

Dynamic Reliability Management for Near-Threshold Dark Silicon Processors *

Invited Paper

Taeyoung Kim[†], Zeyu Sun[§], Chase Cook[§], Jagadeesh Gaddipati[§],
Hai Wang^{**}, Haibao Chen[§] and Sheldon X.-D. Tan[§]

[†]Department of Computer Science and Engineering, University of California, Riverside, CA 92521

[§]Department of Electrical and Computer Engineering, University of California, Riverside, CA 92521

[§] Department of Micro/Nano-electronics, Shanghai Jiao Tong University, Shanghai 200240, China

^{**} School of Microelectronics and Solid-State Electronics, UESTC, Chengdu 610054, China

ABSTRACT

In this article, we propose a new dynamic reliability management (DRM) techniques at the system level for emerging low power dark silicon manycore microprocessors operating in near-threshold region. We mainly consider the electromigration (EM) failures. To leverage the EM recovery effects, which was ignored in the past, at the system-level, we propose a new equivalent DC current model to consider recovery effects for general time-varying current waveforms so that existing compact EM model can be applied. The new equivalent DC current is calculated in two steps: firstly, the equivalent square waveform is calculated so that peak and terminal stresses are matched, secondly, the parameterized equivalent DC current is derived in terms of the parameters of the periodic fitted square waveforms from the first step. The new recovery EM model can allow EM-induced lifetime to be better managed at the system level. The system level energy optimization problem considering EM lifetime subject to power and performance constraints is framed by seeking the best dark silicon cores' voltage and on/off status. The resulting problem is solved by the State-Action-Reward-State-Action (SARSA) reinforcement learning algorithm. Experimental results on a 64-core near-threshold dark silicon processor show that the new equivalent EM DC currents can fully exhibit the recovery effects at the system-level so that trade-off between EM lifetime and energy/performance can be easily made. We further show that the proposed learning-based energy optimization can effectively manage and optimize energy subject to reliability, given power budget and performance limits. When the recovery effects are considered, the new optimization method can achieve 8.6X longer lifetime at the costs of 2.0X more energy and 3.3X more performance degradation.

*This work is supported in part by NSF grant under No. CCF-1255899, in part by NSF grant under No. CCF-1527324, in part by Semiconductor Research Corporation (SRC) grant under No. 2013-TJ-2417 and in part by DARPA grant under No. HR0011-16-2-0009

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICCAD '16, November 07-10, 2016, Austin, TX, USA

© 2016 ACM. ISBN 978-1-4503-4466-1/16/11...\$15.00

DOI: <http://dx.doi.org/10.1145/2966986.2980080>

1. INTRODUCTION

Reliability becomes more and more significant as technology scales into 10nm and below. It was expected that the future chips will show sign of reliability-induced age much faster than the previous generations. Reliability issues have to be addressed in a cross-layer way. As the power density cannot be kept an constant (or start to increase) as technology scales due to the failure of Dennard's scaling [1]. The consequence is the emergence of so-called dark silicon manycore microprocessors, which mean only a percentage of cores can be powered on the chip due to the power and temperature limitations. Recently, architecture researchers have begun focusing on the development of dark-silicon manycore devices with as many as 100 or 1000 cores integrated onto a single die. Such manycore systems pose new challenges and opportunities for power/thermal and reliability management of those chips [2].

To further reduce powers for many applications, ultra-low power designs become necessary. Recent research has led to sub-threshold region where CMOS circuits are found to be capable of operating with a supply voltage of less than 200mV. The theoretical lower limit of V_{dd} has been determined to be 36 mV [3]. But at such low voltages, a leakage power dissipation increases drastically making the reduction in dynamic power insignificant. Also the circuit delay increases rapidly as the supply voltage is scaled down, resulting in decreased operation frequency or performance of the circuits.

For dark silicon manycore processors operating in near threshold voltage, reliability becomes quite significant for the long-term reliability such electromigration. To address the increasing reliability issues, a system-level and run-time approach becomes more appealing. There are some existing works on dynamic reliability managements for dark silicon in the past [4, 5]. These works have been proposed to leverage the dark silicon many-core processors in order to save energy while maintaining performance considering reliability. Run-time management of the heterogeneous dark silicon processors and optimal policy of core status have been addressed. Dynamic voltage frequency scaling method have been employed as energy saving techniques in those works. However, dynamic reliability management for near-threshold dark silicon processors has not been studied.

In this article, we propose a new dynamic reliability management (DRM) techniques at the system level for emerging low power dark silicon manycore microprocessors operating in near threshold region. We mainly consider the electromigration (EM) failures. To leverage the EM recovery effects, which was ignored in the past, at the system level, we propose a new equivalent DC current model to consider re-

covery effects for general time-varying current waveforms so that existing compact EM model can be applied. The new equivalent DC current is calculated in two steps: firstly, the equivalent square waveform is calculated so that peak and terminal stresses are matched, secondly, the parameterized equivalent DC current is derived in terms of the parameters of the fitted periodic square waveforms from the first step. The significance of the new EM current model is that it allows EM recovery effects can be considered at the system level for the first time and thus allow EM-induced lifetime of chips to be better managed at the system level.

