low temperature, a phenomenon called thermomigration (TM), the atomic flux due to TM is generally believed to be much smaller than that due to EM [23]. Therefore, TM is not explicitly modeled in (11). Jonggook *et al.* [24] investigated EM in aluminum (Al) interconnects subject to spatial thermal gradients. They modeled EM from a different approach but yielded an equation similar to ours. Since dual-damascene Cu interconnects have become the mainstream technology in modern VLSI design and have quite different EM characteristics from Al [25], in the following, we focus on EM failure in Cu interconnects.

Various experiments [6], [26] showed that, in Cu interconnect, voids tend to nucleate at the cathode end (near the via), and void growth is the dominant failure process because the critical mechanical stress for void nucleation in Cu is much smaller than that for aluminum. With spatial thermal gradients in the interconnect, it is possible that the location of void nucleation is no longer at the cathode end. However, in this case, void growth tends to be slower than that at the cathode, because there are atomic fluxes both going into and coming from the void in the middle of the interconnect [26]. Bearing these observations in mind and assuming a void-growth dominated failure, we choose a boundary condition for (11) to model void growth at the cathode such that the mechanical stress at the cathode end is zero (stress free at void) and the atomic flux at the other end is zero (complete blockage for atomic flux), or

$$\sigma(x = -l, t) = 0,\ J(x = 0, t) \Rightarrow \left[\frac{\partial\sigma}{\partial x} - \alpha(j)\right]\Bigg|_{x=0} = 0$$

where $x = -l$ is the cathode end. This boundary condition is consistent with that used by Clement [14] to model void growth due to EM. The void size at time t can be approximated by [14]

$$\Delta l \approx \int_{-l}^{0} \frac{-\sigma(x,t)}{B} dx$$

where $\sigma(x,t)$ is the mechanical stress (tensile stress) developed along the interconnect at time t and B is the elastic modulus. Because we are unaware of any closed-form solution for (11) with the previous boundary condition, we use numerical solutions to analyze the impact of thermal gradients on EM.

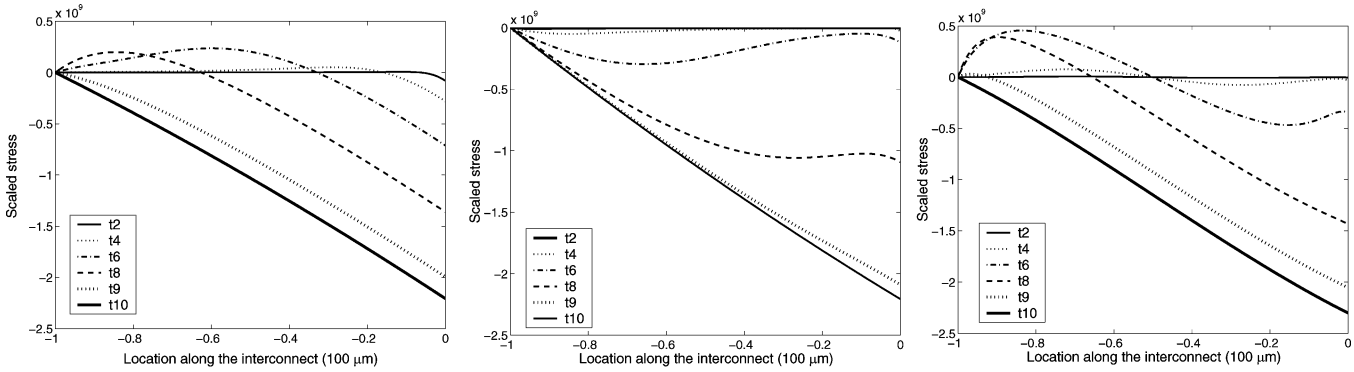

Fig. 8. Stress build-ups at different time points along the interconnect under spatial thermal gradients. (a) Low to high temperature profile. (b) High to low temperature profile. (c) Parabolic temperature profile. Electrons flow from the left to the right, causing compressive stress (negative in the figures) on the right side of the interconnect. The left end (cathode) is stress free to model the growth of a void. The time points ("t2" through "t10") are corresponding to the time points in the plots with the same temperature profile in Fig. 9.

The temperature spatial profile along an interconnect is the combined effects of joule heating and substrate temperature distributions. Fig. 7(a) plots several temperature profiles used in our study and their effects on EM induced void growth. The length of the interconnect is 100 μm, and electrons are assumed to flow from the left end (cathode) to the right end of the interconnect. Though all temperature profiles have the same maximum and minimum temperatures, their void growth differs greatly due to the different thermal gradients along the interconnect Fig. 7(b), resulting quite different failure time. In order to investigate how thermal gradients affect EM induced void growth, we also plot, in Fig. 8, the mechanical stress build-up along the interconnect at different times, with different thermal profiles. In spite of different temperature profiles on the interconnect, in the final EM process stage ("t10" in Fig. 8), a steady stress gradient is built up to counter balance the driving force of electron wind, i.e., $\partial\sigma/\partial x - \alpha(j) = 0$, resulting in voids with comparable saturation sizes Fig. 7(b).

However, as shown in Fig. 8, the kinetic aspects of stress build-up for different temperature profiles are quite different, especially when the relative direction of electron wind and temperature gradients changes. For example, in a "low to high" temperature profile, the temperature increases linearly from the cathode end to the anode end [shown in Fig. 7(a)]. At the early EM stage, as indicated by "t2" and "t4" in Fig. 8(a), the stress gradient near the cathode is negligible. Therefore, the atomic flux at the cathode is only determined by the electron wind at the temperature of that location. Later on in the EM process, the effect of thermal gradients begins to play its role. As shown by "t6" and "t8" in Fig. 8(a), tensile (positive) stress is built up from the cathode end and increases towards the other end, due to the increasing temperature from the cathode end. The resulting stress gradient forms an atomic flux in the same direction as the electron wind. Thus, void growth in the cathode end is enhanced later by the increasing temperature. Similar analysis can be performed for EM stress build-up under different spatial temperature profiles such as those illustrated in Fig. 8(b) and (c) (detailed discussion can be found in [20]). In summary, in the early stage of the EM process, the void growth is largely dependent on the temperature at the cathode, while later on, the void growth is enhanced or retarded depending on the temperature gradient near the cathode. Finally, void growth is suppressed by the back-flow stress gradient just like in the case with a uniform temperature distribution.

