The Efficacy of Error Mitigation Techniques for DRAM Retention Failures

Samira Khan, Donghyuk Lee, Yoongu Kim,
Alaa R. Alameldeen, Chris Wilkerson, and
Onur Mutlu

SAFARI Carnegie Mellon intel

Motivation

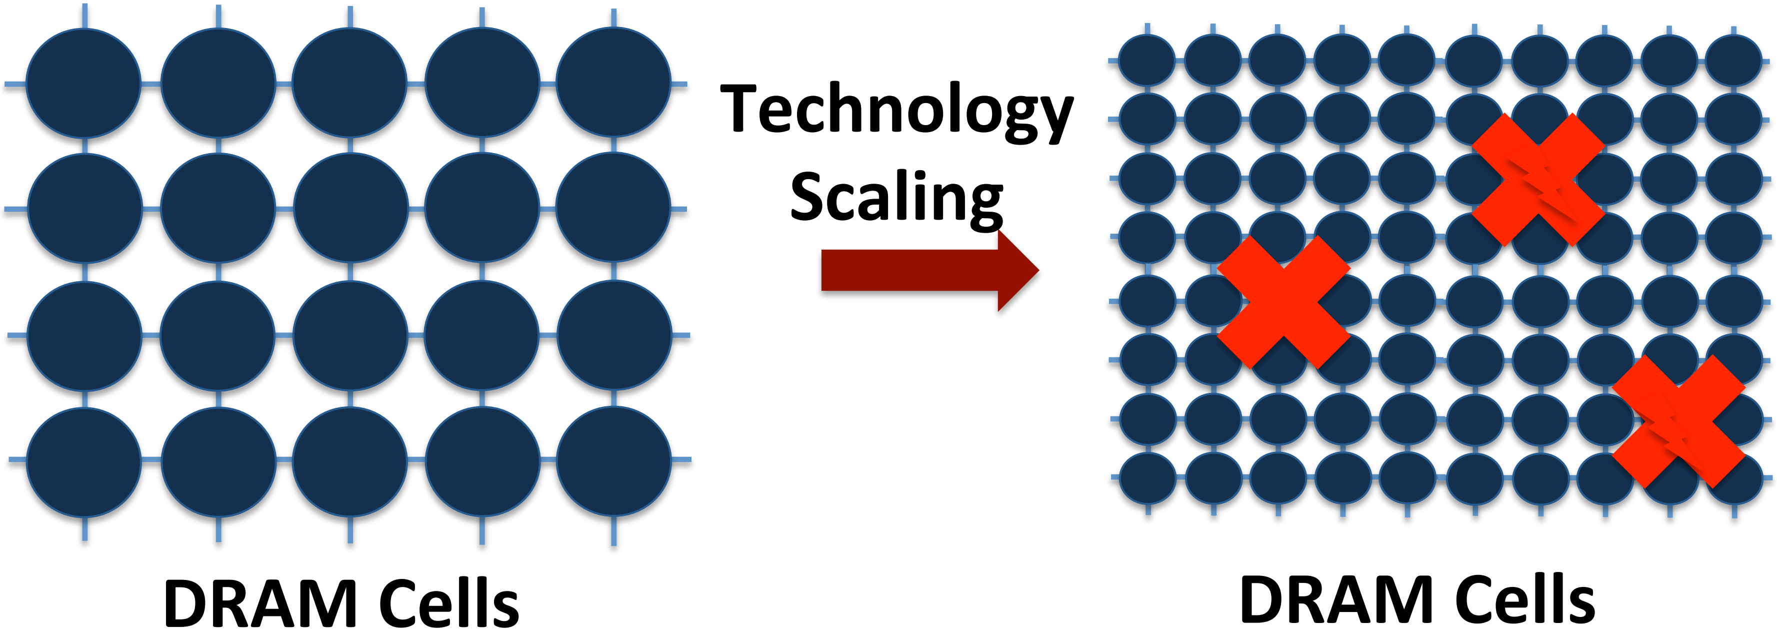

Scaling DRAM cells results in more failures

- Longer manufacture-time tests
- Lower yield
- Higher cost

Vision: Online Profiling

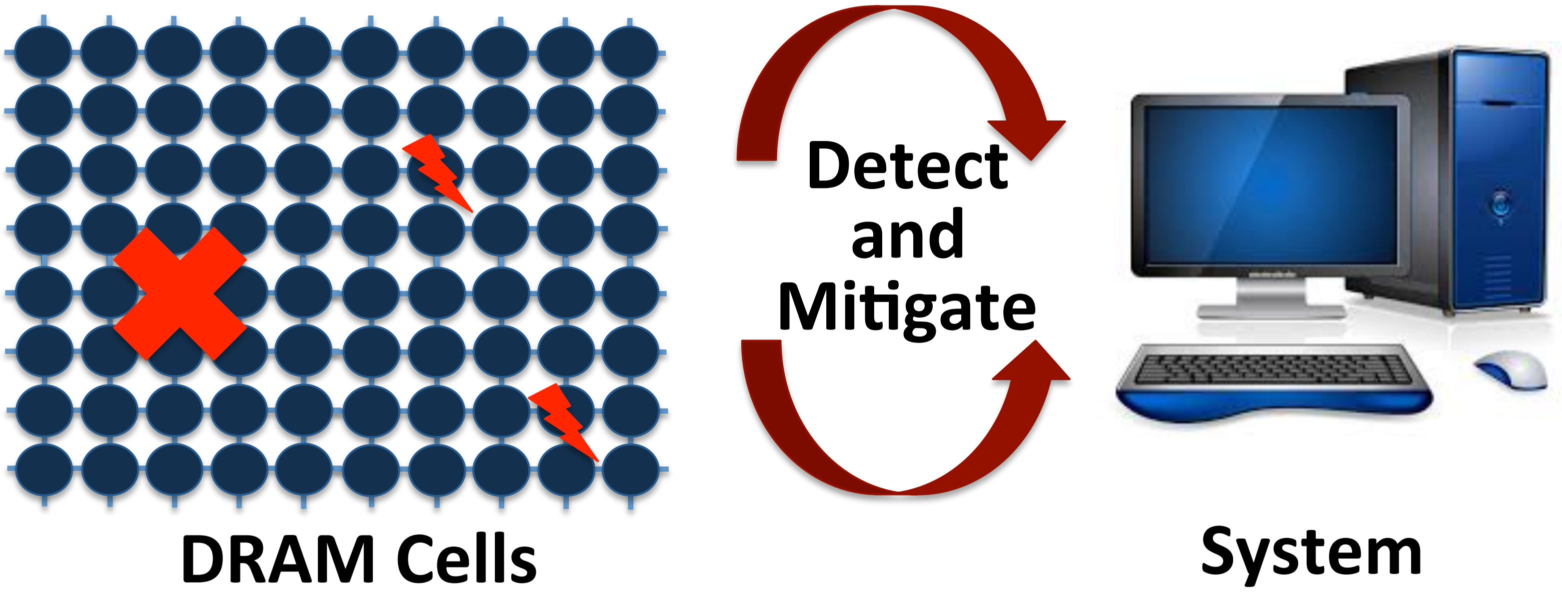

Detect and mitigate errors after the system has become operational

Reduces cost of testing, increases yield, enables scaling

What is the effectiveness of system-level detection and mitigation techniques?

Summary

- ***We*** ***analyze*** **the efficacy of testing, guardbanding, ECC, and recent techniques**
  - **Using experimental data from *real DRAMs***
- ***Key Conclusions***
  - **Testing alone cannot guarantee reliable operation**
  - **A combination of ECC, testing, and guardbanding is more effective**
  - **Testing+ECC-based techniques block memory for significant time → Performance degradation**
- ***We*** ***propose*** **a possible online profiling mechanism**

Outline

- DRAM Scaling Problem
- Online Profiling as a Solution
- Efficacy of System-Level Detection and Mitigation
  - Simple Techniques
  - Recently Proposed Techniques
- Towards an Online Profiling System
- Conclusion

Outline

- **DRAM Scaling Problem**
- **Online Profiling as a Solution**
- **Efficacy of System-Level Detection and Mitigation**
  - **Simple Techniques**
  - **Recently Proposed Techniques**
- **Towards an Online Profiling System**
- **Conclusion**

