# Cross-layer Codesign for Resilient Hardware

Xinfei Guo
*NVIDIA Corporation*
Westborough, MA, USA
xfguo@ieee.org

Optimizing today's computing systems requires a collaborative efforts across the whole system stack starting from transistors, through circuit and architecture up to the whole system and ecosystems. Single layer solution such as technology scaling or architectural innovation need to work together in a more effective way to ensure a lower cost. The key idea of the *cross-layer co-design* is to *divide* a metric optimal task into a set of sub tasks, which can be implemented at different levels of a system stack. Examples of such tasks can be error prediction or detection, security, power optimization, etc. These tasks can be treated as steps that the system follows to handle a particular effect even as they may not occur sequentially. In this talk, I will discuss a series of techniques that were used across the hardware abstractions to address chip aging issues, one of the major device level reliability threats in today's nano-scale integrated circuits.

Nowadays, on a computer chip with size of a fingernail, there are billions of transistors that serve as the smallest computing units. Just as in the biological world, these transistors and their interconnects will age with time, the degradation over time leads slowly but surely to decreased switching speeds, and it can even result in outright circuit failures. It is also unfortunate that aging is becoming much more troublesome for design teams at 10nm and below. A good sign is that many of these aging mechanisms can be recovered slowly, but the existing research only "scratch a surface" on understanding the underlying mechanism due to the difficulties of setting up experiments and capturing the degradations within a reasonable time frame. To fulfill the need for understanding recovery behaviors thoroughly, we designed a set of novel experiments to collect measured results from actual chips (FPGA for transistor aging [1], test chips for interconnect aging [2]), each set of measurement has been carefully designed by considering different combinations of recovery conditions and lasts for more than 3 days. Through extensive experiments, we discovered that recovery can be made active, and the irreversible components can be completely avoided. It was also the first time that we found aging mechanisms follow a *"circadian rhythm"* like pattern - the whole process of chip aging and recovery can be compared directly to the biological world, in which scheduled recovery periods (or healing) are necessary after extensive workouts (or stress), with their athletic performance actually getting even better after the rest periods [3]. We borrowed the idea and successfully applied to the electronic world, thus the equivalent circadian rhythm for an operating chip would start with the active status until the irreversible aging kicks in, then an active recovery period is followed so that the irreversible aging becomes almost unobservable after active recovery even under extreme stress cases. These experimental results provided brand new insights on recovery from aging, such as frequency dependency [4], accelerated and active recovery and long-term vs. short recovery behaviors [3]. Such insights contributed as new experimental evidences for reliability community to create better and more accurate device models for emerging aging effects, they also bring strong implications of designing reliable systems by considering recovery as a main design knob.

With the discovered unique device level aging and recovery behaviors, we further explored on-chip solutions to fully utilize these behaviors especially for applications that require extended lifetime and very low error tolerance. Cross-layer co-design is a closer to optimal way of maximizing reliability by breaking the abstraction layers boundaries across the system stack. At the circuit level, recovery circuit and wearout sensors were proposed to distribute on the aging-critical units, and they are triggered by higher level decisions from scheduler or load balancer. Architecture-level accelerated self-healing solutions can utilize some intrinsic sleep behaviors and heat to recover inactive parts. It can also compensate some of the power overhead introduced by the circuit level recovery solutions. System level scheduling is able to divide the recovery tasks and make the high-level recovery decisions. Distributing the recovery tasks across the system stack can improve performance, and reduce power and area costs by taking advantage of the characteristics of each layer. Overall, active accelerated self-healing shows as a promising technique for solving aging issues and can be introduced as a key design knob for cross-layer resilience during the design process to achieve the optimal resilience effectively.

## REFERENCES

[1] X. Guo, W. Burleson, and M. Stan, "Modeling and experimental demonstration of accelerated self-healing techniques," in *2014 51st ACM/EDAC/IEEE Design Automation Conference (DAC)*. IEEE, 2014, pp. 1–6.

[2] X. Guo and M. R. Stan, "Deep Healing: Ease the BTI and EM Wearout Crisis by Activating Recovery," in *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2017, pp. 184–191.

[3] ——, *Circadian Rhythms for Future Resilient Electronic Systems - Accelerated Active Self-Healing for Integrated Circuits*. Springer, 2020.

[4] ——, "Work hard, sleep well-avoid irreversible ic wearout with proactive rejuvenation," in *2016 21st Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2016, pp. 649–654.