B. Empirical Bounds for Void Growth With Nonuniform Temperature Distribution

In Section III, our analysis reveals that the EM process with time-varying temperature variations can be approximated by an EM process using a constant reliability equivalent temperature $T_{\rm eq}$, as long as $\beta(T_{\rm eq}) = E[\beta(T(t))]$. However, in the case where there is a nonuniform temperature across the interconnect, we cannot find a similar reliability equivalent temperature, due to the difference in the EM kinetics in the different stress build-up stages as shown in Fig. 8. Instead, we try to find two constant temperatures, such that the void growth due to EM with nonuniform temperature can be bounded by the void growth with uniformly distributed temperature equal to these two bounding temperatures, respectively. By providing the bounding temperatures, one can still use Black's equation to evaluate the reliability effects of nonuniform temperature distributions. Table I summarizes our proposed formulas for calculating the EM bounding temperatures (detailed derivation can be found in [20]).

Notice that since void growth at the cathode is only dependent on the atomic flux nearby, the temperature gradient near the cathode plays the major role in determining (enhancing or retarding) the void growth. In our formulas for calculating the bounding temperatures, only temperature distribution along the first half (cathode side) of the interconnect is used. This observation is verified by testing with various temperature profiles (Due to space limitations, we cannot show them all here). The void growth with different temperature profiles as well as those with uniform temperatures are compared in Fig. 9. In these plots, the void growth with spatial thermal gradients is closely bounded by the void growths with the calculated bounding temperatures. Blindly using the average temperature to evaluate the EM lifetime will either overestimate [e.g., Fig. 9(a)] or underestimate [e.g., Fig. 9(b)] the void growth, let alone using the maximum temperature.

TABLE I
PROPOSED BOUNDING TEMPERATURES FOR VOID GROWTH IN AN INTERCONNECT WITH LENGTH l SUBJECT TO A NONUNIFORM TEMPERATURE DISTRIBUTION. $T_{\rm lb}$ IS THE LOWER BOUND. $T_{\rm ub}$ IS THE UPPER BOUND. $T(x)$ IS THE TEMPERATURE PROFILE. $x = -l$ AND $x = 0$ ARE THE LOCATIONS OF THE CATHODE AND THE ANODE, RESPECTIVELY (AS SHOWN IN FIG. 7)

Temperature gradient at cathode	Lower bound temperature	Upper bound temperature
Increasing temperature in the electron wind direction (EM enhanced)	$T_{lb} = T_{cathode}$	$\beta(T_{ub})(\alpha(T_{ub})) = \frac{\int_{-l}^{-\frac{l}{2}} \alpha(T(x))dx}{\int_{-l}^{-\frac{l}{2}} \frac{1}{\beta(T(x))}dx}$
Decreasing temperature in the electron wind direction (EM retarded)	$\beta(T_{lb})(\alpha(T_{lb})) = \frac{\int_{-l}^{-\frac{l}{2}} \alpha(T(x))dx}{\int_{-l}^{-\frac{l}{2}} \frac{1}{\beta(T(x))}dx}$	$T_{ub} = T_{cathode}$

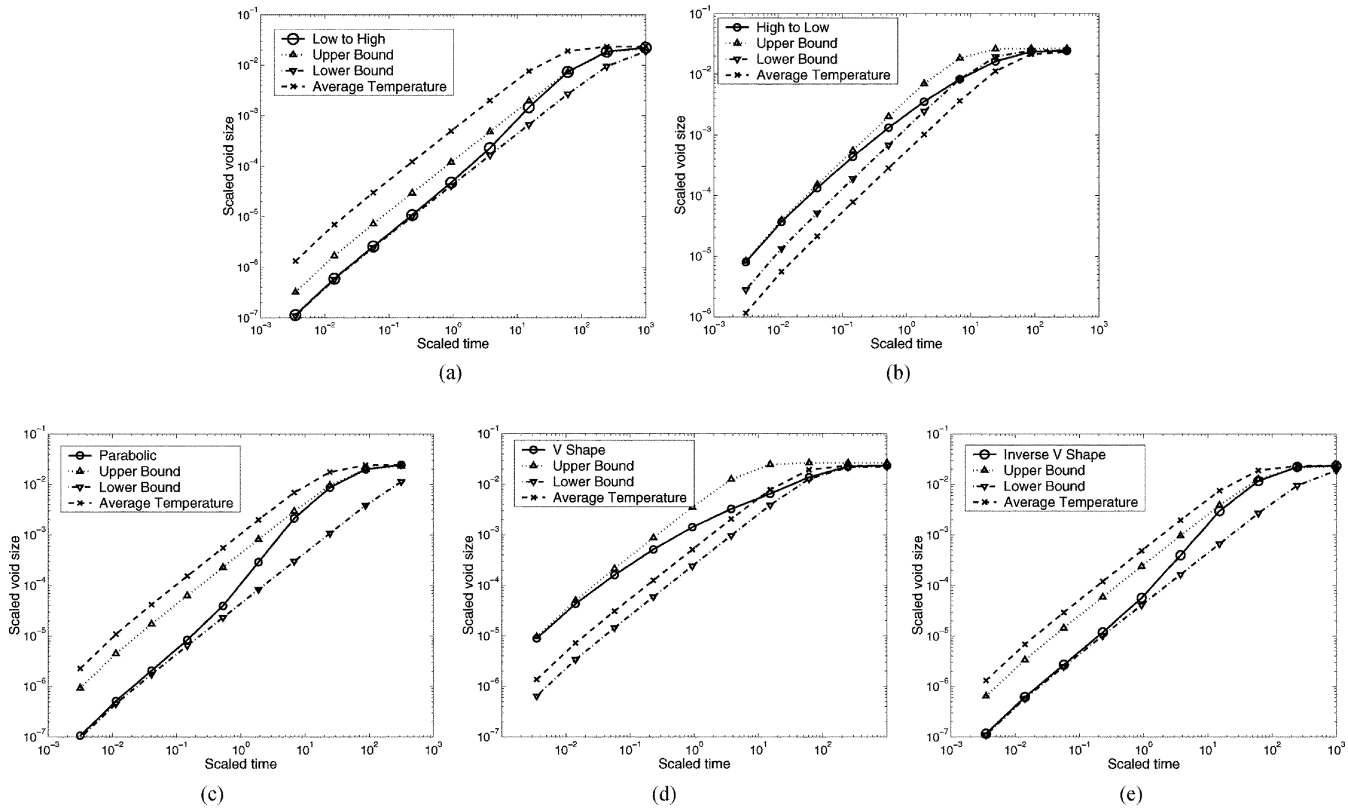

Fig. 9. Void growth with nonuniform temperature distribution is bounded by those with a uniformly distributed temperature. (a) Low to high temperature profile. (b) High to low temperature profile. (c) Parabolic temperature profile. (d) V-shape temperature profile. (e) Inverse V-shape temperature profile.