Retention Failure

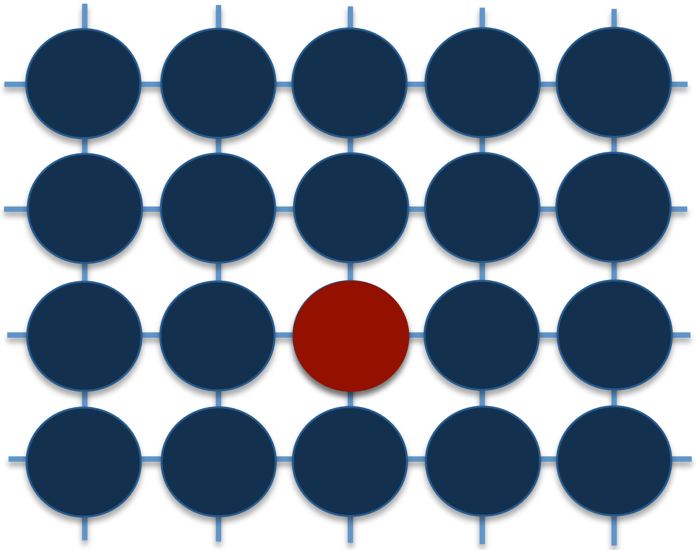

DRAM Cells

Retention Failure

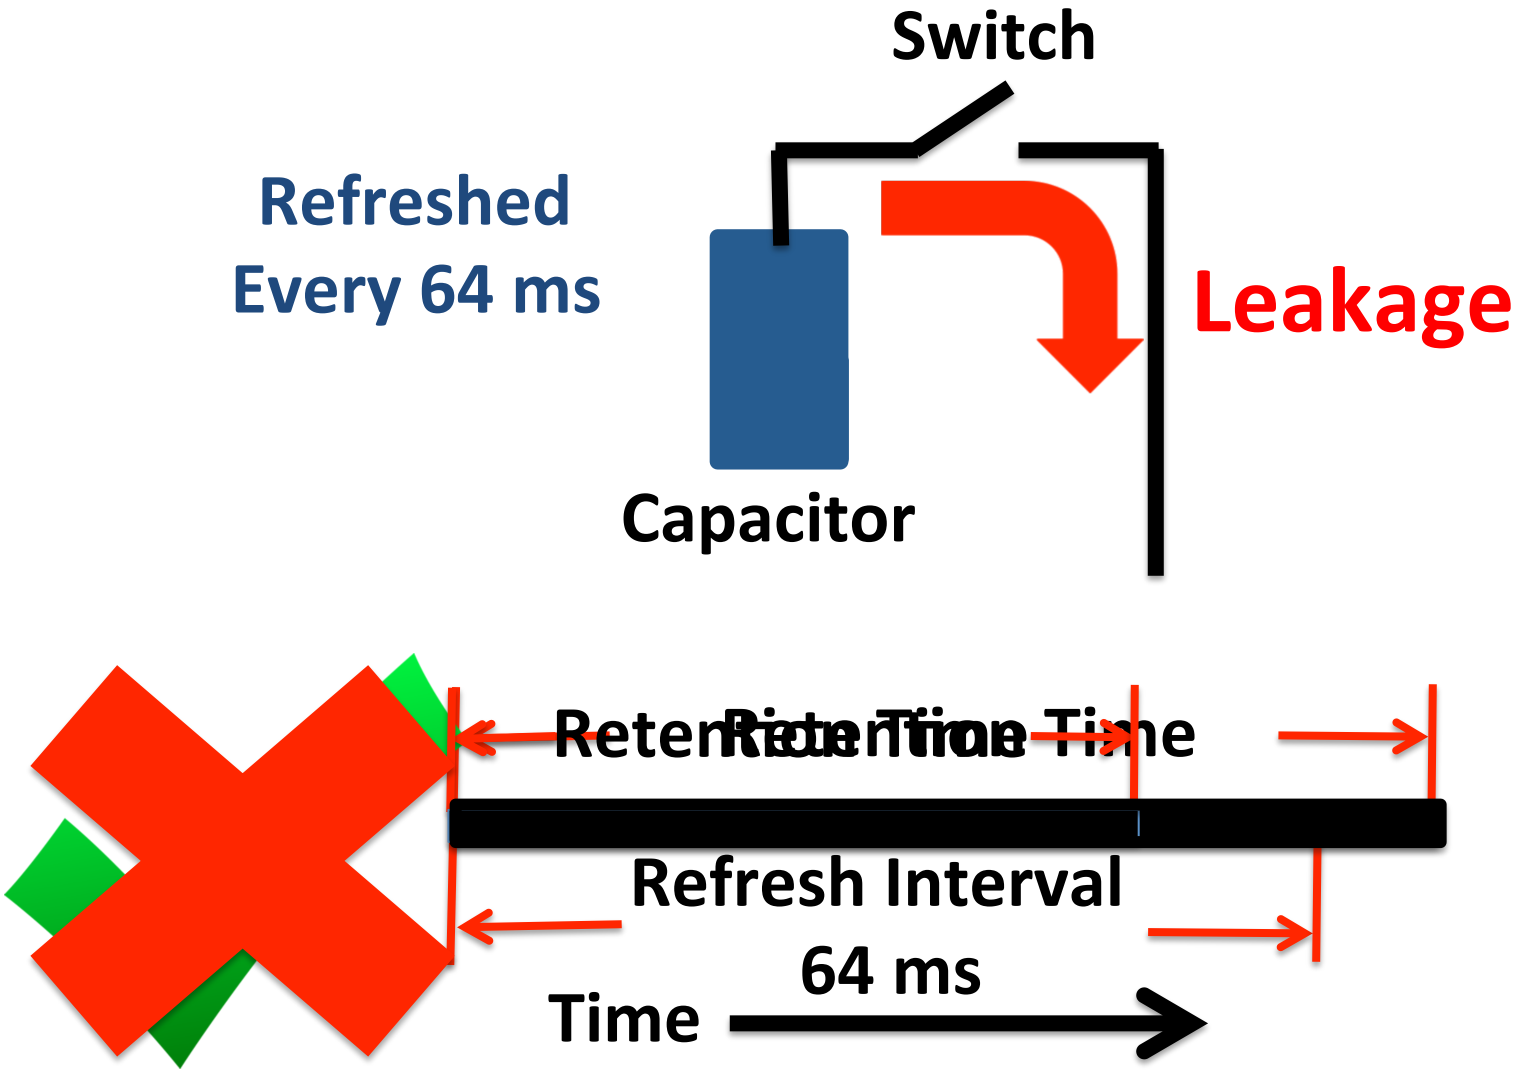

Intermittent Retention Failure

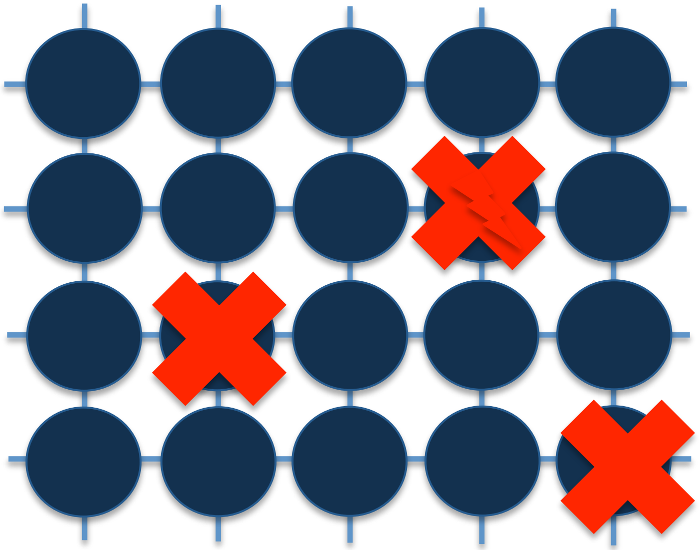

DRAM Cells

- Some retention failures are intermittent
- Two characteristics of intermittent retention failures

