

Cross-layer Codesign for Resilient Hardware



Xinfei Guo, Ph.D.
U of Virginia & NVIDIA Corporation
Westborough, MA

Jan 29th, 2021

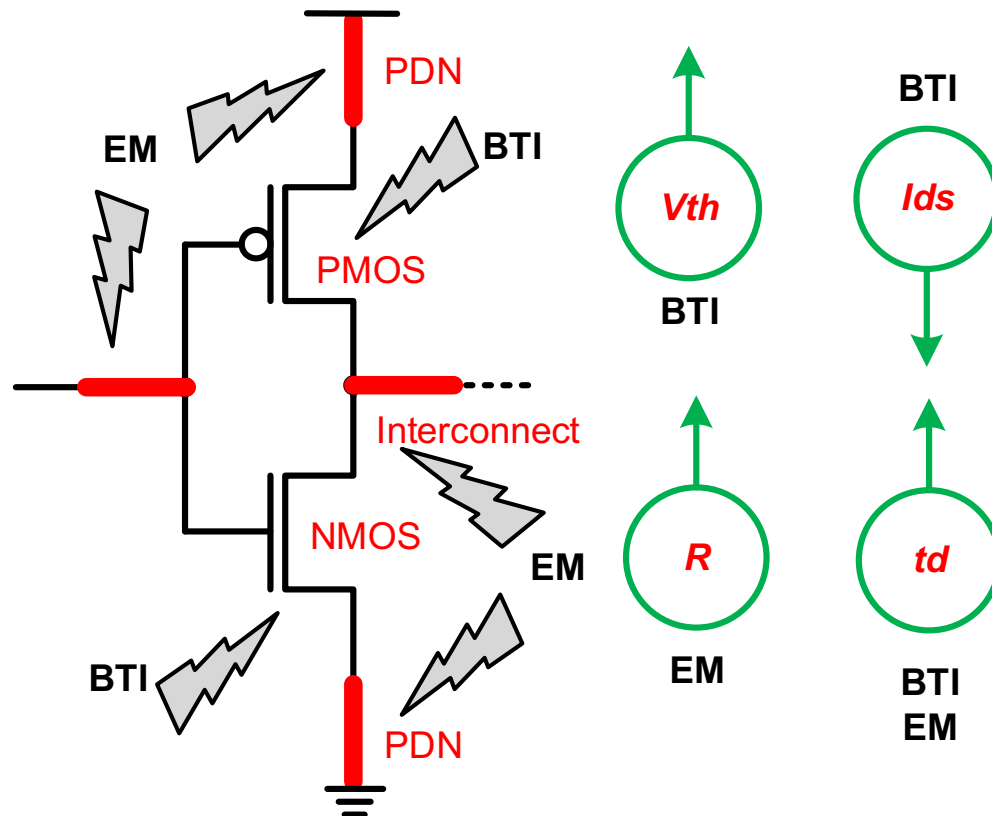
xfguo@ieee.org
www.xinfeigu.com

Outline

- Why?
- What?
- How?
- Now what?

Starting from Transistors – Aging/Wearout

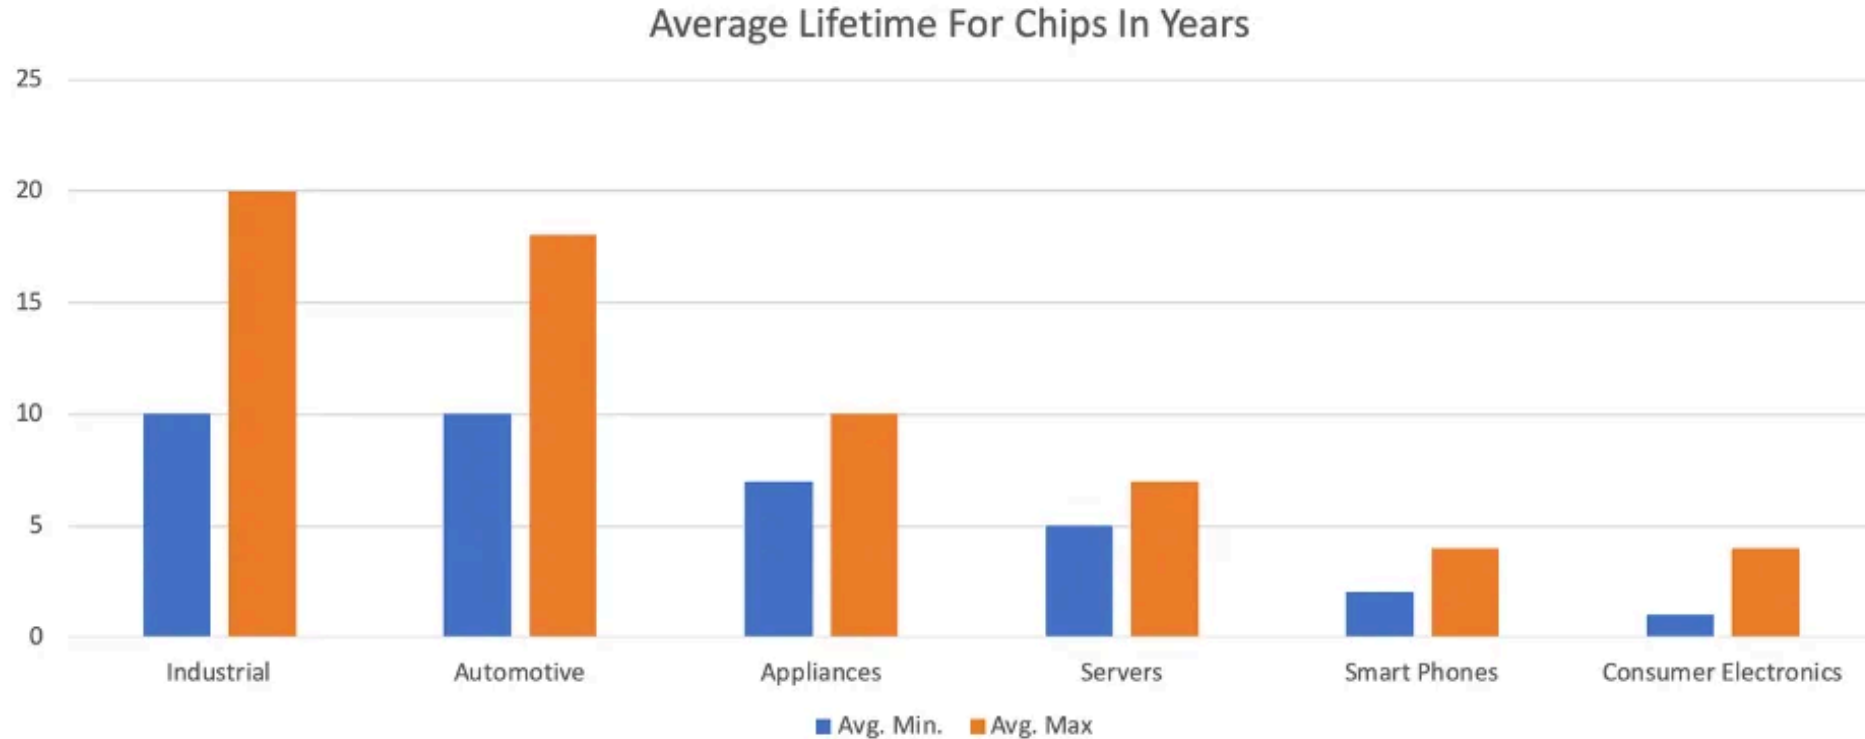
- Time-dependent device degradations
 - Transistors (P/NMOS)
 - Metal layers (Interconnect, PDN)



BTI → Bias Temperature Instability
EM → Electromigration

PDN → Power Delivery Network
V_{th} → Threshold Voltage
I_{ds} → Current
R → Resistance
td → Propagation Delay

Back to applications – lifetime is critical!



- ❑ Longer expected lifetime
- ❑ Extreme environmental conditions
- ❑ Longer run time (higher utilization)
- ❑ Higher susceptibility to aging
- ❑ Replacement costs
- ❑ ...