The system level energy optimization problems considering recovery-aware EM-induced reliability subject to power and performance constraints was framed by seeking the best dark silicon cores' voltage and on/off status. The resulting problem was solved by the State-Action-Reward-State-Action (SARSA) reinforcement learning algorithm. Experimental results on a 64-core near-threshold dark silicon processor show that the new equivalent EM DC currents can fully exhibit the recovery effects at the system level so that trade-off between EM lifetime and energy/performance can be easily made. We further show that the proposed learning-based energy optimization can effectively manage and optimize energy subject to reliability, given power budget and performance limits. When the recovery effects are considered, the new optimization method can achieve 8.6X longer lifetime at the costs of 2.0X more energy and 3.3X more performance degradation.

2. ELECTROMIGRATION MODELING AT SYSTEM LEVELS

2.1 Review of EM physics and physics-based models

EM is the physical phenomenon of the migration of metal atoms along a direction of applied electrical field. Atoms (either lattice atoms or defects/impurities) migrate toward the anode end of metal wire along the trajectory of conducting electrons. Over time, the lasting unidirectional electrical load increases these stresses, as well as the stress gradient along the metal line. In some cases, usually when a line is long, this stress can reach a critical level, resulting in a void nucleation at the cathode and/or hillock formation at the anode end of line.

The EM effects are mainly modeled and heavily used by empirical Black's equation [6] and Blech limit [7]. However, those models are not physics-based and they do not fully consider the predictability for varying stressed conditions and complicated interconnect wire structures. Additionally, they do not address the inherent redundancy in the power grid networks, which are the most vulnerable wires in a chip.

To address those problems, a more physics-based compact EM model has been recently proposed for full-chip reliability analysis [8,9], which is the basis for EM reliability modeled in this work. The EM development process consists of two phases - the nucleation phase and the growth phase. In the first nucleation phase, a closed-form expression to compute the nucleation time (t_{nuc}) is given, which is a function of current density, temperature, and the residual stress of the wire due to thermal and other effects as well as other wire geometry and material parameters. The approximate value of void nucleation time (t_{nuc}) is determined as the instant in time when stress at the cathode end of the line reaches σ_{crit} , which corresponds well to an analytical formulation of t_{nuc} derived from the approximate solution of continuity equations for evolution of vacancy and plated atom concen-

trations (see, for example [10]) in the confined 1D line.

$$t_{nuc} \approx \tau^* e^{\frac{E_V}{kT}} e^{-\frac{f\Omega}{kT}(\sigma_{Res} + \frac{eZ\rho l}{4\Omega}j)}$$

$$\ln \left\{ \frac{\frac{eZ\rho l}{4\Omega}j}{\sigma_{Res} + \frac{eZ\rho l}{4\Omega}j - \sigma_{CR}} \right\} \quad (1)$$

where $\tau^* = \frac{l^2}{D_0} e^{E_D/kT} \frac{kT}{\Omega B}$. Here, j is the current density, T is temperatures, k_B is the Boltzmann's constant, l is the segment length, E_V and E_D are the activation energies of vacancy formation and diffusion, f is the ratio of volumes occupied by vacancy and lattice atom, σ_{crit} is the critical stress needed for the failure precursor nucleation (void/hillock). σ_{Res} is the residual stress of the metal segment from the cooling process and other factors.

The second phase is the void size growth: voids are formed at t_{nuc} and grow at $t > t_{nuc}$. The wire resistance starts to increase over the time in the growth phase [8].

2.2 New equivalent DC current based modeling for EM recovery effects

For EM failures, one of the important phenomenon is that the EM-induced stress can go down when the stressing current becomes small. This effect is called "EM recovery effect" as it represents important transient effects due to time-varying currents. Fig. 1 shows the EM-induced stress changes over time over a periodic current pulse. As we can see the stress can go down significantly. The net effect for such recovery effect is that the lifetime of wire due to EM can be extended significantly as it will take longer time for the stress to reach to the critical stress over time.

However, EM recovery effects were ignored completely in the existing EM models as most of those models assume constant current or current density. To mitigate this problem, a physics-based EM recovery model was proposed recently [11,12] by obtaining an analytical solution of the Korhonen's equation describing the stress evolution kinetics of EM effects. Although the accuracy of this model is high, it is still too complicated for practical use.