1. **Data Pattern Sensitivity**
2. **Variable Retention Time**

❶ Data Pattern Sensitivity

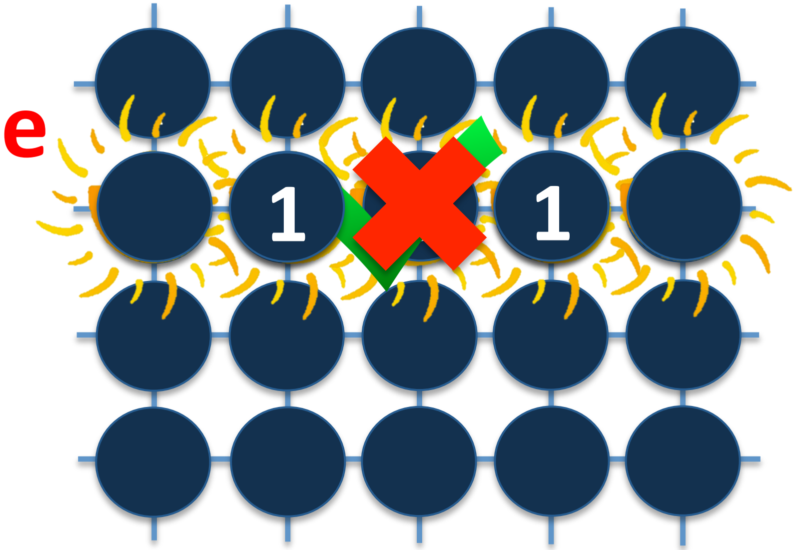

Noise Interference

No Failure

Failure

Some cells can fail depending on the data stored in neighboring cells

❷ Variable Retention Time

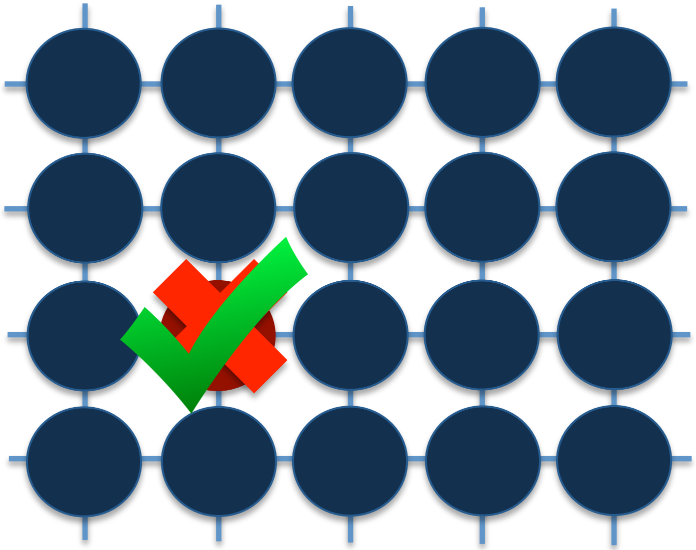

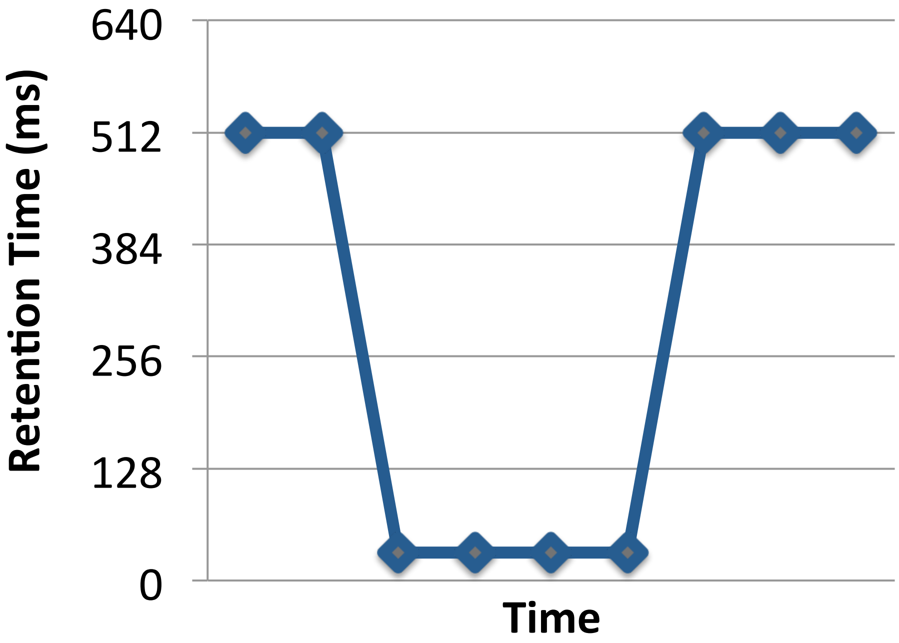

Retention time of some cells change at random points of time

Testing for Retention Failures

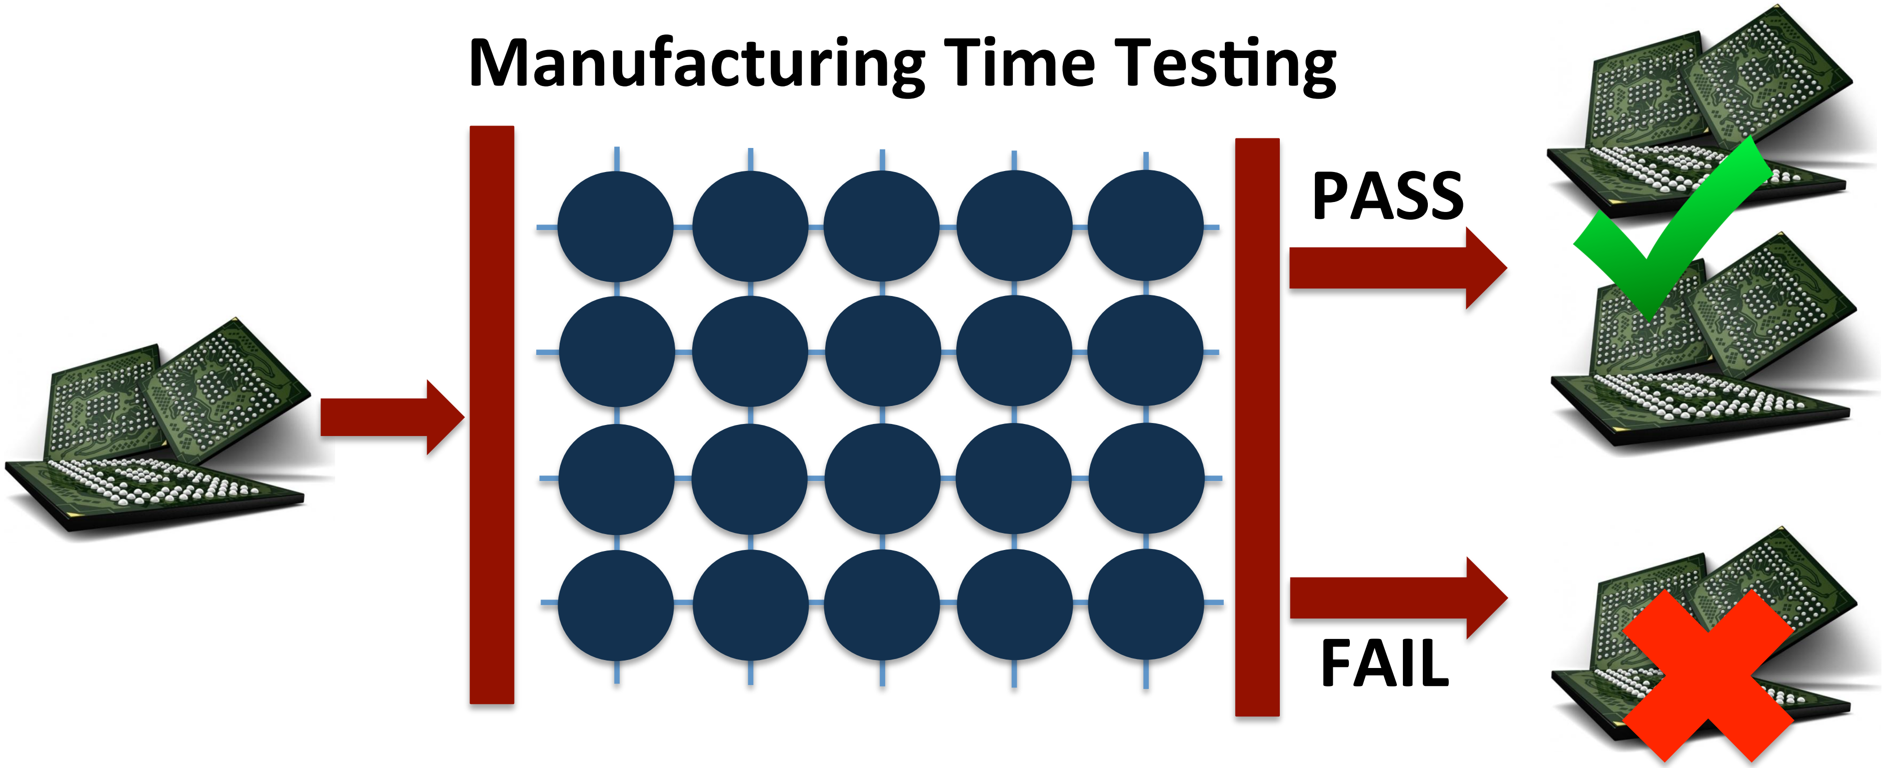
Manufacturing Time Testing

