

System-Level Reliability Modeling for MPSoCs

Yun Xiang[†] Thidapat Chantem[‡] Robert P. Dick[†] X. Sharon Hu[‡] Li Shang[§]

[†] EECS Department
University of Michigan
Ann Arbor, MI 48109

[‡] CSE Department
University of Notre Dame
Notre Dame, IN 46556

[§] ECE Department
University of Colorado
Boulder, CO 80309

{xiangyun@, dickrp@eecs.}umich.edu {tchantem, shu}@nd.edu li.shang@colorado.edu

ABSTRACT

The reliability of multi-processor systems-on-chip (MPSoCs) is affected by several inter-dependent system-level and physical effects. Accurate and fast reliability modeling is a primary challenge in the design and optimization of reliable MPSoCs. This paper presents a reliability modeling framework that integrates device-, component-, and system-level models. This framework contains modules for electromigration, time-dependent dielectric breakdown, stress migration, and variable-amplitude thermal cycling. A new statistical reliability distribution is proposed for accurate characterization of components containing too few devices for an extreme value distribution to be appropriate. A hierarchical system-level survival lattice based Monte Carlo technique is used to estimate the temporal fault distributions of MPSoCs that use arbitrary static and dynamic reliability-enhancing redundancy schemes. Physical process variation, which may have a significant impact on MPSoC reliability, is considered in the model. The proposed modeling technique has 5% average error in mean time to failure and reduces simulation time by nearly 3 orders of magnitude relative to a non-hierarchical Monte Carlo technique.

Categories and Subject Descriptors: C.4 [Performance of Systems]: Reliability, availability, and serviceability; Fault tolerance.

General Terms: Reliability, Performance

1. INTRODUCTION

Integrated circuit (IC) process scaling has resulted in decreasing minimal feature sizes, increasing temperatures, and increasing current densities, all of which increase the rate of wear due to lifetime fault processes such as electromigration, time-dependent dielectric breakdown, stress migration, and

This work was supported in part by SRC under award number 2007-TJ-1589 and in part by NSF under award numbers CCF-0964763, CNS-0347941, and CCF-0702705. We also acknowledge the support of the U.S. Department of Education through a GAANN Fellowship for T. Chantem (award P200A090044).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CODES+ISSS'10, October 24–29, 2010, Scottsdale, Arizona, USA.
Copyright 2010 ACM 978-1-60558-905-3/10/10 ...\$10.00.

thermal cycling. As a result of scaling from a 180 nm to 65 nm process technology, temperature generally increases by 14°C [1]. Many failure mechanisms are exponentially dependent on temperature; a 10–15°C increase in temperature may result in less than half the original lifetime [2].

Reliability analysis and enhancement techniques play key roles in improving IC lifetime, and can be used at various abstraction levels. Considering reliability during system-level design has the potential to yield greater improvements in lifetime at lower costs than considering only lower-level reliability enhancement techniques. System-level decisions, such as component placement, power state control, and functional unit duplication can significantly impact IC temperatures, current density profiles, and redundancies, thus influencing IC reliability.

Empirically developing and evaluating system-level reliability models are challenging tasks due to the number and duration of tests required to determine the system-level temporal fault distributions, and the dependence of these distributions on parameters such as temperature and current density. Though a large amount of work exists on modeling reliability for individual devices, current system-level reliability models have substantial limitations. Coskun et al. [3] and Rosing, Mihic, and De Micheli [4] developed MPSoC reliability models based on exponential component temporal fault distributions. Exponential distributions imply a constant failure rate throughout the MPSoC lifetime. They are appropriate when only early failure is considered, but are insufficient for MPSoCs with long planned lifetimes, for which wear is important. Srinivasan et al. proposed a reliability-aware microprocessor design model [5, 6], which is based on lognormal component temporal fault distributions. Monte Carlo simulation was used to derive system mean time to failure (MTTF) in the presence of system-level redundancy. Gu et al. [7] employed the same assumption about component temporal fault distributions but used Min-Max approximation to estimate the system MTTF. The use of lognormal distributions for components composed of many devices may cause substantial error in system MTTF estimation, as will be shown in Section 7.3. Karl et al. [8] estimated system-level reliability under the assumption that all devices are essential. This approach is accurate but lacks the flexibility to model complex, fault-tolerant designs. Moreover, the increasing number of devices per system makes flat full-chip simulation too time consuming for use in reliability-aware MPSoCs synthesis and optimization.

In this paper, we present a system-level reliability modeling framework that builds on knowledge of device-level temperature- and current-dependent temporal fault distributions and component redundancy to support accurate and efficient system-level reliability estimation. In particular,

the proposed infrastructure supports variable failure rates, variable component temporal fault distribution functions, as well as variable-amplitude and high-frequency thermal cycle modeling. It models components containing devices with lognormal or Weibull device failure distributions, appropriately adjusts component temporal fault distribution functions based on the number of devices in the component, and handles intra- and inter-die process variation. The proposed modeling technique is validated by comparing the estimated system MTTFs with those produced by a (very slow but accurate) full-system device-level Monte Carlo technique. Simulation results show that the proposed model has 5% error on average and decreases simulation time by almost 3 orders of magnitude compared to the full-system device-level Monte Carlo method.

The rest of this article is organized as follows. Section 2 gives a high-level overview of our reliability modeling process and describes the resulting model. Section 3 provides background on device-level failure mechanisms. Section 4 provides equations to calculate temporal failure distributions under time-varying temperatures and describes an accurate thermal cycling model. Section 5 describes a method to derive component-level reliability models from device-level models. Section 6 describes the Monte Carlo technique used to estimate both static and dynamic system-level reliability. Section 7 evaluates the accuracy and efficiency of the proposed modeling technique. Section 8 concludes the paper.

2. MODELING INFRASTRUCTURE

The proposed system-level modeling framework is hierarchical. Three abstraction levels are used: device, component, and system. A device is defined as the smallest discrete circuit element that is vulnerable to failure, e.g., a transistor, a wire, or a via. A component is composed of numerous critical devices; the failure of any device causes the component to fail. A system is a group of components connected in an arbitrary structure; some components may be redundant. Note that functional units with internal redundancy can be modeled within this framework by treating them as multiple redundant components. At the device level, temporal failure mechanisms are characterized by empirical testing. The device-level temporal distributions for many failure processes are well understood: we assume they are known. In contrast, component- and system-level failure distributions are expensive to empirically determine. It is therefore our goal to automatically construct component- and system-level reliability models starting from known device-level models and knowledge of system-level redundancies.

