Multi-Mechanism Reliability Modeling and Management in Dynamic Systems

Eric Karl, *Student Member, IEEE*, David Blaauw, *Senior Member, IEEE*, Dennis Sylvester, *Senior Member, IEEE*, and Trevor Mudge, *Fellow, IEEE*

***Abstract*—Reliability failure mechanisms, such as time-dependent dielectric breakdown (TDDB), electromigration, and negative bias temperature instability (NBTI), have become a key concern in integrated circuit (IC) design. The traditional approach to reliability qualification assumes that the system will operate at maximum performance continuously under worst case voltage and temperature conditions. In reality, due to widely varying environmental conditions and an increased use of dynamic control techniques, such as dynamic voltage scaling and sleep modes, the typical system spends a very small fraction of its operational time at maximum voltage and temperature. In this paper, we show how this results in a reliability "slack" that can be leveraged to provide increased performance during periods of peak processing demand. We develop a novel, real time reliability model based on workload driven conditions. Based on this model, we then propose a new dynamic reliability management (DRM) scheme that results in 20%–35% performance improvement during periods of peak computational demand while ensuring the required reliability lifetime.**

***Index Terms*—Integrated circuit (IC) reliability, reliability management, system-level reliability modeling.**

I. Introduction

TRADITIONAL stress-based reliability qualification techniques, such as the JEDEC JESD-47 Standard [1], qualify designs by stressing sample systems under pessimistic environmental conditions with a zero failure pass/fail criteria. While the traditional approach is an accepted method of ensuring reliability, the limits it places on supply voltage and temperature leave a significant and increasing reliability margin between circuit performance at worst case conditions and at typical conditions. Widely varying environmental conditions linked to portable products combine with dynamic power reduction techniques to exacerbate the limitation of this conventional worst case qualification methodology.

Hence, the need for alternative approaches to ensuring lifetime reliability under dynamic operating conditions is clear. An alternative to stress-based qualification, knowledge-based risk assessment is one alternative to simplistic corner-case stress testing. The knowledge-based approach (a framework for knowledge-based qualification is defined by JEDEC JESD-34 [2]) requires careful characterization and analysis of individual failure modes to assess a reasonable system reliability risk factor given the reliability targets for the system. In this paper, we propose the use of so-called *dynamic reliability management* (DRM), where real-time workloads and thermal information provides accurate inputs to real-time knowledge-based reliability models for projecting the degradation caused by various failure mechanisms. We then use the projected failure probability to control the maximum voltage assigned by a dynamic voltage scaling (DVS) algorithm.

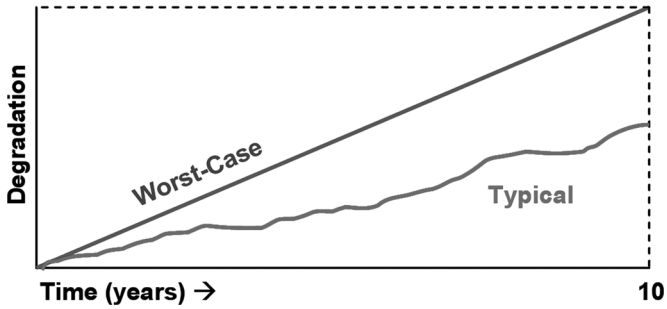

Fig. 1. Reliability degradation over time.

The concept of DRM is conceptually motivated in Fig. 1. The line labeled *worst case profile* shows the accumulated damage due to a failure mechanism, such as oxide breakdown, over a 10-year time span under worst case operating conditions (maximum operating frequency, voltage, and ambient temperature). In traditional analysis, the maximum supply voltage is set such that at the accumulated damage at the 10-year mark results in a failure probability that meets the specified constraint (such as 63.2%, or $1 - e^{-1}$). However, performance traces collected from an actual desktop processor show that the processor spends over 85% of its time in low power or sleep modes where the incurred damage rate is significantly lower. The accumulated damage from such typical usage is shown in Fig. 1 with the line labeled *typical profile* and results in a much lower final damage at the 10-year mark. Hence, the failure probability for this typical usage is well below that of the specification and the maximum allowed operating voltage was unnecessarily constrained resulting in a loss of potential performance.

Since the maximum supply voltage is currently set *a priori* at design time or during post-fabrication testing, it is becoming increasingly difficult to anticipate the actual usage of a part and hence worst case conditions must be assumed. Under DRM, however, it is possible to dynamically monitor the operating voltage and temperature during part operation. With this operating condition history, we propose the use of models to project the expected reliability and to dynamically adjust the maximum

Manuscript received August 22, 2006; revised May 8, 2007. This work was supported in part by the MARCO/FCRP Gigascale Systems Research Center and the National Science Foundation.

The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48105 USA (e-mail: ekarl@umich.edu; blaauw@umich.edu; dennis@eecs.umich.edu; tnm@umich.edu).

Digital Object Identifier 10.1109/TVLSI.2007.915477

1063-8210/$25.00 © 2008 IEEE

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

supply voltage available to a DVS algorithm such that the required reliability constraint is met while delivered processing performance during peak demand is maximized. DRM has the added potential benefit of theoretically allowing a user to select a desired lifetime or degradation envelope. Furthermore, DRM provides designers and architects with the ability to control lifetime reliability independently for individual projects in a manner that is transparent to the manufacturing technology.

The concept of DRM was first introduced in [3] using a sum of failure rates method while considering multiple reliability mechanisms. Lu *et al.* [4] analyzed electro-migration (EM) effects and suggested dynamic thermal management. However, both approaches focus on short time scales that are not indicative of realistic reliability requirements, and more critically, do not propose an actual control system required to obtain performance gains.

In this paper, we explore a DRM framework for digital logic using physics-based failure models for oxide breakdown, EM, thermal cycling, and negative bias temperature instability (NBTI) expressed as incremental damage mechanisms using a linear cumulative damage model referred to as Miner's rule [5]. The reversibility of NBTI damage is modeled and lifetime projection and the recovery effect on achievable system performance are explored. The performance impact of DRM in systems with DVS control techniques is analyzed with a focus on macro-level user collected processor usage profiles rather than traditional benchmark applications. The DRM system sets a maximum supply voltage based on the degradation characteristics modeled during operation and exceeds the nominal supply voltage when possible while meeting the required reliability constraints. With the proposed implementation of a specific control DRM algorithm, this work demonstrates and quantifies the potential performance improvements of DRM utilizing dynamic voltage scaling.

This paper is organized as follows. Section II describes the adapted reliability models used to estimate failure rates under variable conditions. Section III describes the projection of failure rate at a desired lifetime using inputs from the reliability models. The DRM control system that enables operation beyond nominal voltages is presented in Section IV, while the simulation setup and results discussion are presented in Section V. Section VI summarizes the simulation results and highlights the contributions of this paper.

II. Reliability Modeling

In order to implement a real-time dynamic reliability management scheme, we require accurate models that can comprehend dynamic stress behavior with minimal computational expense. High level compact models for oxide breakdown, EM and thermal cycling are addressed in the following sections. The models are adapted from state-of-the-art physics of failure work and applied to real-time DRM. In our proposed approach, we cast all reliability models such that they express wear-out in terms of an accumulated damage or fraction of lifetime consumed. This approach allows simple projections of the failure rate at the desired lifetime and is key to the efficient computation of total failure probability which drives the proposed DRM control method. It also allows the use of degradation dependent models for each reliability mechanism, a capability that is lost when dealing directly with probabilities.

A. Oxide Breakdown

Oxide breakdown, or dielectric breakdown, is a degradation mechanism that results in a low-impedance path through an insulating or dielectric barrier. During normal operation, each electron passing a dielectric barrier has a small probability to enter a high-energy state to tunnel through the insulating layer. Defect paths in the dielectric barrier reduce the energy level required for conduction through the layer, and therefore increase the probability that electrons will travel through the layer.