Manufacturers perform exhaustive testing of DRAM Chips

Chips failing the tests are discarded

DRAM Scaling Problem

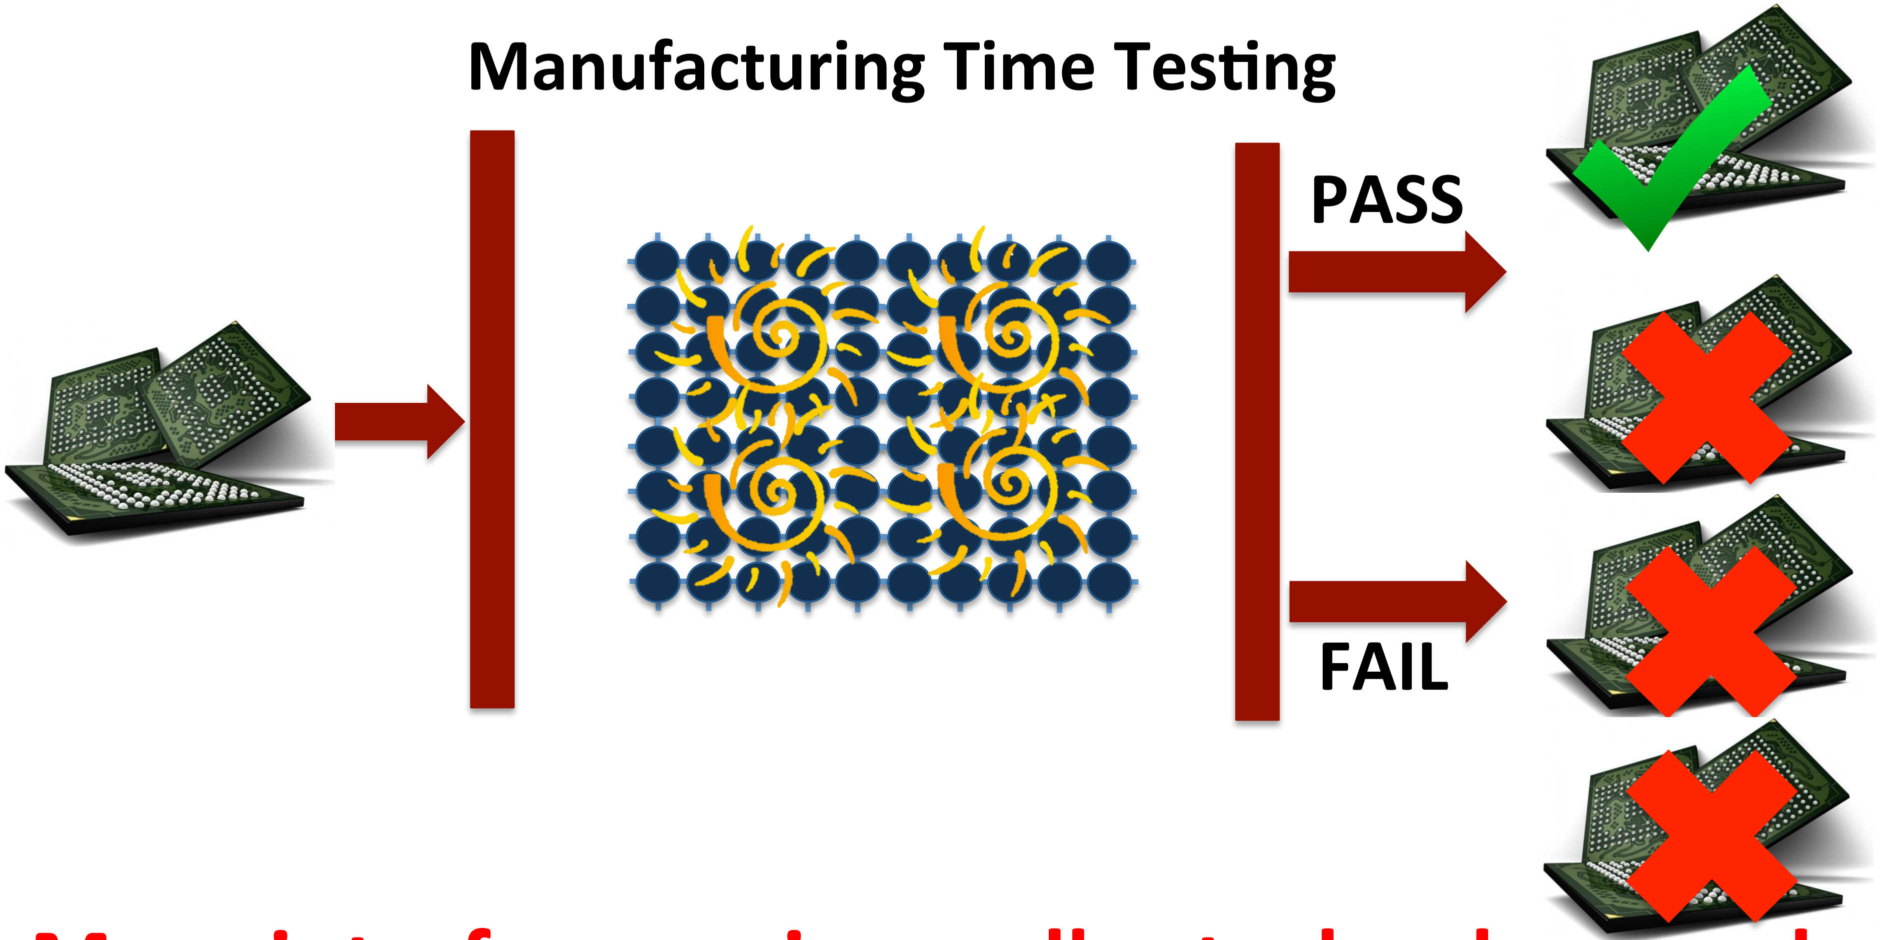

Manufacturing Time Testing

PASS

FAIL

More interference in smaller technology nodes leads to lower yield and higher cost

Outline

- DRAM Scaling Problem
- **Online Profiling as a Solution**
- Efficacy of System-Level Detection and Mitigation
  - Simple Techniques
  - Recently Proposed Techniques
- Towards an Online Profiling System
- Conclusion

