

Power, Thermal, and Reliability Modeling in Nanometer-Scale Microprocessors

David Brooks, Robert P. Dick, Russ Joseph, and Li Shang

Contact Author

Robert P. Dick

dickrp@eecs.northwestern.edu

847-467-2298

Fax: 847-467-4144

L477 EECS Dept.

2145 Sheridan Rd.

Northwestern University

Evanston, IL 60208

Power, Thermal, and Reliability Modeling in Nanometer-Scale Microprocessors

David Brooks, Robert P. Dick, Russ Joseph, and Li Shang

Abstract

Power is the source of the greatest problems facing microprocessor designers. High-power processors rapidly deplete battery energy. Rapid changes in power consumption result in on-chip voltage fluctuations that bring transient errors. High spatial and temporal power densities bring high temperatures, which result in decreased lifetime reliability. High temperatures also increase leakage power consumption, thereby closing a self-reinforcing power-temperature feedback loop. The effects of increasing power consumption, power variation, and power density are expensive and distasteful. The wages of power are bulky short-lived batteries, huge heatsinks, large on-die capacitors, high server electric bills, and unreliable microprocessors. The only alternative is optimizing microprocessor power consumption, temperature, and reliability. However, this depends on accurate modeling and rapid analysis of properties that span different disciplines and levels, from device physics, to numerical methods, to microarchitectural design. This article surveys the most important problems brought directly and indirectly by power consumption, indicates the relationships among them, explains recent methods of efficiently modeling them, and indicates promising directions in the ongoing fight against power.

Keywords: Power consumption, microarchitectural modeling, integrated circuit reliability

1. INTRODUCTION: POWER, TEMPERATURE, AND RELIABILITY

Increasing system integration and performance requirements are dramatically increasing the power consumptions and power densities of high-performance microprocessors. Some microprocessors now consume 100 W. High power consumption introduces challenges to various aspects of microprocessor and computer system design. It increases the cost of cooling and packaging design, reduces system reliability, complicates power supply circuitry design, and reduces battery lifetime. Intensive effort has recently been dedicated to address power-related design problems. Modeling is the essential first step towards design optimization. In this article, we will explain the power, thermal, and reliability modeling problems and survey recent advances in their accurate and efficient analysis.

Figure 1 illustrates the relationships among power consumption, temperature, reliability, and process variation. Shaded boxes indicate attributes that do not significantly influence other attributes in the figure. Dotted boxes explain the processes by which one attribute affects others. Dynamic and leakage power consumption are at the root of the problem. Rapid changes in dynamic power consumption, indicated by a solid box in Figure 1, can result in transient voltage fluctuations as a result of the distributed inductance and resistance in off-chip and on-chip microprocessor power delivery networks. These dI/dt effects, indicated by oval boxes in Figure 1, can change logic combinational path delays, resulting in timing violations; this results in transient faults or requires decreasing the processor frequency. High dynamic power consumption implies high current density, resulting in accelerated wear and permanent faults due to processes such as electromigration. Dynamic power consumption produces heat, which increases microprocessor temperature. The effects of leakage power consumption are similar, with the exception of dI/dt -related problems. In general, leakage power variation is smaller than dynamic power variation. Both dynamic and leakage power consumption decrease battery lifespans in portable devices and increase server electric bills.

Temperature is increased by both dynamic and leakage power. Temperature profiles depend on the temporal and spatial distribution of power and the cooling and packaging solutions used for the microprocessor. High temperature increases charge carrier concentrations, resulting in increased subthreshold leakage power consumption. In addition, it decreases charge carrier mobility, decreasing transistor and interconnect performance, and decreases threshold voltage, increasing transistor performance. Finally, temperature has tremendous impact on the fault processes that

This work was supported in part by the NSF under awards NSF CCF-0048313 and CNS-0347941, in part by Intel, in part by IBM, and in part by the NSERC under Discovery Grant #388694-01.

lead to the majority of microprocessor permanent faults. Modeling and optimizing microprocessor thermal properties is thus essential for reliability, power consumption, and performance.

Process variation has great influence on the other power-related integrated circuit (IC) characteristics explained in this article. It influences transient fault rate via changes to critical timing paths; permanent fault rate via changes to numerous parameters such as wire and oxide dimensions; leakage power consumption via changes in dopant concentration; and dynamic power consumption.

As Figure 1 illustrates, the relationships among power, temperature, and reliability are complex. Controlling power, temperature, and reliability requires modeling and optimization at multiple design stages as well as runtime support. Section 2 explains the power modeling problem. Section 3 defines the thermal analysis problem and describes how temperature profiles may be accurately and efficiently computed based on temporal and spatial power consumption profiles. Section 4 describes a rapid and accurate temperature-aware method of modeling leakage power consumption. Section 5 explains the challenges and methods of modeling the impact of power and temperature on permanent and transient faults within microprocessors. Section 6 explains the impact of process variation on power, performance, and reliability. We summarize and point out a number of open research problems in Section 7.

2. POWER MODELING

Pre-RTL architectural power modeling has gained prominence with the rising importance of energy-constrained portable devices and thermally-constrained high-end machines. Understanding the power characteristics and ramifications of early-stage design decisions is essential because architectural decisions can have a huge impact on overall power-efficiency, and serious mis-steps made at the architecture level are often too difficult to be corrected at later stages in a design project.