Aging – Industry Attentions

Transistor Aging Intensifies At 10/7nm And Below

f t in G+ 80

Device degradation becomes limiting factor in IC scaling, and a significant challenge in advanced SoCs.

JULY 13TH, 2017 - BY: ANN STEFFORA MUTSCHLER

SEMICONDUCTOR ENGINEERING

Home > Manufacturing, Design & Test > Improving Transistor Reliability

MANUFACTURING, DESIGN & TEST

Improving Transistor Reliability

f t in G+ 53 3

...s is being made, but there are no easy answers

TEST, MEASUREMENT & ANALYTICS

Chiplet Reliability Challenges Ahead

46 Shares f 15 t 9 in 20 <

Determining how third-party chiplets will work in complex systems is still a problem.

AUGUST 12TH, 2020 - BY: ED SPERLING

SEMICONDUCTOR ENGINEERING

Home > Manufacturing, Design & Test > The End Of Silicon?

MANUFACTURING, DESIGN & TEST

The End Of Silicon?

f 20 t in 7 G+ 30

Negative bias temperature instability could force chip change course on materials.

MAY 4TH, 2015 - BY: KATHERINE DERBYSHIRE

SEMICONDUCTOR ENGINEERING

Home > Manufacturing, Design & Test > Reliability After Planar Silicon

MANUFACTURING, DESIGN & TEST

Reliability After Planar S

f 11 t in 2 G+ 3

Second of two parts: Why silicon is nearing its end next.

JULY 29TH, 2015 - BY: KATHERINE DERBYSHIRE

Making Chips To Last Their Expected Lifetimes

238 Shares f 42 t 16 in 140 <

Lifecycles can vary greatly for different markets, and by application within those markets.

Chip Aging Accelerates

f 81 t in G+ 76

As advanced-node chips are added into cars, and usage models shift inside of data centers, new questions surface about reliability.

FEBRUARY 14TH, 2018 - BY: ED SPERLING

Aging Problems At 5nm And Below

1.7k Shares f 121 <

Semiconductor aging has moved from being a foundry issue to a user problem. As we get to 5nm and below, vectorless methodologies become too inaccurate.

JUNE 11TH, 2020 - BY: BRIAN BAILEY

Why do we care about aging now?

■ Reliability threat!

- Permanent errors
- Shorten lifetime
- Worsen metrics such as performance, power and area

■ Getting worse with technology scaling!

- Increased power density → Heat
- Increased effective electrical field → More stress
- More components → Require lower single failure rate
- Advanced nodes → New stress and issues: e.g. self-heating

Wire broke due to EM


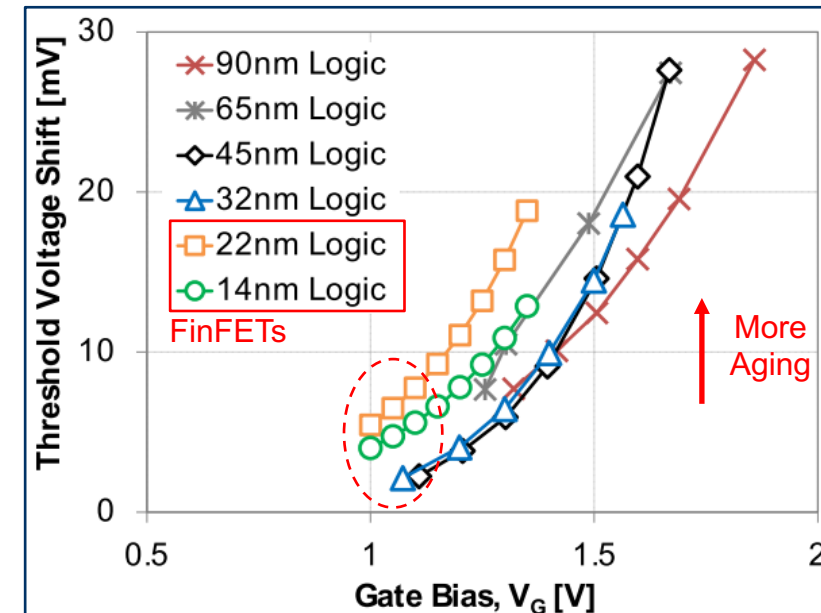


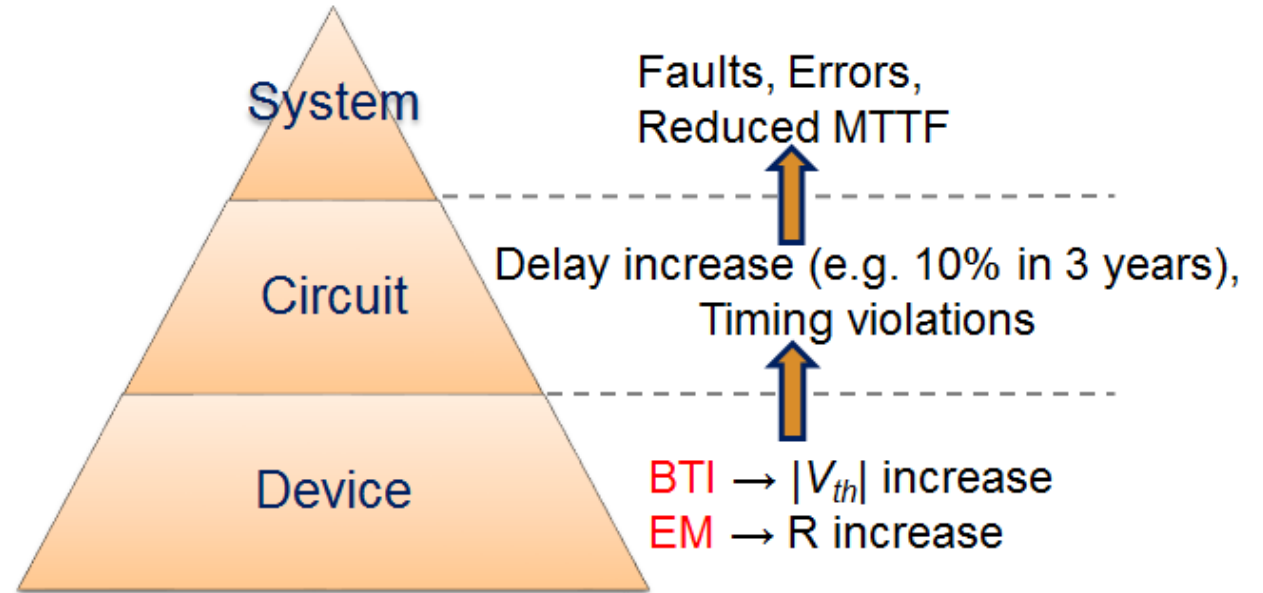
Figure: [N. Cheung et al., UC Berkeley]



Source: [S.M. Ramey, et al. (Intel), 2018]

A cross-layer effect

- Device level
 - Threshold voltage V_{th} increase (BTI)
 - Resistance increase (EM)
- Circuit level
 - Performance degradation
 - Timing failures
 - Leakage power
- Architecture level
 - Failures
 - Errors
- System level
 - MTTF



Traditional solutions

Adding Margins

- Over-estimation
- Under-estimation
- Uncertain operation conditions
- e.g. 10% for a 3-year lifetime constraint

Adapting (Sensing + Actuation)

- The worst case is getting worse
- Aging is unchecked
- Tracking power (over 10 sensors per partition)

Passive Recovery (More idle time)


- Very slow
- Unpredictable
- Permanent part will keep accumulating



Single-layer solution is not adequate to deal with aging any more as it is becoming a bottleneck!

