




Contents lists available at ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

A survey of cross-layer power-reliability tradeoffs in multi and many core systems-on-chip

Ahmed A. Eltawil^a, Michael Engel^b, Bibiche Geuskens^c, Amin Khajeh Djahromi^c, Fadi J. Kurdahi^{a,*}, Peter Marwedel^b, Smail Niar^d, Mazen A.R. Saghir^e

^a Center for Embedded Computer Systems, University of California, Irvine, CA, USA

^b Chair for Embedded Systems, Informatik 12, TU Dortmund, 44221 Dortmund, Germany

^c Intel Labs, Hillsboro, OR, USA

^d LAMIH - University of Valenciennes, ISTV2 UVHC, Campus Mont Houy 59313, Valenciennes Cedex 9, France

^e Texas A&M University at Qatar, Electrical and Computer Engineering Program, P.O. Box 23874, Doha, Qatar

ARTICLE INFO

Article history:

Available online xxx

Keywords:

Multi-core
Many-core
Power
Performance
Reliability
Cross-layer

ABSTRACT

As systems-on-chip increase in complexity, the underlying technology presents us with significant challenges due to increased power consumption as well as decreased reliability. Today, designers must consider building systems that achieve the requisite functionality and performance using components that may be unreliable. In order to do so, it is crucial to understand the close interplay between the different layers of a system: technology, platform, and application. This will enable the most general tradeoff exploration, reaping the most benefits in power, performance and reliability. This paper surveys various cross layer techniques and approaches for power, performance, and reliability tradeoffs are technology, circuit, architecture and application layers.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Multicore platforms are quickly becoming the platform of choice to implement complex Systems-on-Chips (SoCs). The transition to new process technologies has enabled significant on-chip device densities. The shrinking size of transistors has resulted in lower power consumption, thereby narrowing the power gap between programmable and ASIC approaches. In spite of these advantages, many challenges still remain in the design and implementation of specific applications onto multicore systems. One can identify three main challenges facing multicore designers. The first and foremost is power consumption which is on the rise due to the complex algorithms executing on these platforms that demand both a heavy use of computational resources, as well as a large volume of memory and communication. The second challenge is technology related, where scaling is both an enabler and a limiter: it enables unprecedented integration, including the ability to integrate large memories on chip, with the downside being a penalty in leakage power as well as reliability. Finally, the third challenge is cost, driven by a highly competitive marketplace that demands the smallest die size possible. Thus SoC designers are faced with the daunting dilemma of generating high yielding architectures

that integrate vast amounts of logic and memories in a minimum die size with minimum power consumption.

In its current definition, yield indicates a 100% defect free chip, where circuits such as built-in self-test and built-in self-repair are used extensively to guarantee a high yield. Many traditional design approaches have focused on error-free design, and there has been significant research in attempting to guarantee error-free design. However, the International Technology Roadmap for Semiconductors (ITRS) trends clearly show that it becomes economically impractical to insist on a 100% error-free SoC in terms of area and power [1]. Thus, there is a critical need for a radically new approach to designing reliable multicore systems using inherently unreliable components: this approach must necessarily expand the design space across abstraction layers and cross-couple constraints across the circuit, architectural platform and the application abstractions.

Towards that end, one can broadly classify systems in two major categories:

1. Applications that are *inherently error tolerant* such as communications, multimedia and wireless which provide an opportunity to generate a range of acceptable designs for varying amounts of error in the system. For instance, communication and wireless systems have a high level of redundancy introduced at the system level, allowing for a tradeoff between attributes such as bit-error rate (BER) and signal-to-noise ratio (SNR). By

* Corresponding author. Tel.: +1 949 824 8104.

E-mail address: kurdahi@uci.edu (F.J. Kurdahi).

removing the artificial barrier between the system level design and the circuit level implementation, designers can explore an entirely new design space (as shown in Fig. 1) where controlled hardware errors can be treated in a similar manner as “channel errors” thus contributing to the noise floor while still meeting stated system metrics. This scenario presents the most opportunity for innovation by actively exploiting errors across abstraction levels, e.g., aggressive voltage scaling may introduce errors at the circuit/hardware level, but these errors are made visible to, and handled at the system level.