Joule heating in an interconnect usually results in a symmetric temperature distribution with the maximum temperature in the middle, due to the much lower thermal resistance of the vias on both ends. Therefore, the symmetric temperature distributions along the interconnect are of more practical interest. The"parabolic" and "inverse V-shape" temperature profiles shown in Fig. 7(a) are used to represent similar temperature distributions. Interestingly, for these temperature distributions, as indicated by Fig. 9(c) and (e), even the upper bound temperature for void growth is lower than the average temperature. In the EM measurements of Cu interconnects performed by Meyer *et al.* [27], they considered the nonuniform temperature distribution due to self-heating, and tried to fit their measurements with Black's equation by using different temperatures (e.g., maximum, average, weighted average, via (minimum) temperature). They reported that the best fit temperature is strongly weighted to the via temperature. Their findings agree with our analysis presented here.

C. Concurrent Temporal and Spatial Thermal Gradients

So far we have discussed the interconnect lifetime prediction under temporal and spatial temperature distributions separately. In practice, due to circuit activity variations, one might expect the spatial temperature distribution over an interconnect would change over time. The EM diffusion equation [see (2) or (11)] can be extended to capture this situation by assuming that temperature T is a function of both time and interconnect location. However, we cannot obtain a closed-form analytic solution in this highly complex scenario. Instead, we propose to combine the results we have found so far in the cases of temporal gradients only and spatial gradients only to estimate the interconnect

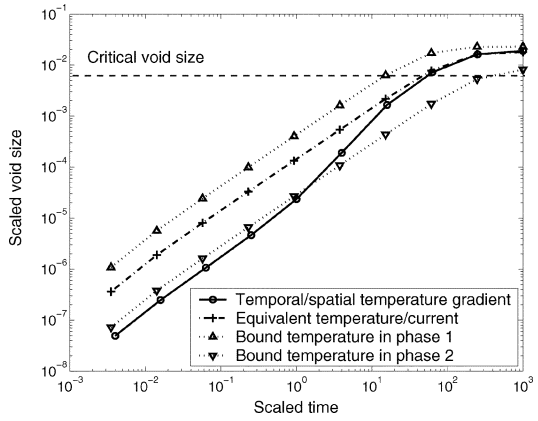

Fig. 10. Void growth subject to both temporal and spatial thermal gradients can be bounded by that using uniform temperature and current.

lifetime subject to both temporal and spatial temperature gradients.

At time t, the temperature profile across an interconnect is denoted by $T(t,x)$ and the current density is $j(t)$. Using the formulas in Table I, we can find the bounding temperature at t, denoted by $T_b(t)$. Applying (9) and (10) for the case of temporal temperature gradients to $T_b(t)$ and $j(t)$, we could find an equivalent uniform temperature and current to approximate the void growth subject to both temporal and spatial temperature gradients. Fig. 10 shows one example. In this example, the interconnect Cu line is subject to two "parabolic" temperature profiles, with each one for half the time (i.e., 50% duty cycle), denoted by "phase 1" and "phase 2" in the figure. We solve (11) numerically with temperature being a function of both time and space, and we plot the void growth. This figure indicates that the void growth subject to the time varying temperature profile can be bounded by that using time invariant uniform temperature and current, as calculated according to the procedures proposed here. As a comparison, we also plot the void growth at the bound temperature of each temperature profile alone. If the critical void size is close to the saturation void size, as shown in the figure, one can use the calculated equivalent temperature and current to estimate the interconnect lifetime subject to both temporal and spatial thermal gradients, using Black's equation [(1)]. We have also tested this finding for other temperature profiles and duty factors, and the results are similar but are not presented here due to space limitations.

VI. Runtime Reliability Aware Thermal Management

The proposed dynamic EM reliability model not only increases the accuracy of reliability estimates, it can also be integrated in the existing IC design flow, allowing designers to reclaim significant design margins from the traditional conservative design approaches. Examples to exploit dynamic reliability model for design time optimizations can be found in [9] and [20]. In this section, we illustrate how our dynamic reliability model can be applied to improve system performance with runtime DRM.

Many DTM techniques [4], [11], [28] have been proposed to ensure that a chip will never operate above some temperature threshold. However, these techniques do not explicitly study the effects of transient behaviors on system reliability, and instead implement a temperature upper bound at the expense of degraded performance. By modeling lifetime as a resource to be consumed over time, we can manipulate chip lifetime directly at runtime.

High temperature limits the circuit performance directly by increasing interconnect resistance and reducing carrier mobility. However, it has been shown that using DTM to compensate the temperature dependency of clock frequency induces very mild performance penalty [4]. On the other hand, Banerjee *et al.* [1] showed that temperature induced reliability issue tends to limit the circuit performance in future technology generations. In the following discussion, we assume that the temperature threshold is set solely for reliability specification, and circuits can operate correctly above this threshold whenever allowed. Although extreme high temperature may cause immediate thermal damage for IC circuits, we study a range of operating temperatures only with longterm reliability impacts (i.e., temperature induced aging). High temperatures causing immediate or unrecoverable damage are assumed to be far above the range of normal operating temperatures studied here (e.g., the temperature used in accelerated EM test is usually around 200°C [29]). A monitoring and feedback mechanism is implemented at runtime to ensure that circuits operate well below such temperatures.

A. Lifetime Banking Opportunities

Fig. 11 depicts the temperature profiles for two different workloads commonly seen in general purpose computing. Fig. 11(a) shows a single program workload and Fig. 11(b) illustrates a multiprogram workload with context switching. In the single program workload, temperature changes over time due to activity variability in the program execution. In the multiprogram workload, besides the execution variations within each program, inter-program thermal differences also affect the overall thermal behavior of the workload. For example, in Fig. 11(b), the workload is composed of one cold program (*applu*) and one hot program (*gcc*). Thus, the temperature fluctuation in Fig. 11(b) is quite different with various context switching intervals. When the actual temperature is under the reliability equivalent temperature, the lifetime is consumed with a slower speed, which allows subsequent execution above the reliability equivalent temperature.