Architectural power simulators are often tightly coupled with existing cycle-level architectural performance models. Microarchitectural utilization information, including bit-level activity in some cases, is combined with energy models for the power consumption of microarchitectural blocks under various usage scenarios (e.g., power consumption with various amounts of activity after applying power or clock-gating techniques). Collecting utilization

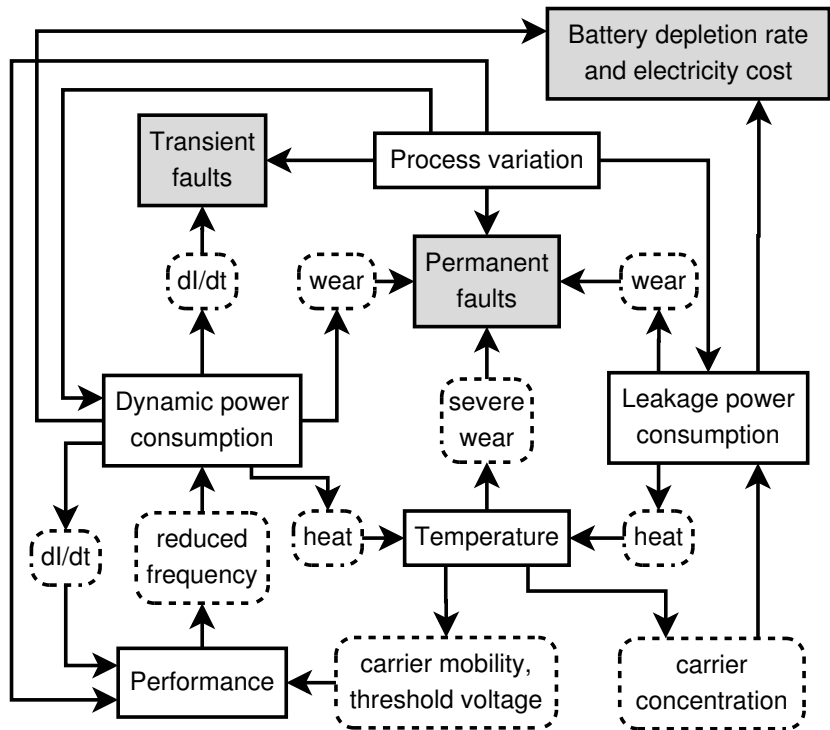


Figure 1. Power and its implications.

information is not particularly challenging, and the rest of this section will focus on the more difficult task of developing power models for microarchitectural blocks. Most existing efforts in this area can be broadly characterized into analytical and empirical techniques for both dynamic and leakage power consumption.

2.A. Analytical Power Models

Analytical models for dynamic power dissipation can be constructed by identifying the key internal nodes of a microarchitectural block and deriving analytical formulas for the capacitance of these nodes. This technique has been widely applied to RAM/CAM structures such as caches, register files, buffers, and queues [1]–[3], and regular combinatorial logic such as decoders, select trees, and arbiters. In contrast, the power of irregular combinatorial logic such as integer and floating-point ALUs is generally estimated using empirical models.

For example, most popular analytical models use Equations 1 and 2 to calculate SRAM array bit line energy, where $C_{bitline}$ is the total stacked bitline capacitance and V_{swing} is the voltage swing of the bitline. The total bitline capacitance includes N stacked access transistor diffusion capacitance, pre-charge circuit capacitance, bitline wire capacitance and, column select circuit capacitance.

$$E = C_{bitline} \times V_{DD} \times V_{swing} \quad (1)$$

$$C_{bitline} = C_{diff_access} \times N + C_{diff_precharge} + C_{diff_column} + C_{wire} \quad (2)$$

Analytical power models provide a fast, flexible method of estimating dynamic power consumption. However, the accuracy of these models suffers for several reasons:

- **Inaccurate Capacitance Estimates:** Node capacitance depends heavily on the actual design layout (including overlap and node-to-node coupling capacitances) and the supply voltage. Thus, it can be difficult to accurately estimate using simple analytical formulas.
- **Direct-Path Current:** Analytical models neglect direct-path current, because this is difficult to estimate as it depends heavily on the circuit design. However, direct-path current can be appreciable in the pre-charge circuitry of certain memory designs.
- **Memory Cell Toggle Power:** Many analytical models neglect the power to toggle memory cell bits when writing to the memory. Although this power is usually small, it can be a fairly significant error source when there are periods of intensive writes and toggles to the memory bits.

Analytical power models for leakage power are also available, based on simple analytical formulas for subthreshold leakage and gate leakage [4], [5]. For example leakage current is commonly estimated as follows:

$$I_{leak} = \beta \cdot e^{b(v_{dd} - V_{dso})} \cdot V_t^2 \cdot (1 - e^{-\frac{V_{dd}}{V_t}}) \cdot e^{-\frac{|V_{th}| - V_{off}}{nV_t}} \quad (3)$$

This leakage model is an approximation and several of the parameters are highly dependent on circuit topology and circuit state, e.g., in Equation 3, parameters I_s and n are empirically derived. Leakage power is also closely related to temperature, requiring tight-coupling between leakage-power and temperature models, as discussed in Section 4.

2.B. Empirical Power Models

Empirical modeling approaches are based on circuit-level power analysis performed on structures used in recently-designed processors. This is the approach used by SimplePower [6] and PowerTimer [7]. PowerTimer’s energy models are based on circuit-level power analysis of all microarchitectural structures in the POWER4 microprocessor. Macro-level circuit power analysis, where several large macros combine to form a microarchitectural structure, determines power for microarchitectural blocks as a function of the activities of individual blocks. These power estimates are then combined with utilization information from performance simulators.

Empirical approaches are accurate in formulating models for design macros that tend to be reused in newer designs with relatively small changes. As such, they are quite useful at exploring the design space of existing microarchitectures that are similar to the original analyzed design. However, the models rely on complete circuit simulation, which is generally slow. In addition, flexibility is a major problem, especially if a new process technology

or a new microarchitecture is being proposed. Simple linear scaling of architectural parameters may not be accurate, and the use of newer technologies often brings even larger errors.