Figure 1: Stress evolution caused by periodical current density

For practical chip design, EM assessment and signoff still uses simple EM models like Black's model [6] or recently proposed, more accurate EM model in (1), which takes constant current and temperature as inputs. In order to consider practical no-DC currents, a simple time-varying equivalent DC current is computed as following,

$$j_{trans,EM,eff} = \frac{1}{P} \left(\int_0^P j_+(t) dt - \gamma \int_0^P |j_-(t)| dt \right); \quad (2)$$

where $j_+(t)$ and $j_-(t)$ are the current densities of the positive and negative phases of the bipolar current, γ is the EM recovery factor, P is the period of the current waveform. When the current density is unidirectional, $j_{trans,EM,eff}$ essentially the time averaged current density. However, using the effective current formula in (2) will create a number of

problems [13]. First of all, the recovery factor depends on the specific current waveforms which is not constant. Also, it ignores important transient effects such as the recovery and peak stress effects. Fig. 2 shows the stress evolutions over time driven by two current waveforms, the actual one and the time-varying equivalent DC current. As we can see, the peak stress due to the actual current waveform can exceed the critical stress while the average current never leads to void nucleation (wire is immortal).




Figure 2: Stress evolutions caused by actual currents and traditional effective DC current

In order to solve the problems of these models, in this paper, we propose a new and novel equivalent DC current method to consider the transient EM recovery effects. The new model is based on first-principle based numerical analysis of EM effects. Here we use nucleation phase to compute the time to failure of a wire as a demonstration of this proposed method. The idea is that for a given EM model, the DC equivalent current will lead to the same time to failure (TTF) computed from the detailed numerical EM analysis of stress diffusion equation. This is better illustrated in Fig. 3(b) in which the periodic current density and a DC current gives the same nucleation time t_{nuc} . Unlike the traditional method which could ignore the case that the peak stress exceeds the critical stress but the equivalent current density never leads to void nucleation, so the transient effects are explicitly taken into account in this method.

That model works well for standard periodic square waveforms with one high current density (j_1) and one low current density (j_2). As shown in Fig. 3(a), j_1 , j_2 , period (P) and duty cycle (D) are used as the variables for the model. Also we find that the temperature (T) is one of the dominant parameters for the equivalent EM DC equivalent current density (j_{em}).

To further derive the parameterized equivalent DC current in terms of two currents, period and duty cycle and temperature, response surface methodology (RSM) [14] is carried out over many different waveforms from measured or detailed numerical analysis information. (3) is the fitted model to obtain equivalent DC current in terms of the five parameters.

$$\begin{aligned}
 j_{em} = & 4.988 * 10^9 - 0.0663 * 10^9 * X_1^2 - 1.114 * 10^9 * X_1 \\
 & * X_2 - 0.9981 * 10^9 * X_1 * X_3 - 0.1390 * 10^9 * X_1 * X_4 \\
 & - 0.3485 * X_1 * X_5 - 0.0315 * 10^9 * X_2^2 - 0.1728 * 10^9 \\
 & * X_3^2 - 0.3461 * 10^9 * X_3 * X_4 + 0.0181 * 10^9 * X_4^2 \\
 & + 0.0934 * 10^9 * X_5^2
 \end{aligned} \quad (3)$$

where

$$\begin{aligned}
 X_1 &= \frac{j_1 - 7.5 * 10^9}{2.5 * 10^9} & X_2 &= \frac{j_2 - 1.75 * 10^9}{2.75 * 10^9} \\
 X_3 &= \frac{D - 70\%}{25\%} & X_4 &= \frac{P - 5.5 * 10^3}{4.5 * 10^3} \\
 X_5 &= \frac{T - 380}{20}
 \end{aligned}$$




Figure 3: (a) Original input driving current density. (b) Calculated EM DC equivalent current density with t_{nuc}

However, this model can only handle regular square waveforms, but for practical cases, the current waveforms are arbitrary. To mitigate this problem, one of the ideas is to convert the arbitrary current waveform to an equivalent square waveform before we apply the aftermentioned parameterized equivalent DC current modeling. In this conversion process, we make sure that the stresses derived by the square waveform and the actual current waveform will match at both the highest peak stress and the final stress (end of period or time) as shown in Fig. 4(b). By matching the two stress points, we can find the two currents j_1 for highest stress point and j_2 for the end of period stress as shown in Fig. 4. During this conversion process, we assume that the given current waveform will repeat itself over time so that it becomes a periodic waveform. This assumption is reasonable as the future current or power of a chip cannot be predicted precisely in general and the recurrent assumption is a good guess.

The other idea is to convert the arbitrary current waveform directly to a DC equivalent current so that the stresses from two waveforms match at the end of period time as shown in Fig. 4(b). But this approach may lead to large errors for time to failure estimation as it ignores the peak stress, which can be significant to determine the time reaching the critical stress (time to failure).