B. DRM Based on Lifetime Banking

In Section III, we derived the lifetime mode for EM subject to dynamic stresses [see (9) and (10)]. Considering void growth limited failures such as those in dual-damascene Cu interconnects where current exponent $n = 1$, we can rewrite MTF by combining (9) and (10) as

$$T_f \propto \frac{1}{E\left[j(t)\left(\frac{\exp\left(\frac{-Q}{kT(t)}\right)}{kT(t)}\right)\right]}.$$

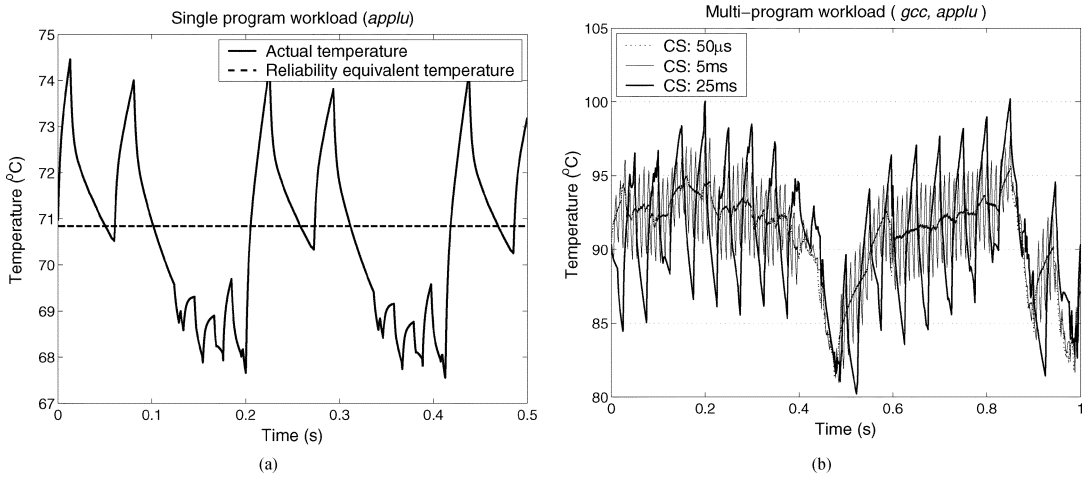

Fig. 11. Temporal temperature variation. (a) Single program workload. (b) Two-program workload with context switching. [12].

Or equivalently, by eliminating the expected-value function, one can express the MTF in an integral form

$$\int_0^{T_f} j(t) \left(\frac{\exp\left(\frac{-Q}{kT(t)}\right)}{kT(t)} \right) dt = D \qquad (12)$$

where D is a constant determined by the structure of the interconnect. Equation (12) models interconnect time to failure (i.e., interconnect lifetime) as a resource consumed by the system over time. Function $r(t) = [j(t)\,(\exp(-Q/kT(t))/kT(t))]$ can be regarded as the consumption rate. In DSM Cu technology, void growth failure (e.g., at vias) is the major EM induced failure mechanism [30], and $r(t)$ can be regarded as the void growth rate (i.e., the atom drift rate at the cathode) in this case.

Equation (12) provides a model to capture the effect of transient behaviors on system lifetime. One interesting case is $j(t) = 0$, which occurs when the system is inactive as commonly seen in systems with nonserver, user-driven workloads. When this happens, the atomic flux becomes zero while the effect of the back-flow diffusion near the cathode created during active periods is worth careful examination. If the inactivity happens in the early stage of the void growth, the back-flow diffusion is negligible and the void simply stops growth during the power-down periods. If the back-flow diffusion is comparable to the normal EM atomic flux, which happens at a very long time after EM begins (e.g., at the order of several years). This back-flow diffusion tends to reduce the void size at the inactivity periods by refilling the void with atoms. However, in order to have significant impact on the void size already formed, this healing process has to last for a duration comparable to the time it took to grow to the current void size, e.g., several years. In normal usage, the inactivity period is usually much less than this time scale. Therefore, the void size is essentially unaffected in inactivity. Our simulations confirm this observation, and more detailed discussion on this aspect is out of the scope of this paper. In summary, the void size remains unchanged during the inactive period if the inactive period is much less than the total active time. Equation (12) accurately models this phenomenon by specifying $r(t) = 0$ for inactivity.

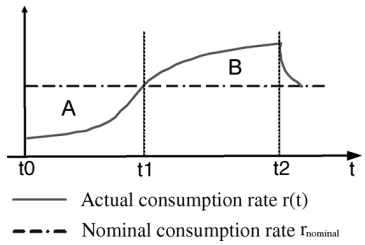

Fig. 12. Illustration of S_DRM technique [12].

Ideally, we would like to monitor the temperature and current for each individual interconnect to build an exact full chip reliability model. In practice, only a limited number of temperature sensors are available on die, and a detailed and complex full chip reliability model is not suitable for runtime management due to the computation overhead. In this study, we use the maximum temperature measured across the chip at runtime, together with the worst case current density specified at design time, to calculate the dynamic consumption rate. This is a conservative but safe approach. Thus, the results obtained in this study provide a lower bound for the potential benefits delivered by the proposed DRM method. Further refinement of the full chip reliability model will be part of future work. In our study, dynamic voltage/frequency scaling is the major throttling mechanism. When DVS is applied, the worst case current density in the IC interconnects should be scaled according to the voltage/frequency setting used. The relationship between current density, supply voltage, and clock frequency can be modeled by transferred charges per clock cycle [31]: $j \propto CV/T = CVf$, where C is the effective capacitance.

When a chip is designed, usually an expected lifetime (e.g., ten years) is specified under some operating conditions (i.e., temperature, current density, etc.). We use r_{nominal} to denote the lifetime consumption rate under the nominal conditions (e.g., reliability constrained temperature threshold). During runtime, we monitor the actual operating conditions regularly, calculate the actual lifetime consumption rate $r(t)$, and compare it with the nominal rate r_{nominal} by calculating $\int (r_{\text{nominal}} - r(t))dt$, which we call the "lifetime banking balance." When $r(t) < r_{\text{nominal}}$, the chip is consuming its lifetime slower than the nominal rate. Thus, the chip's lifetime balance is increased. When $r(t) >$

