

Hardware Reliability Margining for the Dark Silicon Era

Liangzhen Lai

Department of Electrical Engineering
University of California, Los Angeles
e-mail: liangzhen@ucla.edu

Puneet Gupta

Department of Electrical Engineering
University of California, Los Angeles
e-mail: puneet@ee.ucla.edu

Abstract— Hardware reliability margin should be derived from the worst-case aging scenario, which typically occurs when the circuits are operating at peak performance state with the highest operating voltage and frequency. However, as integrated circuits enter the “dark silicon” era, it is impossible to power up all circuits throughout the entire lifetime. Reliability margining in absence of architecture-level power/thermal constraints can be overly pessimistic. In this work, we propose a margining scheme that employs the power/thermal contexts and system management policies to derive the actual worst-case workload pattern for different reliability phenomena. Our experiment results show that at 60% dark ratio, conventional margining approach can overestimate the aging degradation due to EM and BTI by up to 3-7X and 18% respectively. Our margining method is able to eliminate these over-pessimism and results in about 20% delay margin and 40%-60% metal width margin reduction.

I. INTRODUCTION

With ending of perfect Dennard scaling, integrated circuits have entered the “dark silicon” era [10], where chips are fundamentally power- and thermal-constrained. Multi-core design has been widely adopted as the approach to both improve energy efficiency by operating at lower power state and exploiting workload parallelism [14, 18]. Meanwhile, single-thread performance is still desired by the workload that is not parallelizable. Techniques such as “turbo boost” [9] are used to temporarily operate one or several processor cores at higher voltage/frequency under certain power/thermal constraints. This creates a large dynamic range in performance and power that the hardware is designed to run at but maybe only for part of the circuits or only for a limited amount of time. For example, Intel i7-920XM [1, 2] has a base frequency of 2 GHz. It can be boosted to operate at maximum 2.26 GHz with four active cores and to 3.2 GHz with only one core active. This gap can be more pronounced for systems with tighter thermal constraints such as mobile devices [21].

Conventionally, hardware reliability margin is derived from the worst-case aging scenario with possibly pessimistic assumptions. Although sense-and-adapt approach may be used to reduce parametric (e.g., voltage, delay) margin at runtime [11], physical (e.g., wire width or spacing) margin must be determined during hardware design time. Since most circuit wear-out mechanisms (e.g., Negative- and Positive-Bias Temperature Instability (N/PBTI), Hot-Carrier-Injection (HCI), Electromigration (EM) etc.) degrade faster at higher voltage and temperature, the margin is typically derived when the circuits

are operating at peak performance state with the highest voltage, current and temperature.

The dark silicon contexts impose additional constraints so that it is impossible for all circuits to operate at peak performance state throughout the entire lifetime. Therefore, the reliability margin derived by the conventional approach can be overly pessimistic and is expected to be worse with technology scaling and increasing silicon “darkness”. Different terms have been used to describe this “darkness”, such as dark silicon fraction [10], dark area [17], and transistor utilization wall [25]. In this work, we define a term d called “dark ratio”. The definition of d for power-constrained scenario is described in Equation 1, where $d = 0$ means all circuits can be fully powered up. The definition of d for thermal-constrained scenario will be described later in Equation 9.

$$d = 1 - \frac{\text{maximum allowed power}}{\text{maximum possible power}} \quad (1)$$

The over-pessimism may be reduced by correctly accounting for the dark silicon contexts. The main challenge is how to transform the high-level power/thermal (i.e., spatial) constraints into hardware aging (i.e., temporal) profile, especially with mixture of complicated reliability models and system power/thermal management schemes. In this work, we propose a hardware reliability margining methodology (see Fig. 1). It uses an accumulation model to represent the aging degradation dependence on circuit power states. An optimization problem is formulated to search in the space of possible system workload (i.e., power states of each processor cores) patterns that stress the circuits under the power/thermal constraints. By exploiting the properties of workload and management policy, this spatial distribution workload can be transformed into temporal aging profile, which can be used with accumulation model in the optimization to determine the worst-case aging degradation under the power/thermal constraints. The contribution of our work are:

- To the best of our knowledge, this is the first work on reliability margining considering the high-level power/thermal constraints imposed by the dark silicon contexts.
- We formulate and solve the margining problem for both power- and thermal-constrained cases. Experimental results show that at 60% dark ratio, conventional margining approach can results in 3-7X and 18% overestimate of aging degradation due to EM and BTI respectively. Our margining method is able to eliminate these over-pessimism and results in about 20% delay margin and 40%-60% metal width margin reduction.