Both analytical and empirical models have been widely used by microarchitects to explore power-efficient architectures and are a critical building block for temperature and reliability modeling and research. However, it is important for architects to understand the limitations of the modeling infrastructure before attempting a detailed quantitative evaluation. Section 7 highlights some future directions in the area of power modeling that have the potential to address the accuracy and scalability concerns of existing modeling approaches.

3. THERMAL MODELING AND ANALYSIS

IC thermal analysis requires the simulation of thermal conduction from an IC's power sources, i.e., transistors and interconnects, through cooling package layers, to the ambient environment. Modeling thermal conduction is analogous to modeling electrical conduction, with thermal conductance corresponding to electrical conductance, power dissipation corresponding to electrical current, heat capacity corresponding to electrical capacitance, and temperature corresponding to voltage.

The equation governing thermal conduction in chips and cooling packages follows.

$$\rho c \frac{\partial T(\vec{r}, t)}{\partial t} = \nabla \cdot (k(\vec{r}) \nabla T(\vec{r}, t)) + p(\vec{r}, t) \quad (4)$$

In Equation 4, ρ is the material density; c is the mass heat capacity; $T(\vec{r}, t)$ and $k(\vec{r})$ are the temperature and thermal conductivity of the material at position \vec{r} and time t ; and $p(\vec{r}, t)$ is the power density of the heat source.

To solve the above differential equation using numerical methods, a finite difference discretization method is applied to decompose the chip and package into numerous discrete thermal elements, which may be of non-uniform sizes and shapes. Adjacent thermal elements interact via heat diffusion. Each element has a power dissipation, temperature, thermal capacitance, as well as a thermal resistance to adjacent elements. Thus,

$$\mathbf{C}dT(t)/dt + \mathbf{A}T(t) = Pu(t) \quad (5)$$

where the thermal capacitance matrix, \mathbf{C} , is an $[N \times N]$ diagonal matrix; the thermal conductivity matrix, \mathbf{A} , is an $[N \times N]$ sparse matrix; $T(t)$ and P are $[N \times 1]$ temperature and power vectors; and $u(t)$ is the unit step function. Note that in this tutorial, we assume that P is constant within an analysis interval but may vary between analysis interval.

3.A. Steady-State Thermal Analysis

Steady-state thermal analysis characterizes temperature distribution when heat flow does not vary with time. For microprocessors, steady-state thermal analysis is sufficient for applications with stable power profiles or periodically-changing power profiles that cycle quickly.

For steady-state thermal analysis, the left term in Equation 5 expressing temperature variation as function of time, t , is dropped. Thus,

$$\mathbf{A}T = P \longrightarrow T = \mathbf{A}^{-1}P \quad (6)$$

Therefore, given thermal resistance matrix \mathbf{A} and power vector P , the main task of steady-state thermal analysis is to invert matrix \mathbf{A} . The computational complexity of steady-state thermal analysis is thus determined by the size of matrix \mathbf{A} , which is in turn determined by the discretization granularity of chip and cooling package. Accurate IC thermal analysis requires fine-grain discretization. The size of matrix \mathbf{A} is typically large. Matrices may be inverted using direct or iterative methods. Matrix inversion using direct methods, such as LU decomposition, is sufficient for small matrices, e.g., those associated with models in which functional units are modeled with single thermal elements [8]. However, it is intractable for large matrices. Iterative methods permit larger problems to be solved. Recently, multigrid relaxation has been used for steady-state IC thermal analysis [9],[10]. Each relaxation stage is responsible for eliminating errors in the frequency band corresponding to the stage's discretization granularity. Although multigrid methods often permit rapid convergence for linear systems, our recent analysis shows no significant performance improvement from using multigrid relaxation for steady-state IC thermal analysis, compared to an otherwise identical iterative solver.

3.B. Dynamic Thermal Analysis

Dynamic thermal analysis is used to characterize run-time IC thermal profile when transient variations are significant. Time-domain and frequency-domain methods have been applied to this problem.

Time-domain methods use numerical integration to estimate the run-time IC thermal profile. The targeted time interval is partitioned into numerous time steps. If the size of each time step is small enough, IC thermal transition within each time step can be accurately estimated using finite difference temperature approximation function

$$T(T+h) = \frac{h(Pu(t) - AT(t))}{C} + T(t) \quad (7)$$

in which h is the time step size. Numerical integration constitutes a broad family of methods, which are typically classified based on the approximation error bound. Well-known numerical integration methods include Euler's method, the Trapezoidal method, Runge–Kutta methods, and Runge–Kutta–Fehlberg methods. Fourth-order Runge–Kutta methods are widely used due to their accuracy and speed.

Frequency-domain methods estimate run-time IC thermal profile using approximate analytical solutions, thereby avoiding repeated numerical integration. As shown in Equation 8, the Laplace transform is used to translate the dynamic thermal analysis problem into the frequency domain.

$$\begin{aligned} C(s\mathbf{T}(s) - \mathbf{T}(0^-)) &= \mathbf{A}\mathbf{T}(s) + \mathbf{P}/s \\ \mathbf{T}(s) &= -\mathbf{A}^{-1}(\mathbf{I} - s\mathbf{C}\mathbf{A}^{-1})^{-1}(\mathbf{P}/s + \mathbf{C}\mathbf{T}(0^-)) \end{aligned} \quad (8)$$

Expanding $(\mathbf{I} - s\mathbf{C}\mathbf{A}^{-1})^{-1}$ at $s = 0$:

$$\mathbf{T}(s) = -\mathbf{A}^{-1} \left(\sum_{k=0}^{\infty} (s\mathbf{C}\mathbf{A}^{-1})^k (\mathbf{P}/s + \mathbf{C}\mathbf{T}(0^-)) \right) \quad (9)$$

Thus, IC run-time temperature can be accurately represented as an infinite series in the frequency domain. Since the long time-scale impact of high-frequency components is negligible, truncating Equation 9 and ignoring high-frequency components yields an analytical approximation.

