Dynamic reliability management based on resource-based EM modeling for multi-core microprocessors

Taeyoung Kim [a], Zao Liu [b], Sheldon X.-D. Tan [c,*]

[a] Department of Computer Science and Engineering, University of California, Riverside, CA 92521, USA
[b] Xilinx Inc., San Jose, CA 95124, USA
[c] Department of Electrical and Computer Engineering, University of California at Riverside, Riverside, CA 92521, USA

ARTICLE INFO

ABSTRACT

This article presents a new approach for system-level reliability management for multi/many core microprocessors. In the new approach, the electro-migration (EM) induced time to failure (TTF) at the system level is modeled as a resource, which is abstracted, from a recently proposed physics-based EM model, at the chip level. In this model, a single core can spend the TTF resources at different rates specified by the temperature and the related power consumption. As a result, the new resource-based EM model allows more flexible EM-reliability management for multi/many-core systems. As an application of the new model, we propose a novel task migration method to explicitly balance consumption of EM resources for all the cores. The new method aims at equalizing the probability of failure of each core, which will maximize the lifetime of the whole multi/many core system. To more efficiently regulate the lifetime of the multi/many core system, dynamic voltage and frequency scaling (by using different performance states (p-states), which can represent different operating voltages and frequencies) is further employed to compensate for the excessively consumed lifetime of all the cores when the chip is loaded with heavy tasks for a certain period of time. In this way, the TTF of all the cores could be compensated to meet the lifetime requirement, giving the multi/many core system more flexibility to handle heavy task assignment on demand. Experimental results on a 36-core processor platform show that the proposed task migration scheme could balance the lifetime consumption of all the cores, and maintain even-distributed TTF slacks across different cores, while the existing temperature-based task migration schemes lead to diverged TTF of cores. The results also show that the TTF consumption could be easily compensated for by switching to a low p-state, which could not be achieved if TTF consumption diverged across different cores.

1. Introduction

Long-term reliability is becoming a limiting constraint in high-performance microprocessor designs due to the high failure rates in deep submicron and nanoscale devices. The increase in failure rates is caused by high integration levels and higher power densities, which leads to excessive on-chip temperatures. The introduction of new materials, processes and devices, coupled with voltage scaling limitations and increasing power consumption will impose many new reliability challenges. The semiconductor industry faces the challenge of maintaining reliability despite the continued increase in the die size and a number of transistors and the constant scaling of the transistor size for performance [1]. Increasing transistor density and thus power density is causing higher temperatures on chip, resulting in failure accelera-tion. Scaling to smaller transistors increases failure rates by shrinking the thickness of dielectrics. This has led the International Technology Roadmap for Semiconductor (ITRS) to predict the onset of significant reliability problems in the future, and at a pace that has not been seen in the past [2].

Early works mainly focus on dynamic thermal management (DTM) techniques, which have been proposed in the past to keep the temperature to stay below a limit to avoid the temperature-sensitive long-term reliability problems [3–5]. Those techniques, which typically consist of dynamic voltage and frequency scaling (DVFS), task throttling and clock gating, were first developed for single-core microprocessors. Recently, these techniques have been extended for multi-core architectures and multiprocessor system-on-a-chips. They include frequency-control method [6], the combined DVFS and task migration methods

---

[7,8], the predictive control method [9,10], and task migration based methods [9,11–15]. However, all of those techniques fail to explicitly consider the reliability and lifetime directly. As we will show in the work, the thermal-only DVFS can lead to much shorter lifetime for the many-core processors than reliability-aware DVFS methods.

Some initial efforts have been carried out for system level reliability analysis for SoCs (system-on-a-chip). RAMP is the first architecture level tool for modeling the long-term processor reliability of microprocessors at the design stage [16]. The follow-up work by the same authors proposed a dynamic reliability management (DRM) concept by dynamic voltage and frequency scaling (DVFS) [17]. It showed that it was not sufficient to just manage the temperature or power from the reliability perspective. A method in Ref. [18] shows that the power/performance and reliability are intrinsically conflicting metrics and have strong interactions on SoC designs, and proposes a joint policy optimization method. Another dynamic reliability management method was proposed in Ref. [19], in which a simple PID based run-time control was applied to optimize the performance subject to the long-term reliability constraints. Recently, DVFS techniques considering negative bias temperature instability (NBTI) effects were proposed for microprocessors [20]. A supply voltage scheduling technique was proposed for optimizing energy subject to NBTI constraints [21]. Also, reliability resource managements with wear limiting techniques were proposed in Refs. [22–24]. These methods utilized the non-volatile memory behavior that short latency leads to short lifetimes and slow writes provides longer lifetime in non-volatile memory. However, there is no physics detail for their reliability model, and it does not consider the power grid analysis in microprocessors for long-term reliability. Thus, these methods can be only limited in memory technologies for effective wear-leveling.