The rest of the paper is organized as follows: Section II introduces some background of EM and BTI modeling and margining. Section III describes how to derive the required data and model for our methodology. Section IV and Section V present

This work is supported in part by NSF Variability Expedition grant CCF-1029030

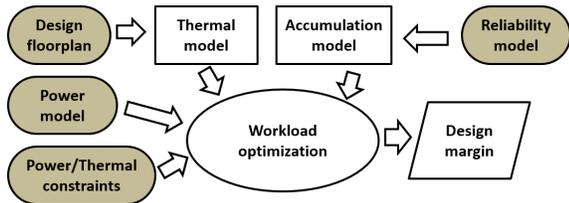


Fig. 1. Margining Methodology Overview.

the problem formulation for reliability margining under power constraints and thermal constraints, respectively. Experimental results and some discussions are presented in Section VI, and Section VII concludes the paper.

II. BACKGROUND

In this work, we focus on two reliability phenomena, EM and BTI, as they require different types of design margin. The basic modeling for EM and BTI are described in Section II.A and II.B. Design margining for EM and BTI is introduced in Section II.C.

A. EM Modeling

Electromigration [3, 13] is a wear-out mechanism on interconnects. Void are created on metal wires due to momentum transfer from the electrons moving in the wires. The Mean Time to Failure (MTTF) of EM can be modeled by the famous Black's equation [5]:

$$MTTF = A_0 J^{-2} e^{E_m/kT} \quad (2)$$

where A_0 is a constant, J is the current density, E_m is the activation energy for the EM mechanism, k is Boltzmann's constant, and T is the temperature.

For non-DC current, Hunter [13] derived an "average" current density model for arbitrary unipolar current. This model is useful for interconnects with unipolar current, such as power distribution networks (PDNs). An effective current density can be derived with Equation 3.

$$J_{eff} = \frac{1}{T} \int_0^T J(t) dt \quad (3)$$

where J_{eff} is the effective current density for EM purpose, T is the period, and $J(t)$ is the current density at time t .

Banerjee etc. [3] derived the model with a recovery factor R for symmetric bipolar (AC) current. So the effective current density J_{eff} is R times the absolute current density. This model is useful for evaluating EM on signal wires, because in digital circuits, signal wires are used to charging up or charging down the same loading capacitance. So the average current density will be proportional to the supply voltage, loading capacitance and switching frequency. This dependence can be derived as Equation 4:

$$J_{eff} = R \cdot \frac{\int |J(t)| dt}{T} = R \cdot \frac{\mu C_{load} V f}{w \cdot h} \quad (4)$$

where μ is the signal switching probability, C_{load} is the loading capacitance, V is the supply voltage, f is the operating frequency, w is the metal width, and h is the metal height.

B. BTI Modeling

Bias Temperature Instability is a wear-out mechanism related to the transistor. Depending on the transistor type, BTI effects include NBTI for PMOS and PBTI for NMOS. Due to

BTI, the device threshold voltage increases when the transistor is turned on (i.e., positively or negatively stressed). When the stress voltage is removed, part of the degradation can be recovered.

The exact BTI mechanism and model is still under debate within the reliability community. There are two major BTI models, reaction-diffusion (RD) model [4, 26] and trapping/detrapping (TD) model [24]. However, most existing analytical BTI models are derived for static operating conditions and may not be directly used for dynamic operating conditions introduced by typical system-level power management mechanisms such as DVFS and power gating. BTI simulators [7] have been used to simulate the degradation under arbitrary stress profile.

C. Margining for EM and BTI

Unlike typical design optimization, hardware margining must be pessimistic to account for all possible or even pathological worst-case scenarios. In many cases, pessimistic assumptions are unavoidable and necessary. In this work, our approach will strictly follow this rule. The pessimism in the assumptions and approaches will be summarized in Section VI.D.

As suggested by Black's equation (see Equation 2), EM margining can be done by increasing the interconnect width in order to reduce the current density. Typically, certain design guidelines are given by the foundry as part of the reliability design rules. For example, in a commercial 45nm design manual, the reliability design rules are given as the current-dependent minimum width requirement for each metal layer. The amount of EM margin can also affect other design metrics such as performance, power and area [15].

In a circuit design, there are different types of interconnect that can be affected by EM. In this work, we consider EM on both local PDN for a processor core and the signal wires inside the processor core. We omit the global PDN because the margin can be directly derived from the power constraints (i.e., dark ratio and total power).