Time-domain and frequency-domain methods are applicable to different modeling scenarios. Time-domain methods rely on numerical integration. Their running times increase with increasing time scale. In addition, since approximation error accumulates, the accuracy of time-domain methods generally degrades as time scales increase. Therefore, time-domain methods are most suitable for short time-scale thermal analysis, i.e., up to a few tens of milliseconds. Frequency-domain methods, in contrast, approximate IC run-time thermal profile using analytical solutions, avoiding the need for numerical integration. However, the one-time computational cost of deriving the analytical approximation is high. Moreover, frequency-domain methods ignore high-frequency components, introducing short time-scale errors. Therefore, frequency-domain methods are more appropriate for long time-scale thermal analysis, i.e., more than a few milliseconds.

3.C. Adaptive Thermal Analysis

The major challenges of numerical IC thermal analysis are high computational complexity and memory usage. High modeling accuracy requires fine-grain discretization, resulting in numerous grid elements. For dynamic thermal analysis using time-domain methods, higher modeling accuracy requires fine spatial and temporal discretization granularities, increasing the computational overhead and memory usage. For dynamic thermal analysis using frequency-domain methods, deriving analytical approximations involves computation and memory intensive numerical operations, such as the inversion and multiplication of large matrices. High thermal element count may hinder or prevent the application of frequency-domain methods to model complicated IC thermal setups.

Recent progress on adaptive numerical modeling methods tackles the high computation complexity and memory usage of steady-state, time-domain, and frequency-domain thermal analysis [10].

Spatial Adaptation: Chip and package thermal gradients exhibit significant spatial variation due to the heterogeneity of thermal conductivity and heat capacity in different chip and cooling package materials, as well as variation in power profiles. The wide distribution of temperature differences across chip and cooling package motivates the design of a thermal gradient sensitive adaptive spatial discretization refinement technique that automatically adjusts the spatial partitioning resolution to maximize thermal modeling efficiency and guarantee modeling accuracy. In this technique, the spatial discretization process is governed by temperature difference constraints. Iterative refinement is conducted in a hierarchical fashion. Regions with high magnitude gradients are recursively partitioned until elements are isothermal and the temperature profile converges. This method uses fine-grain discretization only where necessary for accuracy, thereby improving the efficiency of steady-state and time-domain dynamic thermal analysis, and enabling the use of frequency-domain methods on modeling complicated chips and packages.

Temporal Adaptation: There are two categories of time-domain methods: non-adaptive and adaptive. Non-adaptive methods use the same time step size throughout analysis. Therefore, their performance is bounded by the smallest time step required by any thermal element at any time. Temporally adaptive methods, in contrast, can improve performance substantially without loss of accuracy by adapting time step sizes during dynamic analysis. Temporally adaptive methods can be further classified into two categories: synchronous time marching methods and asynchronous time marching methods. In synchronous time marching methods, such as adaptive fourth-order Runge-Kutta methods, all thermal elements advance in time in lock-step. However, before each step, the step size is adjusted to the minimal size required for accuracy by any thermal element. In contrast, the asynchronous time marching method permits different elements to advance forward in time at different rates. It exploits temporal and spatial differences in temperature variation within chip and cooling packages. Each thermal element uses the largest safe step size for its position and time, instead of being forced to use the same step size as all other elements in the chip and package. By allowing elements to progress forward in time asynchronously, thermal analysis can be dramatically accelerated without loss of accuracy.

4. TEMPERATURE-DEPENDENT LEAKAGE POWER ANALYSIS

Technology scaling is increasing the proportion of power consumption due to leakage; accurate leakage analysis is now important. IC leakage power consumption is a strong function of temperature. Iterative temperature-dependent leakage analysis has been widely used for accurate leakage estimation. This approach was developed based on the following observations. First, subthreshold leakage power consumption increases superlinearly with temperature. This has led to the incorrect assumption that highly-accurate thermal modeling embedded within the power analysis flow is necessary to accurately determine temperature-dependent leakage power consumption. Second, leakage variation per unit temperature change is less than 1, i.e., $dP_{leakage}/dT < 1$. An iterative-based thermal-leakage flow can thus guarantee convergence and correctness. The major challenge of existing thermal-leakage analysis flow is high computational complexity, which significantly increases the running times of both IC thermal analysis and power analysis.

A recent study shows that highly-efficient and accurate temperature-dependent leakage power analysis is possible using coarse-grained thermal modeling [11]. Leakage power mainly depends on IC thermal profile and circuit design style. Despite the non-linear dependence of leakage power on temperature, within the operating temperature ranges of real ICs, using a linear leakage model for individual functional units results in less than 1% error in leakage estimation and permits more than four orders of magnitude reduction in analysis time relative to an approach relying on detailed accurate thermal analysis.

5. POWER AND TEMPERATURE RELATED MICROPROCESSOR RELIABILITY

As shown in Figure 1 power and its byproducts are now a great threat to the reliability of microprocessors. High power consumption translates to high temperatures: a recipe for permanent faults due to electromigration and other failure processes. Rapidly changing power consumption of components supplied by resistive and inductive power delivery networks results in transient voltage fluctuations and therefore transient faults. Soft errors are also becoming an increasingly important source of IC reliability problems [12]. However, the focus of this article is power consumption and most soft errors are orthogonal to power consumption.

5.A. Modeling Permanent Power and Temperature Related Faults

For IC permanent faults, major failure mechanisms include electromigration, thermal cycling, time-dependent dielectric breakdown, and stress migration [13]–[15].

