

HCSS: Designing Reliable Embedded Systems atop of Unreliable Hardware Platform

Yuan Xie (The Pennsylvania State University)

1 The challenges

The advent of more sophisticated embedded systems (such as those in automotive systems) that support more powerful functions, and the reliance on deep sub-micron process technologies for their fabrication, have brought reliability concerns to the forefront. The major three important categories of reliability concerns facing the design of embedded systems are: (1) *Variability* in fabricated circuit primitive parameters, (2) *Signal integrity* issues arising from internal and external noise sources, (3) *Accelerated aging* of the devices. While these concerns also affect other computing system types, embedded systems pose unique challenges.

- First, most of the embedded systems are cost sensitive and often work with limited resources such as smaller memory size or diskless designs. These constraints can make it difficult to apply many traditional computer design methodologies for improving embedded-system reliability. For example, executing multiple redundant versions of the same thread to ensure reliable operation, while common in server environments, can be quite costly in embedded systems.
- Second, providing reliable embedded systems operation while satisfying other stringent constraints such as power consumption and real-time throughput is essential. Consequently, reliability-oriented designs that rely on aggressive redundancy, such as triple-module redundancy, might not be viable from an energy viewpoint. Similarly, approaches such as checkpointing and rollback, used in many server environments, might have no relevance in real-time embedded applications if the recovery happens after a deadline.
- Third, embedded systems often have reduced noise margins thanks to aggressive power optimizations or deployment in harsh environments. For example, the use of sub-threshold circuits in ultralow power environments such as sensor nodes significantly lowers the amount of external charge needed to upset a circuit node. In addition to hardware reliability concerns, the increasing amount of software content in embedded systems also poses a major challenge. Many documented cases of embedded systems failures have been ascribed to software malfunction, including the recent incident of a hybrid car stalling suddenly at highway speeds.

2 Three important reliability factors

This section describes the major three reliability concerns for designing dependable embedded systems on top of unreliable hardware components.

- **Variability.** The challenges in fabricating small feature size transistors have resulted in significant variation in transistor parameters such as channel length, gate-oxide thickness, and threshold voltage across identically designed neighboring transistors and across different identically designed chips. Although designing for worst-case process margins has been used as a traditional option, the degree of variability encountered in the new process technologies and the cost sensitivity make designing for the worst case unacceptable for embedded systems. Addressing these concerns will require beyond traditional memory fault-tolerance techniques. Further, operating the entire memory based on the worst-case process-variation delay would be overly pessimistic and cause performance loss. Consequently, developers need new techniques that can recognize and adapt to the resulting variation. For example, letting the software map latency-critical data to memory portions with shorter access times, while mapping the rest to portions that have higher access times, might possibly leverage the inherent tolerance to slower memory operations for certain data in an application. Developers also need adaptive techniques that can permit the design to tune its frequency based on the encountered process variation.

- **Transient errors.** Transient faults are errors caused by temporary conditions on the chip (such as when power supply noise or interconnect noise exceed a certain threshold) or by external noise (such as soft errors caused by neutron striking the chip). The circuit itself is not damaged even though computational errors are introduced. As supply voltages diminish and feature sizes become smaller in future technologies, soft-error tolerance will become a significant challenge when designing future embedded systems. Further, energy constraints could force embedded systems to incorporate aggressive power optimizations and deploy techniques such as dynamic voltage scaling, which can reduce the amount of critical external charge required to upset the value stored at the circuit node. Additionally, with deeper pipelining, the number of logic stages between latches becomes smaller, increasing the probability of a single-event upset making its way into a latch. Soft errors in memory can be protected by error correcting code, but cheap solutions for logic and latch soft errors are not well understood yet. Approaches such as checkpointing and rollback may not be viable for embedded applications because they might violate real-time constraints.
- **Accelerated aging effect.** Embedded system designs using new process technologies cause higher on-chip temperatures, which result from higher power densities. This significantly accelerates various failure mechanisms, including electromigration, stress migration, time-dependent dielectric breakdown, and thermal cycling, leading to an overall decrease in reliability. Accelerated aging creates a serious risk that devices will fail within an embedded system's warranty period, which poses the key challenge of finding countermeasures that effectively prolong a system's lifetime.

3 Summary

As technology scales, the underlying hardware platform for embedded systems becomes unreliable, due to three major reliability concerns, including *variability*, *transient errors*, and *aging effects*. Designing a dependable embedded system atop a less reliable hardware platform poses great challenges for designers. Cost and energy sensitivity, as well as real-time constraints, make some fault-tolerant techniques unviable for embedded system design. Many techniques to improve reliability can incur performance, energy, or cost penalties. Further, some solutions targeted at a specific failure mechanism could negatively affect other mechanisms. For example, lowering operational voltage can help mitigate thermal problems but increases vulnerability to soft errors. Developers must understand the tradeoffs when designing reliable embedded systems. *Both hardware and software approaches* must be explored to provide a dependable system design. For example, hardware can protect and even correct transient faults at the cost of redundant circuits. Software approaches can also protect/correct these faults, e.g., by instruction duplication or algorithmic design, or even by choosing fault-resilient data representations.

4 Biography

Prof. Yuan Xie is an Assistant Professor in the Computer Science and Engineering Department at the Pennsylvania State University. He received Ph.D. degree from Electrical Engineering Department, Princeton University. Prior to joining Penn State in Fall 2003, he was working for IBM Microelectronics Division's Worldwide Design Center. Dr. Yuan Xie's research interests include Embedded Systems Design, VLSI Design, Computer Architecture. Dr. Xie won the Semiconductor Research Corporation's Inventor Recognition Award in 2002. He has served as the tutorial chair for EMSOFT 2005, and as a TPC member in CASES 2006. He has presented various tutorials on reliable system designs in ASPLOS 2004, ASP-DAC 2005, ISCA 2005, and VLSI/CAD 2006. More information can be found at <http://www.cse.psu.edu/~yuanxie/>