Introducing Accelerated Self-Healing

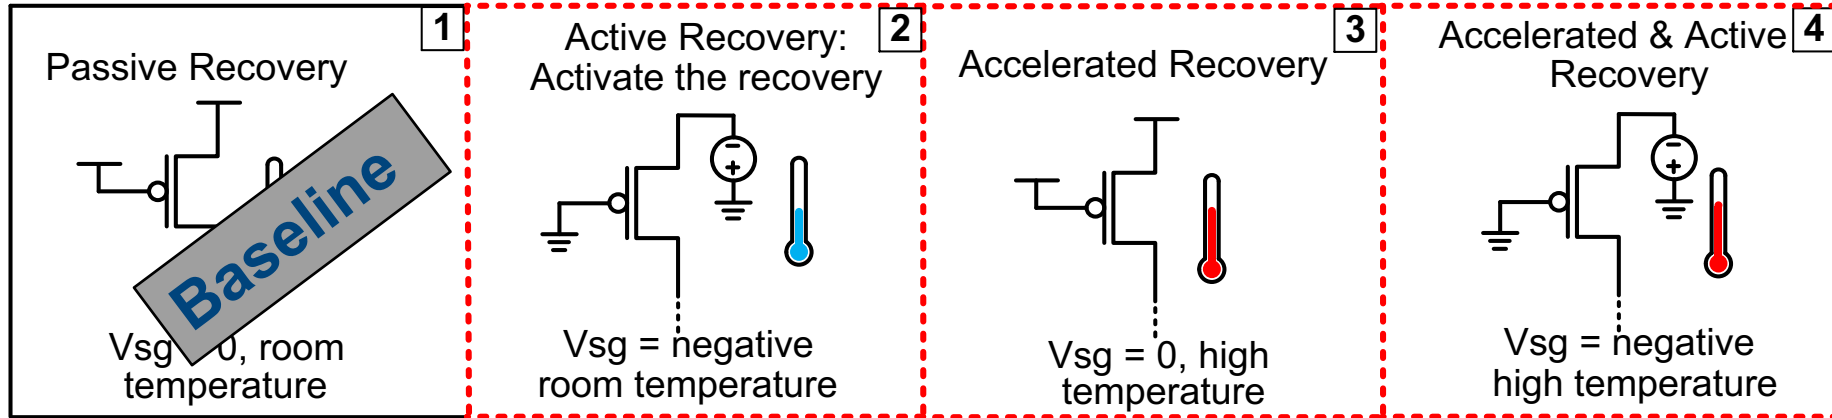
- Fact - Aging is **partially** recoverable under **passive recovery**, but **it is very slow**.
- **Key Idea: Reverse the directions of aging and enable active Recovery**




Accelerated Self-Healing

Key Idea: Recover by reversing the directions of Aging

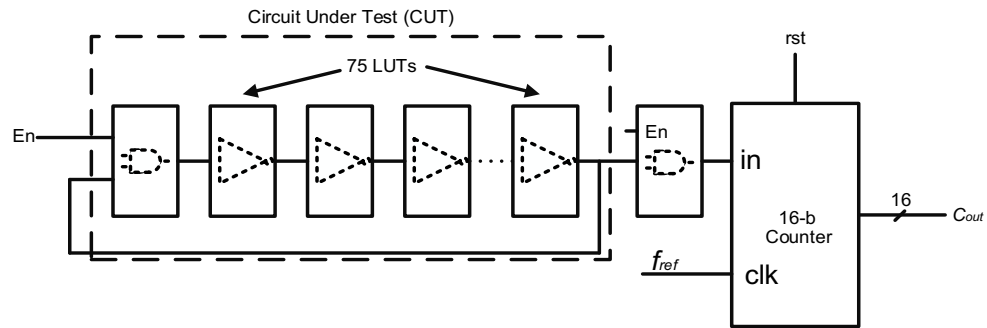
BTI Accelerated Self-Healing




EM Accelerated Self-Healing



Experiments for demonstration




BTI Test Setup




40nm FPGA

EM Test Setup



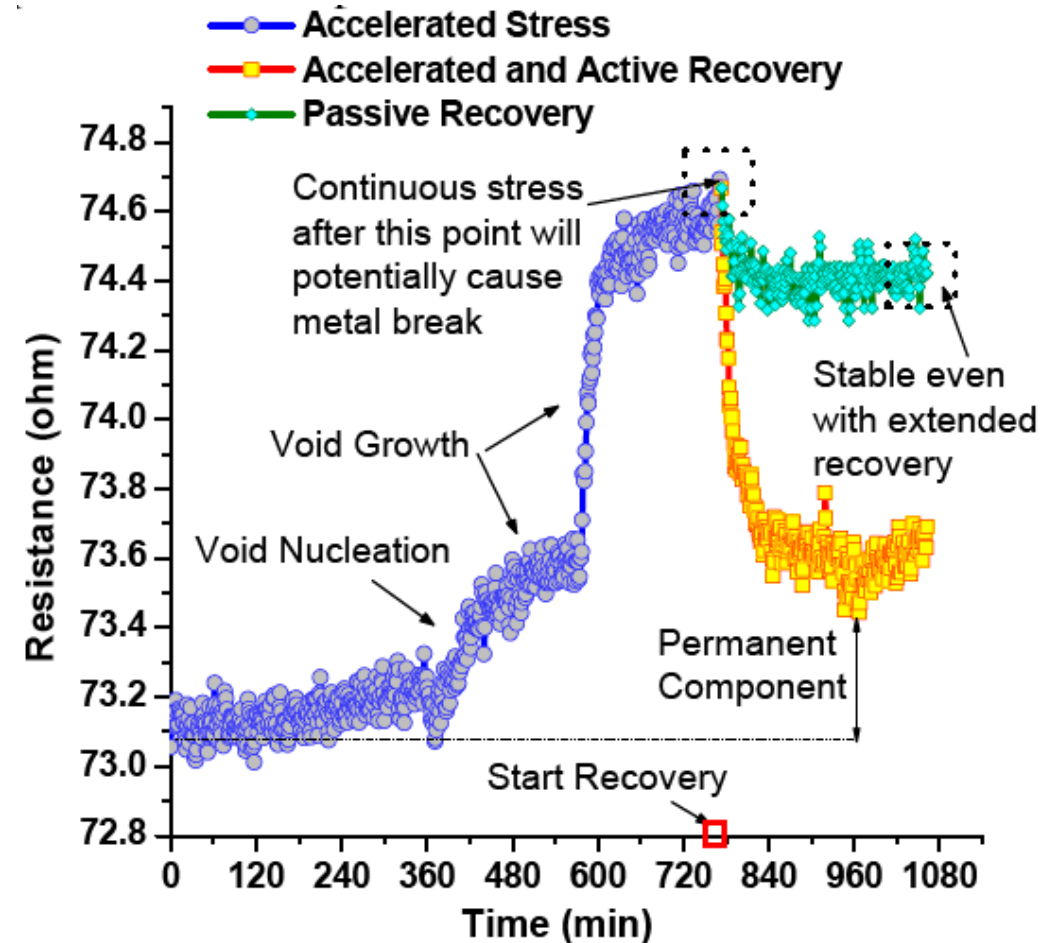
Technology	180nm
Material	Copper
Thickness	0.8um
Length	2.673mm
Width	1.57um
Resistance (@rt)	35.76Ω



On-chip Metal Wires

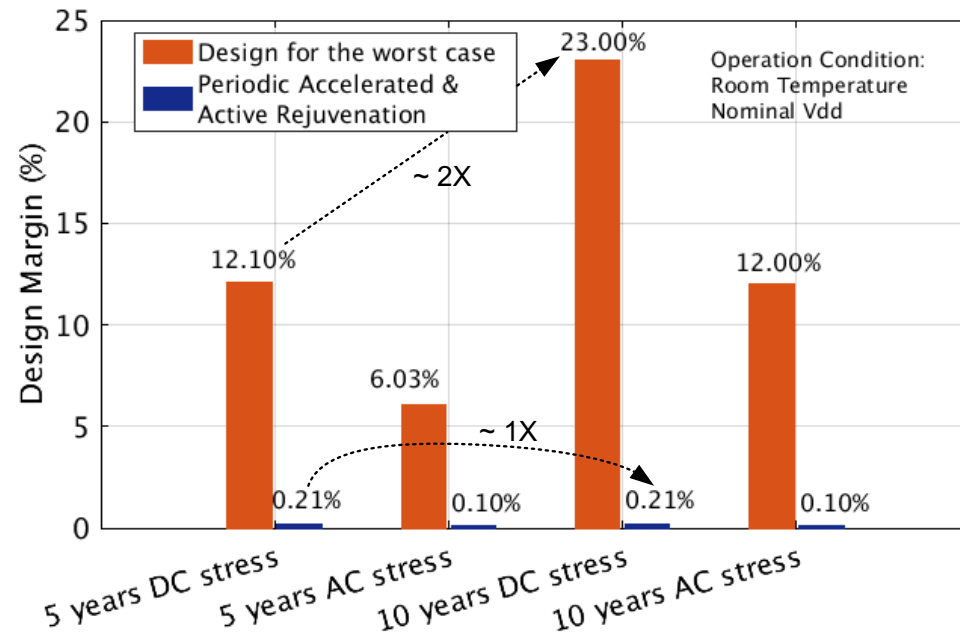
Measurement results summary

- Recovery from aging can be made **active** and be **accelerated**, even the irreversible component can be **fully** eliminated or avoided through various techniques such as **higher temperatures, negative voltages, active vs. sleep ratio ...**
- **What does this mean for chip designers and architects?**
A: Cross-layer Accelerated Self-Healing