Figure 4: (a) Original input driving current density. (b) Calculated EM DC equivalent current density with two methods

To study the accuracy of the two modeling methods, the two-step method (square waveform modeling and RSM fit-

ting and the direct equivalent DC current method) is proposed. We compare stress generated by two step method and the stress given by direct equivalent DC current method against the stress generated by the original current waveform. The results are shown in Fig. 5. As we can see, equivalent square DC current density (two step method) has smaller error compared to the direct DC equivalent method in terms of time to failure estimation. As a result, in this paper, we will use the two-step method to compute the parameterized equivalent DC current.




Figure 5: Comparing the nucleation time of two equivalent methods and original stress

2.3 EM modeling for varying temperature effects

At the system-level EM reliability, the manycore system will run on different tasks under different voltages and frequencies. As a result, its temperature and current densities will change with time. However existing EM models including the new physics-based model can only take a constant temperature. The previous study shows that whole system MTTF or lifetime under different temperature can be approximated by [15]:

$$MTTF_{EM} = \frac{1}{(\sum_{m=1}^n (\Delta t_m \frac{1}{MTTF_m})) / T} \quad (4)$$

where $MTTF_m$ is the actual MTTF (mean time to failure) under the m -th power and temperature settings for Δt_m period, assuming the chip works through n different power and temperature settings and $T = \sum_{m=1}^n \Delta t_m$. Each $MTTF_m$ will be computed based on the EM models discussed in the previous section.

To consider a system-level EM reliability on a manycore dark silicon processor, we use the shortest lifetime among all the cores as the lifetime for all manycore processors [16].

3. NEW LEARNING-BASED RELIABILITY MANAGEMENT FOR NEAR THRESHOLD DARK SILICON FOR EM RECOVERY EFFECTS

3.1 Near threshold dark silicon

Near-Threshold Computing (NTC) has been proposed as a viable solution to overcome the limit of energy efficient computing by using optimal near-threshold voltage between super-threshold and sub-threshold region.

NTC cores are operated at or near their threshold voltage V_{th} . By reducing supply voltage V_{dd} from nominal 1.1 V to 500 mV, a 10X energy efficiency gain can be achieved at the expense of 10X performance degradation [17]. Compared

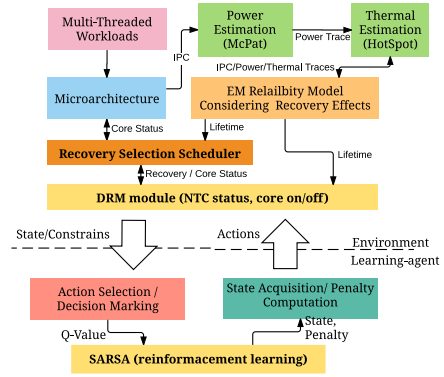


Figure 6: The DRM and NTC Framework

to the sub-threshold region where 20X energy efficiency can be achieved, but the 50X performance degradation due to increased circuit delay is too big a factor to ignore for large-scale applications. Applications with significant standby times benefit greatly from NTC. Memories have to retain their contents even though digital logic is to be powered off. Thus supply voltage scaling results in a significant reduction in leakage power.

On the other hand, NTC is also a promising technique to mitigate the effects of dark silicon as cores can reduce power and temperature under a given power budget, thus, allowing a larger number of cores to be turned on simultaneously at costs of allowed performance losses. Recently, instead of operating the entire cores at either nominal voltage or near-threshold voltage (NTV), voltage islands have been defined such that only partial cores are operated at NTV and the rest is operating at nominal voltage for more flexible trade-off between power and performance [18]. Supply voltage is proportional to the threshold voltage of the transistors in the core. The core with the highest threshold can determine supply voltage for the voltage island. However, the different types of parallel workloads can lead to performance degradation and energy waste. Efficient dynamic management and scheduling to find suitable NTC regions are needed.

In addition to energy and performance, NTC has an effect on reliabilities. NTC may exhibit better long-term EM reliability, as a lower voltage can lead to lower temperature, current density and residual stresses, which are the major factors of EM effects [8]. The NTV, which is a lower supply voltage, can improve EM-induced lifetime of dark silicon processors. However, using NTV for many core can make significant performance issue since NTV still use some cores operated at the nominal voltage and a many core system's reliability can be highly affected by those core's reliability [16].

3.2 Framework of dark silicons in near-threshold computing region

We present the framework for Dynamic Reliability Management (DRM) at NTC region in dark silicon. The DRM framework employs several simulator models (microarchitecture, power, thermal), a policy optimization module, all in conjunction with EM recovery model. Additionally, the DRM has policy optimizer that cores can choose the best NTC policy to maximize energy efficiency while meeting performance limit and power budgets. This work uses a 45nm-based 64-core dark silicon simulation framework with the threshold voltage of $V_{th} = 0.20V$, core On/Off knobs, and NTV capabilities.

