Safety-Critical Systems Reliability

Mark Porter, Medtronic
Glenn Forman, General Electric
Kevin Kemp, Freescale
Claude Moughanni, Freescale
Tony Reipold, Freescale
Xiaowei Zhu, Texas Instruments

As electronic systems become more pervasive in society, the complexity of hardware and software used in applications where malfunctions could cause serious injury or death is also growing. The desire to construct such systems stems from the significant benefits they provide. Implantable pacemakers and defibrillators provide life-support to people with heart conditions that might otherwise prove to be fatal; airbag deployment and traction-stability control in automobiles improve the chance of survival in dangerous accident situations; beam control and placement in radiotherapy suites allow doctors to destroy malignant tumors while minimizing damage to healthy tissue; diagnostic imagery enables the identification and treatment of diseases and injuries that significantly increase beneficial medical outcomes; air transportation is optimized for millions of passengers a day through advanced air traffic control.

Reliability research has become more challenging due to increased system complexity, reduced development timelines, smaller feature sizes, and the demand for new product availability. Accelerated testing can fail to match real-world experience due to component and system miniaturization, and may lead to inferior reliability models. More predictive methods and more resilient architectures are needed.

Life/safety-critical systems that employ next generation computing devices will require increasingly productive and accessible reliability tools, models, and data. In addition, deep-sub-micron CMOS scaling is driving the need for a more advanced understanding of emerging pathogenic IC mechanisms. Reliable computing in the presence of these failure mechanisms will require an affordable, comprehensive immunization against fault propagation with the use of novel cross-layer architectural solutions.

Resilient systems for safety-critical applications must incorporate the benefit of cross-layer improvements for next-generation and future-generation electronic IC components, as well as packaging and interconnections to achieve a total system solution. In many instances, safety-critical designs tend to lag behind the state of the art in IC technology, however, more advanced medical life sustaining, prosthetic limb/vision, and molecular imaging equipment requires the best and highest performance computing available.

This research aims to change the nature of the host response to random and systemic faults, and to better understand the pathogenic role of known and emerging causes. This includes increasing transistor variation due to scaling, soft-error vulnerability, and permanent faults due to device wear-out.

In aerospace and defense, certain device designs have addressed upsets by such methods as hardened-by-design, redundancy, or hot-spares. However, the increased

power, area, and cost of such brute force defenses are counter to economic, commercial, and performance forces that represent the computing industry's more substantial and fundamental market drivers. Such solutions are also less suitable to implantable devices and portable/wearable medical instruments due to their larger size, weight, and power. Ultimately, our effort to address pathogens by newly formed methods of cross-layer immunization will be a balance between costs and benefits. By using cross-layer mechanisms, more affordable and relevant life/safety-critical solutions will ultimately be discovered.

Implantable Medical Devices

The first implantable pacemaker was developed in the late 1950s using solid-state transistors that allowed the devices to be small enough to run on battery power. Over the decades since that time, improvements in available technology to design and build these devices has allowed an ever wider array of therapies to be delivered to patients to treat many additional forms of illness.

The medical device industry has used a deliberate strategy of lagging the leading edge, especially of semiconductor technology. As transistor dimensions have shrunk along the Moore's Law curve, medical devices continued to use geometries that were several generations behind the high performance computer industry. We expect this conservative design principal to continue for two main reasons:

1. Leading edge technologies undergo a yield learning curve that requires a period of time to elapse before defect density has been substantially reduced. Reliability also suffers during this early period due to a higher number of latent defects.
2. As transistor geometries have shrunk to below 0.25μm, off-state leakage currents have also grown. In sub-65nm technologies, gate leakage has become such a significant problem for commercial applications that new materials have been introduced to counter the effect. In the medical device industry, where battery lifetimes are required to meet 10-year longevity, the current drain of advanced CMOS technologies is unsustainable.

Countering this trend is the desire to provide both higher reliability and additional performance. As the complexity of closed loop feedback systems and diagnostic data sets grows, there is a need for advanced technologies to facilitate the boost in performance, without a corresponding increase in energy usage. As part of the solution to address this application space, the need to ensure high reliability among all the components of the system will require new paradigms in design, test, and resiliency.

Automotive Electronics

The last several decades have seen explosive growth in the application of semiconductor electronics in automobiles, already exceeding 70 microcontrollers on some high-end vehicles. Some of these include: engine and power train control to maximize performance and fuel efficiency; driver and passenger information, comfort and entertainment systems; active safety systems such as airbag deployment and seat belt pre-tension devices; and vehicle dynamic safety and control including anti-lock brakes, traction control and electronic stability programs. Newer safety-oriented

applications include adaptive driver assistance systems using radar, active cruise control, lane departure warning and night-vision systems, while future concepts may include electric drive-by-wire steering and braking. The ultimate paradigm for automotive transportation is that of fully autonomous vehicles operating on an intelligent highway with no need for a human driver at all.

Driven by consumer quality expectations, warranty service and recall costs, and their increased use in safety/life-critical applications, automotive electronics are required to meet unprecedented levels of reliability and resiliency against potential failure mechanisms. In addition, these systems are required to operate for 10-20 years in harsh environments including extreme thermal, moisture, vibration and electromagnetic noise conditions.

Challenges

Critical to the continued performance and reliability improvement of safety-critical systems is the ability to abstract the defects and errors that can occur during design and manufacturing into a set of rules that can be used to guide resilient architecture choices.

1. Fault-tolerance/resiliency design choices are difficult to implement in a timely and efficient manner. This reduces both the number and robustness of the candidates and limits tradeoff options. Only the simplest or most critical cases are generally addressed.
2. Error correlation between device physics and architectural effect is poorly characterized and categorized. It is extremely difficult for designers to understand how any particular failure, hard or transient, will expose weaknesses in a given design.
3. Sensors, discrete components, and passives are essential elements for most safety-critical systems and can limit the overall system reliability but are generally ignored in the literature of resilient design. Safety-critical systems must have visibility into defects of all components, not just CMOS, and be prepared to address their failure or misbehavior.
4. It is often unclear how or where to make tradeoffs between the electronic components of a safety-critical system and the application in which it is embedded, where other resiliency techniques may be available. For example, in an automotive environment there exist mechanical back-up systems that can take over for malfunctioning electronics.
5. Translation form concepts to practice is hindered by the lack of a standard model for describing, comparing, and composing resilience techniques. Single-layer resiliency techniques provide complete solutions only in very limited circumstances (e.g. memory ECC); multi-layer protocols built in an ad-hoc fashion are expensive and non-portable between applications or system architectures. Comparing design solutions, especially as proposed by multiple authors in the literature, is difficult or impossible without a complete understanding of all the nuances of every aspect of the design, rendering the information less useful to new architectural designs.

6. Exploring and characterizing resilient/fault-tolerant design techniques and their efficacy in simulation and/or hardware is time-consuming and expensive.

In addition to the need for research into the design resiliency challenges listed above, safety-critical systems will increasingly be required to conform to standards definitions that govern their design, performance validation, and maintenance (e.g. IEC 61508 or ISO26262).

7. It is currently difficult or impractical to integrate safety standard protocols and certifications into design flows. This challenge places a significant cost and effort burden on manufacturers that prevents many companies from adhering to standards that could increase consumer safety. This suggests the need to correlate standards with quantifiable increases in safety and to better align tools, design flows, and automation with standard specifications.