The flow of our multi-level system reliability modeling technique is shown in Figure 1. Device specifications consist of distribution type, shape parameters, as well as reference MTTF and voltage. Component specifications consist of component device counts and process variation parameters. These specifications are used in our modeling frameworks to obtain accurate component temporal failure distributions. The simplified component distribution supports estimation of component reliability at any time given a temperature and voltage time series. Component reliability information, initial wear, still functioning system components, a list of valid states in the system survival lattice, and expected lifetime are used to determine the time of the next component failure. A *survival lattice* is a description of all the fault conditions the system can tolerate [9]. *Expected lifetime* is a random variable generated for each component at the beginning of each Monte Carlo trial; a component fails when

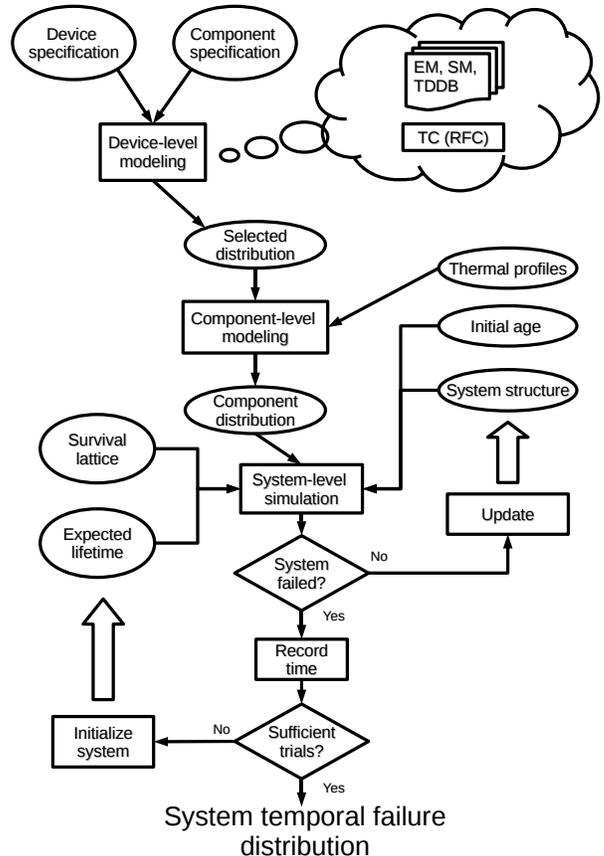


Figure 1: Reliability modeling infrastructure.

its wear exceeds its expected lifetime. If the system survives the failure of a component, the system structure, component wear, and thermal profile are updated. Note that adaptation in workload distribution and power management state can cause thermal and current density profiles to change after the failures of individual components. If, instead, the system fails, the Monte Carlo trial ends.

3. DEVICE-LEVEL FAILURE MECHANISMS BACKGROUND

This section reviews the device lifetime failure mechanisms that are presently dominant for ICs: electromigration, time-dependent dielectric breakdown, stress migration, and thermal cycling. This background material will be useful in explaining the proposed reliability modeling techniques.

1. *Electromigration (EM)* refers to dislocation of metal atoms caused by momentum imparted by electrical current in wires and vias. The MTTF due to electromigration is given by the following equation [10, 11]:

$$MTTF_{EM} = \frac{A_{EM}}{J^n} e^{\frac{E_{aEM}}{\kappa T}}, \quad (1)$$

where A_{EM} is a constant determined by the physical characteristics of the metal interconnect, J is the current density, E_{aEM} is the activation energy for electromigration, n is an empirically determined constant, κ is the Boltzmann constant, and T is the temperature.

2. *Time-dependent dielectric breakdown (TDDB)* refers to deterioration of the gate oxide layer. Gate current due to hot electrons causes defects in the oxide, which eventually form a low-impedance path and cause the transistor to permanently fail. This effect is strongly influenced by temperature, and is increasing 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 [10, 9]:

$$MTTF_{TDDB} = A_{TDDB} \left(\frac{1}{V} \right)^{(a-bT)} e^{\frac{A+B/T+CT}{\kappa T}}, \quad (2)$$

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

3. *Stress migration* is caused by the directionally biased motion of atoms in metal wires due to mechanical stress caused by thermal mismatch among metal and dielectric materials. The MTTF resulting from stress migration is given by the following equation [10]:

$$MTTF_{SM} = A_{SM} |T_0 - T|^{-n} e^{\frac{E_{aSM}}{\kappa T}}, \quad (3)$$

where A_{SM} is a fitting 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_{aSM} is the activation energy.

4. *Thermal cycling* refers to wear caused by thermal stress resulting from mismatched coefficients of thermal expansion for adjacent material layers; run-time temperature variation results in inelastic deformation, eventually leading to failure [12]. The number of cycles to failure (N_{TC}) can be calculated using a modified Coffin-Manson equation [13]:

$$N_{TC} = A_{TC} (\delta T - T_{th})^{-b} e^{\frac{E_{aTC}}{\kappa T_{Max}}}, \quad (4)$$

where A_{TC} is an empirically determined constant, δT is the thermal cycle amplitude, T_{th} is the temperature at which inelastic deformation begins, b is the Coffin-Manson exponent constant, E_{aTC} is the activation energy, and T_{Max} is the maximum temperature during the cycle. Note that δT can change between cycles.

The above empirical MTTF or cycles to failure equations estimate the lifespans of individual devices. However, they cannot be directly used in component- or system-level analysis. Device temporal failure densities are usually characterized using Weibull and lognormal distributions. Experimental results show that TDDB failures usually have Weibull distributions [14] and EM failures have lognormal distributions [15, 16]. The distributions for stress migration and thermal cycling are not known with certainty due to the scarcity of experimental data. We assume that they have Weibull distributions in this work. If future research indicates that they have other distributions, the proposed reliability modeling infrastructure can be easily adjusted. The reliability (i.e., probability of surviving until a particular time) of a device with Weibull distribution can be calculated using the following equation:

$$R(t) = e^{-\left(\frac{t}{\eta}\right)^\beta}, \quad (5)$$

where η is the scale parameter and β is the shape parameter. For devices with lognormal distributions, the reliability can

be calculated using

$$R(t) = \frac{1}{2} - \frac{1}{2} \operatorname{erf} \left(\frac{\ln(t) - \mu}{\sqrt{2\sigma^2}} \right), \quad (6)$$

where μ is the scale parameter, σ is the shape parameter, and erf is the error function.

4. TEMPERATURE VARIATION

Modern ICs typically employ dynamic thermal, power, and performance management techniques such as dynamic voltage and frequency scaling. As a result, reliability-related parameters such as temperature, voltage, and frequency may vary over time. In the rest of the paper, we will focus on modeling temporal and spatial temperature variation since they have strong impacts on most of the important failure mechanisms. The same modeling techniques can be easily applied to other parameters. Modeling temperature variation poses two challenges, (1) steady-state wear estimation and (2) thermal cycling wear estimation, which this section addresses. These methods are used in both device- and component-level modeling steps in our framework.

4.1 Reliability Computation

In the presence of temperature variation, using Equation 5 or 6 to compute reliability can lead to large errors. In this section, we present a method for determining reliability given arbitrary temporal temperature variation. We focus our discussion on device-level modeling. However, the method can be readily applied to component-level modeling.