- 1) **Electromigration** refers to the gradual displacement of the atoms in metal wires caused by electrical current. It leads to voids and hillocks within metal wires that result in open and short circuit failures. The MTTF due to electromigration is given by the following equation:

$$MTTF_{EM} = \frac{A_{EM}}{J^n} e^{\frac{E_{\alpha_{EM}}}{\kappa T}} \quad (10)$$

where A_{EM} is a constant determined by the physical characteristics of the metal interconnect, J is the current density, $E_{\alpha_{EM}}$ is the activation energy of electromigration, n is an empirically-determined constant, and T is the temperature.

- 2) **Thermal cycling** refers to IC fatigue failures caused by thermal mismatch deformation. In chip and package, adjacent material layers, such as copper and low- k dielectric, have different coefficients of thermal expansion. As a result, run-time thermal variation causes fatigue deformation, leading to failures. The MTTF due to thermal cycling is given by the following equation:

$$MTTF_{TC} = \frac{A_{TC}}{(T_{average} - T_{ambient})^q} \quad (11)$$

where A_{TC} is a constant coefficient, $T_{average}$ is the chip average run-time temperature, $T_{ambient}$ is the ambient temperature, and q is the Coffin-Manson exponent constant.

- 3) **Time-dependent dielectric breakdown** refers to deterioration of the gate dielectric layer. This effect is a strong function of temperature, and is becoming increasingly prominent with the reduction of gate-oxide dielectric thickness and non-ideal supply voltage reduction. The MTTF due to time-dependent dielectric breakdown is given by the following equation:

$$MTTF_{TDDDB} = A_{TDDDB} \left(\frac{1}{V} \right)^{(a-bT)} e^{\frac{A+B/T+CT}{\kappa T}} \quad (12)$$

where A_{TDDDB} is a constant, V is the supply voltage, and $a, b, A, B,$ and C are fitting parameters.

- 4) **Stress migration** refers to the mass transportation of metal atoms in metal wires due to mechanical stress caused by thermal mismatch among metal and dielectric materials. The MTTF due to stress migration is given by the following equation:

$$MTTF_{SM} = A_{SM} |T_0 - T|^{-n} e^{\frac{E_{\alpha_{SM}}}{\kappa T}} \quad (13)$$

where A_{SM} is a constant, T_0 is the metal deposition temperature during fabrication, T is the run-time temperature of the metal layer, n is an empirically-determined constant, and $E_{\alpha_{SM}}$ is the activation energy for stress migration.

Equations 10–13 indicate that IC fault processes are strongly influenced by temperature, and in many cases are exponentially dependent on it. As a result, it appears that detailed and accurate thermal modeling is necessary for reliability estimation. Some of these fault processes are also accelerated by increases in other parameters related to power such as current density and voltage.

The most dangerous fault processes in microprocessors accelerate as a result of wear and this complicates reliability modeling and analysis. A number of researchers have assumed that microprocessor fault processes are Poisson processes and used exponential distributions to model them. Exponential distributions are mathematically convenient because they permit the rates of different fault processes operating on different components to be added in order to determine the failure rate of the entire microprocessor. However, they do not model wear, which is generally required for accurate reliability modeling [16]. The lognormal distribution better models prominent microprocessor fault processes because it models the increase in failure rate with increasing time and wear. However, this property complicates system-level modeling. There is no straightforward method of deriving a closed-form expression for the failure rate of a microprocessor composed of numerous components for which lognormal fault process models

are used. Therefore, some microprocessor reliability estimation work assumes exponential fault processes, which may be inaccurate, while other work uses Monte Carlo simulation [17] or techniques based on statistical curve fitting, each of which may be slow. At present, efficient and accurate system-level lifetime reliability estimation is a goal that remains just slightly out of grasp, but toward which research is rapidly progressing.

5.B. Modeling Transient Power-Related Faults

In recent years, interest has been building in microarchitectural support for an increasingly important reliability concern – power supply noise. The circuits on high-performance chips place stringent demands on the power delivery infrastructure responsible for satisfying current demands and maintaining reference voltages. Power supply integrity is essential because even minor deviations in power or ground reference levels can introduce noise or delay into critical signal transitions, leading to unrecoverable error. Noise reduction is difficult due to parasitic inductance in the power delivery system. Whenever the processor current demands vary, the transient load causes voltage fluctuations given by the well known equation $V = LdI/dt$ where L is the value of the inductor and dI/dt is the rate of change of the current. This relationship gives inductive noise its informal name, the dI/dt problem.

Interest in microarchitectural support has grown due to the increasing difficulty of mitigating noise through conventional avenues [18]–[20]. Traditionally, inductive noise could be addressed by reducing effective inductance or adding capacitance on-die or in the package to dampen out the noise. However, the voltage noise tolerance of future processors will decrease at a rate that outstrips our ability to remove the parasitics or add capacitance. In particular, current processor designs use a large amount of on-die capacitance; as much as 15–20% of die area may be devoted to decoupling capacitors. While on-die capacitors are very effective at dampening inductive noise, they consume leakage power and die area because they are implemented as non-switching transistors with large gate capacitances. Furthermore, in future designs absolute voltage tolerance decreases as the supply voltage scales and load transients increase as total power increases. Consequently, the rate of change for load current (dI/dt) will increase and the allowable variation in supply voltage will shrink.

Microarchitectural techniques seek to limit inductive noise by controlling the rate at which load current changes. This is done essentially by smoothing out or altering the current profile of the processor to eliminate problematic load transients. These techniques can reduce the burden on traditional circuit and package solutions for inductive noise.