- Applications that are *stringently error-constrained*, where the error must be detected and corrected at a cost in terms of latency and performance. For instance, consider the cache of a processor in a multicore system: due to process variations, circuit-level techniques that enhance memory performance may result in errors that necessitate changes in the architecture of the circuit to both detect and correct the errors. This approach is in effect changing the statistics of the underlying error mechanisms. Such applications require the design of highly optimized hardware that utilize parallel architectures or time sharing to detect and correct for errors as well as microarchitecture approaches such as hardware shadowing or redundancy.

Thus the ability of the system to handle errors is highly dependent on the statistics of the errors and also on the algorithm running on the hardware, which implies that this has to be a dynamic process, *optimized at design time and managed during run-time*. To be able to extract the most benefit out of this error aware approach, it is important to examine the relationship between (a) the constituent components of an architecture and their vulnerability in terms of power consumption and reliability as a function of the operating conditions, and (b) the needs, assumptions and requirements of the application layers depending on this architecture. The intricate interaction between different controlling mechanisms and their benefits and costs creates an opportunity for finding a global optimum in terms of performance spanning across multiple levels of the design hierarchy.

The ITRS roadmap [1] indicates that embedded memories will dominate the die area in the near future, rising from 71% now, to close to 94% by 2014. The increasing market demand for having larger size memories on chip has flagged the power consumption of the SRAM/Cache as the major portion of chip power consumption. For this reason, we focus many (but not all) of our investigations in this paper on SRAM-related techniques for exploring power-performance-reliability design space exploration.

The remainder of the paper is organized as follows: Section 2 examines the technology layer and the generic concept of variability, Section 3 examines the platform layer at the hardware and micro architectural level, while Section 4 considers the software and compilation perspective. Section 5 considers the application layer. Conclusions are drawn in Section 6.

2. Technology scaling and challenges

Over the past few decades technology scaling has continued to follow Moore’s Law. As this pursuit continues for technologies beyond 22 nm, the decrease in feature size (Fig. 2, [2]) has supported ever increasing on-chip device densities. In order to reap the full benefits of technology scaling, a variety of challenges needs to be managed: increasing process variations, transistor aging variations and exponentially increasing leakage currents. As transistors continue to shrink in size, the limits imposed by leakage currents on transistor threshold voltage make it difficult to continue reducing transistor supply voltage as required by Dennard’s scaling rules [3]. Coupled with increasing transistor counts due to Moore’s Law, the resulting increase in power density has come to be known as the power wall, and is a prime reason for reliability problems that are directly related to increasing die temperatures. These include electromigration, stress migration, electron tunneling, gate-oxide breakdown, time-dependent dielectric breakdown, and thermal cycling, which can all lead to permanent, catastrophic, hard failures. Low transistor supply voltages and noise margins also lead to timing violations and increase the susceptibility of sequential circuit elements and memories to single-event upsets (SEU’s), which are due to atmospheric neutrons and alpha particles. SEU’s flip the values of stored bits but do not otherwise cause permanent damage. Nonetheless, these transient, soft errors can result in logical and timing errors.

To ensure correct functionality, designers will need to rely on careful co-optimization of process, circuit and layout techniques to meet ever challenging performance and power targets. Traditionally, designs have built fixed margins into operating frequency and voltage to ensure error-free operation under worst-case conditions in the presence of variation. This worst-case approach does not consider circuit behavior or implementation details and hence decreases design efficiency. Accurate statistical modeling of variation and its inclusion into timing and power convergence is necessary to recover design circuit margin while preserving pessimism to ensure quality and yield [4].

2.1. Variation sources

The CMOS variation sources can be classified into two groups: historical and emerging variations [5]. Historical variation sources include patterning proximity effects, line-edge and line-width roughness, polish variation, gate oxide thickness variation, fixed charge, defects and traps. These sources continue to require innovative solutions for each subsequent technology node. The emerging variation resources used to have a minor impact, but now present major challenges. Chief among them are random dopant fluctuation, implant and anneal variation, variation associated with strain and gate material granularity.




Fig. 1. (a) Layered system design, (b) traditional power-delay design space, (c) emerging powerdelay-(un)reliability design space.

Download English Version:

<https://daneshyari.com/en/article/10343569>

Download Persian Version:

<https://daneshyari.com/article/10343569>

[Daneshyari.com](https://daneshyari.com)