We start by assuming Weibull device temporal failure distributions (lognormal distributions will be discussed later). Suppose the temperature time series is represented by a sequence of tuples $[(t_1, T_1) \cdots (t_j, T_j)]$, where t_i is the interval duration and T_i is the temperature during the i th time interval. In the absence of process variation, the shape parameter is constant (we relax this assumption in Section 5.1). Therefore, temperature fluctuation only changes the scale parameter, which is η for Weibull distributions and μ for lognormal distributions. For each time interval, the scale parameter is calculated as follows:

$$\eta_i = \frac{MTTF_i}{\Gamma \left(1 + \frac{1}{\beta} \right)}, \quad (7)$$

where $MTTF_i$ is the calculated MTTF for temperature T_i , β is the Weibull shape parameter, and Γ is the gamma function. The change in scale parameter between time intervals causes discontinuity in the temporal failure probability density function. Therefore, for two consecutive time segments t_i and t_{i+1} , when the device temperature changes from T_i to T_{i+1} , the duration of the previous time segment should be adjusted according to the new distribution parameters. At time t when the device temperature changes, the following equation must hold:

$$R_i(t) = R_{i+1}(t'), \quad (8)$$

where $R_i(t)$ is the reliability at time t for scale parameter η_i , and t' is defined as the time resulting in wear equivalent to that at t if the scale parameter were changed to η_{i+1} . Assuming that the reliability at time $t_0 = \sum_{k=1}^{i-1} t_k$ is R , since Equation 8 must be satisfied, at time $t = t_0 + t_i$ we have

$$e^{-\left(-\ln(R)\right)^{\frac{1}{\beta}} + \frac{t_i}{\eta_i}} = e^{-\left(-\ln(R)\right)^{\frac{1}{\beta}} + \frac{t'_i}{\eta_{i+1}}}. \quad (9)$$

By solving Equation 9, we can express t'_i as a function of t_i :

$$t'_i = \frac{\eta_{i+1}}{\eta_i} t_i. \quad (10)$$

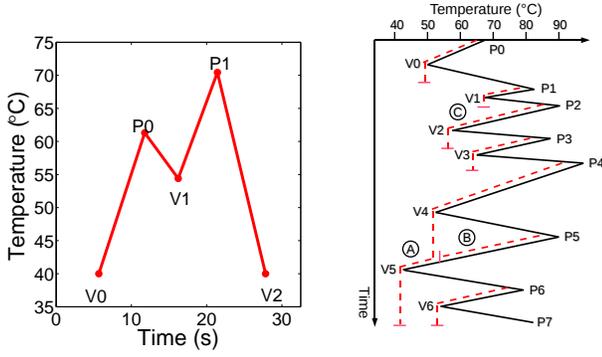


Figure 2: Example (a) temperature time series and (b) application of rainflow counting method.

By applying Equation 10 to the entire temperature time series, the reliability at time $t = \sum_{i=1}^j t_i$ can be calculated using the following equation:

$$R(t) = e^{-\left(\sum_{i=1}^j \frac{t_i}{\eta_i}\right)^\beta}. \quad (11)$$

Given stationary temperature and voltage characteristics, a sequence of sufficient length may be used to represent typical wear conditions. Since the sequence is short relative to the total lifespans of devices, which is typically tens of years, we can calculate the approximate reliability for devices with Weibull temporal failure distributions using the following equation:

$$R_{Weibull}(t) = e^{-\left(\frac{\sum \frac{t_i}{\eta_i}}{\sum t_i} t\right)^\beta}. \quad (12)$$

The same analysis can be used for lognormal distributions.

$$R_{lognormal}(t) = \frac{1}{2} - \frac{1}{2} \operatorname{erf}\left(\frac{\ln(t) + \ln\left(\frac{\sum \frac{t_i}{\sigma^{\mu_i}}}{\sum t_i}\right)}{\sqrt{2\sigma^2}}\right), \quad (13)$$

where σ is the lognormal shape parameter and μ_i is the lognormal scale parameter for the i th time segment. μ_i can be calculated using the following equation:

$$\mu_i = \ln(MTTF_i) - \frac{\sigma^2}{2}. \quad (14)$$

Equations 12 and 13 allow the computation of device reliability given time-varying temperatures. However, they cannot be applied to the thermal cycling failure mechanism, which requires memory for accurate modeling.

4.2 Thermal Cycling Model

A thermal cycle occurs when the temperature starts from some initial value, reaches an extreme point (either a local minimum or a local maximum), and returns to the starting value. The thermal cycling MTTF depends strongly on the peak temperature and cycle amplitude. A cycle can be denoted by a valley-peak (or peak-valley) pair. The peak and valley that form a cycle need not be adjacent to each other.

Consider Figure 2(a), which shows a time series of temperatures in which the P labels denote peaks and V labels denote valleys. There are three cycles: $(V0, P0)$, $(V1, P1)$, and $(V0, P1)$. $(V1, P1)$ and $(V0, P1)$ overlap each other. Counting them both in a naïve way, i.e., double-counting, may lead to underestimation of MTTF. In contrast, if only neighboring valleys and peaks are considered, $(V0, P0)$ and $(V1, P1)$ will be counted. This would result in a $2.7\times$ overestimation

of MTTF due to neglecting the large cycle $(V0, P1)$. Proper estimation of wear due to variable-amplitude thermal cycling is complex.

In previous system-level reliability modeling work, it is typically assumed that the temperature varies slowly over time and has regular pattern that simplify cycle counting by eliminating overlapping cycles [3, 17]. However, in real applications, high-frequency and irregular temperature variation can occur. An accurate thermal cycling model should avoid double-counting cycles and appropriately count overlapping cycles.

4.3 Rainflow Counting Method

Wear due to thermal cycling appears to be primarily caused by inelastic deformation of materials due to temperature changes of connected materials with dissimilar coefficients of thermal expansion, i.e., thermal cycling produces repeated, potentially-inelastic deformation when connected materials expand and contract at different rates. This is analogous to the variable-amplitude stress loading problem [18]. Rainflow counting is the most widely used cycle counting method in material science and is commonly accepted as the best method for estimating the damage of a material due to random loading fluctuation [19].

An example of the rainflow counting algorithm is illustrated in Figure 2(b). Each peak serves as a source of water that flows down along the slope. When a water flow is terminated, a half cycle is counted. A water flow is terminated when it reaches the end of time, e.g., A ($V5, P4$); merges with another flow, e.g., B ($V4, P5$); or comes opposite a maximum more positive than the maximum from which it initiated, e.g., C ($V2, P2$) [20]. Two half cycles with identical amplitudes and peak temperatures but opposite directions form a full cycle. In our application it can be assumed that, although the thermal profile has arbitrarily varying amplitude in short term, its pattern will eventually be repeated given a long enough period, enabling the use of Downing's simplified rainflow counting algorithm [21]. Even if the interval does not repeat, as long as it contains many cycles and subsequent intervals will produce similar amounts of wear, the proposed modeling approach is valid.

The rainflow counting method satisfies the requirement we suggested in Section 4.2. Therefore, we will consider its result as the ground truth value.

The Coffin-Manson equation combined with Miner's rule [18] is used to calculate the resulting total number of cycles to failure for the entire temperature time series. According to Miner's rule, for the i th temperature swing for which the number of cycles to failure is N_i , the damage is $1/N_i$. Therefore, assuming N_{TC} is the mean number of cycles to failure for a given thermal profile that contains m cycles, and the duration of the time series is short compared to the device lifespan, N_{TC} can be approximated as follows:

$$N_{TC} = \frac{m}{\sum_{i=0}^m \frac{1}{N_i}}. \quad (15)$$

Thus, the thermal cycling MTTF can be calculated using the following equation:

$$MTTF_{TC} = \frac{N_{TC} \sum_{k=1}^m t_k}{m}, \quad (16)$$

where $\sum_{k=1}^m t_k$ is the total duration for the given temperature time series.