Each tunneling charge has a small probability of creating a defect when passing through the oxide. This probability of defect generation is the wear-out mechanism for thin dielectric films. When a critical defect density is reached, there is a high probability that a low-impedance defect path exists in the oxide and a runaway current path through the insulating film will develop. The exact microstructure and nature of the defects is not well understood and less than 1% of defect paths ultimately lead to an uncontrolled current path and oxide breakdown. The relationship between charge tunneling through the oxide and the defect density is expressed in (1), where N_{BD} is the defect density, P_{DG} is the probability of defect generation, and I_{tunnel} is the tunneling current, V is the voltage across the oxide, and T is the temperature [6]

$$N_{BD} \approx \int^{t} P_{DG}(V,T) I_{\text{tunnel}}(V,T) dt. \tag{1}$$

A simple simulation methodology for estimating the critical defect density required for a low-impedance defect path was originally developed by Degraeve [7] using a percolation concept. The percolation model places defects of a certain size into a 3-D oxide volume until a path of overlapping defects is created between the top and bottom planes. By running this simulation repeatedly for a given dielectric thickness, one can obtain a probability density function modeling the probability of a defect path related to the defect density. From this PDF, the approximate reliability of a thin-film dielectric is determined. The PDF generated by Monte Carlo simulation of the percolation model is fit to a Weibull distribution and used to calculate the probability of entering the onset of defect-induced oxide breakdown for an individual device.

The tunneling current through a gate oxide is calculated using BSIM4 model equations, yet alternative methods could be employed. The BSIM4 model for gate oxide leakage is well-suited for this calculation due to readily available parameters for most processes and the significant validation efforts to ensure accuracy. The probability of defect generation is a technology-specific term with a increasing exponential trend with increasing supply voltage and an Arrenhius temperature relationship. In this paper, published defect generation relationships from an IBM technology node are used in the simulations [6]. This oxide breakdown model allows an incremental summation of defect density at variable supply voltage and temperature stress conditions. This closed-form, high-level oxide breakdown model

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

is therefore ideal for a real-time DRM system considering dynamic stress conditions.

B. EM

EM is a failure mechanism caused by the movement of metal atoms through wires, creating voids (vacancies) and hillocks (deposits) that force open and short circuits in the surrounding wire networks. The transport phenomenon is primarily caused by electrical current, temperature gradients, and diffusion processes in the conductors. Black's formula [8] is a well-known relationship between the mean time to failure of an interconnect and the current density, temperature, and physical dimensions of the wire as shown in (2)

$$\mathrm{MTF} = AJ^{-n}\exp(E_a/kT). \tag{2}$$

The term A is a constant related to the materials and the geometric structure of the wire and it generally increases with both width and thickness of the structure. J is the current density, E_a is activation energy for atom transport, k is the Boltzmann constant, and T is temperature. The value of n is also a constant that depends on the criterion for EM failure and the treatment of wire-self heating. Typical values of n lie in the range of 1.0 to 2.0 when wire self-heating is considered, although larger values may fit data more accurately when self-heating is not considered. When the criterion for failure is related to a critical void size, a value of n close to 1.0 is used, whereas when considering a critical value of stress, a value of 2.0 is commonly used. The results in this paper are presented using a value of $n = 2.0$.

Ideally a model for a DRM should be expressed as a wear-out mechanism, with a quantifiable stress or damage term that is summed over time. Black's formula is instead expressed as a lifetime estimate based upon a single current density and temperature. We, therefore, use Miner's rule [5] (linear cumulative damage) to estimate the EM lifetime of a conductor by adding the percentage of lifetime consumed during each period of varying stress

$$\sigma_{\mathrm{life}} = \sum_t \frac{\mathrm{MTF}_{\mathrm{ref}}}{\mathrm{MTF}(J,T)}\Delta t. \tag{3}$$

Equation (3) summarizes the adaptation to Black's formula that allows variable stress conditions to be expressed as a percentage of lifetime, σ_{life}. $\mathrm{MTF}_{\mathrm{ref}}$ is a reference value that would be characterized at worst case conditions for the design and $\mathrm{MTF}(J,T)$ is an MTF calculation that is performed with varying current density and temperature averaged over a time window Δt.

A Weibull distribution is again used to convert the percentage of lifetime figure (σ_{life}) to a probability of failure. Due to the scarcity of published distributions of failure relating to EM, the specific parameters are not available and a Weibull curve similar to the oxide breakdown curve is used. This Weibull distribution would need to be characterized for the specific process and geometric structures in the interconnect stack to provide sufficient accuracy for DRM. A self-consistent temperature is calculated for wires in each layer considering the thermal effects of wire resistance and the current density at a given supply voltage [9]. Equation (4) relates the resistance of the wire to the temperature (T_{wire}), thickness of the wire (t), and resistivity of the material (ρ_0 at T_0)

$$R_{\mathrm{wire}} = \frac{\rho_0}{t}\left[1 + \alpha(T_{\mathrm{wire}} - T_0)\right] \tag{4}$$

$$T_{\mathrm{wire}} = T_{\mathrm{sub}} + R_{\mathrm{thermal}}P_{\mathrm{wire}}. \tag{5}$$

Equation (5) demonstrates that the temperature of the wire is a function of the power, which is a function of the resistance R_{wire}. The thermal resistance, R_{thermal} depends upon the layer of the interconnect stack and increases for upper levels of metal which are typically dominated by power and ground wires. The EM modeling in this paper is limited to unidirectional currents (power/ground network) due to the greatly reduced experimental observation of failures in wires with bidirectional current [10]. However, the analysis could be extended to include bidirectional current carrying interconnects as well.

C. Thermal Cycling

Thermal cycling related failures are a growing concern in microelectronic devices as continued scaling has led to rising power densities and temperatures. Systems with power saving techniques (such as DVS or sleep modes) exacerbate the incidence of thermal cycling by modulating the power consumption, and therefore temperature, at a much greater frequency than in a conventional system.

Thermal cycling is a mechanical stress mechanism that is manifested in many locations on an integrated circuit including solder connections and thin-film interfaces. As the temperature of the component materials on a chip changes, the components will expand and contract at differing rates, since most materials will have different thermal coefficients of expansion. The intermolecular bonds in materials will actually change length as they store increased amounts of energy, leading to a change in volume for a component. These changes in volume over time can eventually create adhesion problems between layers, or potentially create shorts or opens in extreme cases. Reducing the number of thermal cycles a system undergoes will decrease the rate at which such mechanical stress abnormalities are observed.

Blish [11] related the number of cycles of thermal fatigue of various materials on a silicon die to the thermal swing via the well-known Coffin–Manson equation [12]

$$N_{\mathrm{cyc}} \approx (\Delta T)^{-m} \tag{6}$$

$$N_{\mathrm{cyc}} = \sum_t \frac{(\Delta T_{\mathrm{ref}})^{-m}}{(\Delta T)^{-m}}. \tag{7}$$

The number of cycles N_{cyc} before breakdown is related to the thermal swing ΔT and a coefficient depending upon the materials involved. In this paper, we consider thin-film cracking damage with a Coffin–Manson exponent of 8.4 [11] and again use Miner's rule to express (6) as a percentage of lifetime. Equation (7) formalizes the use of Miner's rule with the Coffin–Manson equation.

In (7), ΔT_{ref} is a reference thermal swing that the system is designed to withstand and ΔT represents a measured thermal swing that may differ from the reference. Temperature traces are monitored in real-time to detect thermal swings. Equation (7) is used to sum the damage caused by thermal swings to express their contribution in terms of equivalent cycles of a larger swing

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

ΔT_{ref}. Due to the lack of availability of information (related to the difficulty in isolating this effect) regarding the probabilistic behavior of thermal cycling related wearout, some assumptions are made regarding the distribution of failures. A weibull distribution centered at 10 000 cycles of max thermal stress is used to approximate the effect of thermal swings on system reliability. As empirical data becomes available, this approximation should be improved and validated.