The DRM module sets the voltage and frequency policy for the chip. Each core (k) can be assigned a voltage (V_k), nominal ($V_k^{nom} = 1.0v$) or near threshold (V_k^{nt}) which is defined as $0.40v \leq V_k^{nt} < V_k^{nom}$. Additionally, the DRM can turn a core off (V_k^{off}) and then turn a core back on ($V_k^{nom/nt}$). The DRM makes these decisions based on an on-line policy optimization module that employs the SARSA algorithm which is explained later. In the framework, the DRM receives the new policy from the optimization module. It then sets each V_k or turns the core off. Additionally, each core operating frequency (f_k) is affected by V_k , because of this we use 5 as a relation to calculate f_k based on its respective V_k . This ensures that the performance degradation from NTV is reflected in the simulation framework. This policy is then propagated to the architecture, thermal, and power simulators as well as the EM recovery model, and optimizer.

$$f_k \propto \frac{(V_k - V_{th})^2}{V_{th}} \quad (5)$$

The framework uses the Sniper architecture simulator [19] to generate system performance for given workloads on a specified architecture. Parameters (chip floorplan, number of cores, frequency, and cache design) describing the architecture are passed to Sniper. Benchmarks representing the desired workloads are also used as inputs to sniper to simulate the system’s functionality. Sniper then outputs system performance, such as performance characteristics, instruction per cycle, of the chip for each given benchmark run. This is repeated in our experimental setup for several different set of workloads. The whole near threshold dark silicon framework is illustrated in Fig. 6.

Once architectural system performance is generated and transferred to the physical simulators Hotspot [20] and MCPAT [21]. Based on the architecture of the chip, its system performance from Sniper, and the voltage scheduling from the DRM, MCPAT (Multi-Core Power Area and Timing simulator) will generate a power trace for each chip component including each core P_k . Hotspot then uses the chip floorplan in conjunction with the power trace generated by MCPAT, to produce a thermal trace for each chip component and core T_k .

After the power/thermal/voltage characteristics have been generated by the various simulators, the EM Recovery model can use these parameters (P_k , T_k , and V_k) to estimate the chip’s time to failure considering any recovery effects the chip may experience, from V_k^{nt} and V_k^{off} for that given policy.

Lastly, all the information generated, in addition to the current policy enacted by the DRM, are passed to the policy optimizer which will generate a new policy. This new policy will find the best voltage schedule for the various cores to optimize the energy of the chip while meeting MTTF, performance, and power budgets.

3.3 SARSA-based learning algorithm for DRM considering long-term recovery

We can model our DRM problem as a Markov Decision Process (MDP) with states $s(t)$, and actions $a(t)$ where states are the parameters of the framework for the time-step Δt , e.g., $f_k(t)$, $T_k(t)$, $MTTF(t)$, $P_k(t)$, and $V_k(t)$. Actions are defined as changes to these parameters which in our case is the tuning of $V_k(t)$. In our case, our goal is to achieve the best policy that minimizes energy E while meeting all constraints and budgets.

The reinforcement learning algorithm used to optimize the DRM policy is the State Action Reward State Action algo-

rithm, or SARSA, first presented in [22]. SARSA is a combination of Q-learning and the traditional Temporal Difference method (TD) [22]. This algorithm exchanges the greedy updates of Q-learning with a policy driven update that is closer to the TD method. The result is an on-policy reinforcement learning algorithm with faster convergence when compared to Q-learning [22].

The major differences with traditional Q-learning, is that the maximum reward (minimum penalty) for the next state is not used for updating the Q-values. Instead, a new action is selected using the same policy that determined the original action.

The SARSA algorithm works first by populating a Q-table $Q(s(t), a(t))$, where $s(t)$ is a state and $a(t)$ is an action for time-step Δt . It then selects an action from the states using some policy. This action is taken and the penalty $PT(t+1)$ (negative reward) and new state $s(t+1)$ are generated. From this new state, another action $a(t+1)$ is selected from $s(t+1)$. The Q-table is then updated using a penalty function shown below.

$$Q^{t+1}(s(t), a(t)) = Q^t(s(t), a(t)) + \alpha(t) \times (PT(t+1) + \gamma Q(s(t+1), a(t+1)) - Q(s(t), a(t))) \quad (6)$$

Here, α is the learning rate and γ is the discount factor. In our DRM, we employ multiple-constrained penalty function ($PT(t+1)$, [4]) and modify to accommodate each value (EM, power, temperature, and performance) and to also incorporate the power budgets, performance/thermal limits assigned as constraints.