System-Level Online Profiling

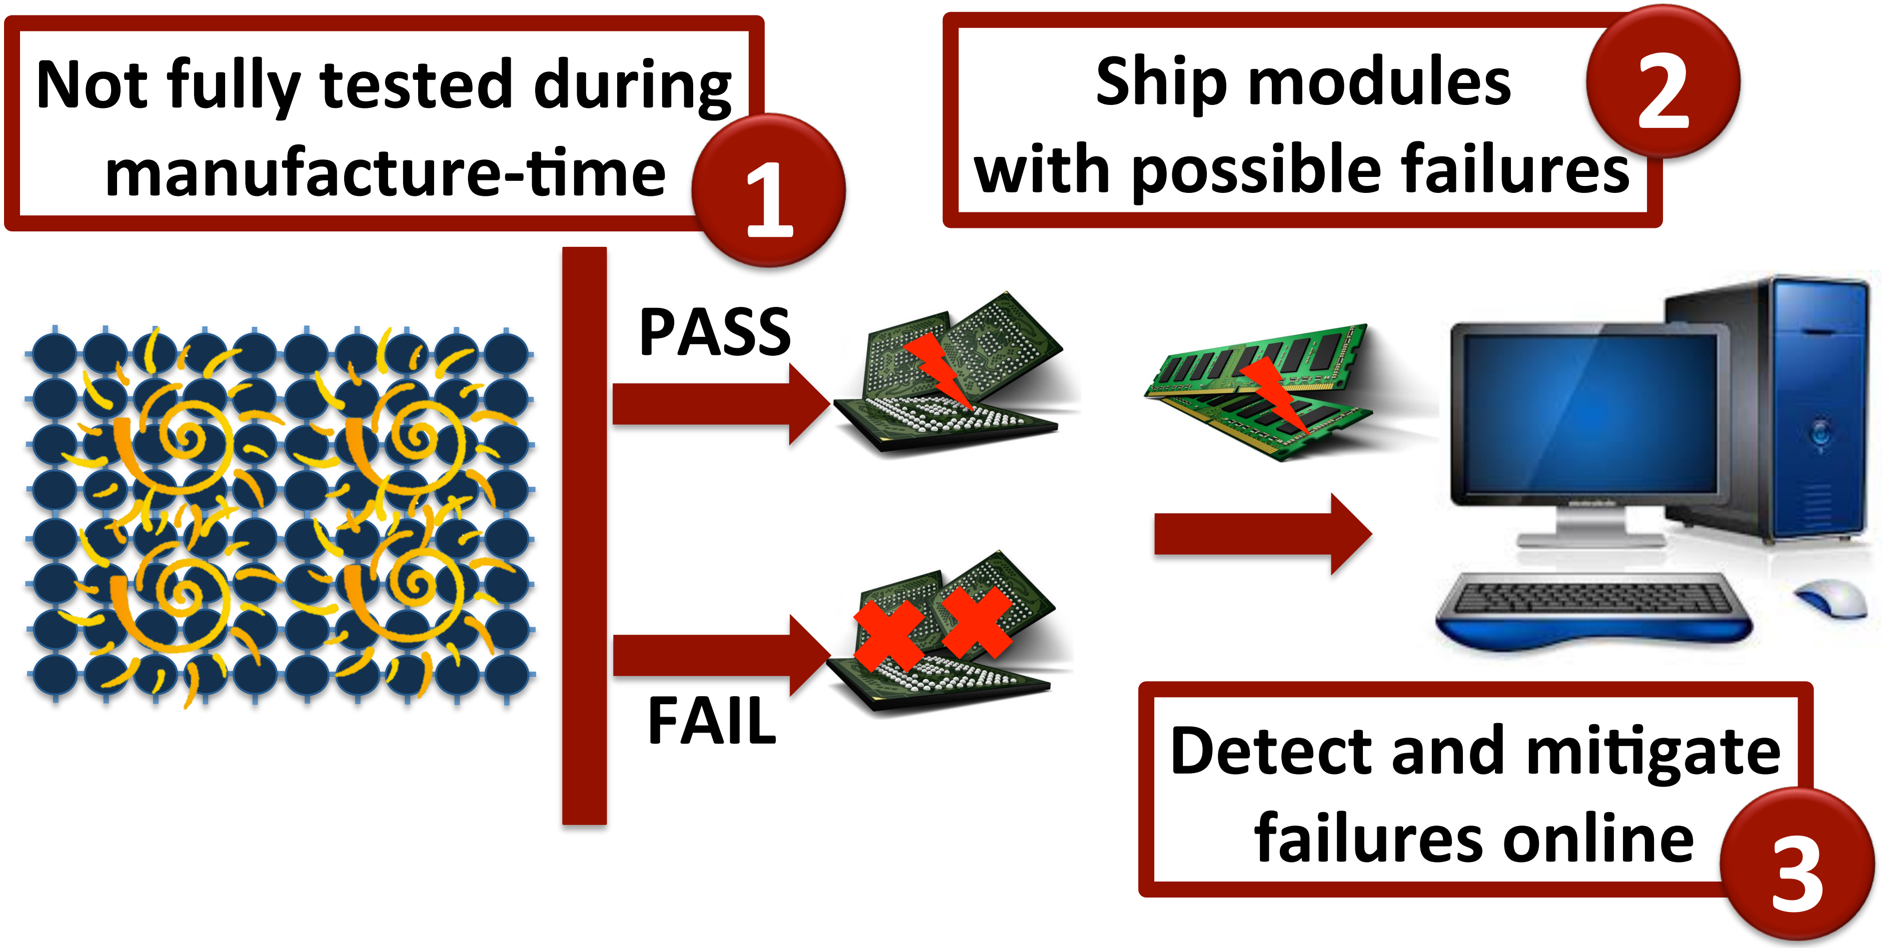

1 Not fully tested during manufacture-time

PASS

FAIL

2 Ship modules with possible failures

3 Detect and mitigate failures online

Increases yield, reduces cost, enables scaling

System-Level Online Profiling

What is the effectiveness of detection and mitigation techniques for retention failures?

Our goal is to analyze the efficacy of

1. Simple Techniques

- **Testing, Guardbanding, ECC**

2. Recently Proposed Techniques

- **ArchShield, RAIDR, SECRET, RAPID, VS-ECC, Hi-ECC**

We analyze the effectiveness of these techniques using experimental data from real DRAM

ArchShield ISCA'13, RAIDR ISCA'12, SECRET ICCD'12, RAPID HPCA'06,
VS-ECC ISCA'11, Hi-ECC ISCA'10

Methodology

FPGA-based testing infrastructure

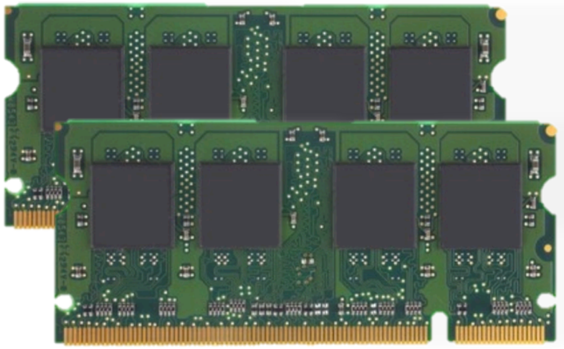

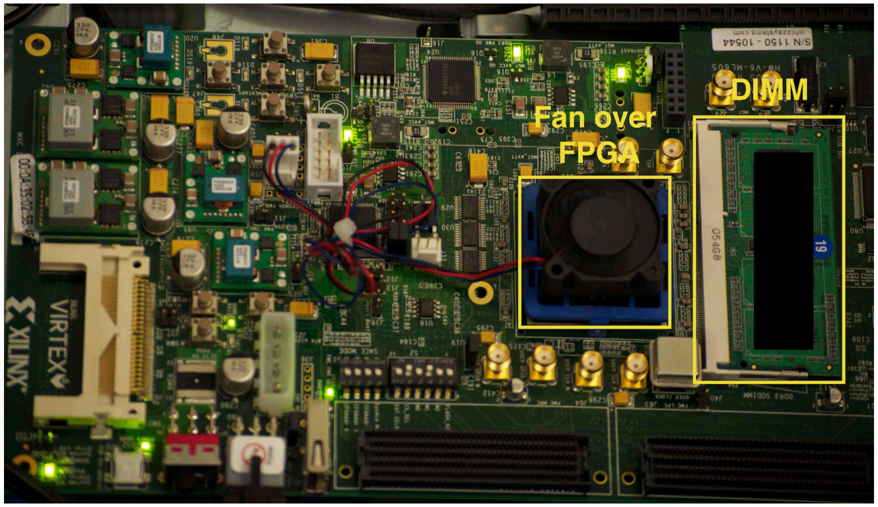

Evaluated 96 chips from three major vendors

Outline

- DRAM Scaling Problem
- Online Profiling as a Solution
- **Efficacy of System-Level Detection and Mitigation**
  - **Simple Techniques**
  - Recently Proposed Techniques
- Towards an Online Profiling System
- Conclusion