D. NBTI

The NBTI leads to shifts in parameter values (V_{TH} and I_{DSAT}) in pMOS devices after extended periods of stress at negative voltages across the gate to channel region. The effect is caused by the dissociation of hydrogen atoms from Si-H bonds that are present near the interface of the dielectric oxide layer and the doped silicon channel region. NBTI is primarily observed at elevated temperatures when the device is biased in the inversion regime, when interaction with holes weakens the Si-H bonds. The dissociation of a hydrogen atom leaves a dangling Si+ bond that serves as an interface trap for free electrons in the surrounding area. The creation of interface states near the oxide-silicon surface leads to a reduction in the effective saturation current (I_{DSAT}) of pMOS devices and is often modeled as an increase in the threshold voltage (V_{TH}) of the MOSFET device.

A unique characteristic of the NBTI effect, is the recovery phenomenon that occurs when the electric field and temperature are relaxed. The dissociated hydrogen atoms return to the oxide interface and anneal many of the interface states that were created during the period of stress, leading to a partial recovery in the saturation current of the pMOS device. Although the exact mechanisms governing the recovery effect are under debate in various physics journals, data collected from several test chips [13] indicates a strong link between the temperature of the device during stress and the extent to which recovery is possible. The model used in the proposed DRM system attempts to capture the dynamics and qualitative nuance of the stress phase and the temperature-limited recovery phase.

The reaction-diffusion (R-D) model [14] is the standard chemistry model for the reaction that governs interface state (N_{it}) creation and annealing. The R-D model is detailed as follows in (8) and (9):

$$\frac{dN_{\mathrm{it}}}{dt} = k_f[N_o - N_{\mathrm{it}}] - k_r N_{\mathrm{it}} N_H^o \tag{8}$$

$$\frac{dN_{\mathrm{it}}}{dt} = D - \frac{dN_H}{dx}|_{x=0} + \frac{\delta}{2}\frac{dN_H}{dt}. \tag{9}$$

Equation (8) represents the reaction-rate component of the R-D model where the forward reaction constant k_f depends on the initial Si-H bond density N_o and the current density of interface states, N_{it}. The reverse reaction is governed by k_r, the reverse reaction rate N_{it}, and the concentration of hydrogen at the interface N_H^o. The diffusion equation models the outflow of hydrogen from the interface and the inflow of hydrogen across the oxide interface of dimension δ. The reaction rate equation dominates initial creation of interface traps, leading to a rapid increase in NBTI damage when stress is applied. After the initial period of damage, the diffusion component of the reaction becomes the limiting factor and subsequent generation of interface traps slows considerably.

The R-D model matches measured circuit degradation well, yet is unsuitable for use in a DRM system due to the computational requirements of solving each iteration numerically. The following model, adapted from Cao [15] uses a piecewise function that is a numerical solution of the R-D model for a hypothetical stress phase and recovery phase. A piecewise function is an excellent candidate for use in a traditional integrated circuit (IC) with a static power supply, or even with a sleep mode IC with well defined stress and recovery phases. In a DRM system, using DVS to actuate the reliability mechanisms, it is difficult to define stress and recovery phases

$$N_{\mathrm{it}} = \sqrt{K_v^2 t_s^{2n} + N_{\mathrm{it0}}^2} \tag{10}$$

$$K_v = A t_{\mathrm{ox}} \sqrt{C_{\mathrm{ox}}(V_{\mathrm{gs}} - V_{\mathrm{th}})}\left[1 - \frac{\alpha V_{\mathrm{ds}}}{(V_{\mathrm{gs}} - V_{\mathrm{th}})}\right] \times e^{(E_{\mathrm{ox}}/E_0)} e^{(-E_a/kT)}. \tag{11}$$

Equations (10) and (11) model the accumulation of interface traps based upon the voltage and temperature stress and the period of time the circuit has been stresses. K_v is the stress factor, t_s is the time under stress, n is a technology dependent factor that is usually around 0.25–0.3 and N_{it0} is the initial concentration of interface traps when the stress period was initiated. The stress factor K_v is a strong function of V_{gs}, and the related E_{ox}, and follows an arrenhius relationship with temperature. If V_{ds} does not equal zero, the NBTI stress is reduced at the drain or source end and results in a lower incidence of hydrogen dissociation. The t^{2n} term in the stress phase equation captures the initial rapid increase in interface traps and the transition to a diffusion limited reaction where trap generation slows

$$N_{\mathrm{it}} = (N_{\mathrm{it0}} - \delta_v)\left[1 - \sqrt{n(t_r - t_0)/t_r}\right]. \tag{12}$$

Equation (12) is the model used by Cao for NBTI recovery when the electric field across the oxide is removed and hydrogen atoms have a probability of annealing the interface traps contributing to NBTI degradation. N_{it0} is the initial interface trap concentration at time t_0, n is a technology dependent variable around 0.35, and t_r is the recovery time elapsed since t_0. This piecewise function that alternates between stress and recovery has some significant drawbacks that required resolution for use in the proposed DRM system.

The stress equation cannot handle varying voltage and temperature due to the reliance on a fixed time component t_s in the formulation. For example, if voltage increases slightly, the term K_v will see an increase and t_s will remain the same, leading to a discontinuity in the calculated interface trap concentration N_{it} that is not present in measured data in the literature. To utilize this stress equation in a DVS system with frequently changing voltage and temperature stress, each time K_v changes, the time t_s must be recalculated. The recalculated t_s is an "effective time" given the previous stress, to maintain continuity in the

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

trap generation curve and correctly model the future trend in trap generation

$$t_{s,\text{eff}} = \left(\frac{N_{\text{it}}^2}{K_{v,\text{new}}^2}\right)^{-2n}. \tag{13}$$

Equation (13) is the simple method used to calculate the "effective" time at stress $K_{v,\text{new}}$ that is required to reach the current interface trap density. This value of t_s is used with the NBTI stress equation until another change in the stress value of K_v is encountered or the model transitions to the recovery function.

Consider the following scenario. NBTI stress begins and the initial N_{it} curve follows the $t^{2\text{n}}$ trend, a brief period of stress relaxation occurs, then the stress is reapplied. In this circumstance, there should be a minor period of recovery, but the overall interface trap concentration should follow a trend line similar to that of an uninterrupted period of stress, as the simulation continues. If the $N_{\text{it}0}$ term is used in (10) following the recovery period and t_s is reset to 0, the final trend line for interface trap generation will greatly exceed that of an uninterrupted period of stress. Intuitively and chemically, there is no evidence that this behavior should occur. Using the "effective" time calculation from (13) and ignoring the $N_{\text{it}0}$ term in (10) will prevent this abnormality that comes up quite frequently in a DVS system modeled by these equations, as seen in (14). In (14), time t_0 is the time that the stress factor $K_{v,\text{new}}$ was recalculated and $N_{\text{it}0}$ was the accumulated interface trap density at that time

$$N_{\text{it}} = \sqrt{K_{v,\text{new}}^2\left[\left(\frac{N_{\text{it}0}^2}{K_{v,\text{new}}^2}\right)^{-2n} + (t - t_o)\right]^{2n}}. \tag{14}$$

Another simple modification to the recovery model is to account for the limited annealing following stress at high temperature. There are no known models for this damage "lock-in" effect of high temperature, but adding a generic function will allow more sophisticated models of this effect to be added as they are verified and come into use. A simple model of the recovery is shown in (15)

$$N_{\text{it}} = N_{\text{it}0}\left(1 - \psi(V_{\text{dd}}, T, N_{\text{it,max}})\sqrt{\frac{n(t_r - t_o)}{t_r}}\right)$$

let

$$\psi(V_{\text{dd}}, T, N_{\text{it,max}}) \approx 0.4. \tag{15}$$

The ψ function is estimated to be around 0.4 for circuits with stress temperatures in the 85 °C–100 °C range from published data in [16]. This effectively represents a limitation to a recovery of 60% of the interface traps that were created during the stress period preceding.