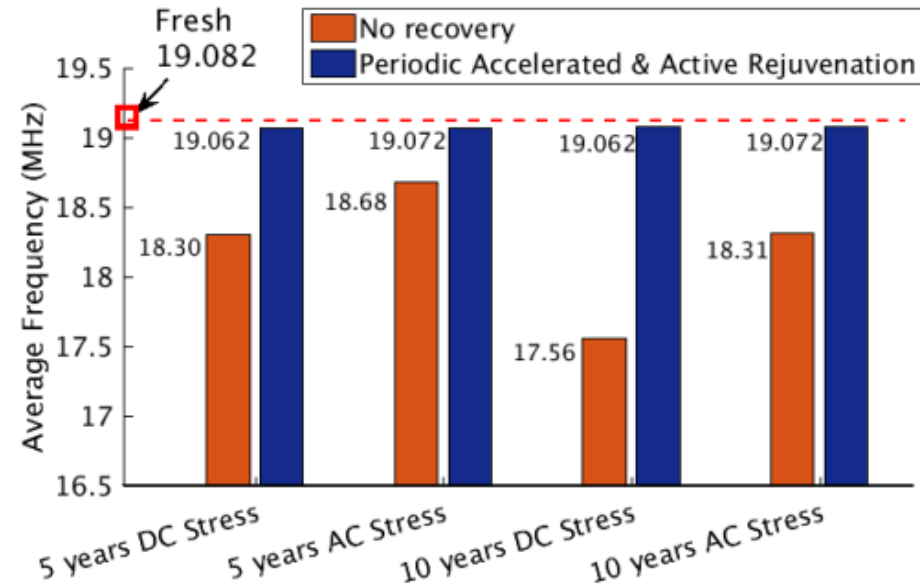


Implication – Metric Improvements

- ❑ >60x reduction of necessary margin for all cases
- ❑ The average performance is close to the fresh during the whole lifetime
- ❑ Both metrics don't scale with the increase of the lifetime constraint

