

Cross Layer Resiliency in Real World

Organizer: Vikas Chandra (ARM US)

Presenters: Subhasish Mitra (Stanford University) Vikas Chandra (ARM US) Chen-Yong Cher (IBM US), Silvia Melitta Mueller (IBM Boeblingen)

Abstract:

Resilience at different design hierarchies will be needed in Complex SoCs to handle failures due to variability, reliability and design errors (logical or electrical). The main reasons for the marginal behavior are sheer design complexity, uncertainties in manufacturing processes, temporal variability and operating conditions. In this session, we will cover the basics of cross layer resiliency and explore the reliability challenges in both embedded processors as well as large scale computing resources.

*Talk 1: Cross-layer resiliency exploration and optimization
(Subhasish Mitra)*

This talk will discuss systematic methodologies for exploring cross-layer resilience, encompassing error detection, correction and recovery techniques, for complex SoCs consisting of processor cores accelerators, and uncore components. The objective is to answer several key questions such as: (i) Given a design, is cross-layer resilience always the best way to protect that design? While cross-layer resilience is generally considered as the right way forward, systematic quantification of cross-layer resilience benefits is required. (ii) What are the right models that link resilience techniques across multiple layers for quick, yet accurate, estimation of costs in terms of power/energy, performance, and area? (iii) What is the proper framework to explore the large space of existing resilience techniques for error detection, correction, and recovery across various abstraction layers?

*Talk 2: Reliability challenges in Embedded processors
(Vikas Chandra)*

Embedded processors are now at the heart of the mobile revolution and have the aspirations to power even high performance data centers. It is of utmost importance to understand the reliability challenges in embedded processors and find ways to tackle them across different layers of design abstraction. In this talk, I will talk about the reliability requirements in embedded processors, the challenges we are facing and our approach to make the design more robust. We will discuss our approaches of measuring wearout in commercial processors as well as efficient design of in-situ monitors to track timing errors.

*Talk 3: Billion chips of trillion transistors: How to make them reliable?
(Chen-Yong Cher, Silvia Melitta Mueller)*

Due to increasing demand for personal devices, high performance computing systems and commercial data centers, microprocessor and main memory designers face numerous challenges in delivering large number of chips at effective cost. While frequency scaling effectively ended, technology scaling continues to provide increasing number of transistors. To effectively utilize these transistors for performance, designers turn to sophisticated and highly integrated chip designs such as multi-core (e.g., Intel i7, IBM POWER7, BlueGene/Q), GPGPU (e.g., NVIDIA Tegra) heterogeneous SoC (e.g., IBM Wirespeed). The increasing demand for chips and transistors presents numerous challenges on reliability, power and manufacturing costs. In large scale HPC systems and data centers, the increasing number of chips also raises per-chip reliability requirement in order to achieve system reliability targets.