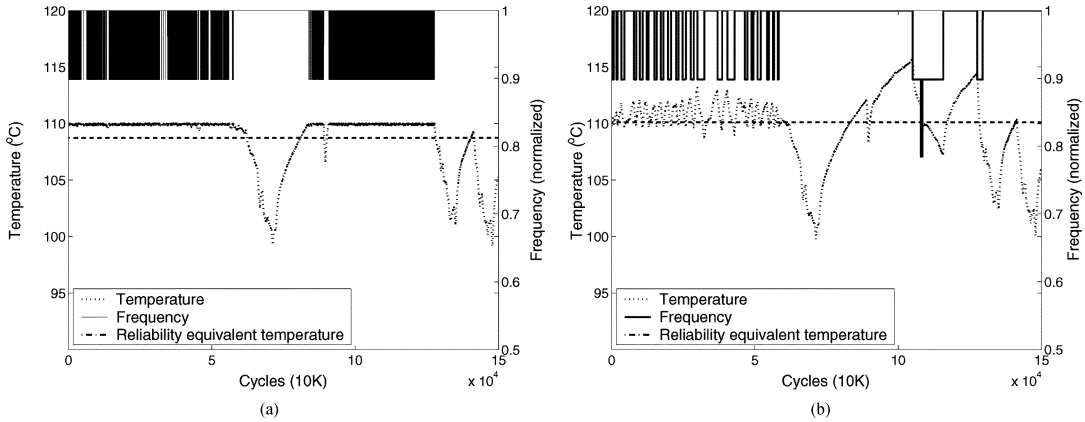

Fig. 13. Temperature and clock frequency profiles in different thermal management techniques for benchmark *gcc*. (a) Conventional DTM (threshold temperature = 110 °C). (b) Reliability banking based DRM (reliability target temperature = 110 °C).

r_{nominal}, the chip is consuming its lifetime faster and reducing the lifetime banking balance. According to (12), as long as the lifetime balance is positive, the expected lifetime will not be shorter than that under the nominal consumption rate r_{nominal}. Fig. 12 illustrates this simple dynamic reliability management (S_DRM) technique. For example, in the interval $[t0, t1]$, the reliability of the chip is banked, while in $[t1, t2]$, the banking balance is consumed. At time instance $t2$, the banking balance becomes less than some threshold, and a cooling mechanism has to be engaged to quickly pull down the lifetime consumption rate to the nominal rate, just as is done in conventional DTM techniques.

Therefore, the difference between conventional DTM and our S_DRM lies in the case where the chip's instantaneous consumption rate is larger than its nominal rate. In DTM, the lifetime consumption rate is never allowed to be larger than the nominal. In S_DRM, before we engage thermal throttling mechanisms, we first check if the chip has a positive lifetime balance. If enough lifetime has been banked, the system can afford to run at a larger than nominal rate. Otherwise, we apply some DTM mechanism to lower the consumption rate, preventing a negative lifetime balance. Since S_DRM only needs to monitor the actual lifetime consumption rate and to update the lifetime banking balance, the computation overhead is comparable to that of DTM.

C. Experimental Setup

We run a set of programs from the Spec2000 benchmark suite on a processor simulator (SimpleScalar [32]) with the characteristics similar to a 0.13-μm Alpha 21364. We simulate each program for a length of five billion instructions, and obtain both dynamic and static (leakage) power traces, which are fed as inputs to a chip-level compact thermal model *Hotspot* [4] for trace-driven simulations. In the trace-driven simulations, we include the idle penalty due to frequency/voltage switching, which is about 10 μs in many practical systems [4]. Furthermore, since leakage power is strongly dependent on temperature, we scale the leakage power trace input dynamically according to the actual temperature obtained during trace-driven simulation, using a voltage/temperature-aware leakage model [33]. Since the *Hotspot* model is highly parameterized, one can easily run experiments on a simulated

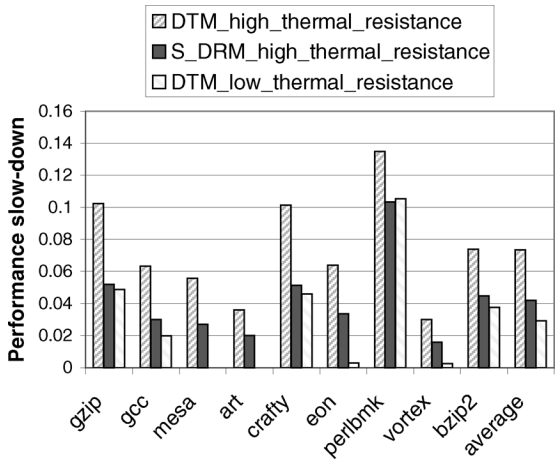

Fig. 14. Performance comparison of DTM and the proposed S_DRM. The results for S_DRM are based on high convection thermal resistance configuration. The results for DTM include two different thermal configurations. [12].

processor with different thermal package settings. In order to obtain meaningful results, one should carefully choose the initial temperature setting for the *Hotspot* model. For each new thermal package setting, we obtain its initial temperatures by repeating the trace-driven simulations until the steady temperatures of the chip are converged, as suggested in [4]. We implement both DTM and S_DRM in the *Hotspot* model and share a common temperature threshold for both runtime management techniques. Both schemes use a feedback controlled dynamic voltage/frequency scaling mechanism to guard the program execution. Since both techniques might slow down the program execution when engaged, we use this performance penalty as a metric to compare these two techniques.

D. Single-Program Workload

Fig. 13 plots the dynamic process for both conventional DTM and the proposed S_DRM techniques for benchmark *gcc*. The feedback controller in DTM effectively clamps the temperature within the target temperature (110 °C) by oscillating the clock frequency between 1.0 and 0.9, resulting in a reliability equivalent temperature less than the target temperature, and causing unnecessarily frequent clock throttling Fig. 13(a). On the other

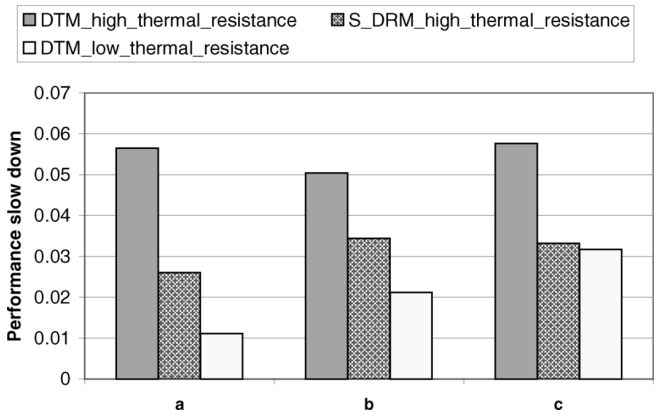

Fig. 15. Average performance comparison of DTM and DRM on a multiprogram workload with different context-switch intervals [(a) 50 μs, (b) 5 ms, and (c) 25 ms]. [12].