5. COMPONENT-LEVEL MODELING

The failure of any device within a component causes the component to fail. We assume the temperature variation

among devices within a component is negligible. Large components with internally varying temperature can be accurately modeled by treating them as multiple components. Devices are grouped into components to reduce the computational complexity compared to performing complete Monte Carlo simulation on every device in the IC. This section describes our approach to obtain component temporal failure distributions based on device-level information.

5.1 Weibull Device Temporal Failure Distributions with Process Variation

At component level, we consider the impact of intra-die process variation. Inter-die process variation is handled at the system level. In Section 4.1, a constant shape parameter β was assumed. However, this assumption does not hold in the presence of process variation. For example, the gate-oxide thickness may vary across the chip. It is reported that the Weibull shape parameter for TDDDB changes linearly with the gate-oxide thickness [14]; given the same defect generation rate, a thinner gate increases the failure probabilities of gates with few defects, thus increasing the variance of the probability density function, i.e., decreasing the shape parameter β . If we assume that the shape parameter has a Gaussian distribution $\phi(\mu, \sigma^2)$, where μ is the mean and σ^2 is the variance, component reliability can be expressed as follows [22]:

$$R(t) = e^{-n \int_{-\infty}^{\infty} \phi(\mu, \sigma^2) \left(\frac{t}{\eta}\right)^{\mu} dx} = e^{-n \left(\frac{t}{\eta}\right)^{\mu - \frac{\sigma^2}{2}} \ln\left(\frac{t}{\eta}\right)} \quad (17)$$

Combining Equations 12 and 17 yields

$$R(t) = e^{-n \left(\frac{\sum \frac{t_i}{\eta_i} t}{\sum t_i}\right)^{\mu - \frac{\sigma^2}{2}} \ln\left(\frac{\sum \frac{t_i}{\eta_i} t}{\sum t_i}\right)} \quad (18)$$

Equation 18 allows us to compute the component reliability considering intra-die process variation. For thermal cycling, Equation 17 is sufficient, since we have already taken temporal temperature variation into account with rainflow cycle counting. Note that if the process variation distribution deviates significantly from the Gaussian distribution, Equation 17 no longer holds. In that case, Equation 17 should be re-derived or the approximation method described in Section 5.3 should be used.

5.2 Lognormal Device Temporal Failure Distributions with Process Variation

Given that the shape parameter μ of a lognormal temporal failure distribution has its own distribution due to process variation, the reliability of a component with n device can be calculated using the following equation:

$$R(t) = e^{n \int_{-\infty}^{\infty} f(x) \ln\left(\frac{1}{2} - \frac{1}{2} \operatorname{erf}\left(\frac{\ln(t) - \mu}{\sqrt{2x^2}}\right)\right) dx}, \quad (19)$$

where $f(x)$ is the probability density function of μ . This equation is too complex to be directly used for system-level reliability estimation. Simpler component-level reliability functions can be used to reduce the computational cost of system-level reliability estimation. Therefore, our goal is to find a distribution that can approximate the ground truth component distribution, which is the limiting distribution for the minimum lifetime of a collection of an arbitrary number of devices with identical and independent lifetime distributions.

Statisticians have previously considered a related problem. Gumbel (1958) showed that for any well-behaved ini-

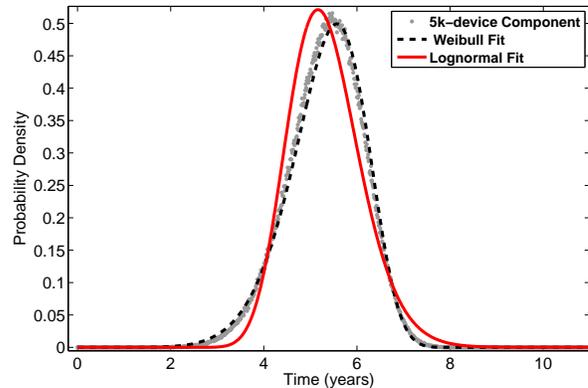


Figure 3: Weibull and lognormal fitting results for a 5,000-device component with process variation.

tial distribution (i.e., one that is continuous and has an inverse), only a few models are needed [23]. The relevant model depends on whether the extreme value of interest is the maximum or minimum and on whether observations are bounded from above or below. In the limit, this conclusion is independent of the distributions of individual components, as stated by the Extreme Value Theory [23]. Since finding the minimum is our objective, the Weibull distribution is appropriate.

Figure 3 shows the ground-truth probability density of a component containing 5,000 devices, each with a lognormal temporal failure distribution. The shape parameter μ is assumed to be randomly distributed in range [0.3, 0.7] [24] to model process variation. Lognormal and Weibull distributions are used to fit the ground-truth distribution. The results show that even under significant process variation, the Weibull distribution provides an accurate approximation. Note that the shape parameter of the fitted lognormal component distribution differs from that of the device distribution. The shape parameters for both the fitted lognormal and fitted Weibull distributions change with device count, as explained in Section 7.3. If a direct derivation is infeasible and the component contains many devices, it is appropriate to use a fitted Weibull distribution for the component.

5.3 Generalized Lognormal Extreme Value Distribution (GLEV)

The lognormal distribution is appropriate for modeling many device-level failure processes but is inaccurate for components containing many devices. The Weibull distribution is appropriate for modeling failure processes in components containing many devices but is decreasingly accurate as the number of devices decreases. It is often useful to model the reliabilities of small regions containing relatively few devices in order to ensure that environmental characteristics such as voltage and temperature are constant. Selecting one of the two distributions depending on the number of components would improve accuracy, but is not a perfect solution. For example, for a system containing 500 components, both lognormal and Weibull distributions have nearly 10% error in time to 1% failure (1-TTF).

We use a distribution that is capable of accurately modeling the impact of lognormal device failure processes in components containing an arbitrary number of devices, which we refer to as the generalized lognormal extreme value (GLEV)

distribution. The probability density function of the proposed distribution is given by the following equation:

$$\begin{aligned}
f_{GLEV}(t) &= w \cdot f_{Weibull}(t) + (1-w)f_{lognormal}(t) \\
&= w \left(\frac{\beta}{\eta} \left(\frac{t}{\eta} \right)^{\beta-1} e^{-\left(\frac{t}{\eta}\right)^\beta} \right) + \\
&\quad (1-w) \left(\frac{1}{t\sigma\sqrt{2\pi}} e^{-\frac{(\ln(t)-\mu)^2}{2\sigma^2}} \right), \quad (20)
\end{aligned}$$

where w , the GLEV parameter, is used to characterize the relative contribution of its corresponding distribution. The values of β and σ in Equation 20 are determined by the number of devices in the component via curve fitting. The values of η and μ are directly calculated based on the MTTF. w , which has range $[0,1]$, is also determined based on the number of devices. When the component has only one device, $w = 0$, the GLEV distribution is lognormal. When the number of devices approaches infinity, $w \rightarrow 1$ and the GLEV distribution approaches the Weibull distribution. It is possible that there are more accurate non-linear combination methods to model the component distribution but our analysis shows that this easy-to-use linear combination is highly accurate even for the most challenging components to model, i.e., those with a moderate number of components.