The final challenge in implementing the piecewise function modeling NBTI degradation and annealing in a DVS system, is defining what a recovery phase is and when is the circuit in a stress phase. Particularly important is the transition from recovery to stress phases following a slight reduction in stress. When the system is binary, either at maximum voltage or in sleep mode, the definition is simple. However, the task of defining the NBTI degradation when the system experiences

```
Calculate K_v
If (K_v >> K_v,prev){
    CASE 1: Increasing Stress
    Calculate T_eff, stress
    Calculate N_it, N_it,perm
}
Else If (K_v << K_v,prev) {
    CASE 3: Initiate Recovery
    Save T_current
    Calculate N_it (N_it,perm influences recovery)
    Calculate T_eff,stress
}
Else {
    CASE 2: Continuation
    If (Recovery){  // Case 2.3
            Calculate N_it,recov (using recovery model)
            Calculate N_it,stress (using T_eff,stress)
            If( N_it,stress >= N_it,recov ) { Enter Damage Mode }
            Else{ N_it = N_it,recov }
    }
    Else{ // Case 2.1
            Calculate N_it, N_it,perm
    }
}
```

Fig. 2. Piecewise NBTI modeling function.

a 300 mV reduction in supply voltage or a 40 °C reduction in temperature is addressed in Fig. 2. When the NBTI model is in the recovery state and the stress is unchanged, both an updated interface traps value using the recovery phase model and the stress phase model are calculated using the current conditions. This models the annealing effect of recovery and the simultaneous creation of new traps, allowing a shift to the stress phase of the model when the $N_{\text{it,stress}}$ value exceeds the $N_{\text{it,recov}}$ value. These modifications to Cao's NBTI model allow it to be used effectively within a DVS-DRM framework.

One final step to projecting NBTI damage over a varying workload is needed, since the computation overhead for simulating yearlong traces of behavior is prohibitive. The previous reliability mechanism models, especially oxide breakdown and thermal cycling are particularly amenable to projection, since the defect density and thermal cycles can be directly summed to create a lifetime estimate. This allows shorter application traces to be characterized with the detailed models and a lifetime trace to be constructed by superposing the results from the short traces to construct a longer trace. With the N_{it} value in the NBTI model, the generation is nonlinear and a simple summation does not suffice for projecting shorter traces to a lifetime value. The approximation used in this paper is translating the N_{it} value from a trace to an "effective" time, in (13), of stress considering nominal voltage and temperature, and summing the damage in the time domain to create a reasonable estimate of lifetime degradation due to the NBTI mechanism.

III. System Level Modeling

This section presents an efficient approach to calculating system-level probability of failure that can be tailored to the desired level of detail. In this paper, a single failure due to

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

any reliability mechanism for any structure on the chip is a sufficient condition to declare that the chip has failed. In reality, individual dielectric breakdown events or EM voiding effects may not induce total system failure and certain components in a design (i.e., memory) may have built-in redundancy. However, this assumption will not significantly alter the conclusions reached.

The models outlined in Section II calculate actual probabilities of failure for individual oxides and wires, considering the historical stress pattern of voltage and temperature in a dynamic system. Since the primary parameters that cause correlation between these failure mechanisms (voltage, temperature) are directly used in the calculation of accumulated stress over the simulated time span, the correlation between these failure mechanisms is naturally considered. This allows system-level probability calculations using probabilistic independence and greatly simplifies the mathematical formulation, as now described.

In order to derive the total projected system failure probability at the end of lifetime t_{life}, we perform two tasks: 1) based on the existing stress history and accumulated damage for a particular failure mechanism and device at the current time t_1, we project the probability of failure for that failure mechanism and device at t_{life}; and 2) we combine the failure probabilities for all considered failure mechanisms and devices. We discuss each step in the following.

For each mechanism discussed in Section II, some concept of the degradation over time is maintained, typically expressed as a damage value. For oxide breakdown, the relevant metric is an estimate of defect density in a typical oxide layer, and NBTI damage is tabulated as the density of interface traps near the oxide-silicon interface N_{it}. Thermal cycling damage is counted in cycles normalized to the maximum expected thermal swing, and EM is defined as the projected MTF due to average current and stress conditions for lack of a suitable damage variable.

In the system, the damage (D_1) at time t_1 is extrapolated to the damage (D_{life}) at time t_{life} based on history information about the rate of damage up to time t_1 using the following simple linear extrapolation:

$$D_{\text{life}} = D_1 \frac{t_{\text{life}}}{t_1}. \tag{16}$$

Equation (16) accounts for environmental conditions and workload history intrinsically, providing a lifetime projection that is tailored to the exact stress conditions historically experienced on the chip. The model implicitly assumes that the future is similar to the past. However, we show in Section V that, even under use profiles that display significant shifts over time, the proposed DRM algorithm provides stable control. Given the projected accumulated damage at t_{life}, the probability of failure of an individual structure is then calculated using a cumulative distribution function for the relevant reliability mechanism characterized for the given process technology.

Equation (16) is useful for scaling the damage done in the near-linear damage model equations, such as oxide breakdown, thermal cycling and the basic EM model described in Section II, but it is inaccurate for a nonlinear mechanism like NBTI. To extrapolate the NBTI damage, the relationship between time and damage is implemented from the actual NBTI model itself to provide decent projections of actual NBTI damage

$$D_{\text{life,NBTI}} = D_1 \left(\frac{t_{\text{life}}}{t_1} \right)^{0.25}. \tag{17}$$

Equation (17) is very similar to (16), where the damage, D_1, at time t_1 is scaled to a lifetime damage prediction at t_{life} using the $(t_{\text{life}}/t_1)^{0.25}$ relationship from the NBTI model presented in Section II.

Individual device reliability projections are used to compute a chip-level reliability projection across all devices and all failure mechanisms using the following expression:

$$\begin{aligned}
(1 - P_{\text{ox}}) &= \prod_{b} \prod_{d} \left(1 - P'_{\text{ox}-b}\right) \\
(1 - P_{\text{EM}}) &= \prod_{b} \prod_{l} \prod_{n} \left(1 - P'_{\text{EM}-b-l}\right) \\
(1 - P_{\text{cyc}}) &= \prod_{b} \left(1 - P'_{\text{cyc}-b}\right) \\
(1 - P_{\text{NBTI}}) &= \min \left(1 - P'_{\text{NBTI}-b}\right).
\end{aligned} \tag{18}$$

P_{ox} is the probability of oxide failure, P_{EM} is the probability of an EM failure, and P_{cyc} is the probability of a thermal cycling failure. Oxide breakdown failure probability is calculated based on the number of devices per functional unit (decibels) with a specific failure rate for a device from each individual functional unit ($P'_{\text{ox}-b}$). Electro-migration failure rate is projected from the individual failure rate ($P'_{\text{EM}-b-l}$) across the number of wires (n), in each layer (l) and in each block (b). Thermal cycling failure is calculated as a component from each functional unit, since separate blocks undergo vastly different temperature traces. In future work, thermal cycling will consider the temperature gradients between functional units for this projection. NBTI failure probability is calculated differently from the other mechanisms, since the nature of an NBTI failure is much different (circuit timing failure versus fundamental device/material failure). NBTI reduction in saturation current of pMOS devices is tracked at the block level and the minimum probability of correct operation for any block is used to represent the NBTI failure contribution, PNBTI. The total chip failure rate is estimated by using the contributions of each failure mechanism in (19)

$$P_{\text{failure}} = 1 - ((1 - P_{\text{NBTI}})(1 - P_{\text{ox}})(1 - P_{\text{EM}})(1 - P_{\text{cyc}})). \tag{19}$$

Equation (19) is the combination of the failure rates due to each individual mechanism contributing to overall failure, P_{failure}. The simplicity of the chip failure rate calculation allows it to be used directly to drive a DRM control algorithm, which is described in Section IV.