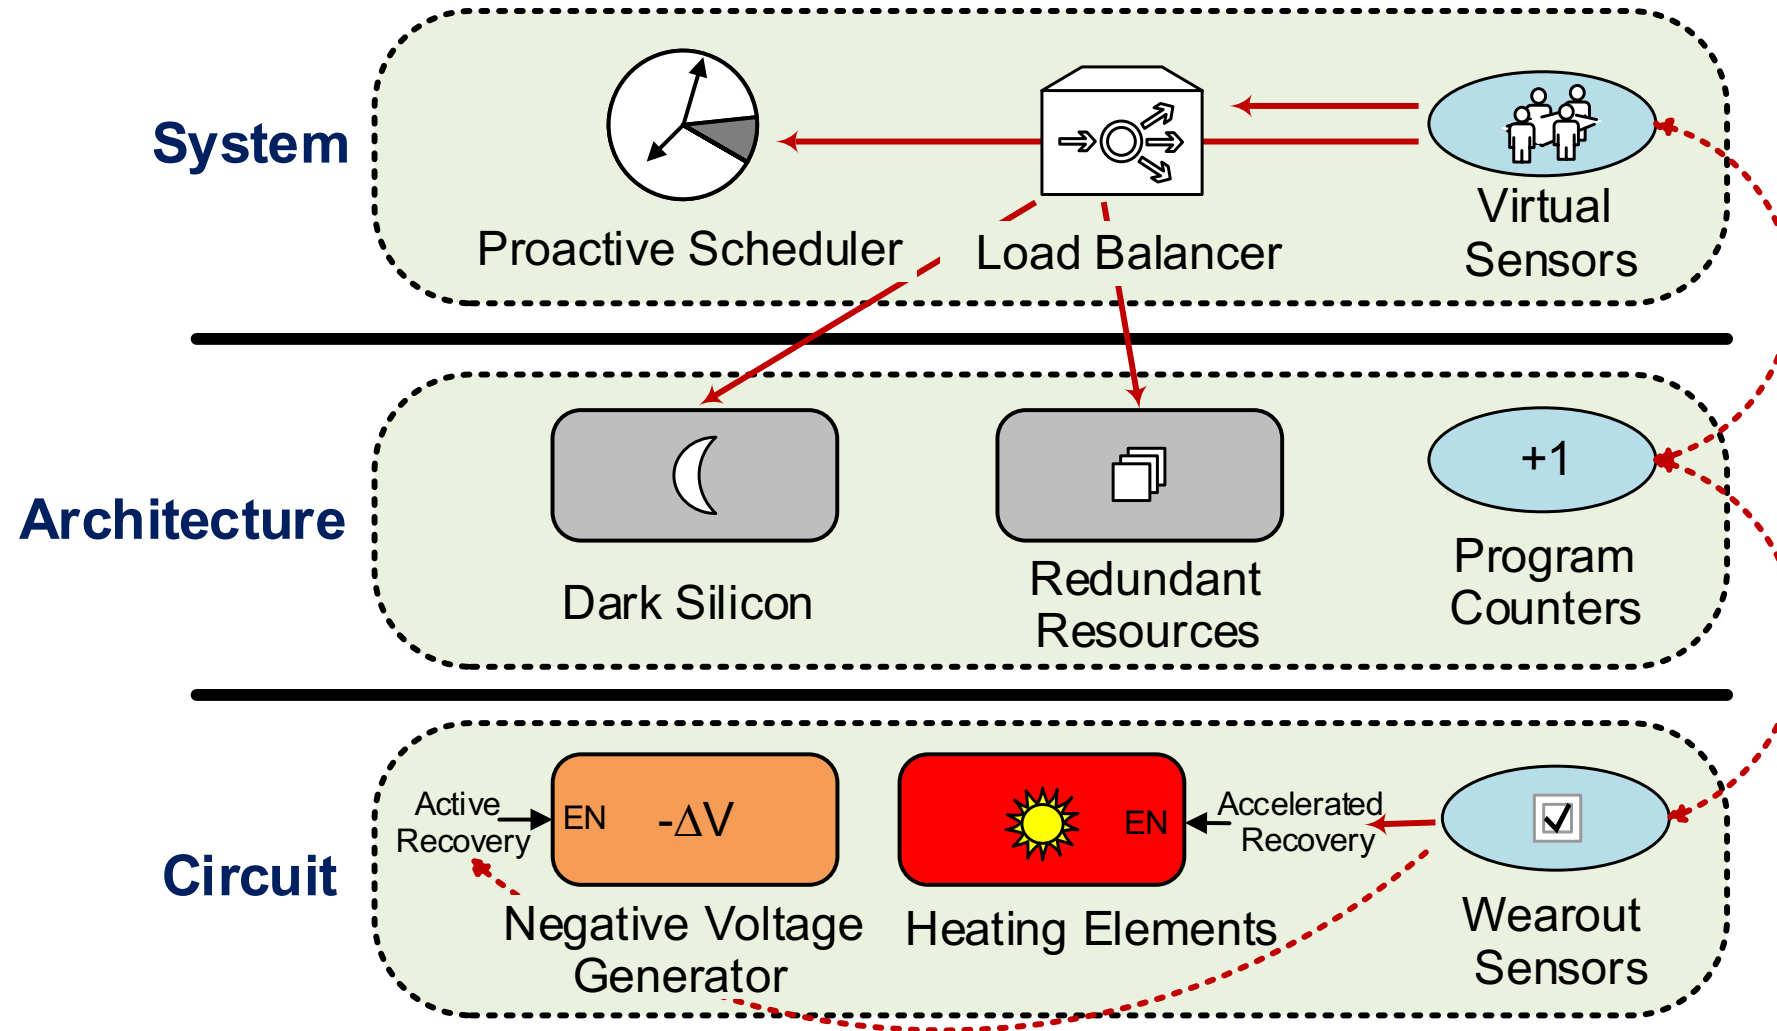


Reduction of Necessary Design Margin

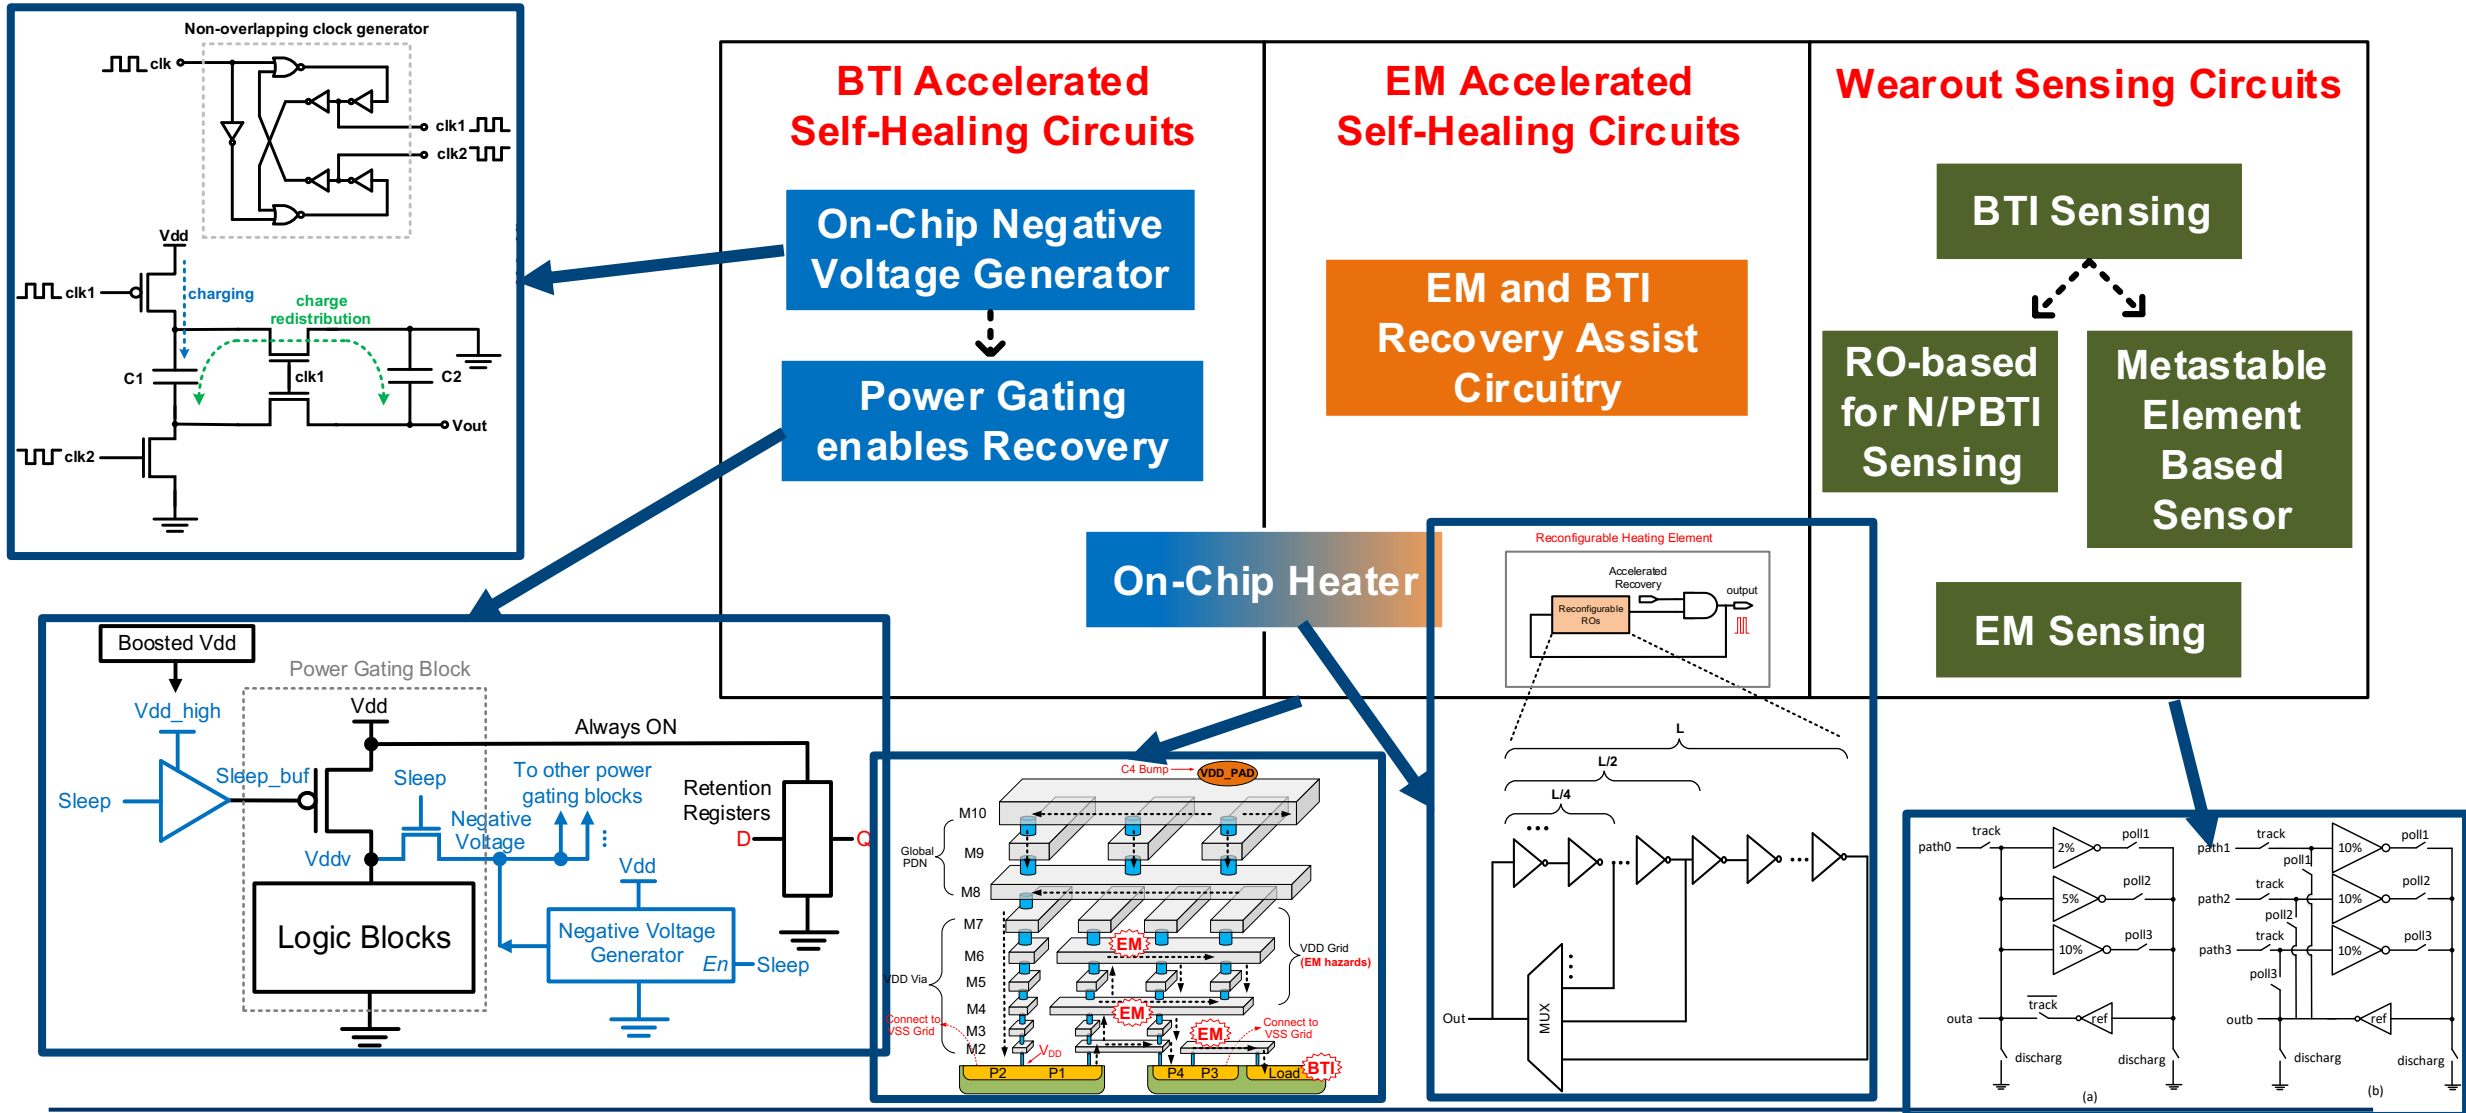


Average Performance Improvement

Cross-layer Accelerated Self-Healing



Circuit Components for Self-Healing



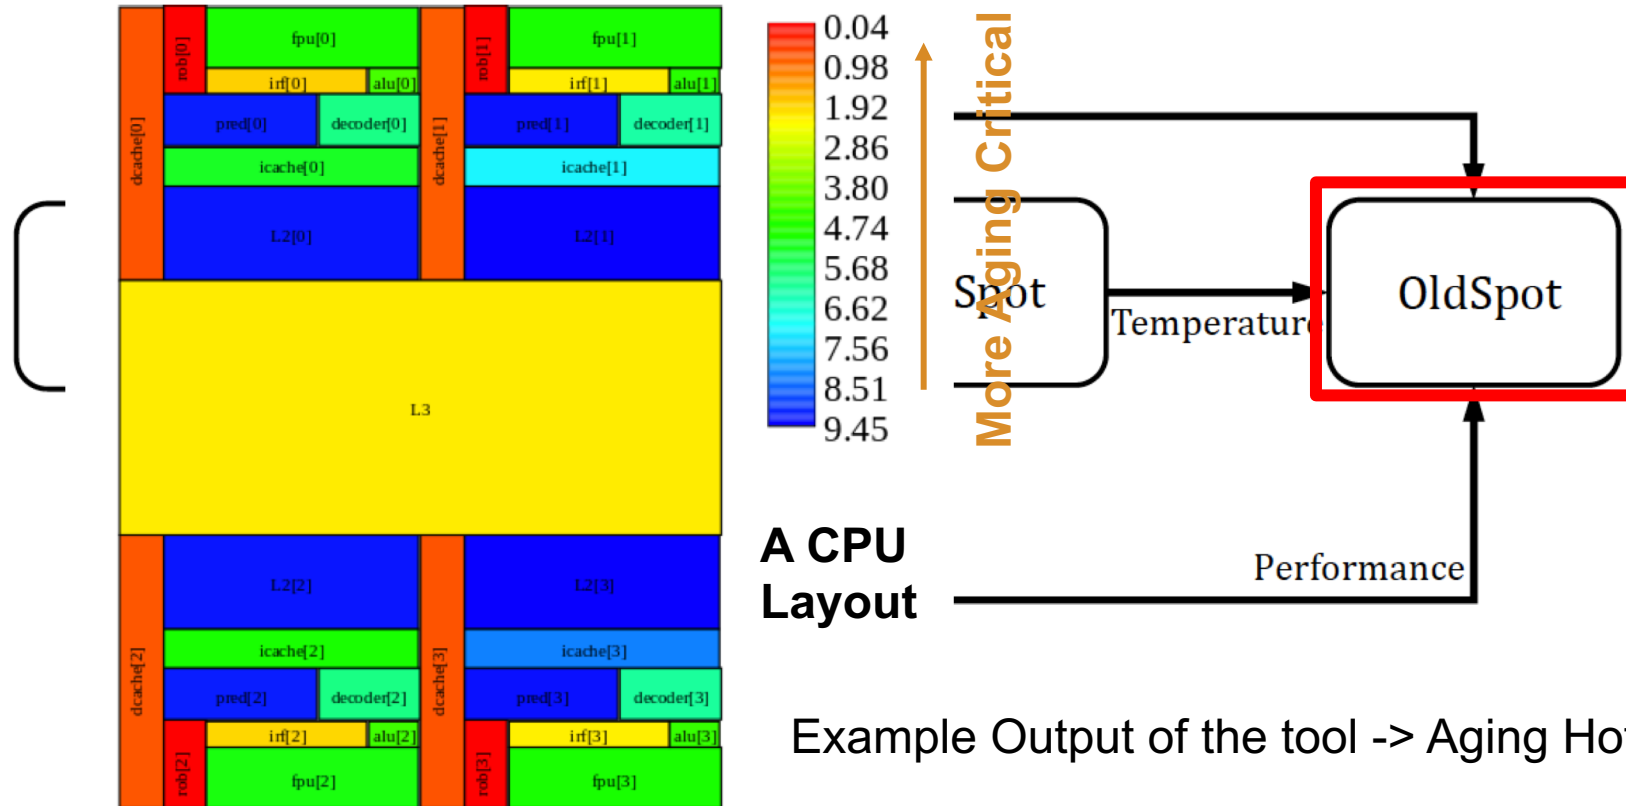
Costs

- Area ↓ Power ↓ Extra Heat ↓
 - Optimal ways of distributing circuit IPs in a large system
 - Avoid unwanted heat
 - Trigger only when necessary

Design Name	Leakage Power	Dynamic Power	Area	Performance
Neg. Voltage Generator	68.85nW	64.47uW	4300um ²	>66.7MHz
On-Chip Heater	16.8nW	75uW	16um ²	-
Multi-mode Recovery Circuit	-	-	58.24um ²	Wakeup time ~170ns

Are there any other opportunities beyond circuit level?

Architectural Simulation Framework for Architecture Level Exploration – “OldSpot”