BTI-induced threshold degradation can affect both logic and memory and result in reduced read/write margin [19] and increased logic delay. Since SRAM operating conditions are usually more stable, in this work, we focus on the BTI-induced delay degradation for logic.

There can be two approaches to margin for BTI-induced delay degradation. One way is estimating the end-of-lifetime delay degradation for the entire circuit and directly applying it as additional delay margin (i.e., reducing clock frequency) or voltage margin (i.e., increasing supply voltage). The other way is to estimate the process corners for the aged devices and performing timing sign-off accordingly (i.e., increase the design efforts such as using larger cells). In this work, we will use the former approach to evaluate the BTI margin, as the margin in latter approach depends on the specific design/constraints.

III. POWER, THERMAL AND ACCUMULATION MODELS

An overview of our proposed approach is shown in Fig. 1. The shaded blocks are assumed to be known at hardware design time. In this section, we will focus on the modeling of power, thermal and reliability. Unless otherwise defined, all notations used in Section III, IV and V are listed in Table I. Matrix transpose is represented with a prime symbol, i.e., transpose of \mathbf{x} is represented as \mathbf{x}' .

A. Accumulation Model

Most reliability models are derived for static scenarios. However, a dynamic reliability model is required to study the worst-case aging scenarios, which can be a mixture of different power

TABLE I
GLOSSARY OF TERMINOLOGY

Term	Description
N	Total number of processor cores
n	Index for processor cores
M	Total number of power states
m	Index for power state
P^{max}	Maximum allowed power
T^{max}	Maximum allowed temperature
T^{amb}	Ambient temperature
d	Dark ratio
\mathbf{x}	$M \times 1$ vector where x_m is the number of cores at power state m
\mathbf{P}	$M \times 1$ vector where P_m is the power consumption of a processor core at power state m
\mathbf{F}	$M \times 1$ vector where F_m is the operating frequency of a processor core at power state m
\mathbf{V}	$M \times 1$ vector where V_m is the supply voltage of a processor core at power state m
\mathbf{T}	$N \times 1$ vector where T_n is the temperature of processor core n
\mathbf{T}^{bak}	$N \times 1$ vector where T_n^{bak} is the background temperature of processor core n
\mathbf{A}	$N \times N$ matrix where $A(n_1, n_2)$ is the temperature increase of core n_1 due to power consumed by core n_2
\mathbf{S}	$N \times M$ matrix where $S(n, m)$ is the proportion of time that core n spends in power state m
$f(\mathbf{x})$	The accumulation model for a processor core spending x_m/M lifetime at each power state m

states. For example, with the same power budget, the circuits can be operating at 0.8V constantly or alternating between 0.7V and 0.9V. We need to be able to compare these different workload patterns in order to determine the worst case.

Most reliability-related phenomena are temperature-dependent. Intuitively, lower power states are usually associated with lower temperature. So we may be able to use lower temperature for calculating the degradation. But this cannot be guaranteed considering the following scenario that the processor core has already been heated up by other heavy workload before entering a lower power state. This history dependence makes it difficult to make any pessimistic assumption about temperature other than using the maximum temperature.

In our proposed methodology, reliability model is used to characterize an accumulation model (see Fig. 1). This helps isolate the model-dependence and makes our methodology more general for different wear-out mechanisms. The accumulation model is a function to calculate the pessimistic estimation of end-of-life degradation under a given workload pattern. The input to the accumulation model is the fraction of circuit lifetime spent in different power states (i.e., \mathbf{x}).

Some reliability models themselves suggest the corresponding accumulation model. For example, the current density for local PDN is proportional to the power consumption. Therefore the EM accumulation model can be derived based on Equation 3 as:

$$f(\mathbf{x}) \propto \mathbf{x}' \cdot \mathbf{P} \quad (5)$$

The EM accumulation model for signal wires can be derived based on Equation 4 as:

$$f(\mathbf{x}) \propto \mathbf{x}' \cdot (\mathbf{F} \circ \mathbf{V}) \quad (6)$$

where \circ represents Hadamard product (i.e., element-wise multiplication).

For other reliability phenomena like BTI, applying the analytical models directly for dynamic cases may lead to mislead-

ing results [7]. In this work, we will derive the accumulation model with a physics-based RD simulator [7].