In order to provide long-term shut-off time for leveraging EM recovery effects, our DRM have a recovery selection scheduler. It is a periodic scheduler on the top of SARSA for long-term recovery cycle ($T_{recovery}$), which is the required time for recovery effects. The selected cores can be turned off long-term recovery cycle as recovery effects needs some long-term as seen in Subsection 2.2. Every $T_{recovery}$ cycle, we use greedy-based selection algorithm by EM-induced lifetime evaluation and find the worst lifetime core set below the certain lifetime threshold ($EM_{threshold}$), then SARSA will work only for the cores except recovering cores for $T_{recovery}$. After each long-term cycle, we find new long-term recovery core set.

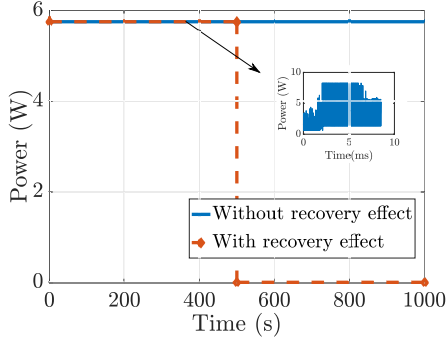
The proposed new energy optimization algorithm in the near-threshold dark silicon framework has been implemented in Python 2.7.9 with the numerical libraries (Numpy 1.9.2 and Scipy 0.15.1). For near-threshold dark silicon framework, we used the architectural simulator (Sniper 6.1), power estimator (McPAT 1.0.32), and thermal simulator (Hotspot 5.02 [23]) to estimate recovery-considered EM-induced lifetime on top of the new physics-based EM model [8]. In the proposed framework as shown in Fig 6, each simulator module is connected with a plugin connector, so that one simulator’s result can dynamically provide the other’s inputs. The learning-based SARSA method and recovery selection scheduler have been implemented for our dynamic reliability management (DRM).

The DRM module assigns each core as the near-threshold or nominal voltage, and calculates its frequency from 5. The near threshold voltages as 0.45V, and a nominal voltage of 1.2V are defined [24]. Additionally, each core can be turned off completely, or turned back on for dark-silicon. Once the simulation runs, the optimization module can then send a new policy to the DRM which it will use to schedule the core voltages. Our energy optimization method is validated with a near-threshold 64-core processor on the SPLASH-2 multi-threaded benchmarks.

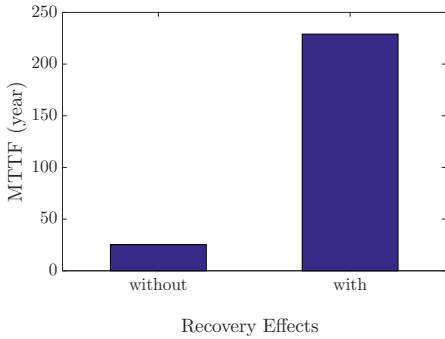
4. NUMERICAL RESULTS

Table 1: Results for dynamic reliability management for 64-core near-threshold dark silicon

64-Core	DRM (No recovery baseline)	All NTV cores	Half-NTV/Half-Dark	Half-Nominal/Half-Dark
Performance (s/tasks)	0.014	0.23	0.48	0.03
EM lifetime (year)	9.03	221	221	7.1
Energy (Joule/tasks)	0.942	14.54	29.05	2.07



(a)



(b)

Figure 7: (a) two cases of power traces from proposed framework and (b) and the resulting MTTF without/with recovery effects

4.1 Evaluation of the lifetime impacts considering EM recovery effects

In order to evaluate system-level EM-induced lifetime considering recovery effects, a single core long-term task example case from our framework is shown here. Our proposed framework, shown in Fig. 6, can properly manage and control for both power-on and shut-down of each core, so we can significantly extend the system-level reliability by leveraging the EM recovery effects introduced in subsection 2.2. We present two different simple power traces and calculate EM-induced MTTF (mean time to failure) in Fig. 7. In this example, the time period is 1000 seconds and switch off for 500 seconds that can be recognized as a sufficient period for recovery effect. The original power traces (5.72 W) is converted to an equivalent power, which equals to 2.122W using our recovery model. Our duty cycle for turning-off is 50% in

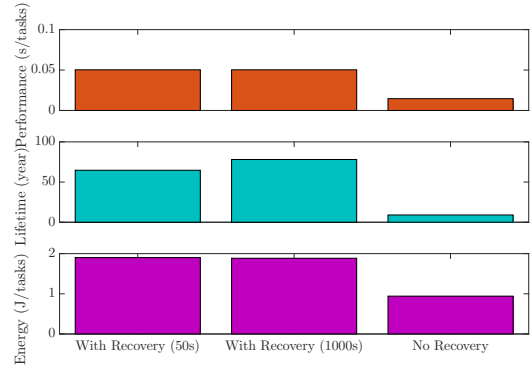


Figure 8: Performance, energy and EM-induced lifetime from proposed DRM considering recovery effects for three cases (1) Recovery effects with $T_{recovery} = 50s$ (first column) (2) Recovery effects with $T_{recovery} = 1000s$ (the second column) (3) Only DRM without recovery effects (the third column)

the recovery case. As a result, it leads to 50% performance degradation. On the other hand, the EM-induced lifetime considering recovery effect is 9.16X higher than that of the original case without recovery effect case, which is quite significant.