The cumulative distribution function of the GLEV distribution follows:

$$\begin{aligned}
F_{GLEV}(t) &= w \left(1 - e^{-\left(\frac{t}{\eta}\right)^\beta} \right) + \\
&\quad (1-w) \left(\frac{1}{2} + \frac{1}{2} \operatorname{erf} \left[\frac{\ln(t) - \mu}{\sqrt{2}\sigma} \right] \right). \quad (21)
\end{aligned}$$

Once a GLEV distribution for a component of interest is derived, it may be used in system reliability analysis.

6. SYSTEM LEVEL MODELING

A system contains multiple components, which need not be critical. A system may also contain hot or cold spares in an arbitrary structure [9]. Hot spares operate concurrently with the protected component(s); cold spares are used only after the protected component(s) fails. In this section, we will discuss methods to estimate static or dynamic system reliability.

6.1 Static System Reliability Estimation

Static reliability is the precomputed time-dependent survival probability of a system. It is an appropriate design-time optimization metric for systems with pre-planned, static responses to component failures. System-level hot and cold spares can make the direct calculation of system-level static reliability intractable. For example, assuming a system with hot spares, whenever a non-critical component fails, task migration and/or voltage and frequency scaling is performed to meet the performance requirements, which might change the thermal profile of the system. As a result, the probability density function of the components become discontinuous and depends on the accumulated wear when partial failure occurs. To directly calculate the system MTTF, we must traverse all states that would permit survival, which requires solving multi-convolution integrals. The situation is even more complicated when inter-die process variation is considered. Therefore, alternative methods are necessary.

If the number of system-level components is limited, it is possible to use Monte Carlo trials for reliability evaluation. We use a survival lattice to handle arbitrary system structures. Each state in the survival lattice represents a

unique set of operating components with which the system can still meet its requirements. States are associated with pre-computed thermal profiles to accelerate reliability analysis. With enough Monte Carlo trials, the temporal failure distribution for any structure of components can be determined. For failure mechanisms with Weibull device-level temporal failure distributions, a uniformly distributed random number u in range $[0, 1]$ is generated for each component. This number represents expected lifetime during system-level simulation. The expected time to failure for each component can be obtained using the following equation:

$$1 - e^{-n \left(\frac{\sum \frac{t_i}{\eta_i}}{\sum t_i} \right)^\mu - \frac{\sigma^2}{2} \ln \left(\frac{\sum \frac{t_i}{\eta_i}}{\sum t_i} \right)} = u. \quad (22)$$

There exist two solutions for t from Equation 22, only one of which is valid. Since u represents the failure probability, it should increase monotonically with t . Thus, the valid solution $t(u)$ satisfies $\frac{dt(u)}{du} > 0$. Therefore,

$$t = \frac{\sum t_i}{\sum \frac{t_i}{\eta_i}} e^{\frac{\mu - \sqrt{\mu^2 + 2\sigma \ln \left(-\frac{\ln(1-u)}{n} \right)}}{\sigma^2}}. \quad (23)$$

The same procedure can be applied to failure mechanisms with lognormal device-level temporal failure distributions. Components with multiple failure mechanisms are modeled as multiple components connected in series, each of which is subject to a single failure mechanism.

We take inter-die process variation into consideration during Monte Carlo simulation. We assume that the two parameters μ and σ in Equation 18 have Gaussian distributions with $\mu \sim \phi(\mu_1, \sigma_1^2)$ and $\sigma \sim \phi(\mu_2, \sigma_2^2)$ across all dies. To model this inter-die process variation, two normal random variables with means μ_1 and μ_2 and standard variations σ_1 and σ_2 are generated at the beginning of each trial and are used to replace μ and σ in Equation 23. The expected lifetime is also generated. The system-level failure time is determined using the survival lattice. The probability of system survival up to time t can be calculated as the percentage of trials for which system survives longer than t .

6.2 Dynamic System Reliability Estimation

Systems dynamically adapt to faults in ways that change the system thermal profile, e.g., task rescheduling and voltage and frequency scaling. Reliability up to time t refers to the overall survival probability in time interval $[0, t]$. However, on-line dynamic measurement-based adaptation techniques already know that the system has survived thus far with probability 1. Therefore, performing optimization based on the static reliability estimate would yield suboptimal solutions. Instead, conditional probabilities should be used. We propose to use the dynamic reliability of the system at time $t + t_{interval}$, where $t_{interval}$ is the time interval from t to the next time point when dynamic optimization can be performed, given the reliability numbers of its components at time t .

Our model calculates the dynamic reliability using the following equation:

$$R_{system}(t + t_{interval} | R_{C_1}, R_{C_2} \cdots R_{C_n}), \quad (24)$$

where R_{C_i} is the reliability of component i at time t . Direct calculation of the dynamic system reliability requires consideration of arbitrary system redundancy structures, for which Monte Carlo trials are used. In Monte Carlo simulation of dynamic systems, at the beginning of each trial the component ages are the estimated ages obtained in the previous

step. Expected lifetime is no longer uniformly distributed in range $[0, 1]$. Instead it is uniformly distributed in range $[F, 1]$, where F is component’s cumulative failure probability at the beginning of the Monte Carlo simulation.

7. RESULTS

To evaluate the proposed system-level reliability modeling framework, we have performed several simulation-based studies. Specifically, we compare the accuracy of the rainflow-based method used in device- and component-level reliability analysis with other cycle counting methods used in existing work. We also evaluate the accuracy of the reliability results obtained by applying our proposed failure distribution modeling technique to components with device counts ranging from 1 to 50,000. Finally, we evaluate the accuracy and efficiency of our multi-level reliability modeling technique. Before presenting these results, we first briefly describe the setup used in our simulation study.

7.1 Experimental Setup

Our test cases are based on MPSoCs consisting of 4 or 16 homogeneous cores. Each core is based on the Alpha 21264 processor, with a maximum power consumption of 120 W at 4 GHz [25]. All cores are DVFS-enabled and have 7 normalized discrete speed levels similar to those for the Intel Core Duo [26]: 0.462, 0.615, 0.692, 0.769, 0.846, 0.923, and 1. A total of 10 task sets consisting of 15 tasks each are randomly generated. The system load ranges from 60% to 80%. Periodic real-time tasks are used, each of which has a worst-case execution time, period, and deadline. Task periods are uniformly distributed between 50 s and 300 s. The worst-case execution time of a task is uniformly distributed between 30% and 60% of its period. Task deadlines and periods are equal. For each task set, tasks are statically assigned to cores using the largest-task-first algorithm [27]. The earliest-deadline-first (EDF) scheduling algorithm is used [28]. The actual execution time of each task is uniformly distributed in the range of 50% to 100% of the worst-case execution time of that task. The simulation is run for each task set (benchmark) for a duration of 86,400 s (1 day). Each simulation results in a power profile for each core. Dynamic thermal analysis was performed using HotSpot [29]. The generated thermal profiles were used as input to our reliability modeling framework (see Figure 1).

For device-level lognormal and Weibull distributions, the shape parameters are set to 0.5 and 5, respectively [14]. The reference device MTTF is assumed to be 120 years at 60 °C and the supply voltage is 1.8 V. Intra-die and inter-die process variation cause a total of 10% deviation from mean value for the shape parameters [30]. One-million trials are used for Monte Carlo simulation. After one million trials, increasing the number of trials by 1,000 consistently leads to less than 0.0005% improvement in the resulting MTTF errors.