We use two steps to derive the accumulation model. The first step is to identify the ordering of different power states that will result in the worst degradation. Based on the simulator, the worst BTI degradation happens when power states are applied in increasing order of stress voltages. Therefore, for a given power state distribution (\mathbf{x}), there is only one degradation value and it is pessimistic.

The second step is obtaining the accumulation function through interpolation or fitting. In this work, we first pick a set of power state distribution samples \mathbf{x} . These samples are generated with fraction of lifetime spent in each power state being multiple of 0.1. To avoid the low activity region where the BTI model shows high non-linearity, we limit the total fraction of idle time to be less than 0.3. Therefore, we will only report the results with $d < 0.7$.

Assuming $g(\mathbf{x})$ is the accumulation function by the simulator, our accumulation function is fitted with a linear function, i.e., $f(\mathbf{x}) = \mathbf{c}' \cdot \mathbf{x} + d$ with \mathbf{c} and d being the fitting coefficients, as the results of the following optimization:

$$\begin{aligned} & \text{minimize:} && b \\ & \text{subject to:} && \forall \mathbf{x} \quad 0 \leq f(\mathbf{x}) - g(\mathbf{x}) \leq b \end{aligned} \quad (7)$$

We validate this accumulation function against the results from the simulator. The maximum overestimation is less than 0.5mV for about 10mV threshold voltage degradation.

B. Power and Thermal Model

For chips in the dark silicon era, there can be two types of constraints, i.e., limit on the instantaneous power consumption and limit on the temperature. For example, Intel Sandy Bridge [20] has total package power control and responsiveness via dynamic turbo. The former one controls the the entire package's total power consumption, while the latter one allows the total power consumption to be temporarily higher (e.g., 1.2-1.3X [20]) than the thermal design power (TDP).

Since the margin is determined at hardware design time, detailed power models may not be available. But the hardware designer should have estimated power for each hardware power states. In this work, we assume that the power consumption at each power state is available for reliability margining purpose. The model should include the power consumption for all possible power states, including different voltage states (i.e., P-state) and idle states (i.e., C-state). These power consumption estimate should also be optimistic (i.e., with smaller values) so that the derived margin will be pessimistic.

The thermal model should give the static state temperature under a given power profile for a fixed floorplan and package. In this work, we will use a thermal simulator, HotSpot [12, 22], which takes the floorplan and power profile trace and can calculate both static state and transient temperature. By using the simulator, we can obtain the static state temperature increase of one processor core due to power consumed by the other cores (i.e., \mathbf{A}). Based on the thermal-electrical duality, the superposition principle can be applied here. Therefore, the temperature can be computed with the values in \mathbf{A} without invoking the simulation in the optimization. We have also validated this superposition principle and the thermal model using the simulator.

IV. RELIABILITY MARGINING WITH POWER CONSTRAINTS

In this section, we describe the problem formulation for system with power constraints. The interpretation of man-

agement policy and power constraints are presented in Section IV.A. The optimization formulation is described in Section IV.B.

A. Management Policy and Power Constraints

For circuits in the dark silicon era, certain management mechanisms and policies are essential to enforce the power/thermal constraints. These policies can affect the system behavior and thus the reliability margin. In some case, a bad policy can result in the same degradation as what conventional reliability approach indicates. For example, a policy with a fixed priority in scheduling workload can overload the first core, resulting in spending entire lifetime at the highest power state.

In this work, we assume a fair round-robin policy that iterates the scheduling priority between all processor cores. The iterating frequency can range from once per hour to once per week. We argue that this is a reasonable, effective, and possibly pessimistic policy due to the following reasons:

- This policy is an open-loop policy that does not rely on any sensing or monitoring capabilities, and thus is easy to implement.
- Given a typical hardware life time of multiple years and scheduling window of a few milliseconds, the number and length of priority iteration is sufficient to redress the workload imbalance between different cores.
- More sophisticated policies, e.g., close-loop sense-and-adapt policies or more advanced management policies [6, 16, 23], may improve the degradation. This makes our assumed policy pessimistic, which is acceptable for margining purposes.

With such management policy, the worst-case workload will be iterated among all processor cores. So the power states of all processor cores at a given time will be iterated through each processor core along its lifetime. Therefore, the temporal distribution of power states for one processor core is equivalent to the spatial distribution of power states among all cores. This property will be exploited in our problem formulation.