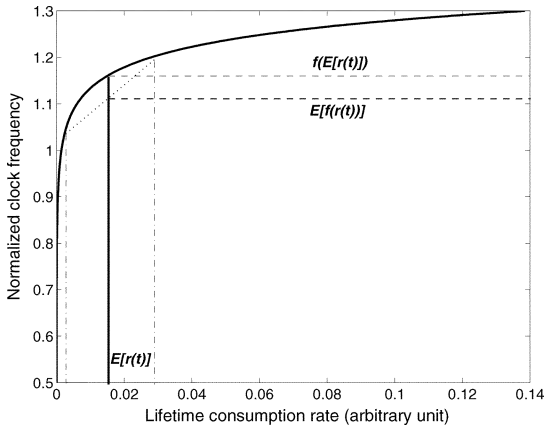

Fig. 16. Relation between performance and lifetime consumption rate.

hand, our S_DRM technique can exploit reliability banking opportunities during the cool phase, and delay the engagement of throttling, while maintaining the specified reliability budget, as proven by the reliability equivalent temperature shown in Fig. 13(b).

Fig. 14 shows the performance penalty for both DTM and S_DRM with the same thermal configuration. Only those benchmarks subject to performance penalties due to runtime management are shown here. As clearly indicated in the figure, performance penalty with the S_DRM scheme is always less than that with DTM scheme, when the thermal configuration is the same. On average, the S_DRM technique reduces the performance penalty by about 40% of that due to DTM (from 7% to 4%). Also shown in the figure is the performance of DTM with a more expensive thermal package whose convection thermal resistance is only one-third of the other's. Fig. 14 shows that, on average, S_DRM with a higher thermal resistance can achieve performance very close to that of DTM with a lower thermal resistance. In other words, if the tolerable performance lost is fixed, the application of S_DRM allows the usage of a much cheaper thermal package than that required by the conventional DTM technique.

E. Multiprogram Workload

Another interesting program execution scenario is a workload of multiple programs context switched among them. When both hot and cold benchmarks are executed together, the average operating temperature should be between the individual benchmarks' operating temperatures. Fig. 11(b) plots the temperature profile of a hybrid workload composed of *gcc* and *applu*, with different context-switch time intervals.

As one expects, the smaller the context-switch interval, the less the temperature fluctuation, with the thermal package of the chip working as if a low-pass filter. When the context-switch interval is increased, individual benchmarks can show their hot/cold properties, and the temperature variation in the workload becomes obvious. Fig. 15 shows the performance penalties of DTM and S_DRM for this multiprogram workload with different context-switch intervals. We observe a similar trend as shown in the single-program workload. S_DRM outperforms DTM with the same thermal package configurations. As the context-switch interval increases, the performance of S_DRM becomes closer to that of DTM with a three-fold smaller convection thermal resistance.

F. Server Workload

The disadvantages of traditional DTM technique become more obvious in server systems such as web servers, in which hot phases usually imply an increased number of service requests and the engagement of cooling mechanisms then exacerbate the QoS degradation by the server. Server workloads have been observed to exhibit low-frequency variations, often determined by national events or by time of day [34]. This allows a more proactive optimization. To study this, we develop a synthetic workload by concatenating two Spec2000 benchmarks *applu* and *gcc* (see [12] and [20] for details). Our synthetic workload has a distinct cool phase and a hot phase. From an average user's point-of-view, the QoS provided by the server is largely dependent on its performance in the hot phase, as most requests are made during that time. Therefore, in our study, we use the performance in the hot phase as our performance metric for comparison.

The application of the S_DRM technique to server workloads is straightforward, just like in the context-switched multiprogram workload studied previously. However, due to the concave curve between the reliability consumption rate and performance as shown in Fig. 16, the best strategy for reliability management is to spend the lifetime balance evenly during the hot phase (see [20] for detail). Therefore, we proposed a P_DRM (profile-based dynamic reliability management) technique for server workloads. P_DRM is a natural extension of our S_DRM with the awareness of the optimal operating points in the hot phase. When the server is running in the cool phase, P_DRM works the same way as S_DRM with lifetime banked in the cool phase. When the server enters the hot phase, P_DRM calculates a new nominal lifetime consumption rate based on the lifetime balance and the estimated duration of the hot phase (obtained through profiling). Then P_DRM acts just like S_DRM, with the new calculated nominal consumption rate. In some cases, we might not be able to predict the hot phase duration accurately. However, the inaccuracy only affects the performance optimality

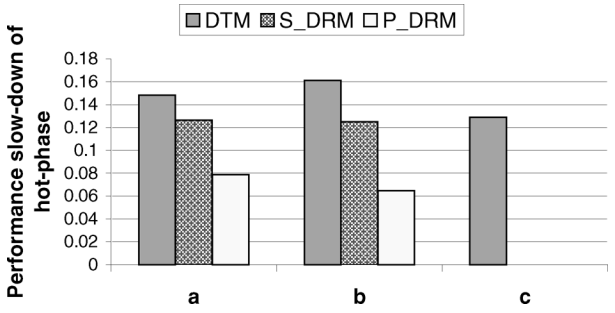

Fig. 17. Performance comparison of different runtime management techniques on the synthetic workload with different duty cycles of the cool phase: (a) 0.5, (b) 0.6, and (c) 0.75. [12].

without harming the reliability budget. That is because our technique always tracks the actual reliability consumption rate and compares it with the nominal lifetime consumption.

The performance of different runtime management techniques for server workloads is shown in Fig. 17. Both DRM techniques outperform DTM, and P_DRM performs the best. P_DRM can fully exploit the banking benefits of the cool phase. For example, when the cool phase occupies 60% of the total time [i.e., as indicated by Fig. 17(b)], P_DRM can reduce the performance penalty from 16% (DTM) to only 6% (or equivalently, the execution speed of the hot phase is increased by P_DRM by about 9.5% over DTM). Interestingly, for the case when the cool phase occupies 75% of the total time [i.e., in Fig. 17(c)], no performance slowdown is incurred for both DRM techniques, because the reliability equivalent temperature for that workload is less than the reliability nominal temperature. Thus, in that case, the lifetime balance banked in the cool phase is enough to support the full speed execution in the hot phase, while DTM clamps the hot phase temperature to the reliability temperature, resulting in about a 13% performance penalty in the hot phase. Further analysis using an analytic model reveals that the "sweet spots" for P_DRM lie in the case where the duty cycle of the cool phase is more than 50% and the temperature difference between cool and hot phases is more than 20 °C, and more than 5% of performance speed-up can be expected [12], [20].