4.2 Evaluation of the DRM for near-threshold dark silicon processors

To evaluate the proposed learning-based energy optimization method in Section 3, we show the comparison results of performance, EM-induced lifetime, energy consumption on processing 64 multi-threaded tasks on a 64-core near-threshold dark silicon processor. Our experiment uses performance and energy metrics as $s/tasks$ and $J/tasks$, which are total execution time and energy consumption for selected 64 multi-threaded tasks (16 CHOLESKYs, 16 RADIXs, 16 RAYTRACEs, 16 VOLRENDs) on our framework. The results are shown in Table 1. We compare our DRM case (without recovery effects) with all-NTV cores (64 cores are near-threshold voltages(NTV)), half-NTV/half-dark (32 cores are NTV and 32 cores are turned off), and half-nominal/half-dark (32 cores are nominal and 32 cores are turned off) cases. The optimization results show the energy consumption improvements for the given budget constraints (power budget is 250W, performance limit 5s per given tasks, and EM-induced lifetime limit is more than 5 years). For all-NTV and half NTV/dark cases, EM lifetimes are significantly improved, but their performance results are 16X and 35X slower than the DRM baseline result. On the other hand, for only half nominal/half dark case, energy and per-

formance are quite improved but EM lifetime is relatively short. Therefore, our DRM result effectively finds better EM lifetime (9.03 years) with significantly high performance with the lowest energy.

4.3 DRM considering recovery effects

Finally, we evaluate our proposed DRM method considering the EM recovery effects. As seen in Fig. 8, we evaluate two different recovery cycles ($T_{recovery} = 50s$ and $1000s$), so all the cores' MTTFs are periodically evaluated for every $T_{recovery}$ to determine which core needs to be turned off for the whole period when the core is below the certain lifetime threshold, $EM_{threshold}$. For our experiment, we set 8 years as the $EM_{threshold}$ in the recovery selection scheduler. As we can see, both DRM cases with recovery effects can significantly improve our EM-induced lifetime (8.6X longer compared to the baseline results, which are shown in the third column in Fig. 8). However, the costs are the 2.0X more energy consumption ($1.9J/tasks$ vs $0.94j/tasks$) and 3.3X performance degradations ($0.05s/tasks$ vs $0.014s/tasks$). But this is still a better trade off for the higher EM lifetime (64.7 years and 78.1 years for $T_{recovery} = 50s$ and $1000s$, respectively) compared with the baseline case in Fig. 8.

5. CONCLUSION

In this paper, we proposed a new dynamic reliability management (DRM) techniques for emerging near-threshold dark silicon manycore microprocessors considering electromigration (EM) reliability. To leverage the EM recovery effects, which was ignored in the past, at the system level, we propose a new equivalent DC current model to consider recovery effects for general time-varying current waveforms so that existing compact EM model can be applied. The new EM current model allows EM recovery effects to be effectively considered at the system level for the first time. To leverage the EM recovery effects, we considered the energy optimization problem for dark silicon manycore processors with Near-Threshold Voltage (NTV) capabilities considering EM reliability. We showed that the on-chip power consumptions have different impact on reliability. The resulting optimization problem was solved with State-Action-Reward-State-Action (SARSA) reinforcement learning algorithm to optimize the near-threshold dark silicon cores' voltage policy to minimize energy considering reliability. Experimental results on a 64-core near-threshold dark silicon processor showed that the new equivalent EM DC currents was able to fully exhibit the recovery effects at the system level so that trade-off between EM lifetime and energy/performance were easily made. We further showed that the proposed learning-based energy optimization can effectively manage and optimize energy subject to reliability, given power budget and performance limits. When the recovery effects were considered, the new optimization method was able to achieve 8.6X longer lifetime at the costs of 2.0X more energy and 3.3X more performance degradation.