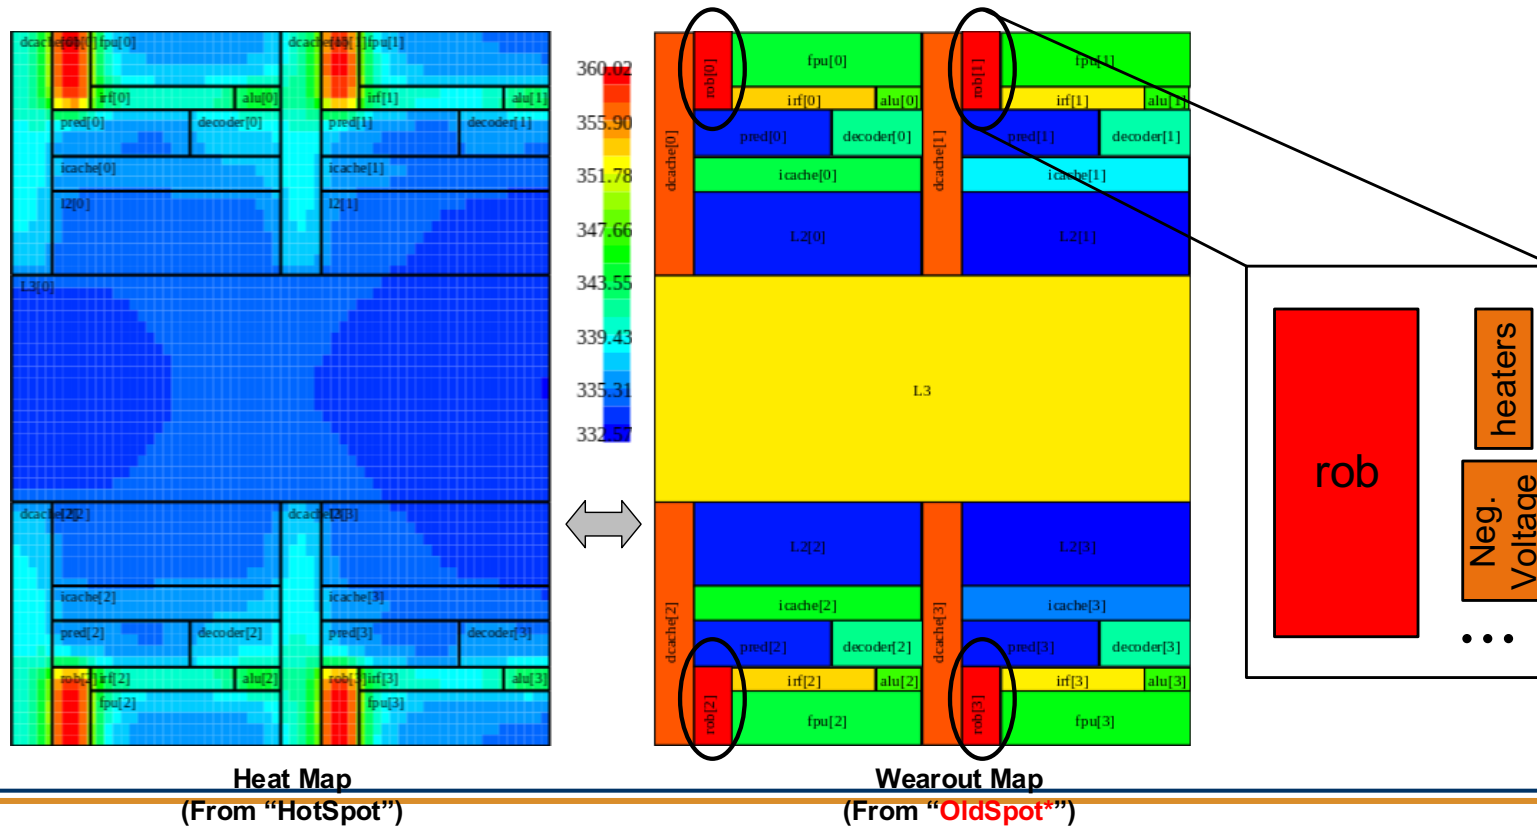


Example Output of the tool -> Aging HotSpot!

<https://github.com/hplp/oldspot>


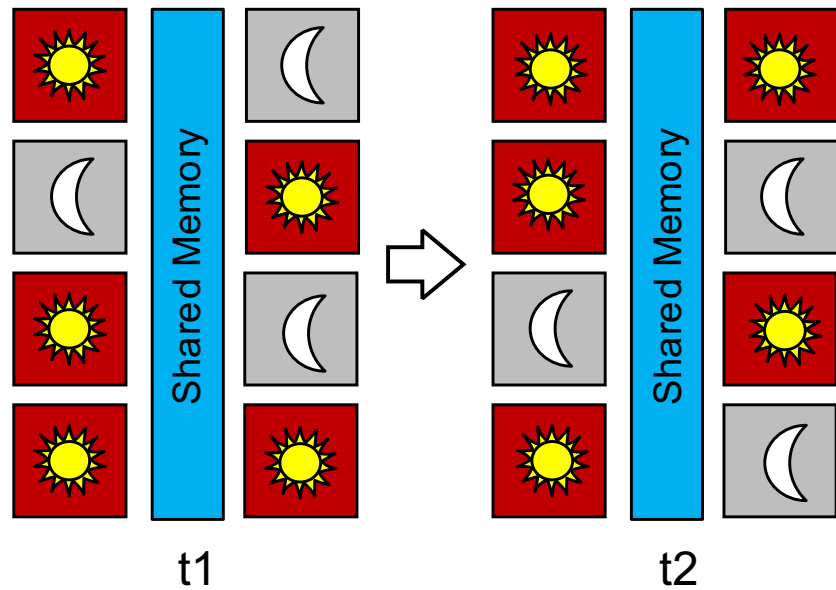
Unit-level Accelerated Self-Healing

- Goal
 - Less area and power overhead
- Solution
 - Placing self-healing IPs **only** for aging-critical units



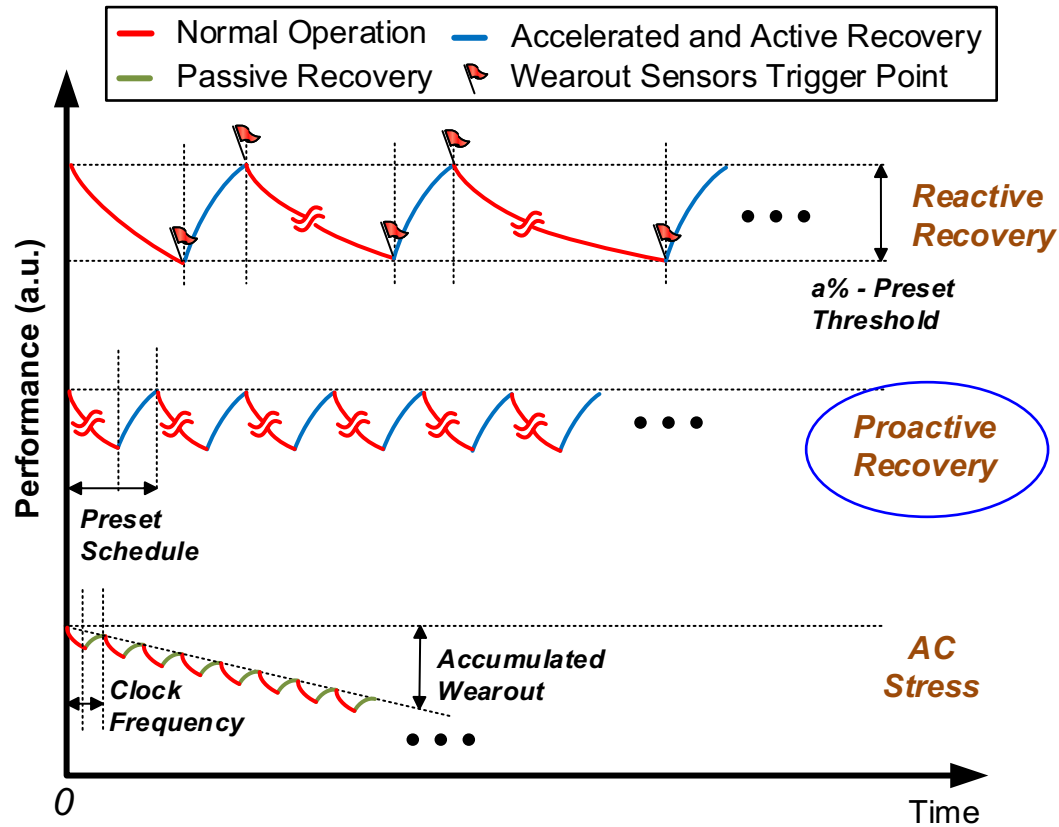
Utilize Intrinsic Heat

- Goal
 - Avoid power overhead for generating extra heat
- Solution
 - Take advantage of dark silicon or core redundancy
 - Utilize intrinsic sleep behaviors