B. Problem Formulation

The power constraint limits the total power of all processor cores. In the power model, processor power is represented as its power state. So the total power consumption can be calculated from the number of cores in each power state (i.e., \mathbf{x} should be integer variables). As discussed earlier, this spatial distribution can be converted into the temporal distribution of power states along the lifetime of one processor core. So the optimization should maximize the resulting degradation with power state distribution of \mathbf{x} . Therefore, the problem with power constraints can be formulated as the integer optimization problem described in Equation 8.

$$\begin{aligned} & \text{maximize:} && f(\mathbf{x}) \\ & \text{subject to:} && x_m \geq 0 \\ & && \mathbf{x}' \cdot \mathbf{1} \leq N \\ & && \mathbf{P}' \cdot \mathbf{x} \leq P^{max} \end{aligned} \quad (8)$$

V. RELIABILITY MARGINING WITH THERMAL CONSTRAINTS

In this section, we describe the problem formulation for system with thermal constraints. The interpretation of thermal constraints is presented in Section V.A. The optimization formulation is described in Section V.B.

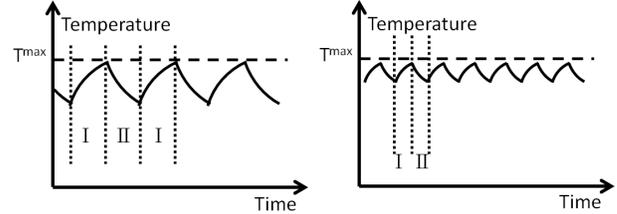


Fig. 2. Example of two workload cases. Both cases are alternating their power states between I and II while keeping the temperature under the thermal limit.

A. Interpreting the Thermal Constraints

Unlike power constraints, thermal constraints do not directly impact the system instantaneous behavior. Circuits can temporarily operate at high power state if the temperature is below the limit. Intuitively, there can be two possible pathological scenarios that may maximize the degradation. The first one stresses the circuits in a two-phase manner which cools down the circuits with lower power state and then heats up with high power state. The second one stresses the circuits in a constant manner so that the temperature is maintained or has small fluctuations right below the temperature limit.

We argue that the second scenario will always result in worse degradation through an example shown in Fig. 2. Both cases have the same alternating power states (i.e., power state I and II), but with different alternating frequencies. Since the average temperature of the second case is higher than the first case, by the thermal model, the heat dissipated from the second case will be more than the first case. Therefore, the second case can spend more time in higher power state and thus results in more wear-out degradation.

So with the thermal constraints, a “constant” stress will result in more degradation. We also define the dark ratio d for this thermal-constrained scenario in Equation 9. The term $T^{max} - T^{amb}$ is the maximum temperature increase allowed for a given ambient temperature. The term $\max(\mathbf{A} \cdot \mathbf{1} \cdot \max(\mathbf{P}))$ is the temperature increase when all processor cores are fully powered up.

$$d = 1 - \frac{T^{max} - T^{amb}}{\max(\mathbf{A} \cdot \mathbf{1} \cdot \max(\mathbf{P}))} \quad (9)$$

B. Formulation

As discussed earlier, the thermal constraints can be represented as the limit on the average power, which can be represented as the fraction of time spent in each power states for all processor cores (i.e., continuous variables \mathbf{S}). The static temperature is calculated based on this average power consumption and the pre-characterized temperature-sensitivity matrix \mathbf{A} . The problem with thermal constraints can be formulated as an optimization problem described in Equation 10.

$$\begin{aligned} & \text{maximize:} && f(\mathbf{S}' \cdot \mathbf{1}) \\ & \text{subject to:} && \forall m, n \quad S(m, n) \geq 0 \\ & && \mathbf{S} \cdot \mathbf{1} \leq \mathbf{1} \\ & && \mathbf{A} \cdot \mathbf{S} \cdot \mathbf{1} \leq T^{max} \times \mathbf{1} - \mathbf{T}^{bak} \end{aligned} \quad (10)$$

VI. EXPERIMENTAL RESULTS AND DISCUSSION

We will demonstrate our proposed architecture-assisted reliability margining methodology for both power- and thermal-constrained scenarios. The experiment setup is described in

Section VI.A. The results for EM margining and BTI margining are presented in Section VI.B and VI.C respectively. Some of the results and the pessimism in our approaches are discussed in Section VI.D.

A. Experiment Setup

The power model is based on a commercial processor benchmark and a commercial sub-32nm process technology and libraries. The processor benchmark has 32KB L1 instruction cache, 32KB L1 data cache and 256KB L2 cache. The power values are derived from the synthesized, placed and routed results of the processor benchmark. Libraries are characterized at different supply voltages to calculate the operating frequency and power. Five different voltage states are used with supply voltage ranging from 0.6V to 0.9V. So the total number of available states is six including the state when the core is shut down or power gated. The SRAM power is derived from the memory compiler results. The power and thermal effects of other un-core components can be accounted for by reducing the power budget (i.e., P_{max}) or adding into the background temperature (i.e., T_{bak}).