7.2 Thermal Cycling Results

The mechanism used to count thermal cycles has a significant impact on MTTF estimation accuracy. Given that the rainflow counting method is a widely accepted means to count variable-loading cycles in the material science field, it is important to see whether applying it to the device- and component-level reliability models in ICs leads to MTTF values that are significantly different from those obtained by methods typically used in existing MPSoC reliability analysis and optimization work. Unfortunately, existing MPSoC reliability modeling publications provide little detail on the actual thermal cycling counting methods used, except that

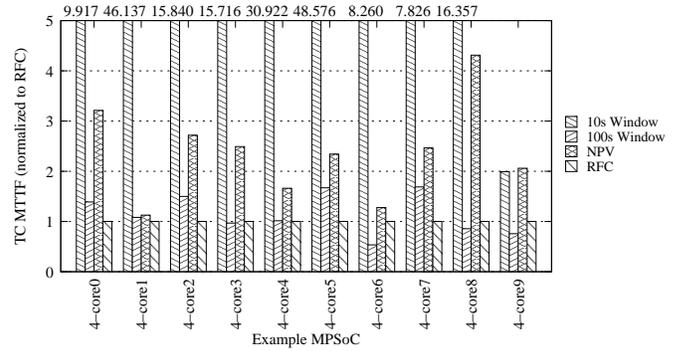


Figure 4: Cycle counting result for core 0 normalized to rainflow counting method.

they all assume fixed-amplitude thermal cycles. We therefore consider two counting methods: fixed-window size based counting and peak-valley pair based counting. Specifically, we have studied the following three setups.

1. *10s Window*: In each window of 10 s, set the thermal cycling range to the maximum temperature difference seen in that window. When computing the number of cycles, any cycle that contains the maximum temperature difference is counted.
2. *100s Window*: This counting method is the same as the *10s Window* except that the window size is increased to 100 s.
3. *NVP*: This method considers every neighboring peak and valley pair to be a cycle. It records both the peak-to-valley and valley-to-peak swings.

Since the accuracies of different counting methods depend on the pattern of the thermal profiles but are independent of the number of components in the system, we evaluate the counting methods on the 4-core MPSoC benchmark only. Each of the thermal cycle counting methods is applied to the thermal profiles generated by the 4-core benchmarks and the MTTF of each core due to thermal cycling is obtained. The thermal profiles contain dynamic and random high-frequency temporal temperature variations. The simulation results are shown in Figure 4. The 10 s Window overestimates thermal cycling MTTF for almost all of the benchmarks. Although a small window size allows the high-frequency cycles to be captured, it breaks down the large and slow cycles into small segments, thus underestimating the impact of the most important cycles on MTTF. The 100 s window size method counts the large global cycles, allowing an estimate closer to the rainflow counting method. However, neglecting high-frequency cycles still results in an average error of 46.01% and a maximum error of 268.10% relative to rainflow counting when all 4 cores are considered. The accuracy of the NVP method is between that of the 10 s Window method and 100 s Window method.

Note that the purpose of this evaluation is to show the impact of choosing different thermal cycling models on MTTF estimation. These results show that for thermal profiles that contain variable-amplitude thermal cycles, different counting methods yield significantly different results. Therefore choosing the appropriate counting method is important for thermal cycling MTTF estimation accuracy. Based on research in the material science community, rainflow counting appears to be the most appropriate cycle counting method.

Table 1: Comparisons of Different Failure Distributions

Component count	Lognormal ($\sigma = 0.5$)			Lognormal			Weibull			GLEV		
	σ	MSE	1-TTF (%)	σ	MSE	1-TTF (%)	β	MSE	1-TTF (%)	w	MSE	1-TTF (%)
1	0.500	1.90e-06	0.06	0.500	1.90e-06	0.06	2.367	4.17e-05	41.44	0.001	1.90e-06	0.06
10	0.500	4.48e-04	37.53	0.291	2.26e-05	10.28	3.927	4.53e-05	22.46	0.400	4.53e-06	5.14
50	0.500	1.37e-03	46.09	0.225	7.76e-05	13.07	5.038	4.09e-05	14.60	0.588	5.70e-06	5.58
100	0.500	1.91e-03	48.37	0.205	1.10e-04	13.67	5.487	3.97e-05	12.30	0.637	5.93e-06	5.10
500	0.500	3.48e-03	52.18	0.175	2.08e-04	13.68	6.419	3.91e-05	9.06	0.713	6.56e-06	4.30
1,000	0.500	4.30e-03	53.37	0.165	2.98e-04	13.57	6.801	3.44e-05	7.96	0.733	5.40e-06	3.78
5,000	0.500	6.58e-03	55.49	0.147	4.03e-04	13.52	7.626	3.99e-05	6.01	0.777	7.43e-06	2.93
10,000	0.500	7.73e-03	56.22	0.140	4.80e-04	13.43	7.967	3.94e-05	5.35	0.794	7.80e-06	2.64
50,000	0.500	1.08e-02	57.66	0.129	6.70e-04	12.93	8.680	4.16e-05	4.41	0.817	8.36e-06	2.20

7.3 Component-Level Approximated Distribution Evaluation

We characterize the accuracies of temporal failure distribution functions for components with device counts ranging from 1 to 50,000. We compare lognormal distribution with fixed $\sigma = 0.5$, lognormal distribution with fitted σ , Weibull, and the proposed GLEV distribution. The shape parameters of all the distributions are obtained using curve fitting. The GLEV parameter w is also obtained through curve fitting, starting from the fitted parameters of the corresponding lognormal and Weibull distributions. The GLEV distribution should therefore have at least as high accuracy as the lognormal and Weibull distributions.

The component distribution evaluation results are shown in Table 1. For all cases, the GLEV distribution gives the best result both in terms of mean square error (MSE) and time to one percent failure error (1-TTF). It reduces the MSE by $5\times$ compared to the minimum MSE of all the other distributions. Its 1-TTF error is consistently half that of the second-best distribution. The results also indicate that a fixed $\sigma = 0.5$ lognormal temporal failure density function results in substantial error, often 50%. This error increases with device count, making it difficult to compensate for. The fitted lognormal distribution is adequate for components with very few devices but inaccurate for components with many devices. In contrast, the Weibull distribution is accurate for components containing many devices but inaccurate for components with few devices. However, neither lognormal nor Weibull distributions work well for components with a moderate number of devices. The GLEV distribution is consistently accurate regardless of device count. The GLEV parameter w increases with number of devices, indicating an increased weight on the Weibull distribution component. This observation is consistent with the Extreme Value Theory and our previous analysis.

Figure 5 illustrates fitted lognormal, Weibull, GLEV, and ground truth curves for a 500-component system. For this example, the lognormal distribution has the greatest error. However, the relative qualities of lognormal and Weibull distributions would be reversed for a system with relatively few components. The GLEV distributions well approximates the ground truth distribution.

