

# Key Reliability Challenges for Consumer Electronics<sup>1</sup>

Over the last several decades, the exponential rate of improvement in both transistor performance and device density known as Moore's Law has been widely lauded as the reason for the growth and success of the consumer electronics and computing industry. However, this interpretation is not quite correct, in that it neglects the fact that users do not see fabrication improvements directly, but instead see improvements in the performance and capabilities of the systems that engineers are able to design using improved fabrication technology. It is these improvements in system performance and capabilities that drive new applications, new products, and users' desire to upgrade their systems well before they physically fail.

This distinction may seem minor, but it motivates the key challenge facing computer electronics/computing over the next 10-20 years: the need to develop designs that tolerate increasing rates of device-level variation, unpredictability, and failures without devoting so much silicon area and power to tolerating these effects that the rate of improvement in user-visible performance over time decreases significantly. Put in a more qualitative fashion, each new "generation" of fabrication technology approximately doubles the number of transistors that can be built on a chip of a given size, but this increase in device density comes at the cost of increased variation and error rates. If too many of the new transistors that a given fabrication process provides must be devoted to tolerating the process' increases in variation and error rates, products built in that fabrication process will not deliver enough improvement in user-level performance or capabilities as compared to products built in the previous fabrication process to motivate consumers to purchase them. If this happens, the growth engine of the electronics/computing industry will stall, because the industry relies on the profits from products implemented in each fabrication technology to fund the development of the next generation.

While errors in computation have been an issue for consumer systems since at least the 1970s, when researchers began to study the rates of alpha particle-induced soft errors in DRAMs, two trends argue that designers will need to shift from the current model of error correction in consumer electronics, which applies individual correction mechanisms to the structures that see the highest error rates, to a more system-level model in which the entire system considers the possibility of errors and variation. First, increasing rates of errors and variation are making it increasingly difficult to deliver sufficient reliability through a collection of mechanisms that tolerate individual causes of errors, such as ECC bits on memory arrays, increasing both the number of mechanisms required by a design and the hardware cost of each mechanism.

Second, designs are increasingly becoming limited by a chip's power budget instead of by the number of transistors that can be fabricated in a given amount of chip area, motivating the desire to reduce the "guard bands" on a chip's power supply and clock rate. Current designs operate at power supply/clock rate combinations that are significantly lower than their peak capabilities in order to ensure that they will continue to have very low error rates even when operated under worst-case conditions and/or at the end of their product lifetimes. In contrast, designs that are able to detect and correct timing errors are often able to operate at significantly (30%) more efficient power supply/clock rate combinations when implemented in conventional CMOS. This

---

<sup>1</sup> "Consumer Electronics" are defined here as computing/electronic systems that don't fit into one of the categories "life-critical systems," "aerospace systems," "infrastructure," or "large-scale." The category is, obviously, vague. In particular, it's hard to draw an exact line between the largest "consumer" computing system and the smallest "large-scale" system.

power/performance benefit from reducing guard bands is expected to increase as fabrication processes scale, due to increasing device variation and sensitivity. Similarly, the efficiency advantages of error-tolerant designs become even higher when implemented in near-threshold-voltage CMOS, because of the higher performance variations seen in such designs.

Because of these and other constraints, we believe that consumer electronic/computing systems will need to adopt full-system approaches to reliability, in which multiple levels of the system stack collaborate to detect, tolerate, and adapt to errors and variation. Developing these approaches will require research at all levels of the system stack and vastly-increased communication and collaboration between researchers at different levels in the stack. To facilitate, guide, and support this research, we have identified the following high-level research focus areas for reliable consumer computing:

- 1) Models and abstractions for errors and variation: Current research tends to focus on solutions to specific physical causes of errors and variation (SEUs, NBTI, etc.), leading to a profusion of extremely-specialized techniques. Developing a small set of abstract categories of errors/variations and showing that a wide range of physical effects can be coerced into one or more of those categories would simplify system design and analysis. Further, it would encourage and support the creation of clean communication interfaces between layers in the system stack by abstracting away some of the less-important details of the physical causes of errors and variation.
- 2) A general framework for multi-level reliability/resilience: One of the difficulties facing researchers is the lack of a “standard” architecture/system stack for a resilient system into which they can insert new techniques. Similar to the way the availability of a “conventional” model of superscalar architecture and tools to model such architectures allowed computer architects to develop new micro-architectural mechanisms without having to re-implement entire processor designs, the availability of a framework for reliable systems and one or more exemplar designs would make it possible for researchers to focus on individual aspects of reliable system design and to compare their techniques to other approaches with some confidence that they are doing a fair comparison. Having such a framework/toolset would also increase smaller institutions’ ability to contribute to reliability research, by lowering the “barrier of entry” required to generate useful results.
- 3) Testing/verification strategies for reliable systems: Systems that tolerate/correct errors promise to increase fabrication yields by providing correct operation in the face of a small number of fabrication defects, but their very resilience makes it difficult to test them at fabrication time by increasing the number of logic paths that must be tested and by hiding defects. As resilient designs become commonplace, it will no longer be sufficient to merely determine whether or not a chip is functionally correct at fabrication time. Instead, it will be necessary to characterize both functional correctness and the amount of “safety margin” remaining in terms of the chip’s ability to tolerate in-field failures and errors before a chip is declared ready to ship, and it will be necessary to do so without significantly increasing test time and cost, which is already a significant issue in the electronics industry. Similarly, it will be necessary to develop in-field diagnostics and testing techniques that can monitor a system’s state as it ages in order to tolerate changes in circuit behavior and predict when a chip’s reliability will drop below the requirements of the system due to accumulated errors and variation.
- 4) Improved recovery/rollback mechanisms: While this is a subset of the general reliability framework, recovery and rollback in consumer-scale systems is one of the least-studied aspects of the reliability space, and will need significant attention in order

to avoid spending excessive circuitry and power handling infrequent events. In particular, it is likely that future systems will incorporate a hierarchy of recovery mechanisms with different costs and capabilities, such as pipeline squashing, checkpointing at different levels in the memory hierarchy, and infrequent checkpointing to non-volatile storage.

- 5) Lightweight detection: The costs of recovery and rollback are closely tied to the latency of a system's error detection mechanisms. Low-latency error detection significantly reduces the cost of recovering from errors by reducing the amount of work that must be "undone" to restore the system to a state before the error occurred. To maximize efficiency, future systems will require a variety of error-detection mechanisms that are optimized to minimize both error detection latency and overhead. These mechanisms should work with the recovery/rollback mechanism by determining not only what has happened when an error occurs but also bounding the amount of time that has passed since the error occurred, allowing the system to select a recovery/rollback mechanism that minimizes the cost of recovering from the error.
- 6) Interfaces and abstractions for reliable system-on-chip design: While microprocessor architectures receive a great deal of attention, more and more product designs are being done using a system-on-chip approach, and this trend is expected to continue as improvements in device density allow larger portions of a system to be integrated onto a single chip. In particular, systems designed by connecting multiple pre-designed circuit blocks through standardized interface architectures and a small amount of custom logic are becoming extremely common. As it becomes important to consider reliability at all points in the product spectrum, it will become necessary to develop standardized interfaces for reliability, error notification, retry, and reconfiguration in SOC designs.
- 7) Scalable approaches to and abstractions for reliability: Consumer electronics are extremely sensitive to the costs of providing reliability because they compete in a marketplace in which performance (or performance per unit power or cost) is critical, and errors are relatively rare. In contrast, other portions of the industry (aerospace, HPC, etc.) have higher error rates and greater error impact (a crashed airplane as compared to the need to reboot a laptop), but are less-sensitive to cost and overhead. Scalable reliability techniques, which allow system integrators to trade off overhead against system reliability, might make it possible to use the same designs in both consumer and high-reliability products, allowing the high-reliability portions of the industry to benefit from the sales volumes of the consumer electronics industry.