IV. DRM System

Dynamic reliability management is implemented in this work using dynamic voltage scaling, which selects clock frequency and supply voltage pairs based upon workload demand and reliability model feedback. The scope of the DRM in this paper includes digital logic blocks degrading from EM, NBTI, oxide

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

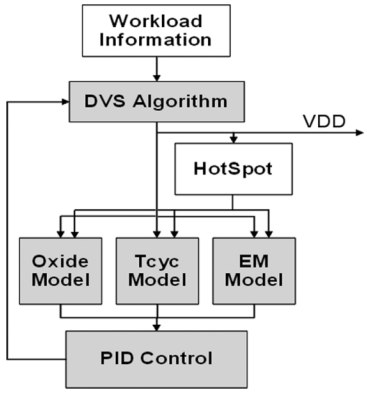

Fig. 3. DRM system block diagram.

breakdown, and thermal cycling. The framework can be extended to include other wear-out mechanisms, such as channel hot carrier effects, or to consider the impact on analog, input/output (I/O) circuits, or even packaging degradation with a similar modeling approach. For this paper, DVS is an ideal control scheme for managing reliability concerns, since oxide breakdown and electromigration are both strongly voltage dependent and reductions in supply voltage greatly reduce the effect of these wear-out mechanisms. Thermal cycling is intuitively exacerbated by the increased variation in power consumption in a DVS system, yet thermal swings can be indirectly limited by capping the absolute maximum supply voltage, which limits the maximum temperature. If necessary, it is possible to further address reliability degradation due to temperature cycling by limiting the rate of voltage change in the DVS algorithm. In our analysis, however, this was found to be unnecessary.

Fig. 3 details the organization of the DRM system that is implemented for maximizing the peak performance of a microprocessor system. Processor utilization traces are used to generate voltage/frequency traces for the DVS microprocessor. The selection of a voltage/frequency pair is converted to a block-based power consumption value that is derived from Wattch [17] application traces. Thermal information for each block is calculated using HotSpot [18] in the simulation flow, but can be replaced by a thermal sensor in an actual silicon implementation. The combination of voltage and thermal information is supplied to the four reliability mechanism models described in Section II and the output of each model is combined to generate the probability of chip failure at the desired lifetime as described in Section III. Chip failure probability information is then used in a proportional–integral–derivative (PID) control algorithm to set a maximum allowable voltage used in the DVS voltage assignment step.

A PID-based control algorithm is proposed as the key mechanism to provide maximum improvement in peak circuit performance when necessary, without affecting steady-state performance or comprising reliability.

Equation (20) describes the behavior of a PID control system, where $e(t)$ is an error signal and $v(t)$ is the output being controlled

$$v(t) = v(t - t_0) + P\left[e(t) + R\int e(t)dt + D\frac{\partial e(t)}{\partial t}\right]. \quad (20)$$

P is the proportional gain, R is the reset or integral gain, and D is the derivative gain. In general, proportional gain controls the response time of the controller, integral gain corrects for offset, and the derivative gain limits overshoot in the error term. In the proposed DRM system, $e(t)$ is the probability of system failure projected to the lifetime, t_{life}, and $v(t)$ is the maximum voltage available to the DVS algorithm.

The DRM system described is a discrete, nonlinear, time-varying control system. Most of the reliability models are inherently nonlinear with stress input and any models with a time-dependence or recovery mechanism (NBTI) are also time-varying, preventing a straightforward expression of the transfer function of the system. Therefore, it is unfeasible to present a general proof of system stability under all conditions. The results section presents some evidence of system stability under dramatic workload shifts using the impulse response of the system. An alternative to a closed-form proof of system stability could be achieved by fitting a mathematical model of the system response for a given implementation to collected data. Given this model, and the general PID equation, a proof of stability should be possible. Since the system is discrete in reality, (21) reflects the modifications made to the theoretical model of PID control in (20)

$$v(t)=v(t-1)+P\left[e(t)+R\sum_{-\infty}^{t}e(k)+D\left(e(t)-e(t-1)\right)\right]. \quad (21)$$

Tuning of the control algorithm is dependent upon the desired response time and the length of time to correct offset issues. Overshoot, or selecting a maximum voltage that is too large, leads to a number of negative side-effects in a DVS system. To compensate for the excessive amount of wear-out damage, the algorithm will reduce the clock frequency below nominal which could limit performance in subsequent time periods. In severe cases of poorly set proportional gain, oscillation between high and low voltages is observed in a classic case of an unstable feedback loop. The integral gain plays a very small and inconsequential role in this system and the control system could be reduced to a P-D system without much impact on the selected voltages. Setting a relatively high derivative gain (on the order of proportional gain) delivers near-optimal control performance in the proposed DRM system by allowing a decent response time to processing demand requests, yet minimizing overshoot and undershoot when correcting the voltage setting.

In order to evaluate the effectiveness of the PID control system in achieving gains in processor frequency by using the available "reliability slack," the figure of merit, peak performance improvement (PPI) is defined. PPI is a measure of the relative improvement in processor frequency possible when the system is operating at its peak demand. Essentially, it is a measure of how far the frequency can be "overclocked" during peak usage to deliver critical results. PPI is not a measure of overall system speedup, since it does not include any information regarding the proportion of total calculation time the processor may spend at the elevated voltage/frequency pairing, however, applications in parallel processing systems can be limited by the peak performance of a critical thread. Many parallel algorithms require global or semi-global reductions or calculations

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

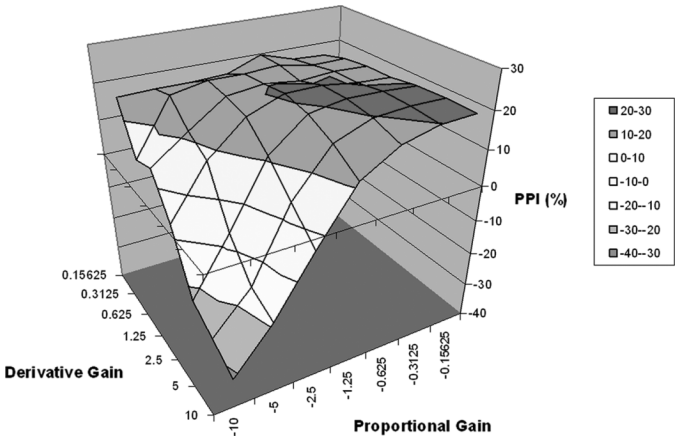

Fig. 4. PID control gain versus peak performance improvement.

that often function in a sequential fashion that benefits greatly from the potential increased single processor performance offered by the proposed system. Synchronization in parallel applications is another bottleneck that can potentially benefit from elevated peak performance to minimize the contention for data. Any system that involves user input and interaction can also potentially benefit from the peak performance gains available from the DRM-DVS concept. Although the PPI figure is not a measure of total system performance, it is certainly a metric of interest in many applications.

Fig. 4 shows the improvement in peak performance (PPI) using the PID control system with varying values of proportional gain and derivative gain. Tuning the control system is somewhat dependent upon the variation in processor demand seen in the workload traces. In this plot, the average PPI is plotted for each set of values over a compilation of varying workload traces. An explanation of the workload trace collection is included in Section V.

An actual implementation of the DRM system could take several forms, from purely hardware driven to purely software controlled. Several factors point in the direction of a predominantly software-oriented approach with limited hardware support. 1) Calculations to update the projected lifetime of the system and update the maximum assignable DVS voltage are needed infrequently, as the timescale of degradation is significantly greater than the timescale of computation. 2) Updating the models for a given system after development will be significantly easier for a system designed in software. Hardware components may consist of distributed sensors and a communication network, or temporary memory-mapped storage registers to maintain information on voltage and temperature history or the output of the sensors. Given infrequent updates, the overhead of the DRM system should be minimal in terms of performance when implemented in software, particularly when run on systems with significant "sleep" time. Area overhead should also be minimal for a software controlled system, with reasonable amounts of sensors. Assuming flexibility in placement of the sensors (placement in available whitespace), and routing, no more than 1%–2% area overhead should be incurred. The remaining factor is the overhead of the DVS system, which has been implemented in an existing industrial processor [19] and shown to have no significant area overhead.