The thermal model (i.e., the values in **A**) is pre-characterized using HotSpot [12, 22]. The chip floorplan size is configured as 15mm x 15mm. The ambient temperature is set as 300 K. Other parameters are set the same as default HotSpot settings. Accumulation model is derived using the BTI simulator [7]. The BTI simulator results are scaled to match the built-in NBTI and PBTI models in the process library. We report the BTI degradation (i.e., ΔV_{th}) as the sum of NBTI and PBTI degradation (i.e., $\Delta V_{thn} + \Delta V_{thp}$).

In the experiments, we will compare against conventional margining approaches, which derives the margin assuming the core is operating at its highest power state along the entire lifetime (i.e., equivalent to the cases when $d = 0$). In order to test the applicability and scalability of our proposed approaches, we consider the case of 4, 16, 64 and 256 cores which are placed in 2x2, 4x4, 8x8 and 16x16 arrays respectively. The power/thermal constraints are determined by the dark ratio d defined in Equation 1 and 9.

B. EM Results

As mentioned earlier in Section III.A, we will apply our margining methodology for both local PDN and signal wires. The accumulation function for unipolar and bipolar current is applied as discussed in Section III.A. For both the power-constrained and thermal-constrained scenarios, we obtain the value of P^{max} and T^{max} by varying the value of dark ratio d . The optimization results of Equation 8 and 10 are reported and normalized with respect to the conventional margining approaches (i.e., the case when $d = 0$).

The results of local PDN under power constraints are plotted in Fig. 3. Compared to our margining approach, conventional approach gives about 3X underestimate of the MTTF at 60% dark ratio, which is equivalent to about unnecessary 40% metal width margin. The results of signal wires under power constraints are plotted in Fig. 4. The MTTF underestimate by conventional margining approach is about 7X, which is equivalent to 60% redundant metal width margin.

We also record the worst-case power state results generated by our methodology. For local PDN, the worst-case happens when some of the cores are at the highest power state and other cores are idle. For signal wires, the worst-case happens when all cores are active and at some intermediate power states. This also suggests different preferences in power management schemes for EM purposes.

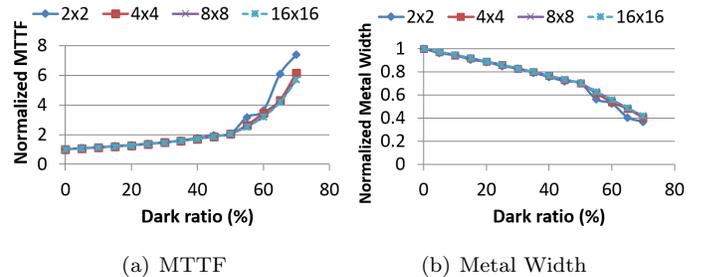


Fig. 3. EM results with power constraints for local power distribution network.

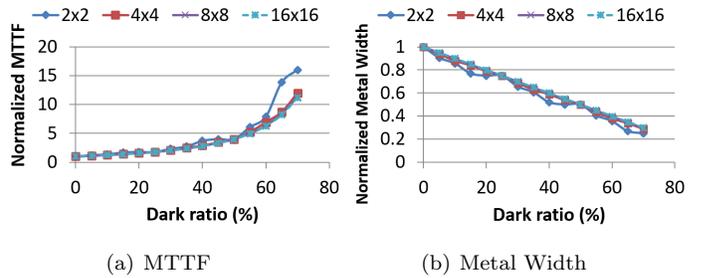


Fig. 4. EM results with power constraints for signal wires.

The results of same scenarios under thermal constraints are plotted in Fig. 5 and Fig. 6. Compared to the results under power constraints, the curves are smoother. This is because the power-constrained problem is an integer problem, while the thermal-constrained problem is continuous. The over-pessimism of conventional margining approach is similar to the cases with power constraints. At 60% dark ratio, our approaches can result in 3-6X less MTTF overestimation and 40%-60% metal width margin reduction.