Efficacy of Simple Techniques

1. Testing
2. Guardbanding
3. Error Correcting Code

① Testing

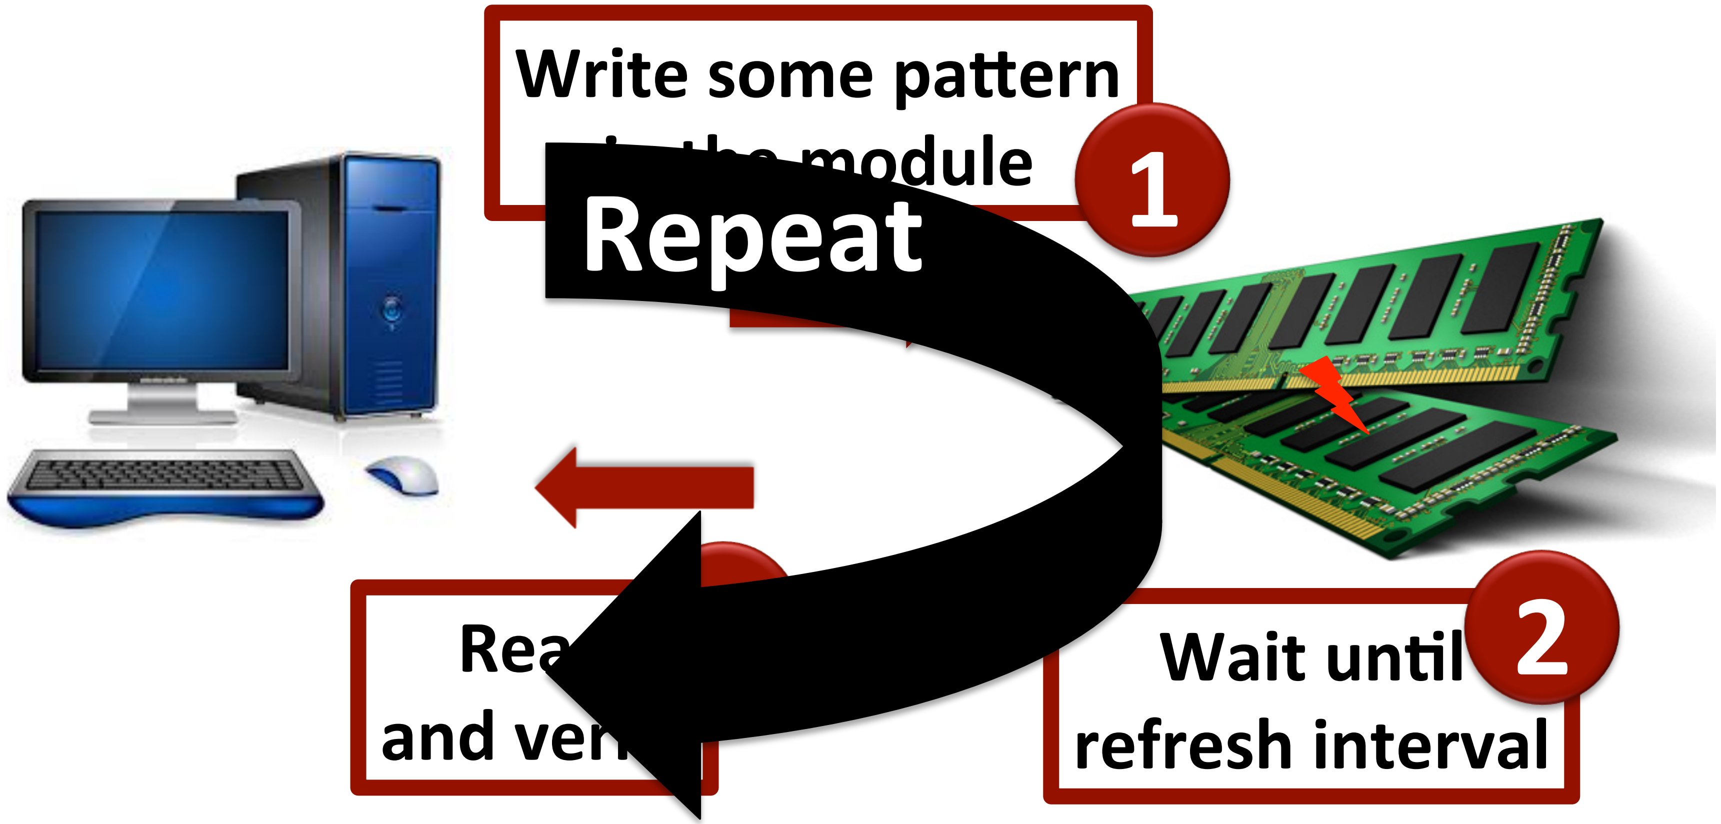

Write some pattern in the module

1

Repeat

Wait until refresh interval

2

Rea

and ver

Test each module with different patterns for many rounds

Zeros (0000), Ones (1111), Tens (1010), Fives (0101), Random

Efficacy of Testing

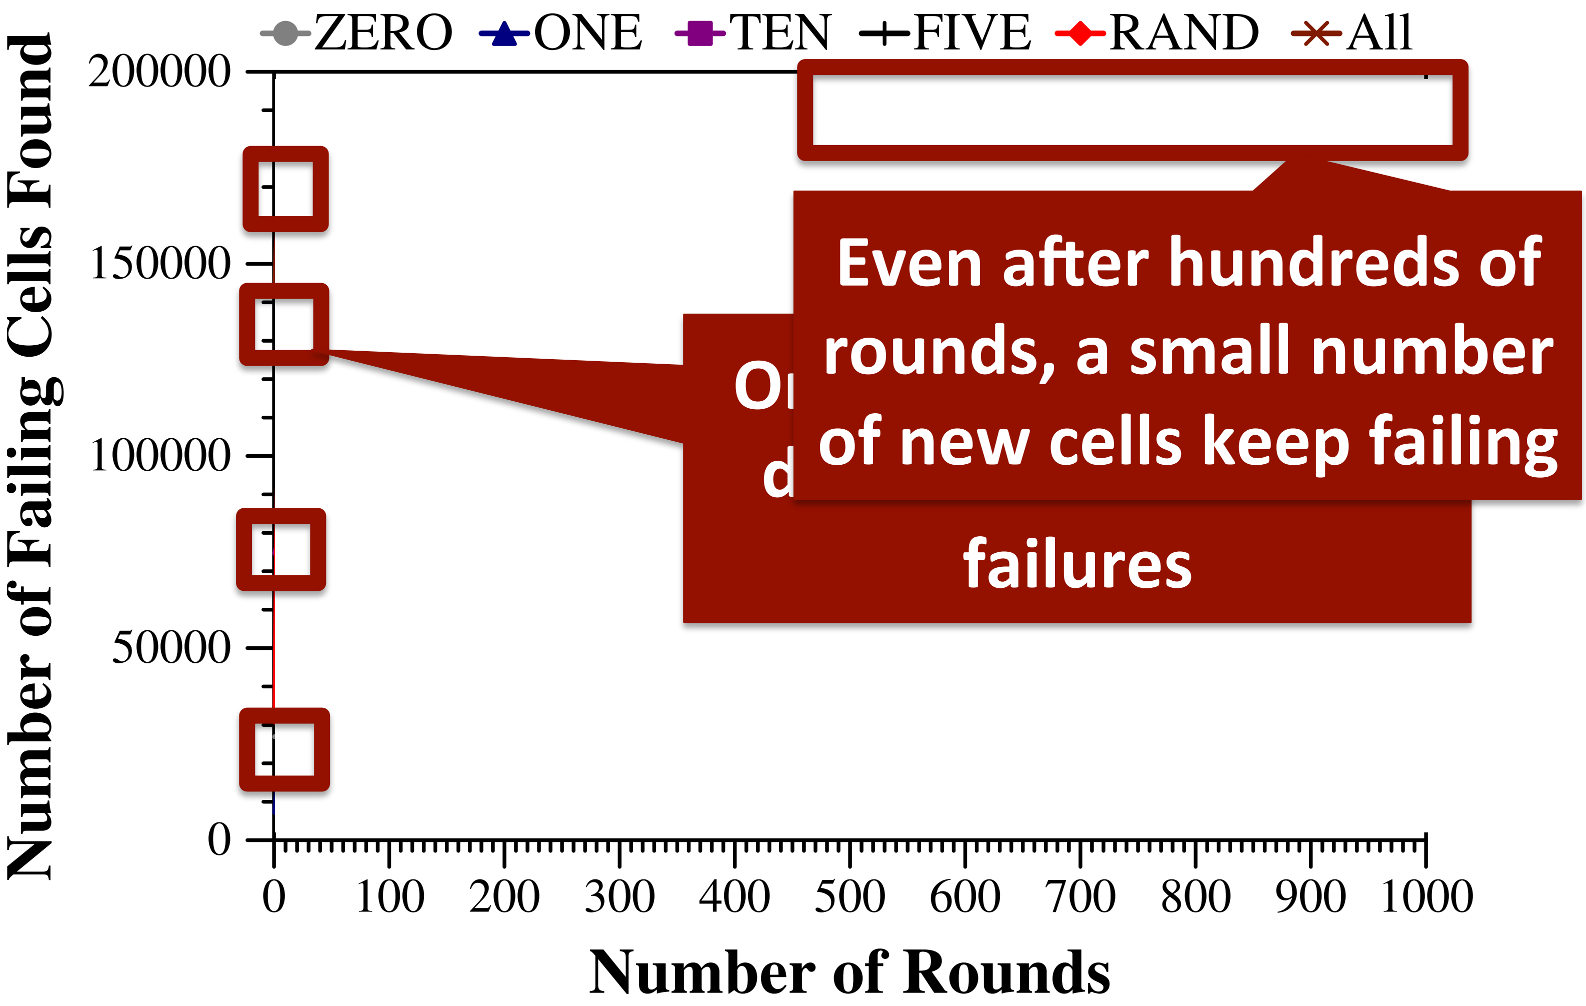

O

d

failures

Even after hundreds of rounds, a small number of new cells keep failing

Testing alone cannot detect all possible failures

② Guardbanding

- Adding a safety-margin on the refresh interval
- Can avoid VRT failures

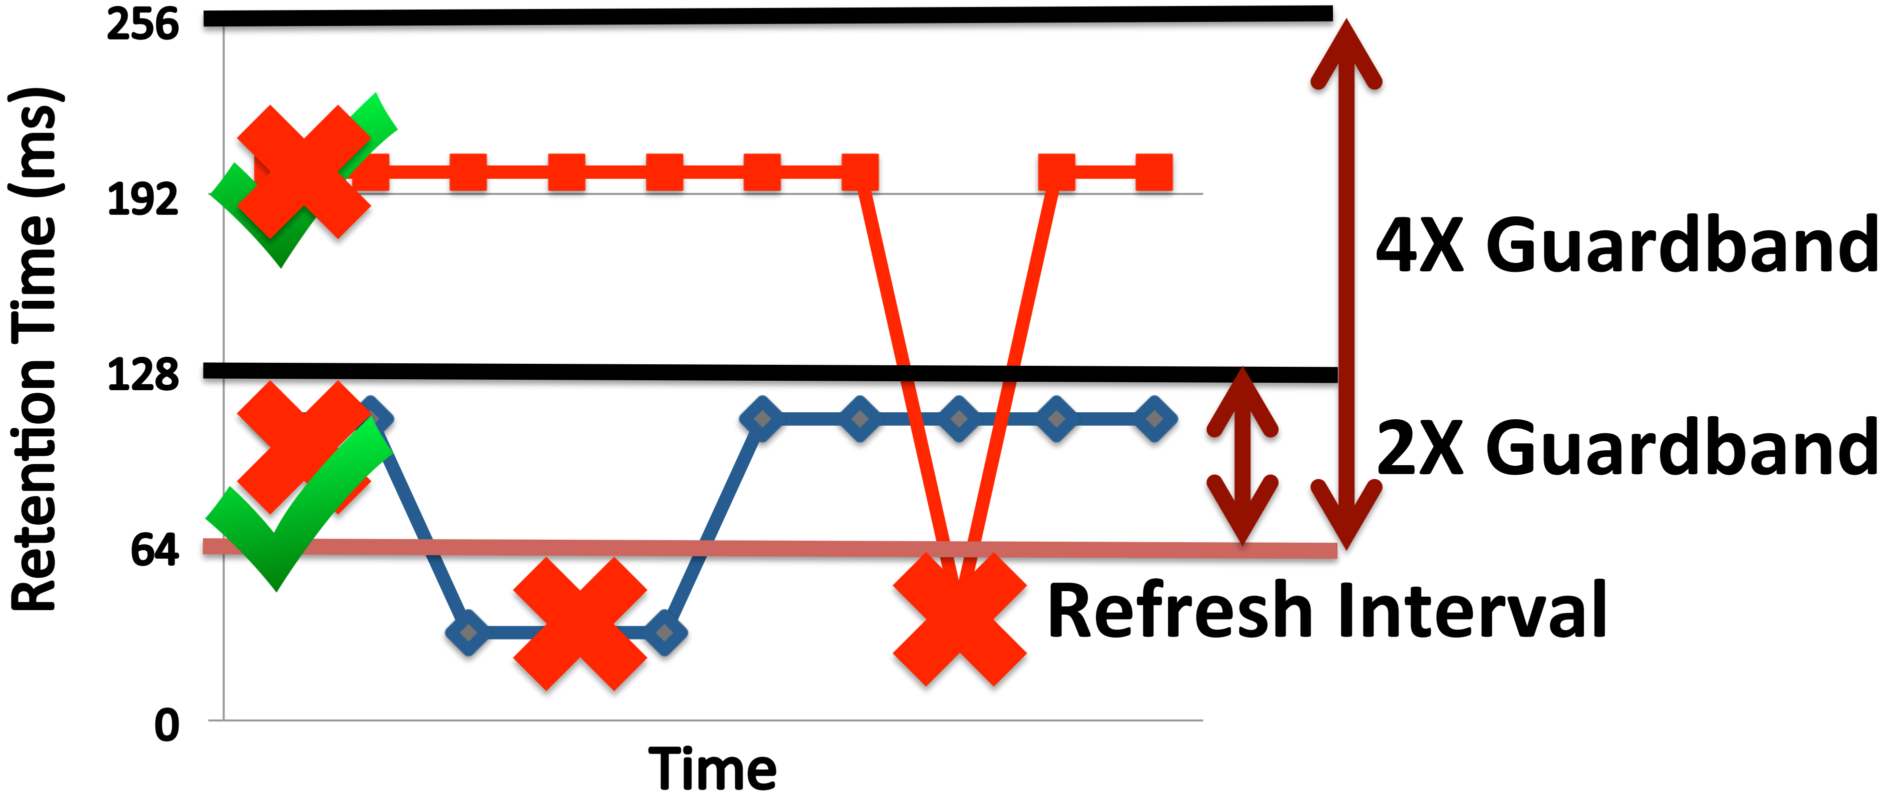

Effectiveness of guardbanding depends on the difference between retention times of a cell