TABLE I
SIMULATION TECHNOLOGY SPECIFICATION AND SELECTED MODEL PARAMETERS OF INTEREST

Symbol	Quantity	Value
L_{drawn}	channel length	130 nm
V_{th0}	device threshold voltage	250 mV
VDD_{nom}	nominal supply voltage	1.2 V
T_{ox}	oxide thickness	1.8 nm
W_L	wire width (local)	140 nm
T_L	wire thickness (local)	350 nm
W_G	wire width (global)	450 nm
T_G	wire thickness (global)	1200 nm
ρ_0	wire resistivity (em model)	1.68×10^{-8} ohm-m
T_0	wire resistivity reference temperature	293.15 K
K_{OX}	wire thermal conductivity	0.25 W/K-m
n_{EM}	technology constant (em model)	2.00
A_{TFC}	thermal cycling constant	5.05×10^{21}
α_{TFC}	thermal cycling constant	-0.33
m_{TFC}	thermal cycling constant	8.4
n_{NBTI}	NBTI constant (stress/recovery)	0.25 / 0.35
A_{NBTI}	NBTI constant	1.80×10^{7}

V. RESULTS AND DISCUSSION

Workload data from several desktop computers was collected over several months to provide realistic processor utilization information with a wide-range of system behavior. A processor layout similar to the Alpha 21264 is used with process parameters based on 130-nm industrial models, which are summarized in Table I. The hypothetical processor is divided into 15 sub-blocks representing individual functional units on the chip (i.e., arithmetic logic unit (ALU), memory, decode unit). Initial power estimates are generated by Wattch and used with the processor utilization data collected to generate workload-based voltage, frequency, and power traces. The PID controller assigns voltage-frequency pairs based upon the requested performance and the reliability state of the system. HotSpot 2.0 is used to calculate temperatures for each functional unit in the design using power numbers adjusted according to the selected supply voltage. The PID controller updates the maximum voltage every 50 μs in the short-time limit simulations presented and every hour in 10-year lifetime simulations.

The technology specification utilized for simulation is an aggressive 130-nm technology based on values from several industrial models, predictive technology models, the SIA roadmap and available figures from the literature on the relevant reliability mechanisms. The relative impact of each mechanism is a strong factor in the simulation results. Given the values used in this study, oxide breakdown was the dominant mechanism in the results presented for the work with oxide breakdown, EM in power wires and thermal cycling. The EM model had a moderate impact on results, with thermal cycling showing a minimal impact (when considering the on-chip component, thin film cracking). When considering the NBTI effect, it had a similar magnitude to the oxide breakdown mechanism, using aggressive figures for degradation. The NBTI effect in this paper scaled

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

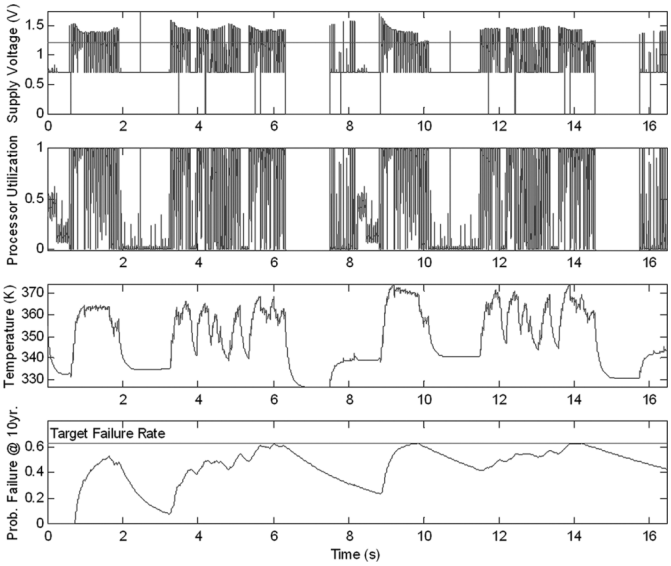

Fig. 5. DRM operation for workload C630 (16 s).

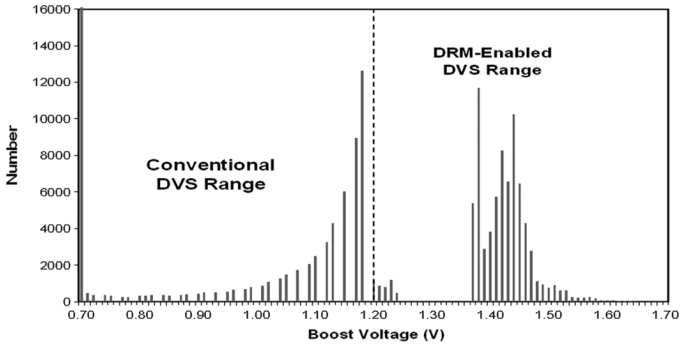

Fig. 6. Voltage histogram for workload in Fig. 5.

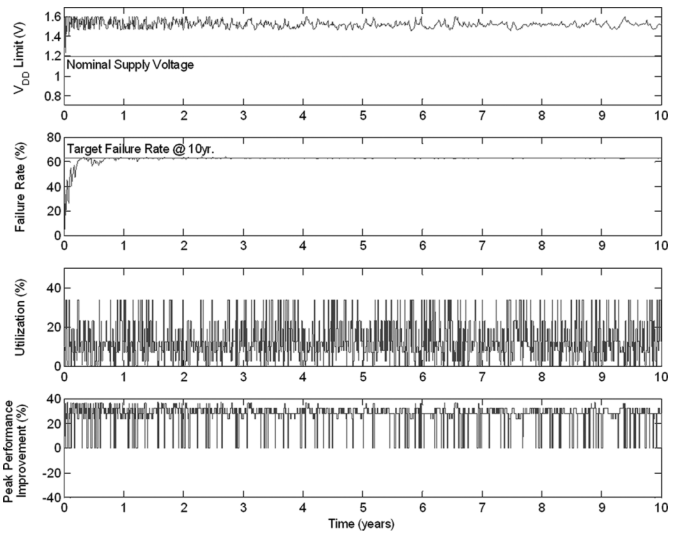

Fig. 7. Ten-year reliability simulation.

up to 15%–17% reduction in circuit speed, which is somewhat higher than figures from more recent published models, which cite 8%–9% expected degradation [20].

A. Results (Without NBTI Model)

A figure of merit to quantify the performance gains available with DRM that will be used throughout the results discussion is "peak performance improvement." This figure is a measure of the improvement in attainable frequency (%) during periods of *peak CPU demand*. This is a convenient measure of how well DRM provides additional performance when it is needed most by the user or application.

The traces in Fig. 5 are of a DRM simulation over a time span of only 16 s, allowing a detailed look at the interplay between the voltage assignment, workload, temperature, and the projected failure rate at the 10-year lifetime. The horizontal line across the supply voltage trace is the nominal supply voltage 1.2 V and the line across the failure rate curve is the target failure rate of 63.2% at 10 years. The plot clearly shows the increase in projected failure rate during periods of high supply voltage and temperature across this high activity profile. Longer simulations result in a much smoother failure rate projection curve, as the slope of the damage projection becomes more stable over time.

The histogram in Fig. 6 shows the frequency of different voltage assignments in the conventional voltage supply range and the boosted DRM voltage range. The distribution is for the Alpha 21264 system running the workload that was plotted in detail in Fig. 5. The plot is bimodal since all tasks that require peak performance are executed at the maximum supply voltage allowed by the PID controller. Although there is no voltage limit upon the system, Fig. 6 shows no data points beyond 1.7 V and the majority of the boost voltage usage occurs below 1.5 V. In an actual implementation of a DRM system, an upper bound on voltage could therefore be placed at 1.5 V to accommodate power distribution or voltage regulator limitations.

Fig. 7 displays the 10-year performance of the DRM control algorithm over a randomized selection of 10 representative 1-h workloads collected from an actual desktop machine. The randomization process selects a workload and then selects a random duration for that workload (ranging from 1 h to 2 weeks), which the 1-hr trace is then repeated to fill that duration. The $V_{\rm DD}$ limit graph in each section of the plot represents the upper limit placed upon the DVS algorithm by the DRM mechanism, not the actual voltage during the entire trace. The PID controller does an excellent job of maintaining the target error rate over the long lifetime simulation. Although the nature of the calculation of error rate prevents any straightforward analysis on the stability of the control algorithm, it is possible to provide evidence of stability with extreme inputs to the system.