Impact of Frequency

The severity of inductive noise depends not only on the magnitude of transient current, but also on the frequency range over which it changes. Conventional power supply systems consist of a complex network of die, package, and motherboard capacitances, inductances, and resistances. These elements are excited to different degrees by specific current variation patterns. Specifically, mid-frequency noise in the range of 50–200 MHz and high-frequency noise near the processor clock rate have achieved the most attention in literature. Mid-frequency noise is associated with package inductance that reacts with die and package capacitors whenever the processor current varies within the problematic frequency range. Variations at these critical frequencies are problematic because they allow resonance to build in the power delivery network, producing violent swings in supply voltage. High-frequency noise produces sharp localized fluctuations in on-chip supply voltage. These high-frequency noise patterns arise when processor execution resources have abrupt changes in current demand that cannot be adequately serviced by neighboring on-chip or in-package capacitors and solder bump connections to the off-chip network. We focus on mid-frequency noise as it provides the most opportunity for architectural solutions.

Mid-Frequency Models

Researchers in the electronic packaging community have characterized the overall response of the power delivery system as a second order linear system [21], which can essentially be represented as a single lumped RLC underdamped network [22]. In this simple model, the R, L, and C values do not correspond to any specific physical elements in the real system, rather they are effective values that consider the composite effects of all elements in the network.

The dominant characteristic of mid-frequency models is resonance. The most problematic load profile is a pulse pattern at the natural frequency of the RLC network, namely $f_c = 1/2\pi\sqrt{LC}$. In current processors this translates

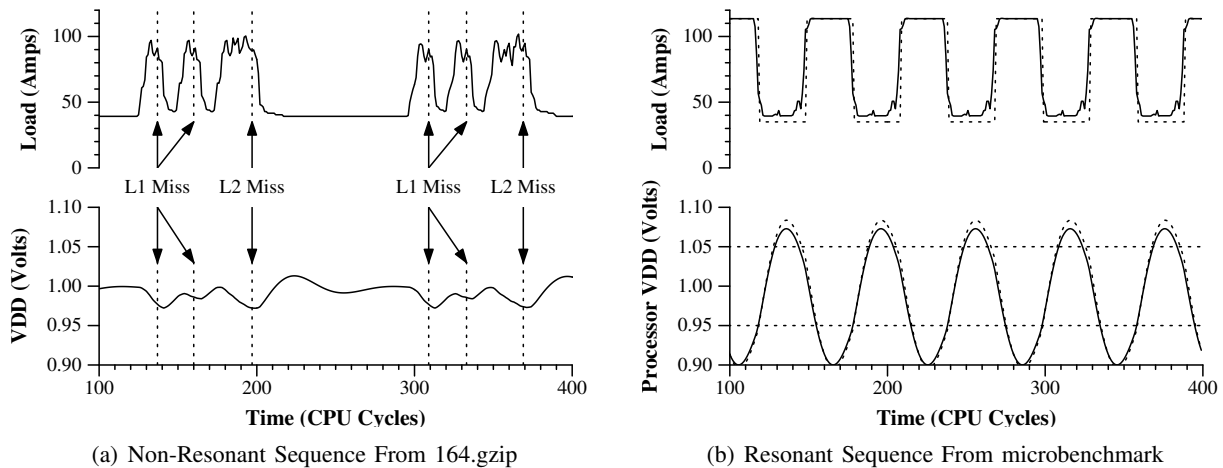


Figure 2. Noise impact of different current consumption patterns: (a) an execution sequence from 164.gzip, which does not hit a resonant frequency and (b) a trace from a dI/dt microbenchmark, which features severe ILP variation and mimics a resonant pulse.

to alternating patterns of on-off activity on the order of many tens to about one hundred processor cycles. This phenomena is illustrated in Figure 2, which shows the impact of a non-resonant current load and a resonant current pulse train, with the latter having a much more significant noise impact. Architecture level phenomena such as cache misses and ILP variations can produce bursty current behavior at this frequency range if they are sequenced in an unfortunate manner.

Microarchitectural simulators with support for cycle-by-cycle power estimates can model the impact of mid-frequency inductive noise. In essence, the total current load of the processor can be assumed to be its instantaneous power divided by the *ideal supply voltage*. Note that in an actual circuit, a real supply voltage droop would cause current draw to decrease, making division by a constant value a conservative estimate. Also note that dividing by varying voltage would be unnecessarily pessimistic. This instantaneous current can be related to varying supply voltage through (1) filter techniques that use convolution to relate a history of current consumption to voltage [18], [19], [23] or (2) circuit simulation that directly models the RLC network [24], [25]. In either case, models can be extended to capture higher order power delivery system effects through use of a more complex convolution filter or a detailed circuit model.

6. MODELING PROCESS VARIATION

Parameter variation is an unavoidable consequence of continued technology scaling. The impact of random and systematic variation on physical factors such as gate length and interconnect spacing will have a profound impact on performance and power consumption. In current designs, foundry-induced physical deviations already produce significant die-to-die variation. In particular, industry data for a high-performance processor in a 180 nm technology shows that individual dies produced with the same fabrication equipment can have as much as a 30% die-to-die frequency variation and a $20\times$ leakage power variation [26]. ITRS predicts that manufacturing variability will have an increasing prominence in future designs.

Current architecture-level models for process variation have focused on effective v_t deviations and attempt to relate physical uncertainty in this highly influential parameter to performance and power. This is done through a mixture of analytical and empirical transistor models, circuit structure characteristics, and statistical principles that relate physical uncertainty to architecturally visible characteristics.

Variation Under Statistical Distributions

Physical variation at the transistor level can be modeled by statistical relationships. A combination of analytical models and Monte Carlo simulation are often used to give modeled transistor parameters the appropriate statistical distributions. Dopant ion concentration is normally modeled as a Gaussian distribution. Under this model, there is no correlation between any pair of transistors in the design. On the other hand, gate length is known to have strong spatial

correlation properties [27]. Due to localized imperfections in lithography, neighboring transistors are likely to have similar gate length deviations. As the distance between transistors increases, correlation decreases linearly [27]. This is most frequently analyzed using Monte Carlo simulation. The die is modeled as a collection of grid sections. All transistors within a section are assumed to have identical parameter variation. Several methods, including hierarchical methods [28] as well as convolution kernels [29], have been used to assign parameter deviations to neighboring blocks. An alternative to Monte Carlo simulation is to model systematic components of leakage variation through an empirically derived gate length deviation model [30]. Note that this approach is useful for average case studies, but cannot capture probabilistic aspects of leakage variation.