C. BTI Results

We also apply our methodology for BTI margining. A linear accumulation function is fitted by using the RD simulator as described in Section III.A. The results for Vth degradation are plotted in Fig. 7. At 60% dark ratio, conventional margining approaches result in about 3mV overestimate of the threshold voltage degradation, which is about 18% of the total degradation. We also apply this degradation difference on a buffer chain and use SPICE to study the delay impacts. The around 16 mV threshold degradation requires up to 14% margin on the total delay at the lowest supply voltage (i.e., 0.6V). At 60% dark ratio, our margining approach is able to reduce these delay margin values by about 20%, i.e., about 3% total delay margin reduction. The margin reduction can imply even more saving in the final design if we consider the interaction between aging margin and design flow [8].

Among all the results, the number of processor cores does

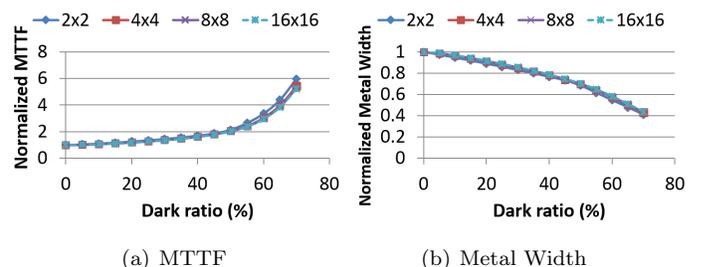


Fig. 5. EM results with thermal constraints for local power network.

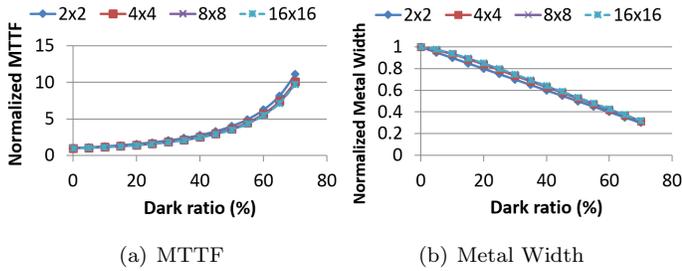


Fig. 6. EM results with thermal constraints for signal wires.

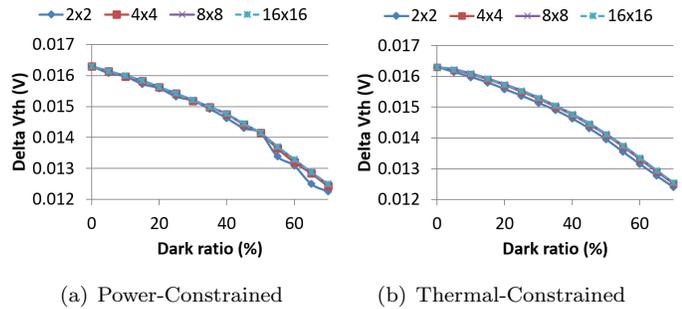


Fig. 7. BTI results of Vth degradation with power constraints and thermal constraints.

not significantly affect the results. The only exception is the power-constrained cases with high dark ratio values. This is caused by the discrete power states and integer variables in power-constrained problem.

D. Pessimism in Our Approach

Due to the nature of design margining, the approach has to guarantee its pessimism against all possible pathological cases. Throughout our margining methodology, possible pessimism is introduced. We summarize and discuss some of them here.

- As discussed in Section III.A, due to the history effects of temperature, we assume the worst-case temperature for all power states.
- As discussed in Section III.A, the accumulation model, if derived from the simulator, can result in pessimism. For example, the BTI accumulation model always assumes the worst-case ordering of power states.
- As discussed in Section III.B, the power and thermal models used in this approach are pessimistic, e.g., the use of best-case power values under process, temperature and workload variations. This will result in optimistic estimation of the number of active cores and therefore result in pessimistic estimation of the margin.

VII. CONCLUSION

In this work, we propose an architecture-assisted hardware reliability margining methodology for chips in the dark silicon era. By exploiting the properties of workload and system management policies, we formulate the margining as an optimization problem. Experimental results show that at 60% dark ratio, our method can reduce the over-pessimism in conventional margining approach for EM and BTI by 3-7X and 18% respectively, which is equivalent to 40%-60% reduction in metal width margin and 20% reduction in delay margin. Our ongoing work is looking at the margining problem for heterogeneous multi-core systems.