Fig. 8 represents the response of the control algorithm to a sudden change in workload at the 6-year period, representing a pessimistic scenario for the proposed system. Tuning the controller involves a tradeoff between response time and stability. In Fig. 8, there is a very small undershoot on the voltage limit, demonstrating the ability of the control algorithm to respond to sudden changes in the workload. The response time of the system can be improved at the expense of the magnitude of the undershoot which could result in unnecessary performance throttling where the $V_{\rm dd}$ limit drop below the nominal supply voltage.

Peak performance gains over the 10-year workloads ranged from 20% to 35% compared to a nominal DVS controlled

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

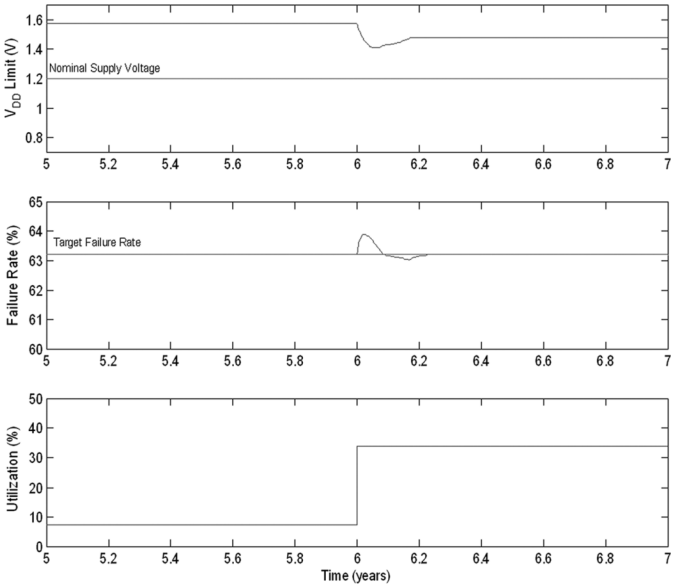

Fig. 8. PID controller impulse response.

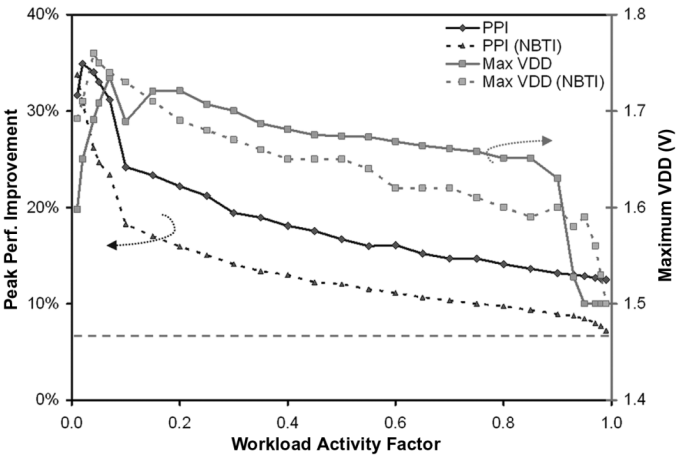

Fig. 9. Peak performance and max supply voltage versus workload activity considering NBTI.

system. Specifically, the profile in Fig. 7 shows a 26.7% peak performance improvement. In cases where the system is operating below the maximum operating temperature, significant overall performance improvements are possible. In Fig. 9, with a design rated to a maximum on-die temperature of 125 °C, an ambient temperature of 60 °C allows 12.5% overall performance improvement before considering the workload.

B. Results (With NBTI Model)

In Fig. 9, the peak performance improvement is plotted versus the workload activity factor. The workload profile for this plot is constructed with oscillations between 100%–0% utilization at a period of 5 min with a variable duty cycle that equals the activity factor. For extremely inactive systems, the voltage may be boosted dramatically above the nominal voltage, delivering a maximum of 34% peak performance improvement during periods of peak CPU demand. The sharp roll off in performance gains as activity factor is increased is related to the higher temperatures that are reached as the chip spends longer periods of

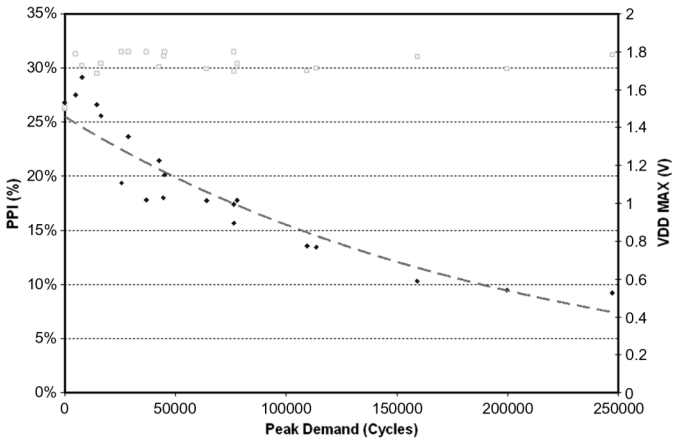

Fig. 10. Peak demand versus PPI (with NBTI).

time above nominal supply voltage. As the activity factor of the workload approaches 1.0, the performance gains are reduced to 12.5%. At this point, the benefits of DRM are derived not from periods of low voltage stress in the workload, but from a lower operating temperature than specified in the worst case reliability corner.

The simulated chip is designed to operate at nominal voltage (1.2 V) and a maximum on-die temperature of 125 °C for 10 years. In this simulation, with an ambient temperature of 60 °C, a 12.5% performance improvement is possible even with an extremely pessimistic workload. This plot demonstrates 15%–20% peak performance improvement for workload profiles below a 100% maximum performance. Collected usage profiles from actual desktop machines indicate workload activity factors between 0.10–0.15 are typical for user-driven systems for which peak performance improvement is approximately 25%. The overall performance improvement when considering NBTI is also in the range of 0.0%–2.6% depending upon the workload, and gains up to 7.5% were observed when operating significantly below the maximum on-die operating temperature.

Providing boosted supply voltages above nominal voltage may require a greater number of pads devoted to the power and ground network to handle the additional current associated with a higher clock frequency/voltage pairing. While the absolute cost in terms of area or packaging is difficult to quantify in a general sense, it is helpful to consider the effects on peak power consumption and lifetime energy consumption for systems with DRM implementations. Energy consumption over the lifetime was found to track the performance gains closely, and is relatively unaffected by a change in the maximum allowable voltage. Approximately 20.1% additional energy consumption is required to obtain peak performance gains of 20.68% over the lifetime of the chip under the workload of Fig. 5. Peak power increases were found to be somewhat larger but naturally tend to be short in duration due to the feedback from the PID controller.

In Fig. 10, the number of cycles with peak demand requested by the workload is linked to the attainable PPI on the left axis. Similar to the workload activity plot in Fig. 9, there is a strong dependence on peak demand and the PPI, however, this data is

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

collected from the workload traces from actual desktop profiles rather than the synthetic benchmark used to generate Fig. 9. The traces plotted here are 250 000 cycles in length, so data points on the left are near 0% workload activity and the right is near 100% workload activity. The maximum assigned voltage at each point is roughly constant at 1.8 V, yet this value is far above the typical value seen in the DVS trace. The high assignment of 1.8 V is largely an artifact of the PID controller error function beginning the trace uninitialized leading to an exaggerated voltage assignment during the first period of peak demand.

An exponential curve is fitted to the simulated data in Fig. 10, showing a rough exponential trend between 10%–30% PPI depending upon the number of samples of peak demand witnessed in the simulated workloads. With proper characterization, the PPI attainable for a given system may be predicted with decent accuracy *a priori*, allowing some potential high-level optimizations in a multiprocessing environment This is an idea that is currently being explored by our current work and future plans involving multiprocessing systems and DRM.