VII. Conclusion

This paper presented an analysis of interconnect EM failures subject to temporal and spatial thermal gradients. For EM under time-varying stresses (temperature/current), we proposed a dynamic reliability model, which returns reliability equivalent temperatures/currents. For EM under nonuniform temperature distributions, we obtained close bounding temperatures to estimate the actual lifetime. Therefore, the commonly used Black's equation is still applicable by using our constant reliability equivalent temperatures. Our analysis reveals that blindly using the maximum or average temperature to evaluate EM lifetime is inappropriate. Our results not only increase the accuracy of reliability estimates but enable designers to more aggressively explore the design space and to reclaim the design margin imposed by less accurate, more pessimistic models. Existing constant-temperature models require designers to observe a static worst case temperature limit, but the analysis presented here enables temperature-aware designers to evaluate the system reliability using runtime information, thus, increasing the confidence about the actual behavior of the system.

As an application example, we detailed the use of the temperature variability and lifetime resource models to develop novel DRM techniques that reduce the performance penalties associated with existing DTM techniques while maintaining the required expected IC reliability lifetime. When the operating temperature is below a nominal temperature (i.e., the threshold temperature used in DTM techniques), lifetime is being consumed at a slower than nominal rate, effectively banking lifetime for future consumption. A positive lifetime balance allows the nominal temperature to be exceeded for some time (thus, consuming lifetime at a faster than nominal rate) instead of automatically engaging DTM and unnecessarily suffering the associated performance penalties. Simulations with general-purpose computing workloads as well as synthetic server workloads were performed and revealed that our reliability banking based DRM techniques provide performance improvements over traditional threshold-based DTM without sacrificing expected lifetime, or allows the usage of cheaper thermal package without sacrificing performance. Although the DRM experiments presented here do not explicitly study the scenarios with long periods of inactivity, which are commonly seen in nonserver, user-driven workloads, our lifetime banking techniques can be applied in a straightforward way, because the dynamic reliability model [(12)] also holds true in these situations. Consequently, much better performance gains would be expected because more lifetime can be banked during those inactivity periods.

References

[1] K. Banerjee and A. Mehrotra, "Global (interconnect) warming," *IEEE Circuits Device Mag.*, vol. 17, no. 5, pp. 16–32, Sep. 2001.

[2] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankanarayanan, and D. Tarjan, "Temperature-aware microarchitecture," in *Proc. 30th Int. Symp. Comput. Arch.*, 2003, pp. 2–13.

[3] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers, "The case for lifetime reliability-aware microprocessors," in *Proc. 31st Int. Symp. Comput. Arch.*, 2004, pp. 276–287.

[4] K. Skadron, K. Sankaranarayanan, S. Velusamy, D. Tarjan, M. Stan, and W. Huang, "Temperature-aware microarchitecture: Modeling and implementation," *ACM Trans. Arch. Code Optimization*, vol. 1, no. 1, pp. 94–125, Mar. 2004.

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

[6] M. Hauschildt, M. Gall, S. Thraasher, P. Justison, L. Michaelson, H. Kawasaki, and P. S. Ho, "Statistical analysis of electromigration lifetimes and void evolution for cu interconnects," in *Proc. Mater., Technol. Reliab. Adv. Interconnects*, 2004, pp. 379–384.

[7] J. R. Black, "Mass transport of aluminum by momentum exchange with conducting electrons," in *Proc. IEEE Int. Rel. Phys. Symp.*, 1967, pp. 148–159.

[8] Y.-J. Park, V. K. Andleigh, and C. V. Thompson, "Simulations of stress evolution and the current density scaling of electromigration-induced failure times in pure and alloyed interconnects," *J. Appl. Phys.*, vol. 85, no. 7, pp. 3546–3555, Apr. 1999.

[9] Z. Lu, W. Huang, J. Lach, M. Stan, and K. Skadron, "Interconnect lifetime prediction under dynamic stress for reliability-aware design," in *Proc. Int. Conf. Comput.-Aided Des.*, 2004, pp. 327–334.

[10] S. M. Alam, G. C. Lip, C. V. Thompson, and D. E. Troxel, "Circuit level reliability analysis of Cu interconnects," in *Proc. 5th Int. Symp. Quality Electron. Des. (ISQED)*, 2004, pp. 238–243.

[11] D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," in *Proc. 7th Int. Symp. High-Performance Comput. Arch. (HPCA-7)*, 2001, pp. 171–182.