Although some early efforts, the research on system or architecture-level reliability analysis and optimization is still in its early stage. For electromigration related reliability effects, these efforts are based on the simple semi-empirical Black's equation [25], enabling estimation of the mean time to failure (MTTF) of interconnect wires and simplified series of wires in the scope of the constant failure rate models [16,26].

$$\text{MTTF} = A j^{-n} e^{\frac{E_a}{kT}} \tag{1}$$

Here, j is the current density, k is the Boltzmann's constant; T is the absolute temperature; E_a is the EM activation energy. The main drawback of Black's model is that it cannot correctly predict the MTTF for wires stressed in quite different conditions as it is not physics-based. It also ignores the impacts from stress migrations in the form of existing residual stress. Also, the traditional pessimistic predictions derived with (1) are based on current assessments where an EM-induced failure rate of the individual segment is considered as a measure of EM induced reliability in the extreme end, an MTTF of the weakest segment is accepted as a measure for the chip lifetime. This results in very conservative EM-aware current density design rules for the leading-edge technology nodes. To mitigate the mentioned problems for full-chip EM assessment and signoff techniques, a novel approach and technique for physics-based EM assessment for VLSI circuits have been proposed recently [27–29]. The proposed technique mainly focuses on the power grid networks, which is more vulnerable to EM failure than the signal nets. It considers EM-induced degradation (IR-drop increase due to wire resistance changes) as a parametric failure instead of single wire failure. The new method can also consider multi-segment interconnect tree in practical power grid networks. Additionally, the model takes into account thermal impacts and other process-induced residual stresses, which is ignored in the traditional Black-Blech based approaches.

This article presents a new approach for system-level reliability management on multi/many-core microprocessors. The new system EM model is based on the recently proposed physics-based EM models for EM assessment at the chip level [27]. In this model, EM life time is treated as a resource, a single core can spend the TTF resources at different rates specified by the temperature and the related power con-

sumption. As a result, the new resource-based EM models allow more flexible EM-reliability management for multi/many-core systems. As an application for the new resource based reliability model, we propose a new run-time reliability management technique. It consists of two optimization knobs. The first one is a new task migration method, which explicitly balances consumption of EM resources for all the cores. The new method aims at the equal chance of failure of these cores, which will maximize the lifetime of the whole multi/many core system. To more efficiently regulate the lifetime of the multi/many core system, DVFS technique (by using different performance states (p-states), which can represent different operating voltages and frequencies) can be employed to reduce power of the systems to meet lifetime requirement of the cores. In this way, it allows compensation of the excessively consumed life time of all the cores when the chip is loaded with heavy tasks for a certain period of time and the new method essentially can have more flexibility to handle heavy task assignment on demand.

We remark that the proposed system level resource based reliability models and the online management methods can also be applied to other long term reliability effects such as TDDB (time dependent dielectric breakdown), NBTI (negative bias temperature instability) and hot carriers etc. as these reliability effects can also be modeled as consumable resources and their consumption rates are also very sensitive to temperature (thus power). As a result, those long-term effects can be considered at the same time if they all can be modeled as resources, which then can be easily managed and optimized at the system level. But in this paper, we only focus on the reliability management for EM effects as an application of the proposed resource models.

Experimental results on a 36-core processor platform show that the proposed task migration scheme can balance the lifetime consumption of all the cores, and maintain evenly-distributed EM slacks across different cores, while two existing temperature-based task migration schemes lead to diverged TTF consumption. The results also show that the TTF consumption could be easily compensated for by switching to a low performance state (p-state), which could not be achieved if TTF consumption diverged across different cores.

This paper is organized as follows. Section 2 reviews the reliability-aware task migration problem. Section 3 reviews the recently proposed physics-based EM models. Then we introduce the new resource based EM modeling at the system level and present the new dynamic reliability management technique based on the new resource base EM models in Section 4. Section 5 presents the results on a 36-core multi-core systems. Last, Section 6 concludes.

2. Problem formulation and reliability modeling

As long-term reliabilities and performance are intrinsically conflicting factors, one has to consider them jointly at system level optimization as shown in the existing works [16,18,19]. To consider and model the reliability effects due to various failure mechanisms, many existing works use the so-called sum of failure rate model (SOFR) [30] to compute MTTF of a whole system from its components. SOFR model consists of the competing risk model which estimates the failure rate of each component and the series model which estimates the failure rate of the system based on failure rates of each component. Such SOFR models only work when the following conditions are met: first, each failure mechanism proceeds independently (they do not affect each other). Second, the whole system fails when the first of its component fails. For EM-related reliability on a practical power grid network neither conditions are met. In this work, we try to mitigate those problems by using a new physics-based EM model and redundancy-aware analysis techniques to compute the TTF of a power grid for given current sources, supply voltages and temperatures.

For multi-core microprocessors, the optimization could be achieved through proper reliability management of resources and tasks. In this work, we treat MTTF as a reliability resource that could be consumed and controlled during task executions. For optimization