REFERENCES

- [1] Intel core i7-920xm processor extreme edition. <http://ark.intel.com/products/43126>.
- [2] Intel turbo boost. http://en.wikipedia.org/wiki/Intel_Turbo_Boost.
- [3] K. Banerjee et al. Coupled analysis of electromigration reliability and performance in ulsi signal nets. In *Proc. IEEE/ACM International Conference on Computer-Aided Design*, pages 158–164. IEEE Press, 2001.
- [4] S. Bhardwaj et al. Predictive modeling of the nbti effect for reliable design. In *IEEE Custom Integrated Circuits Conference*, pages 189–192. IEEE, 2006.
- [5] J. R. Black. Electromigration: a brief survey and some recent results. *Electron Devices, IEEE Transactions on*, 16(4):338–347, 1969.
- [6] A. Calimera et al. Nbti-aware power gating for concurrent leakage and aging optimization. pages 127–132. ACM, 2009.
- [7] T.-B. Chan et al. On the efficacy of nbti mitigation techniques. In *IEEE/ACM Design, Automation and Test in Europe*, pages 1–6. IEEE, 2011.
- [8] T.-B. Chan et al. Impact of adaptive voltage scaling on aging-aware signoff. In *IEEE/ACM Design, Automation and Test in Europe*, pages 1683–1688. IEEE, 2013.
- [9] J. Charles et al. Evaluation of the intel core i7 turbo boost feature. In *Workload Characterization, 2009. IISWC 2009. IEEE International Symposium on*, pages 188–197. IEEE, 2009.
- [10] H. Esmaeilzadeh et al. Dark silicon and the end of multicore scaling. In *International Symposium on Computer Architecture*, pages 365–376. IEEE, 2011.
- [11] P. Gupta et al. Underdesigned and opportunistic computing in presence of hardware variability. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2012. Keynote Paper.
- [12] W. Huang et al. Hotspot: A compact thermal modeling methodology for early-stage vlsi design. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 14(5):501–513, 2006.
- [13] W. R. Hunter. Self-consistent solutions for allowed interconnect current density. ii. application to design guidelines. *Electron Devices, IEEE Transactions on*, 44(2):310–316, 1997.
- [14] C. Isci et al. An analysis of efficient multi-core global power management policies: Maximizing performance for a given power budget. In *Microarchitecture, Annual IEEE/ACM International Symposium on*, pages 347–358. IEEE Computer Society, 2006.
- [15] A. B. Kahng et al. On potential design impacts of electromigration awareness. In *Proc. Asia and South Pacific Design Automation Conference*, pages 527–532. IEEE, 2013.
- [16] U. R. Karpuzcu et al. The bubblewrap many-core: popping cores for sequential acceleration. In *Microarchitecture, IEEE/ACM International Symposium on*, pages 447–458. IEEE, 2009.
- [17] F. Kriebel et al. Aser: Adaptive soft error resilience for reliability-heterogeneous processors in the dark silicon era. In *Proc. ACM/IEEE Design Automation Conference*, pages 1–6. IEEE, 2014.
- [18] R. Kumar et al. Single-isa heterogeneous multi-core architectures: The potential for processor power reduction. In *Microarchitecture, Annual IEEE/ACM International Symposium on*, pages 81–92. IEEE, 2003.
- [19] S. V. Kumar et al. Impact of nbti on sram read stability and design for reliability. In *IEEE International Symposium on Quality Electronic Design*, pages 6–pp. IEEE, 2006.
- [20] E. Rotem et al. Power-management architecture of the intel microarchitecture code-named sandy bridge. *IEEE Micro*, (2):20–27, 2012.
- [21] K. Sekar. Power and thermal challenges in mobile devices. In *Proceedings of the 19th annual international conference on Mobile computing & networking*, pages 363–368. ACM, 2013.
- [22] K. Skadron et al. Temperature-aware microarchitecture. In *ACM SIGARCH Computer Architecture News*, volume 31, pages 2–13. ACM, 2003.
- [23] J. Srinivasan et al. Lifetime reliability: Toward an architectural solution. *Micro, IEEE*, 25(3):70–80, 2005.
- [24] J. Velamala et al. Statistical aging under dynamic voltage scaling: A logarithmic model approach. In *IEEE Custom Integrated Circuits Conference*, pages 1–4, 2012.
- [25] G. Venkatesh et al. Conservation cores: reducing the energy of mature computations. In *ACM SIGARCH Computer Architecture News*, volume 38, pages 205–218. ACM, 2010.
- [26] W. Wang et al. An integrated modeling paradigm of circuit reliability for 65nm cmos technology. In *IEEE Custom Integrated Circuits Conference*, pages 511–514. IEEE, 2007.