Scheduling for Recovery

- Goal
 - Recover effectively only when necessary



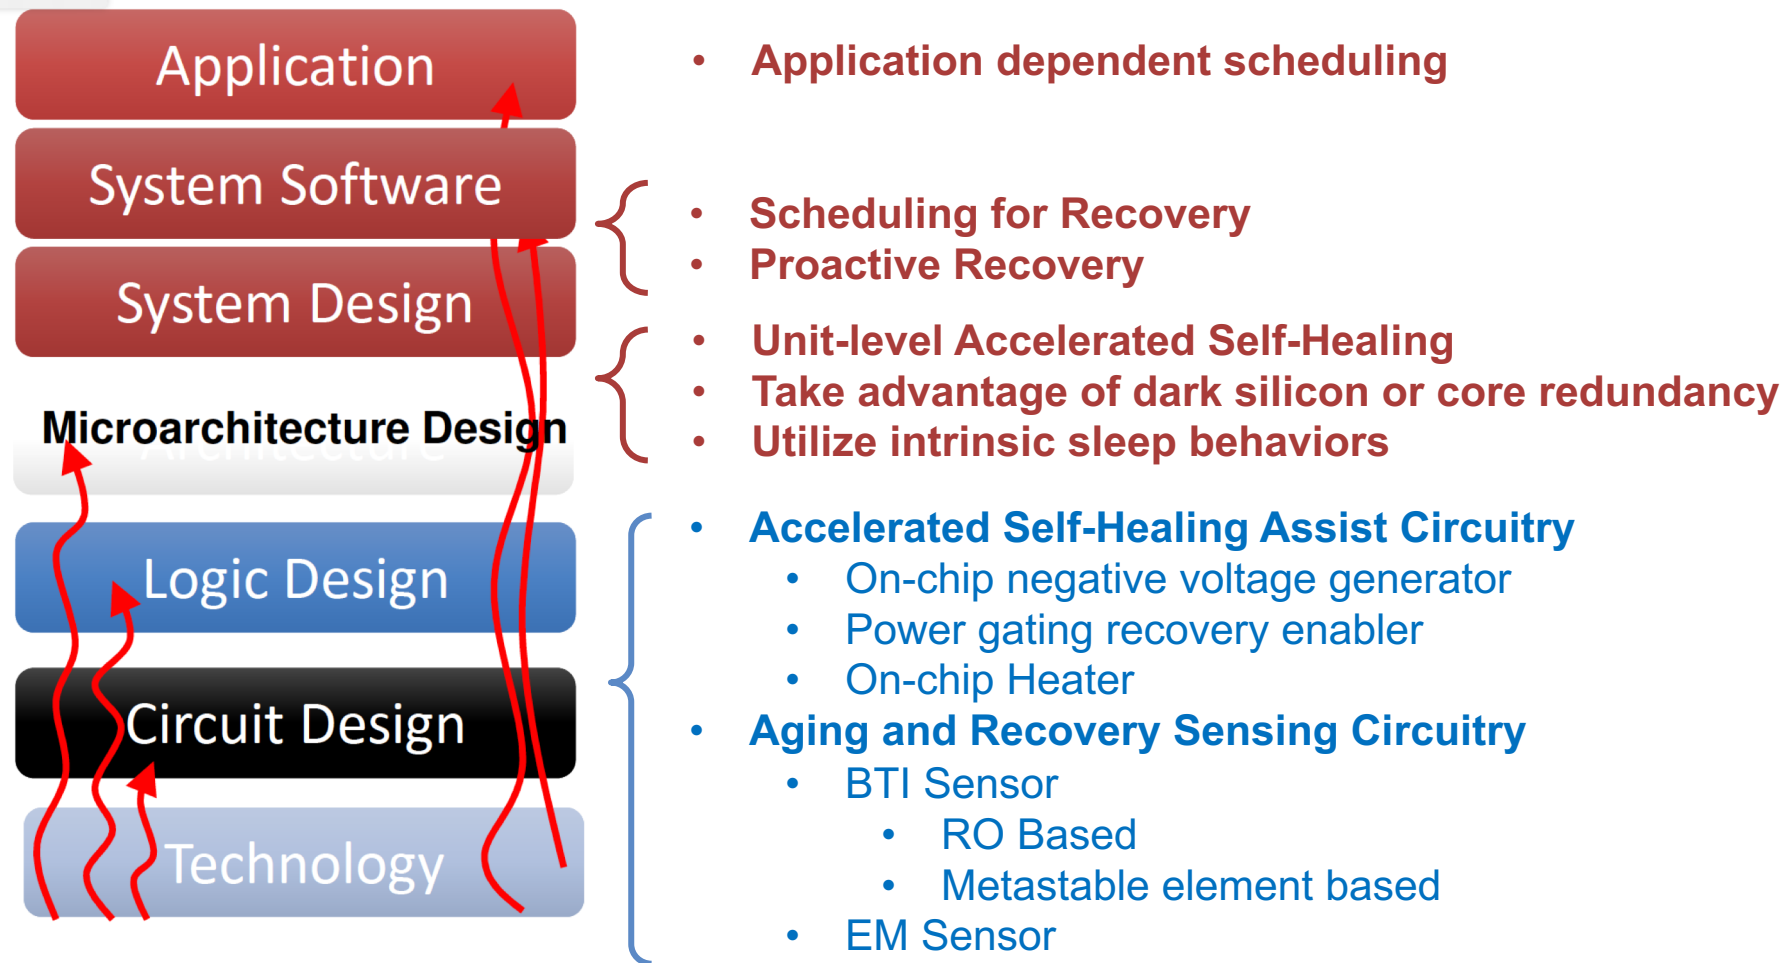
Proactive Recovery




Full recovery time after 12-hour constant stress under normal condition

Application-dependent Scheduling

Putting it All Together: **CLASH** - Cross-layer Accelerated Self-Healing System



CLASH System – Hardware view



What is the key benefit of doing cross-layer codesign here?

Before...

- Margin (e.g. **10 – 20 %**)
- Track and Adaptation (Track during the entire lifetime)
- Passive recovery (**<20%** recovery percentage)

Constraints (Close to data center or IoT application cases)

- 10 year lifetime constraint
- Under DC stress
- Operating at room temperature
- Nominal Vdd

Accelerated Self-Healing

- Margin (**0.21%**)
- Only track the reversible part (**~ 8X** tracking power reduction)
- As high as **100%** recovery rate
- Cross-layer implementation minimizes the cost
- Recovery-driven design method

*All numbers are based on the model prediction which is calibrated by measurements.

Sleep for rejuvenating and healing neurons?

Subscribe

Latest Issues

SCIENTIFIC
AMERICAN®

Cart 0 Sign In | Stay

COVID-19 VIRUS THE SCIENCES MIND HEALTH

Read

Lack of Sleep

We are just starting to investigate an additional benefit of artificial sleep in our simulations. Often, a few neurons in a simulated network fail to function at all when a simulation is started. We have found that applying artificial sleep states seems to reset idle neurons to ensure they become functioning components in the network.

Some types of artificial intelligence could start to hallucinate if they don't get enough rest, just as humans do

By Garrett Kenyon on December 5, 2020

<https://www.scientificamerican.com/article/lack-of-sleep-could-be-a-problem-for-ais/>

BARC 2021


Key takeaways

- Device level behaviors will have a lasting impact to all upper layers
- Cross-layer codesign is an essential way of enlarging the search space
- Challenges co-exist with opportunities
 - Infrastructures
 - Transparency
 - Design cycle
 - ...



References

1. X. Guo, W. Burleson, M. Stan, "Modeling and Experimental Demonstration of Accelerated Self-Healing Techniques," ACM/IEEE Design Automation Conference (DAC), 2014.
2. X. Guo, M. Stan, "Work hard, sleep well - Avoid irreversible IC wearout with proactive rejuvenation," ACM/IEEE Asia and South Pacific Design Automation Conference (ASP-DAC), 2016.
3. X. Guo, M. Stan, "Deep Healing: Ease the BTI and EM Wearout Crisis by Activating Recovery," International Conference on Dependable Systems and Networks (DSN), 2017.
4. X. Guo, M. Stan, "Implications of Accelerated Self-Healing as a Key Design Knob for Cross-Layer Resilience", INTEGRATION, the VLSI journal (VLSI), vol. 56, pp. 167-180, 2017.
5. A. Roelke, X. Guo, M. Stan, "OldSpot: A Pre-RTL Model for Aging and Lifetime Optimization," ICCD, 2018.
6. X. Guo, V. Verma, P. Guerrero, M. Stan, "When things get older - Exploring Circuit Aging in IoT Applications," International Symposium on Quality Electronic Design (ISQED), 2018.
7. X. Guo, M. Stan, "Circadian Rhythms for Future Resilient Electronic Systems - Accelerated Active Self-Healing for Integrated Circuits", Springer, 2020.



Stay Safe and Healthy!



xfguo@ieee.org