Relating v_t Variation to Microarchitecture

Performance and power variations at the transistor level are dominated by effective threshold voltage (v_t) which is in turn determined by (1) drawn gate length and (2) dopant concentration [31]. Process variations in either of these physical factors will cause deviations in V_t and lead to conflicting impacts on performance and leakage power [26], [32]. Lower effective threshold voltages caused by short gate lengths or over-doping decrease delay in a roughly linear fashion but lead to an exponential increase in subthreshold leakage current. Deviations that increase v_t reduce leakage with the penalty of increased delay. While dynamic power can be affected by v_t variation, its impact on leakage power dominates and has been the primary focus of power models for parameter variation.

At the gate level, the effects of varying threshold voltage can be modeled analytically or empirically. Analytical models use device models to relate physical variation in critical parameters, namely gate length and concentration of dopant ions to calculate effective threshold voltage. Based on an effective threshold voltage, the leakage current can be computed using Equation 3.

Likewise, threshold's effect on transistor delay can be modeled by the well known Alpha power law:

$$t_d \propto \frac{V_{DD}}{V_{DD} - V_{TH}^\alpha} \quad (14)$$

Together, these equations govern the relative impact of process variation at the transistor level and form the basis for analytical models. For leakage, existing architectural models provide baseline leakage power on a component basis. The analytical models project how leakage scales. For performance, the impact on effective v_t has can also be related to architectural structure. Within each pipeline stage, the number of critical paths determines the minimum frequency of the stage [32], [33]. In essence, a large number of critical paths increases the probability that one of the paths will be slow and hence decreases the clock frequency. Architectural models for critical path count can be derived from circuit structure [30]. In particular, HDL descriptions of processor designs [34] can be used to count critical paths in each stage. SPICE based empirical models offer another alternative. Detailed simulation on either complete circuits or representative sections can determine the impact on power and performance.

7. FUTURE DIRECTIONS AND CONCLUSION

Continued technology scaling and emerging directions in the design will change the power, thermal, reliability, and process variation modeling problems. We now summarize problems that will likely be encountered in the near future for each of these modeling domains.

- **Power:** The move to nanoscale technologies will cause traditional circuit scaling theory to break down, exacerbating the difficulties associated with analytical dynamic and leakage power modeling. Furthermore, asymmetric and heterogeneous chip-level multiprocessors (CMPs) will lead to an increased diversity in microarchitectures and accelerator cores, limiting the utility of empirical modeling approaches. As more designers explore “many-core” and heterogenous core systems, power modeling tools will need to incorporate better estimates for interconnect and combinatorial logic structures. To address these challenges, future power modeling approaches may rely on a mixture of analytical and empirical modeling techniques, seeking to leverage the advantages of both techniques.
- **Thermal:** Increasing IC integration and power density will bring new thermal modeling challenges. Moreover, the changes that new fabrication processes, materials, and cooling solutions bring to thermal models will require advances in analysis to permit efficiency and speed. For instance, the thermal properties of stacked

3-D ICs will differ greatly from conventional 2-D ICs, and high power density remains one of the major 3-D integration challenges. Increasing microprocessor power density will require novel cooling solutions such as microchannel cooling or nanotube thermal vias, requiring multi-resolution modeling down to the micro-meter or nano-meter scale.

- **Reliability:** Most system-level reliability models are derived using statistical techniques based on empirical models. These statistical approximations can benefit greatly from calibration and validation. Technology scaling further increases the importance and challenges of reliability modeling. Reduced feature size increases the vulnerability of individual devices and interconnects. This will require models that capture the system-level impact of nano-scale components while permitting efficient analysis. However, increasing integration density complicates the development of accurate microarchitectural reliability models. Future CMP architectures also present additional opportunities for per-core power-down and DVFS, leading to additional dI/dt noise challenges.
- **Process Variation:** Architects will need to improve the accessibility of their models and adjust them to keep pace with manufacturing trends. In addition, architects will have to develop unified models that better capture couplings between power, performance, and reliability aspects of variation. While Monte Carlo variation analysis is relatively easy to apply for power/performance studies, it does not offer the same intuition and insights that analytic models can. Mathematical models that are easily parameterizable but can capture probabilistic characteristics such as mean, variance, and skew could be extremely beneficial.

Power-related design challenges have become a critically important topic for computer architects and system designers. Microarchitectural design requires detailed models of power-related phenomena. As advanced design techniques and fabrication process changes reveal new power-related phenomena, power, thermal, reliability, and process variation models will require ongoing improvement.