We also evaluate the accuracy of the Weibull distribution for components composed of lognormal-distribution devices with process variation. We assume that the shape parameters of the devices are randomly varied with σ uniformly distributed in the range $[0.3, 0.7]$ [24]. Table 2 shows the MSE and 1-TTF errors for fitted lognormal and Weibull distributions when used to model a system consisting of 1–

50,000 components with varied shape and scale parameters. The results indicate that when the number of devices in a component is large, a single Weibull distribution is sufficient, even in the presence of process variation. The parameters of the corresponding Weibull distribution can be obtained from Monte Carlo simulation and curve fitting. However, if the device count is moderate, using the GLEV distribution is necessary to avoid significant errors in early life.

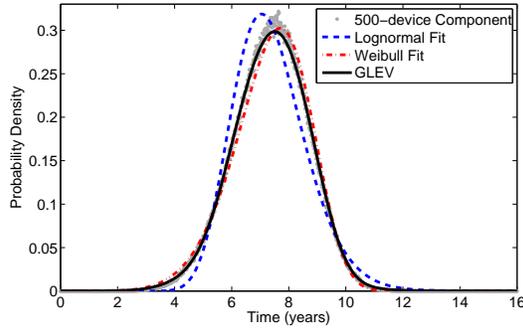
7.4 Comparison with Alternative Models

We evaluate the accuracy and speed of the multi-level system reliability model and compare it with other system-level models described in previous work. To make the comparison fair, we use the same rainflow counting method in all cases. Our objective is to compare the performance of different system-level modeling techniques, not that of the thermal cycling model, which was evaluated in Section 7.2. The following system-level reliability models are used for comparison.

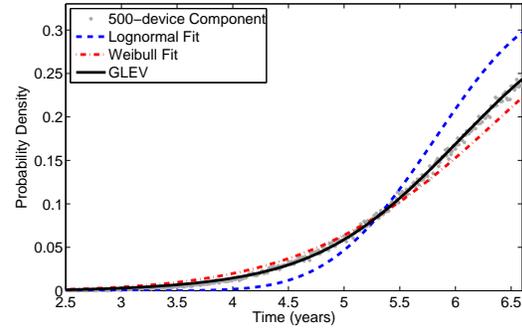
1. Exponential: This model assumes the exponential device-level failure distribution and uses technique called sum of failure rate (SOFR) to compute the total failure rate in a system without spares [3].
2. Log-fixed: This model uses a lognormal distribution for each component with shape parameter $\sigma = 0.5$. At system level it uses Monte Carlo simulation to estimate MTTF [5].
3. Log-simple: This model also assumes lognormal distributions at component level. But instead of performing Monte Carlo simulation, it uses a Min-Max approach to calculate the system MTTF [7].
4. Ground truth: This is the ground truth distribution obtained by performing Monte Carlo simulation on every individual device in the system.

In this experiment, we assume that each component contains 5,000 devices. Although larger device counts are also of interest, the computation time to generate the ground truth distribution by performing Monte Carlo simulation considering each device for the entire system is prohibitive. Note that this is only a concern when comparing reliability modeling techniques; in practice, the ground truth distribution would not be computed. Moreover, the quality of the proposed hierarchical model improves for components with larger device counts, as described in Section 7.3.

Figure 6 shows the ground truth system MTTF values and the normalized results for other system reliability models.



(a) PDFs for the complete lifetime.

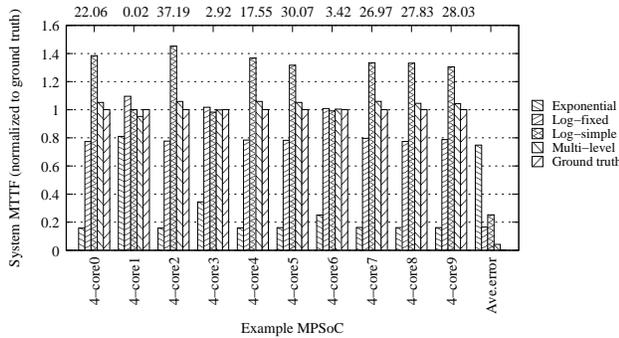


(b) PDFs for the early lifetime.

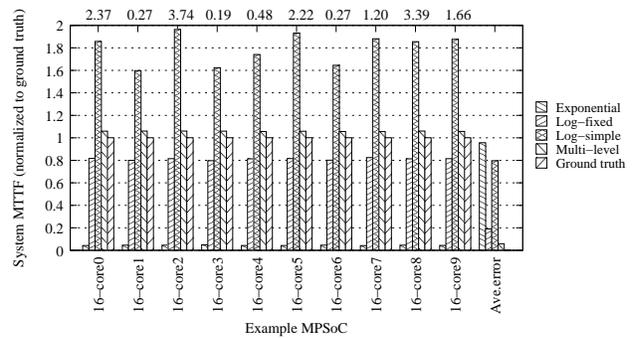
Figure 5: PDFs of the lognormal, Weibull, and GLEV distributions for 500 component system.

Table 2: Comparisons of Different Failure Distributions with Process Variation

Component count	Lognormal			Weibull			GLEV		
	σ	MSE	1-TTF (%)	β	MSE	1-TTF (%)	w	MSE	1-TTF (%)
1	0.364	9.90e-07	0.04	3.165	1.50e-05	34.99	0.002	9.90e-07	0.03
10	0.325	5.77e-05	20.58	3.565	1.45e-05	17.19	0.699	4.69e-06	8.92
50	0.296	1.05e-04	20.16	3.874	2.09e-05	15.68	0.713	4.68e-06	8.33
100	0.285	1.42e-04	21.38	4.028	2.08e-05	13.76	0.735	4.69e-06	7.18
500	0.248	2.44e-04	20.83	4.590	2.97e-05	10.88	0.760	5.93e-06	5.56
1,000	0.236	2.98e-04	20.40	4.787	3.44e-05	10.12	0.763	5.40e-06	5.14
5,000	0.214	4.96e-04	19.86	5.300	3.90e-05	8.19	0.797	7.18e-06	4.28
10,000	0.205	5.94e-04	19.61	5.519	4.37e-05	7.46	0.802	8.06e-06	3.79
50,000	0.188	8.70e-04	18.59	5.984	5.62e-05	6.49	0.812	1.02e-05	3.25



(a) System MTTF for the 4-core MPSoC benchmarks.



(b) System MTTF for the 16-core MPSoC benchmarks.

Figure 6: System-level MTTF comparisons. The MTTF value of the ground truth distribution (in years) is shown at the top. Results of other models are normalized to the ground truth value.

Some benchmarks have relatively small MTTFs due to large thermal cycle amplitudes and high peak temperatures. The proposed hierarchical model has the best performance of all the models for all the example problems. For the 4-core benchmarks, the average error is $3.93\times$ better than that of the second-best modeling approach and the maximum error is $3.85\times$ better. The average error for the hierarchical model is 4.2% and the maximum error is 5.9%. For the 16-core benchmarks, the average error is 5.8% and the maximum error is 6.0%. The average error is $3.26\times$ better than the best alternative. Overall, the average error is improved by $3.60\times$

compared to a technique using a fixed-parameter lognormal distribution for components. Note that the 16-core examples generally have a larger spatial temperature variation than the 4-core benchmarks.

Table 3 shows the results of comparing the running times of the models that use Monte Carlo simulation. The running time of the hierarchical model is related to the system component count and the size of the survival lattice. If the number of components is equal to the number of devices, the running time would degrade to that of the ground truth simulation.