Efficacy of Guardbanding

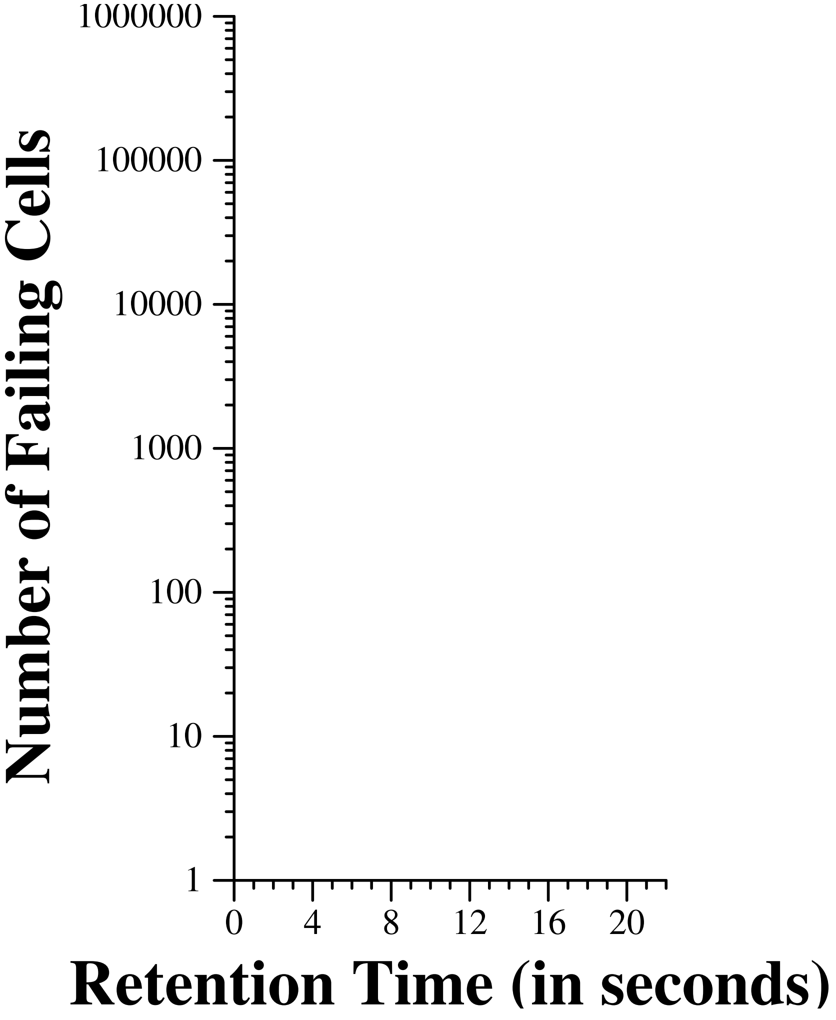

Efficacy of Guardbanding

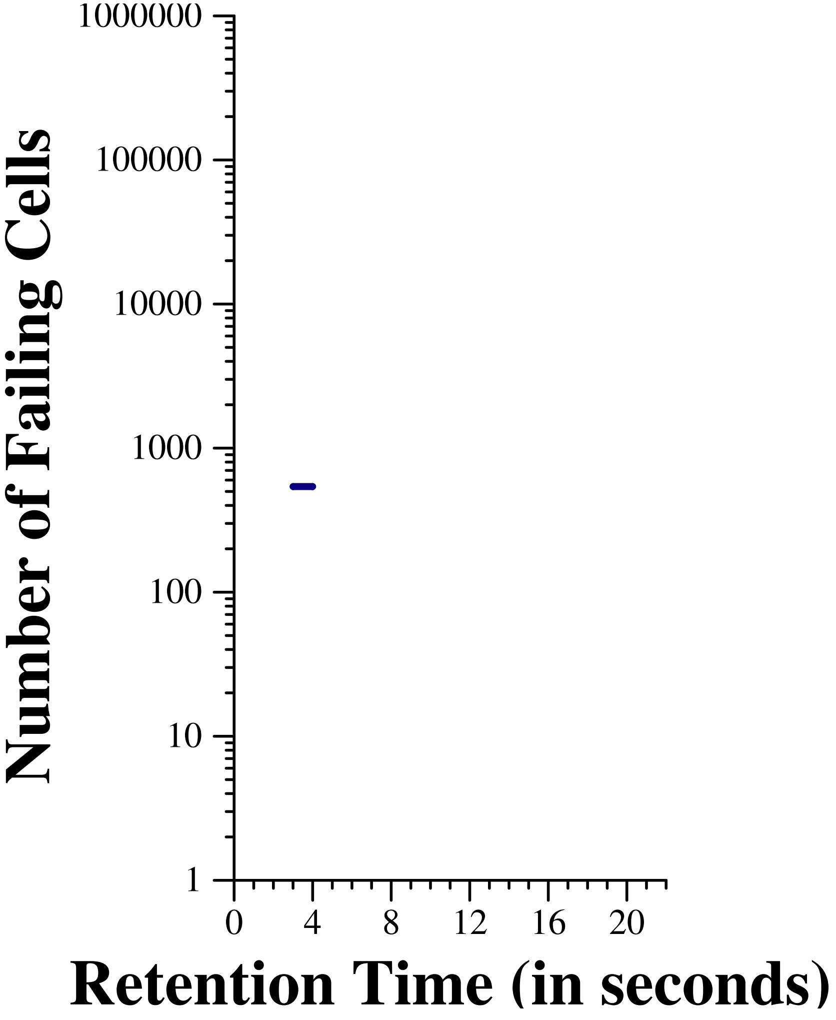

Efficacy of Guardbanding

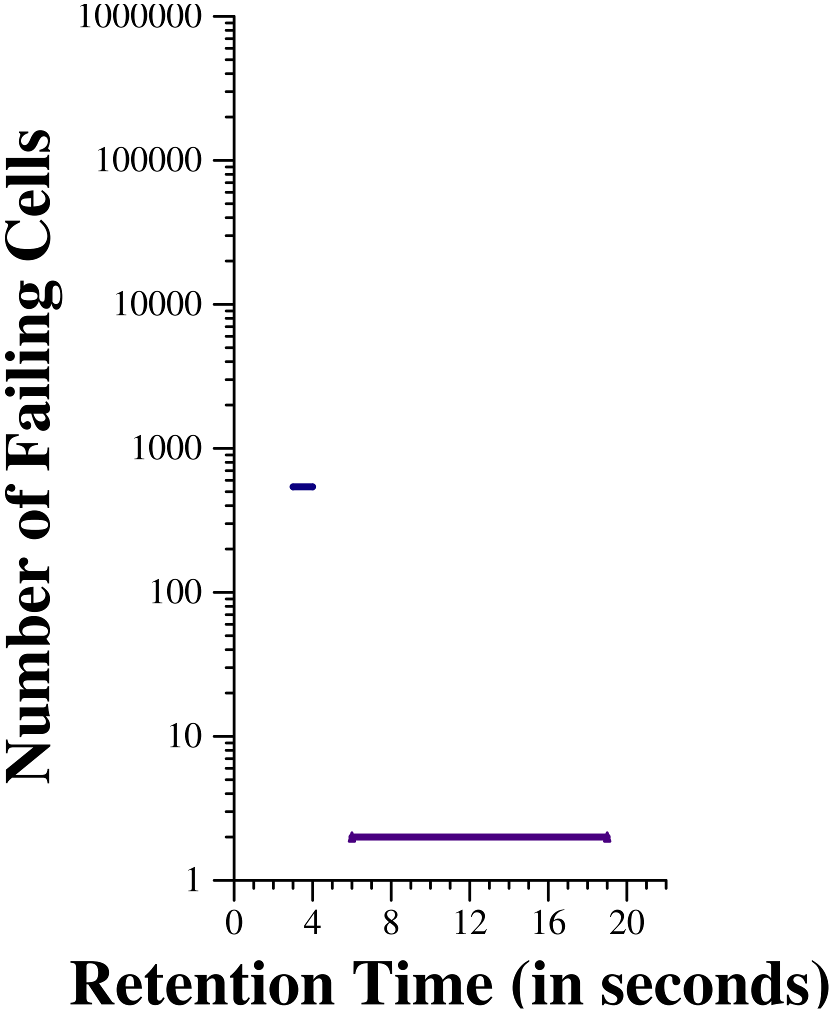

Efficacy of Guardbanding

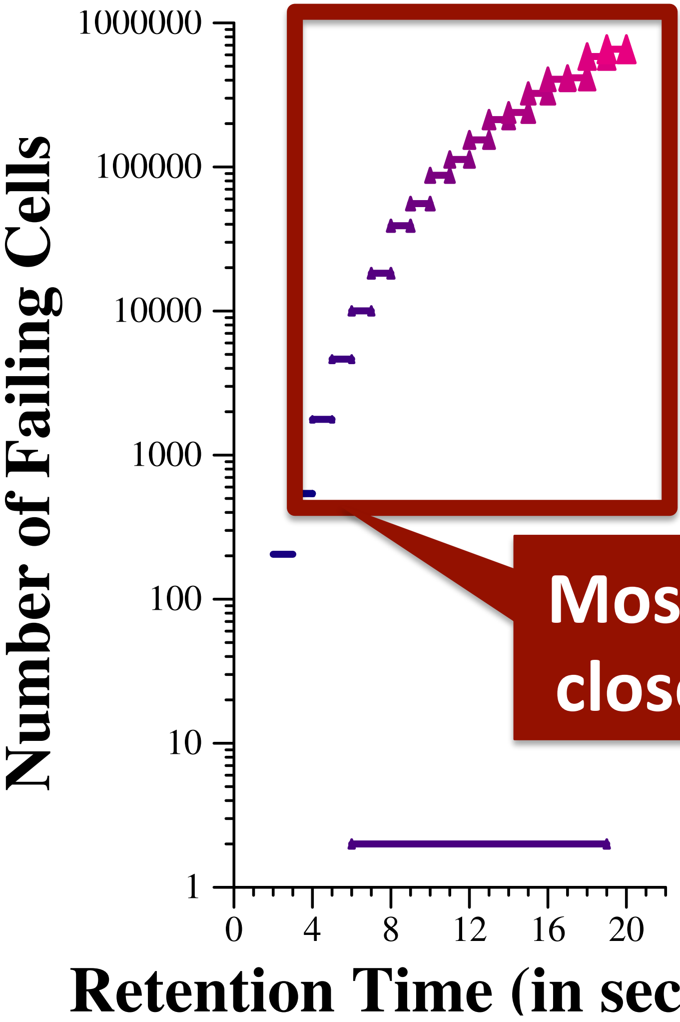

Most of the cells exhibit closeby retention times

Efficacy of Guardbanding