VI. Conclusion and Future Work

A framework for implementing dynamic reliability management is presented, including a rigorous model for failure rate prediction under four common failure mechanisms and a PID-based control system that balances increased throughput in peak-demand periods with the remaining reliability lifetime. Workload and processor utilization information collected over months for typical users is used to quantify achievable gains in peak processor performance. On-chip real-time reliability monitoring allows supply voltages to be boosted beyond nominal values set during worst case profiling and qualification, enabling maximum responsiveness during periods of critical computational demand. Despite minimal overall performance gains typically in the range of 0%–2.6%, we observe typical peak performance gains of 20%–35% over a variety of real-world workloads and lifetime usage profiles, without exceeding the specified lifetime budget. Considering the impact of NBTI reduces the achievable gains by 8%–10% in simulation, yet still allows a peak performance gain of 15%–25% over the typical range of workload activity of 0.05–0.20. The design of low-overhead on-chip sensors (temperature, tunneling current, etc.) is a primary focus in our future work in order to enable silicon implementations and validation of the modeling-based work.

References

[1] *Stress-Test-Driven Qualification of Integrated Circuits: JESD47E*, JEDEC Standard, JEDEC Solid State Technology Association, 2007.

[2] *Failure-Mechanism-Drive Reliability Qualification of Silicon Devices*, JEDEC Standard, JEDEC Solid State Technology Association, 1993.

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

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

[5] M. A. Miner, "Cumulative damage in fatigue," *J. Appl. Mechan.*, vol. 67, pp. A159–A164, 1945.

[6] J. H. Stathis, "Physical and predictive models of ultra thin oxide reliability in CMOS devices and circuits," *IEEE Trans. Device Mater. Reliab.*, vol. 1, no. 1, pp. 43–59, Mar. 2001.

[7] R. Degraeve *et al.*, "A consistent model for intrinsic breakdown in ultra-thin oxides," in *Proc. Int. Electron Devices Meet.*, Dec. 1995, pp. 863–866.

[8] J. R. Black, "Electromigration failure modes for aluminum metallization in semiconductor devices," *Proc. IEEE*, vol. 57, no. 9, pp. 1587–1593, Sep. 1969.

[9] M. R. Casu, M. Graziano, G. Masera, G. Piccinini, and M. Zamboni, "An electromigration and thermal model of power wires for a priori high-level reliability prediction," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 12, no. 4, pp. 349–358, Apr. 2004.

[10] A. Chandrakasan, Ed., *Design of High-Performance Microprocessor Circuits*. Piscataway, NJ: IEEE Press, 2001, pp. 437–438.

[11] R. C. Blish, II, "Thermal cycling and thermal shock failure rate modeling," in *Proc. 35th Int. Reliab. Phys. Symp.*, Apr. 1997, pp. 110–117.

[12] W. Nelson, *Accelerated Testing: Statistical Models, Test Plans and Data Analyses*. New York: Wiley, 1990.

[13] M. Ershov *et al.*, "Transient effects and characterization methodology of negative bias temperature instability in PMOS transistors," in *Proc. 41st Int. Reliab. Phys. Symp.*, Apr. 2003, pp. 606–607.

[14] M. A. Alam and S. Mahapatra, "A comprehensive model of PMOS NBTI degradation," *Microelectron. Reliab.*, vol. 45, pp. 71–81, 2005.

[15] R. Vattikonda, Y. K. Cao, and W. Wang, "Modeling and minimization of the PMOS NBTI effect for robust nanometer design," in *Proc. 43rd Des. Autom. Conf.*, Jul. 2006, pp. 1047–1052.

[16] S. Rangan, N. Mielke, and E. C. C. Yeh, "Universal recovery behavior of negative bias temperature instability," in *Dig. Int. Electron Devices Meet.*, Dec. 2003, pp. 14.3.1–14.3.4.

[17] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: A framework for architectural-level power analysis and optimizations," in *Proc. 27th Int. Symp. Comput. Arch.*, 2000, pp. 83–94.

[18] K. Skadrony, M. Stanz, M. Barcellaz, A. Dwarkaz, W. Huangz, Y. Liy, Y. Maz, A. Naiduy, D. Parikhy, P. Rez, G. Rosez, K. Sankaranarayanany, R. Suryanarayanz, S. Velusamyy, H. Zhangz, and Y. Zhangz, "HotSpot: Techniques for modeling thermal effects at the processor-architecture level," in *Proc. Int. Workshop THERMal Investigations ICs Syst.*, 2002.

[19] K. J. Nowka, G. D. Carpenter, E. W. MacDonald, H. C. Ngo, B. C. Brock, K. I. Ishii, T. Y. Nguyen, and J. L. Burns, "A 32-bit PowerPC system-on-a-chip with support for dynamic voltage scaling and dynamic frequency scaling," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, pp. 1441–1447, Nov. 2002.

[20] S. V. Kumar, C. H. Kim, and S. S. Sapatnekar, "An analytical model for negative bias temperature instability," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Des.*, Nov. 2006, pp. 493–496.

Eric Karl (S'02) received the B.S. degree in electrical engineering *summa cum laude* and the M.S. degree in electrical engineering from the University of Michigan, Ann Arbor, in 2002 and 2004, respectively, where he is currently pursuing the Ph.D. degree in electrical engineering.

During his Ph.D. dissertation work, he held research and development internship positions with Intel Research, IBM Research, Sun Microsystems, and General Motors. His current research interests include integrated circuit reliability and low-power circuit design and automation.

Mr. Karl was a recipient of the EECS Senior Scholar award for outstanding scholastic achievement during his undergraduate coursework.

David Blaauw (SM'07) received the B.S. degree in physics and computer science from Duke University, Durham, NC, in 1986, and the Ph.D. degree in computer science from the University of Illinois, Urbana, in 1991.

Since August 2001, he has been an Associate Professor with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor. Until August 2001, he worked with Motorola, Inc., Austin, TX, where he was the Manager of the High Performance Design Technology Group. His

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.

work has focused on VLSI design and CAD with particular emphasis on circuit design and optimization for high performance and low power applications. He was the Technical Program Chair and General Chair with the International Symposium on Low Power Electronic and Design and was the Technical Program Cochair and member of the Executive Committee the ACM/IEEE Design Automation Conference.

Mr. Blaauw is currently a member of the ISSCC Technical Program Committee.

Dennis Sylvester (S'95–M'00–SM'04) received the B.S. degree in electrical engineering *summa cum laude* from the University of Michigan, Ann Arbor, and the Ph.D. degree in electrical engineering from the University of California, Berkeley, in 1999.

He is currently an Associate Professor with the Department of Electrical Engineering and Computer Science, University of Michigan. He previously held research staff positions with the Advanced Technology Group of Synopsys, Mountain View, CA, Hewlett-Packard Laboratories, Palo Alto, CA, and a Visiting Professorship with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore.

Dr. Sylvester was a recipient of an NSF CAREER Award, the Beatrice Winner Award at ISSCC, an IBM Faculty Award, an SRC Inventor Recognition Award, several Best Paper Awards and nominations, the ACM SIGDA Outstanding New Faculty Award, and the University of Michigan Henry Russel Award for distinguished scholarship. His dissertation research was recognized with the David J. Sakrison Memorial Prize as the most outstanding research in the UC-Berkeley EECS department. He has served on the technical program committee of numerous design automation and circuit design conferences, the steering committee of the ACM/IEEE International Symposium on Physical Design, and was General Chair for the 2005 ACM/IEEE Workshop on Timing Issues in the Synthesis and Specification of Digital Systems (TAU). He is currently an Associate Editor for IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS and previously served as an Associate Editor for IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS. He is a member of ACM, American Society of Engineering Education, and Eta Kappa Nu.

Trevor Mudge (F'95) received the Ph.D. degree in computer science from the University of Illinois, Urbana.

Since then he has been with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor. He was named the Bredt Professor of Engineering after a ten year term as Director of the Advanced Computer Architecture Laboratory—a group of a dozen faculty and 80 graduate students. He is the author of numerous papers on computer architecture, programming languages, VLSI design, and computer vision. He has also chaired 33 theses in these areas.

Prof. Mudge is a member of the ACM, the IET, and the British Computer Society.

Authorized licensed use limited to: IEEE Xplore. Downloaded on November 14, 2008 at 14:23 from IEEE Xplore. Restrictions apply.