Table 3: Run Time Comparison

Device count	Running Time (s)		
	Log-fixed	Multi-level	Ground truth
1	5.71	13.05	15.11
5	5.72	13.09	62.66
10	5.66	13.06	121.74
50	5.68	13.17	574.33
100	5.65	13.10	1137.60
500	5.72	13.15	5631.06
1,000	5.65	12.96	10963.70

8. CONCLUSION

In this work, we have presented a hierarchical MPSoC reliability modeling framework. We use direct calculation to construct component-level temporal failure models from device-level models and perform Monte Carlo simulation to determine system-level reliability. Our model considers intra-die and inter-die process variation at component and system level. We also implement the rainflow cycle counting algorithm, which enables accurate modeling of variable-amplitude thermal cycling. Experimental results show that the hierarchical model improves on the accuracy of models based on simplifying assumptions about thermal cycling and temporal failure distributions, resulting in 3.60× reduction in average error compared to the second best alternative. The proposed hierarchical model has only 5% average error in MTTF compared to ground truth distribution but runs hundreds of times faster for moderate-size components.

9. REFERENCES

- [1] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers, "Lifetime reliability: toward an architectural solution," *IEEE Micro*, vol. 25, no. 3, pp. 70–80, May 2005.
- [2] R. Viswanath, V. Wakharkar, A. Watwe, and V. Lebonheur, "Thermal performance challenges from silicon to systems," *Intel Technology J.*, vol. 4, no. 3, pp. 1–16, Aug. 2000.
- [3] A. K. Coskun, T. S. Rosing, K. Mihic, G. D. Micheli, and Y. Leblebici, "Analysis and optimization of MPSoC reliability," *J. Low Power Electronics*, vol. 2, no. 1, pp. 56–69, Apr. 2006.
- [4] T. S. Rosing, K. Mihic, and G. D. Micheli, "Power and reliability management of SoCs," *IEEE Trans. VLSI Systems*, vol. 15, no. 4, pp. 391–403, Apr. 2007.
- [5] J. Srinivasan, S. Adve, P. Bose, and J. Rivers, "The impact of technology scaling on lifetime reliability," in *Proc. Int. Conf. Dependable Systems and Networks*, June 2004, pp. 177–186.
- [6] J. Srinivasan, S. V. Adve, P. Bose, J. Rivers, and C. K. Hu, "RAMP: a model for reliability aware microprocessor design," IBM Research Report, Tech. Rep., 2003.
- [7] Z. P. Gu, C. Zhu, L. Shang, and R. P. Dick, "Application-specific MPSoC reliability optimization," *IEEE Trans. VLSI Systems*, vol. 16, no. 5, pp. 603–608, May 2008.
- [8] E. Karl, D. Blaauw, D. Sylvester, and T. Mudge, "Reliability modeling and management in dynamic microprocessor-based systems," in *Proc. Design Automation Conf.*, July 2006, pp. 1057–1060.
- [9] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers, "Exploiting structural duplication for lifetime reliability enhancement," in *Proc. Int. Symp. Computer Architecture*, June 2005, pp. 520–531.
- [10] "Failure mechanisms and models for semiconductor devices," Joint Electron Device Engineering Council, Tech. Rep., Aug. 2003, JEP 122-B.
- [11] J. R. Black, "Electromigration—a brief survey and some recent results," *IEEE Trans. Electron Devices*, vol. 16, no. 4, pp. 338–347, Apr. 1969.
- [12] Y. Zhang, M. L. Dunn, K. Gall, J. W. Elam, and S. M. George, "Suppression of inelastic deformation of nanocoated thin film microstructures," *J. of Applied Physics*, vol. 95, no. 12, pp. 8216–8225, June 2004.
- [13] M. Ciappa, F. Carbognani, and W. Fichtner, "Lifetime prediction and design of reliability tests for high-power devices in automotive applications," *IEEE Trans. Device and Materials Reliability*, vol. 3, no. 4, pp. 191–196, Dec. 2003.
- [14] R. Degraeve, G. Groeseneken, R. Bellens, M. Depas, and H. Maes, "A consistent model for the thickness dependence of intrinsic breakdown in ultra-thin oxides," in *Proc. Int. Electron Devices Meeting*, Dec. 1995, pp. 863–866.
- [15] M. Gall, C. Capasso, D. Jawarani, R. Hernandez, H. Kawasaki, and P. S. Ho, "Statistical analysis of early failures in electromigration," *J. of Applied Physics*, vol. 90, no. 2, pp. 732–740, July 2001.
- [16] J. R. Lloyd and J. Kitchin, "The electromigration failure distribution: the fine-line case," *J. of Applied Physics*, vol. 69, no. 4, pp. 2117–2127, Feb. 1991.
- [17] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers, "The case for lifetime reliability-aware microprocessors," in *Proc. Int. Symp. Computer Architecture*, June 2004, pp. 276–287.
- [18] N. E. Dowling, *Mechanical Behavior of Materials*, 3rd ed. Pearson/Prentice Hall, NJ, 2007.
- [19] I. Rychlik, "Note on cycle counts in irregular loads," *Fatigue & Fracture of Engineering Materials & Structures*, vol. 16, no. 4, pp. 377–390, Apr. 1993.
- [20] I. Rychlik, "A new definition of the rainflow cycle counting method," *Int. J. of Fatigue*, vol. 9, no. 2, pp. 119–121, Jan. 1987.
- [21] S. D. Downing and D. F. Socie, "Simple rainflow counting algorithms," *Int. J. of Fatigue*, vol. 4, no. 1, pp. 31–40, Jan. 1983.
- [22] K. Chopra, C. Zhuo, D. Blaauw, and D. Sylvester, "A statistical approach for full-chip gate-oxide reliability analysis," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2008, pp. 698–705.
- [23] NIST, "Assessing product reliability, NIST/SEMATECH e-handbook of statistical methods," <http://www.itl.nist.gov/div898/handbook/>.
- [24] K. D. Lee and P. S. Ho, "Statistical study for electromigration reliability in dual-damascene Cu interconnects," *IEEE Trans. Device and Materials Reliability*, vol. 4, no. 2, pp. 237–245, June 2004.
- [25] R. Rao and S. Vrudhula, "Performance optimal processor throttling under thermal constraints," in *Int. Conf. on Compilers, Architecture, and Synthesis for Embedded Systems*, Oct. 2007, pp. 257–266.
- [26] "Intel Core Duo processor and Intel Core Solo processor on 65 nm process datasheet," <http://www.intel.com/design/mobile/datashts/309221.htm>.
- [27] C.-Y. Yang, J.-J. Chen, and T.-W. Kuo, "An approximation algorithm for energy-efficient scheduling on a chip multiprocessor," in *Proc. Design, Automation & Test in Europe Conf.*, Mar. 2005, pp. 468–473.
- [28] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *J. of the ACM*, vol. 20, no. 1, pp. 46–61, Jan. 1973.
- [29] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. Stan, "HotSpot: a compact thermal modeling methodology for early-stage VLSI design," *IEEE Trans. VLSI Systems*, vol. 14, no. 5, pp. 501–524, May 2006.
- [30] A. Agarwal, D. Blaauw, and V. Zolotov, "Statistical timing analysis for intra-die process variations with spatial correlations," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2003, pp. 900–907.