[12] Z. Lu, J. Lach, M. Stan, and K. Skadron, "Improved thermal management with reliability banking," *IEEE Micro*, vol. 25, no. 6, pp. 40–49, Nov./Dec. 2005.
[13] A. Ramakrishnan and M. G. Pecht, "A life consumption monitoring methodology for electronic systems," *IEEE Trans. Comput. Packag. Technol.*, vol. 26, no. 3, pp. 625–634, Sep. 2003.
[14] J. Clement, "Electromigration modeling for integrated circuit interconnect reliability analysis," *IEEE Trans. Device Mater. Reliab.*, vol. 1, no. 1, pp. 33–42, Mar. 2001.
[15] M. E. Sarychev, Y. V. Zhitnikov, L. Borucki, C.-L. Liu, and T. M. Markhviladze, "General model for mechanical stress evolution during electromigration," *J. Appl. Phys.*, vol. 86, no. 6, pp. 3068–75, Sep. 1999.
[16] J. Clement, "Reliability analysis for encapsulated interconnect lines under DC and pulsed DC current using a continuum electromigration transport model," *J. Appl. Phys.*, vol. 82, no. 12, pp. 5991–6000, Dec. 1997.
[17] M. A. Korhonen, P. Bøgesen, K. N. Tu, and C.-Y. Li, "Stress evolution due to electromigration in confined metal lines," *J. Appl. Phys.*, vol. 73, no. 8, pp. 3790–3799, Apr. 1993.
[18] W. R. Hunter, "Self-consistent solutions for allowed interconnect current density-part II: Application to design guidelines," *IEEE Trans. Electron Devices*, vol. 44, no. 2, pp. 310–316, Feb. 1997.
[19] L. M. Ting, J. S. May, W. R. Hunter, and J. W. McPherson, "Ac electromigration characterization and modeling of multilayered interconnects," in *Proc. Int. Reliab. Phys. Symp.*, 1993, pp. 311–316.
[20] Z. Lu, W. Huang, M. Stan, K. Skadron, and J. Lach, "Interconnect lifetime prediction with temporal and spatial temperature gradients for reliability-aware design and runtime management: Modeling and applications," Dept. Comput. Sci., Univ. Virginia, Tech. Rep. CS-2006-23, 2006, .
[21] A. H. Ajami, K. Banerjee, and M. Pedram, "Modeling and analysis of nonuniform substrate temperature effects on global ULSI interconnects," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 24, no. 6, pp. 849–861, Jun. 2005.
[22] W. Guo, Z. Li, H. Zhu, W. Zhang, Y. Ji, Y. Sun, and G. Shen, "Temperature gradient impact on electromigration failure in VLSI metalization," in *Proc. 14th IEEE SEMI-THERM Symp.*, 1998, pp. 122–127.
[23] H. V. Nguyen, C. Salm, B. Krabbenborg, K. Weide-Zaage, J. Bisschop, A. J. Mouthaan, and F. G. Kuper, "Effect of thermal gradients on the electromigration lifetime in power electronics," in *Proc. 42nd Annu. Int. Reliab. Phys. Symp. (IRPS)*, 2004, pp. 619–620.
[24] K. Jonggook, V. C. Tyree, and C. R. Crowell, "Temperature gradient effects in electromigration using an extended transition probability model and temperature gradient free tests," in *Proc. IEEE Integr. Reliab. Workshop (IRW)*, 1999, pp. 24–40.
[25] E. T. Ogawa, K.-D. Lee, V. A. Blaschke, and P. S. Ho, "Electromigration reliability issues in dual-damascene Cu interconnections," *IEEE Trans. Reliab.*, vol. 51, no. 4, pp. 403–419, Dec. 2002.
[26] Z.-S. Choi, C. L. Gan, F. Wei, C. V. Thompson, J. H. Lee, K. L. Pey, and W. K. Choi, "Fatal void size comparisons in via-below and via-above Cu dual-damascene interconnects," in *Proc. Mater., Technol. Reliab. Adv. Interconnects*, 2004, pp. 373–378.
[27] W. K. Meyer, R. Solanki, and D. Evans, "Near-threshold electromigration of damascene copper on tin barrier," in *Proc. Mater., Technol. Reliab. Adv. Interconnects*, 2003, pp. 145–150.
[28] J. Srinivasan and S. V. Adve, "Predictive dynamic thermal management for multimedia applications," in *Proc. 17th Annu. ACM Int. Conf. Supercomput.*, 2003, pp. 109–120.
[29] R. A. Wachnik, R. G. Filippi, T. M. Shaw, and P. C. Lin, "Practical benefits of the electromigration short-length effect, including a new design rule methodology and an electromigration resistant power grid with enhanced wireability," in *Symp. VLSI Technol. Dig.*, 2000, pp. 220–221.
[30] J. A. Davis and J. D. Meindl, Eds., *Interconnect Technology and Design for Gigascale Integration*. Norwell, MA: Kluwer Academic, 2003.
[31] D. T. Blaauw, C. Oh, V. Zolotov, and A. Dasgupta, "Static electromigration analysis for on-chip signal interconnects," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 22, no. 1, pp. 39–48, Jan. 2003.
[32] D. Burger and T. M. Austin, "The simplescalar toolset, version 2.0," *Comput. Arch. News*, vol. 25, no. 3, pp. 13–25, Jun. 1997.
[33] Y. Zhang, D. Parikh, K. Sankaranarayanan, K. Skadron, and M. Stan, "Hotleakage: A temperature-aware model of subthreshold and gate leakage for architects," Dept. Comput. Sci., Univ. Virginia, Tech. Rep. CS-2003-05, 2003.
[34] E. N. Elnozahy, M. Kistler, and R. Rajamony, "Energy conservation policies for web servers," in *Proc. 4th USENIX Symp. Internet Technol. Syst.*, 2003, pp. 99–112.

Zhijian Lu (S'01) received the B.E. degree in control engineering from Tsinghua University, Beijing, China, in 1999, and the Ph.D. degree in electrical engineering from the University of Virginia, Charlottesville, in 2007.

He is currently with the Marvell Technology Group, Santa Clara, CA. His research interests include power- and temperature-aware computing, software/hardware co-design and optimization, and system-level design automation.

Dr. Lu was a recipient of the Best Student Paper Award at SHAMAN 2002.

Wei Huang (M'07) received the B.S. degree in electrical engineering from the University of Science and Technology of China, Hefei, China, and the Ph.D. degree in electrical engineering from the University of Virginia, Charlottesville, in 2007.

He is currently a Postdoctoral Researcher at the Department of Computer Science, University of Virginia. He held an internship with IBM Research, NY, in 2004. His research interests include VLSI circuit and architecture design with considerations of thermal, variability, and reliability issues.

Dr. Huang was a recipient of the Best Student Paper Award at ISCA in 2003.

Mircea R. Stan (SM'98) is an Associate Professor in the Electrical and Computer Engineering Department, University of Virginia, Charlottesville. He is teaching and researching the areas of high-performance and low-power VLSI, temperature-aware circuits and architecture, embedded systems, and nanoelectronics.

Prof. Stan was a recipient of the National Science Foundation CAREER Award in 1997 and was coauthor on Best Paper Awards at GLSVLSI 2006, ISCA 2003, and SHAMAN 2002. He was TPC chair of ISLPED 2005 and NanoNets 2007, and general chair of GLSVLSI 2003 and ISLPED 2006. He is a member of the Association for Computing Machinery (ACM), IET, Eta Kappa Nu, Phi Kappa Phi, and Sigma Xi.

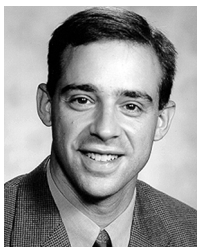

Kevin Skadron (SM'05) received the Ph.D. degree from Princeton University, Princeton, NJ, in 1999.

Currently, he is an Associate Professor in the Computer Science Department, University of Virginia, Charlottesville. His research interests focus on multicore design in the context of physical limits such as reliability, temperature, and parameter variations. He is a co-founder and Associate Editor-In-Chief of IEEE COMPUTER ARCHITECTURE LETTERS. He recently served as program or general co-chair of PACT 2006, MICRO 2004, and PACT 2002, and helped create the Workshop on Temperature-Aware Design.

Dr. Skadron was coauthor of Best Student Paper Awards at ISCA 2003, RTSS 2003, and SHAMAN 2002.

John Lach (SM'05) is an Associate Professor in the Electrical and Computer Engineering Department, University of Virginia, Charlottesville. His primary research interests include IC design techniques, dynamically adaptable and real time embedded systems, fault and defect tolerance, safety-critical system design and analysis, general purpose and application specific processor design, and wearable technologies for biomedical applications. He is an Associate Editor for IEEE TRANSACTIONS ON COMPUTERS and IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS. He was General Chair of GLSVLSI 2005 and is on several conference Technical Program Committees.

Dr. Lach was a recipient of several teaching and faculty awards and he is a member of the Association for Computing Machinery (ACM) and Eta Kappa Nu.