REFERENCES

- [1] P. Shivakumar and N. P. Jouppi, "CACTI 3.0: An integrated cache timing, power, and area model," Western Research Lab 2001/2, Tech. Rep., Aug. 2001.
- [2] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: A framework for architectural-level power analysis and optimizations," in *Proc. Int. Symp. Computer Architecture*, June 2000, pp. 83–94.
- [3] N. S. Kim, et al., "Microarchitectural power modeling techniques for deep sub-micron microprocessors," in *Proceedings of the 2004 international symposium on Low power electronics and design*, 2004, pp. 212–217.
- [4] J. A. Butts and G. S. Sohi, "A static power model for architects," in *Proc. Int. Symp. Microarchitecture*, Dec. 2000, pp. 191–201.
- [5] Y. Zhang, et al., "HotLeakage: A temperature-aware model of subthreshold and gate leakage for architects," Univ. of Virginia, Tech. Rep., May 2003, CS-2003-05.
- [6] W. Ye, et al., "The design and use of simplepower: a cycle-accurate energy estimation tool," in *Proc. Design Automation Conf.*, June 2000, pp. 340–345.
- [7] D. Brooks, et al., "New methodology for early-stage, microarchitecture-level power-performance analysis of microprocessors," *IBM J. Res. Dev.*, vol. 47, no. 5–6, pp. 653–670, 2003.
- [8] K. Skadron, et al., "Temperature-aware microarchitecture," in *Proc. Int. Symp. Computer Architecture*, June 2003, pp. 2–13.
- [9] P. Li, et al., "Efficient full-chip thermal modeling and analysis," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2004, pp. 319–326.
- [10] Y. Yang, et al., "Adaptive multi-domain thermal modeling and analysis for integrated circuit synthesis and design," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2006, pp. 575–582.
- [11] Y. Liu, et al., "Accurate temperature-dependent integrated circuit leakage power estimation is easy," in *Proc. Design, Automation & Test in Europe Conf.*, Mar. 2007, to appear.
- [12] S. S. Mukherjee, et al., "A systematic methodology to compute the architectural vulnerability factors for a high-performance microprocessor," in *Proc. Int. Symp. Microarchitecture*, Dec. 2003, pp. 29–40.
- [13] Joint Electron Device Engineering Council, "Failure mechanisms and models for semiconductor devices," in *JEDEC Publication JEP 122-B*, Aug. 2003.

- [14] J. Srinivasan, et al., "Exploiting structural duplication for lifetime reliability enhancement," in *Proc. Int. Symp. Computer Architecture*, June 2005, pp. 520–531.
- [15] E. Karl, et al., "Reliability modeling and management in dynamic microprocessor-based systems," in *Proc. Design Automation Conf.*, June 2006.
- [16] B. Schroeder and G. A. Gibson, "A large-scale study of failures in high-performance computing systems," in *Proc. Int. Conf. Dependable Systems and Networks*, June 2006, pp. 249–258.
- [17] J. Srinivasan, et al., "The impact of technology scaling on lifetime reliability," in *Proc. International Conf. Dependable Systems and Networks*, June 2004, pp. 177–186.
- [18] W. El-Essawy and D. Albonesi, "Mitigating inductive noise in smt processors," in *International Symposium on Low Power Electronics and Design (ISLPED)*, August 2004.
- [19] R. Joseph, D. Brooks, and M. Martonosi, "Control techniques to eliminate voltage emergencies in high performance processors," in *Proceedings of the 9th International Symposium on High Performance Computer Architecture (HPCA-9)*, February 2003.
- [20] M. D. Powell and T. N. Vijaykumar, "Exploiting resonant behavior to reduce inductive noise," in *Proceedings of 31st International Symposium on Computer Architecture (ISCA-31)*, June 2004.
- [21] D. J. Herrell and B. Beker, "Modeling of power distribution systems for high-performance microprocessors," *IEEE Transactions on Advanced Packaging*, vol. 22, no. 3, pp. 240–248, August 1999.
- [22] I. Kantorovich, et al., "Measurement of low impedance on chip power supply loop," *IEEE Trans. on Advanced Packaging*, vol. 27, no. 1, Feb. 2004.
- [23] E. Grochowski, D. Ayers, and V. Tiwari, "Microarchitectural di/dt control," *IEEE Design and Test of Computers*, vol. 20, no. 3, pp. 40–47, May/June 2003.
- [24] M. D. Powell and T. N. Vijaykumar, "Pipeline muffling and a priori current ramping: Architectural techniques to reduce high-frequency inductive noise," in *Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED)*, August 2003.
- [25] M. S. Gupta, et al., "Understanding voltage variations in chip multiprocessors using a distributed power-delivery network," in *Design Automation and Test In Europe (DATE 2007)*, 2007, pp. 201–204.
- [26] S. Borkar et al., "Parameter variations and impact on circuits and microarchitecture," in *Proc. of the 40th DAC*, 2003.
- [27] P. Friedberg, et al., "Modeling within-die spatial correlation effects for process-design co-optimization," in *Proc. of the 6th Int. Symp. on Quality Electronic Design*, 2005.
- [28] A. Agarwal, et al., "Path-based statistical timing analysis considering inter and intra-die correlations," in *Proc. of TAU*, 2002.
- [29] K. Meng, et al., "Modeling and characterizing power variability in multicore architectures," in *IEEE Symposium on Analysis of Software and Systems (ISPASS)*, April 2007.
- [30] K. S. Eric Humenay, David Tarjan, "Impact of parameter variations on multicore chips," in *Workshop on Architectural Support for Gigascale Integration 2006 (ASGI in conjunction with ISCA-33)*, June 2006.
- [31] S. R. Nassif, "Modeling and forecasting of manufacturing variations," in *Proceeding of the Fifth International Workshop Statistical Metrology*, 2000.
- [32] D. Marculescu and E. Talpes, "Energy awareness and uncertainty in microarchitecture-level design," *IEEE Micro*, vol. 25, pp. 64–76, Sept.-Oct. 2005.
- [33] K. A. Bowman, S. G. Duvall, and J. D. Meindl, "Impact of die-to-die and within-die parameter fluctuations on the maximum clock frequency distribution for gigascale integration," *IEEE Journal of Solid State Circuits*, vol. 37, pp. 183–190, Feb. 2002.
- [34] B. F. Romanescu, S. Ozev, and D. J. Sorin, "Quantifying the impact of process variability on microprocessor behavior," in *Workshop on Architectural Reliability (WAR in conjunction with MICRO-39)*, December 2006.