6. REFERENCES

- [1] R. Dennard, F. Gaensslen, H. Yu, V. Rideout, E. Bassous, and A. LeBlanc, "Design of ion-implanted mosfet's with very small physical dimensions," *IEEE Journal of Solid-State Circuits*, vol. 9, pp. 256–268, October 1974.
- [2] H. Esmailzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," in *Proceedings of the 38th Annual International Symposium on Computer Architecture*, ISCA '11, (New York, NY, USA), pp. 365–376, ACM, 2011.
- [3] R. M. Swanson and J. D. Meindl, "Ion-implanted complementary mos transistors in low-voltage circuits," *IEEE Journal of Solid-State Circuits*, vol. 7, pp. 146–153, Apr 1972.
- [4] T. Kim, X. Huang, V. S. H.-B. Chen, and S. X.-D. Tan, "Learning-based dynamic reliability management for dark silicon processor considering EM effects," in *Proc. Design, Automation and Test In Europe. (DATE)*, Mar. 2016.
- [5] M. Shafique, S. Garg, T. Mitra, S. Parameswaran, and J. Henkel, "Dark silicon as a challenge for hardware/software co-design: Invited special session paper," in *Proceedings of the 2014 International Conference on Hardware/Software Codesign and System Synthesis*, CODES '14, (New York, NY, USA), pp. 13:1–13:10, ACM, 2014.
- [6] J. R. Black, "Electromigration-A Brief Survey and Some Recent Results," *IEEE Trans. on Electron Devices*, vol. 16, no. 4, pp. 338–347, 1969.
- [7] I. A. Blech, "Electromigration in Thin Aluminum Films on Titanium Nitride," *Journal of Applied Physics*, vol. 47, no. 4, pp. 1203–1208, 1976.
- [8] X. Huang, T. Yu, V. Sukharev, and S. X.-D. Tan, "Physics-based electromigration assessment for power grid networks," in *Proc. Design Automation Conf. (DAC)*, June 2014.
- [9] V. Sukharev, "Beyond Black's Equation: Full-Chip EM/SM Assessment in 3D IC Stack," *Microelectronic Engineering*, vol. 120, pp. 99–105, 2014.
- [10] V. Sukharev, A. Kteyan, E. Zschech, and W. D. Nix, "Microstructure Effect on EM-Induced Degradations in Dual Inlaid Copper Interconnects," *IEEE Transactions on Device and Materials Reliability*, vol. 9, no. 1, pp. 87–97, 2009.
- [11] V. Sukharev, X. Huang, and S. X.-D. Tan, "Electromigration Induced Stress Evolution Under Alternate Current and Pulse Current Loads," *Journal of Applied Physics*, vol. 118, pp. 034504–1–034504–10, 2015.
- [12] X. Huang, V. Sukharev, T. Kim, and S. X.-D. Tan, "Electromigration recovery modeling and analysis under time-dependent current and temperature stressing," in *Proc. Asia South Pacific Design Automation Conf. (ASPDAC)*, pp. 244–249, 2016.
- [13] K.-D. Lee, "Electromigration Recovery and Short Lead Effect under Bipolar- and Unipolar-Pulse Current," in *IEEE International Reliability Physics Symposium (IRPS)*, pp. 6B.3.1–6B.3.4, April 2012.
- [14] R. H. Myers and D. C. Montgomery, *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*. Wiley-Interscience, 2002.
- [15] 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. on Computer Aided Design (ICCAD)*, pp. 327–334, IEEE, November 2004.
- [16] A. Das, A. Kumar, and B. Veeravalli, "Reliability-driven task mapping for lifetime extension of networks-on-chip based multiprocessor systems," in *Proceedings of the Conference on Design, Automation and Test in Europe, DATE '13*, (San Jose, CA, USA), pp. 689–694, EDA Consortium, 2013.
- [17] R. G. Dreslinski, M. Wieckowski, D. Blaauw, D. Sylvester, and T. Mudge, "Near-threshold computing: Reclaiming moore's law through energy efficient integrated circuits," *Proceedings of the IEEE*, vol. 98, pp. 253–266, Feb 2010.
- [18] C. Silvano, G. Palermo, S. Xydis, and I. Stamelakos, "Voltage island management in near threshold manycore architectures to mitigate dark silicon," in *2014 Design, Automation Test in Europe Conference Exhibition (DATE)*, pp. 1–6, March 2014.
- [19] T. E. Carlson, W. Heirman, and L. Eeckhout, "Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulations," in *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pp. 52:1–52:12, Nov. 2011.
- [20] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. R. Stan, "HotSpot: A compact thermal modeling methodology for early-stage VLSI design," *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, vol. 14, pp. 501–513, May 2006.
- [21] S. Li, J. H. Ahn, R. D. Strong, B. J. B. D. M. Tullsen, and N. P. Jouppi, "Mcpat: an integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, pp. 469–480, ACM, 2009.
- [22] G. A. Rummery and M. Niranjan, *On-Line Q-learning using connectionist systems*. University of Cambridge, Department of Engineering, 1994.
- [23] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware microarchitecture," in *International Symposium on Computer Architecture*, pp. 2–13, 2003.
- [24] H. Kaul, M. Anders, S. Hsu, A. Agarwal, R. Krishnamurthy, and S. Borkar, "Near-threshold voltage (ntv) design: Opportunities and challenges," in *Proceedings of the 49th Annual Design Automation Conference, DAC '12*, (New York, NY, USA), pp. 1153–1158, ACM, 2012.