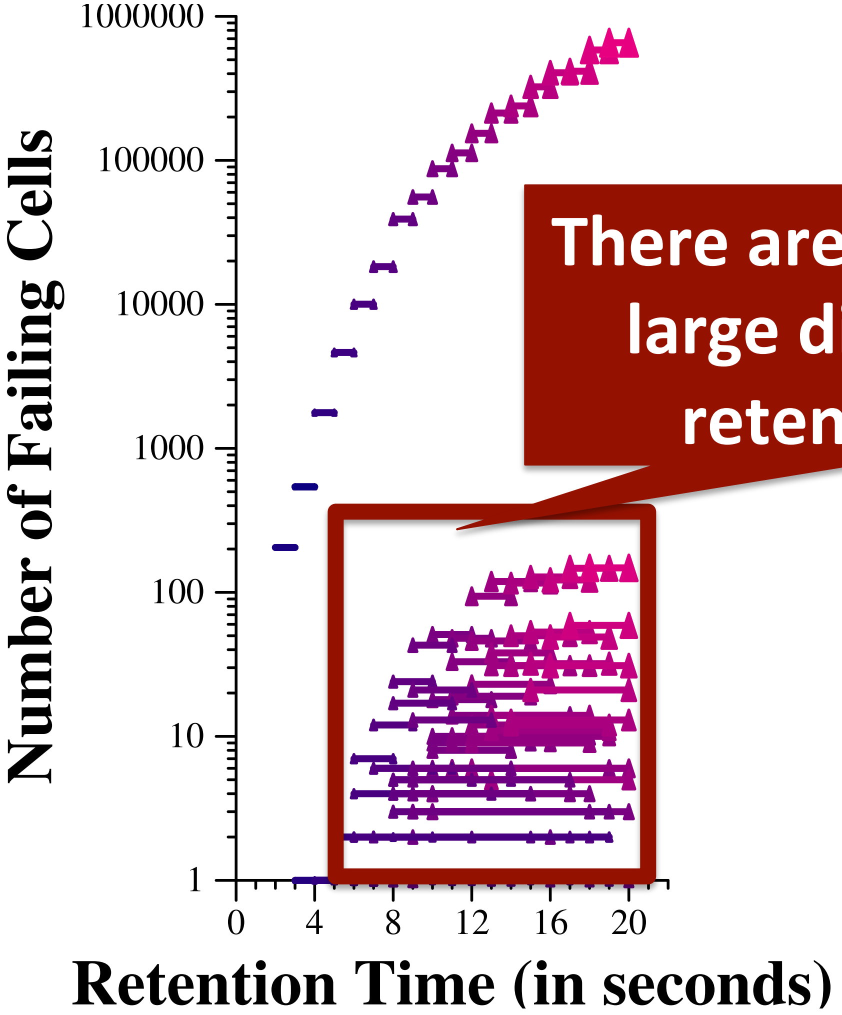

There are few cells with large differences in retention times

Even a large guardband (5X) cannot detect 5-15% of the intermittently failing cells

❸ Error Correcting Code

- Error Correcting Code (ECC)
  - **Additional information to detect error and correct data**

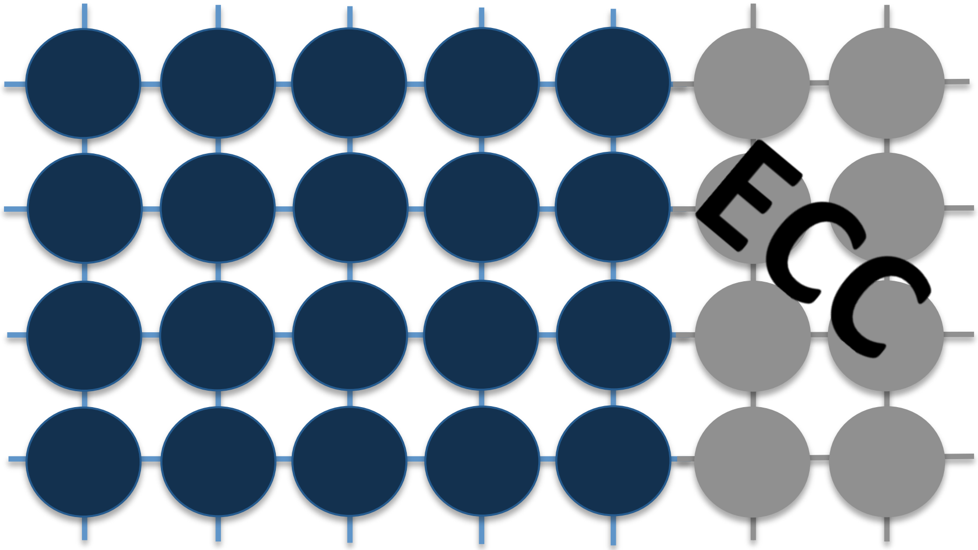

Effectiveness of ECC

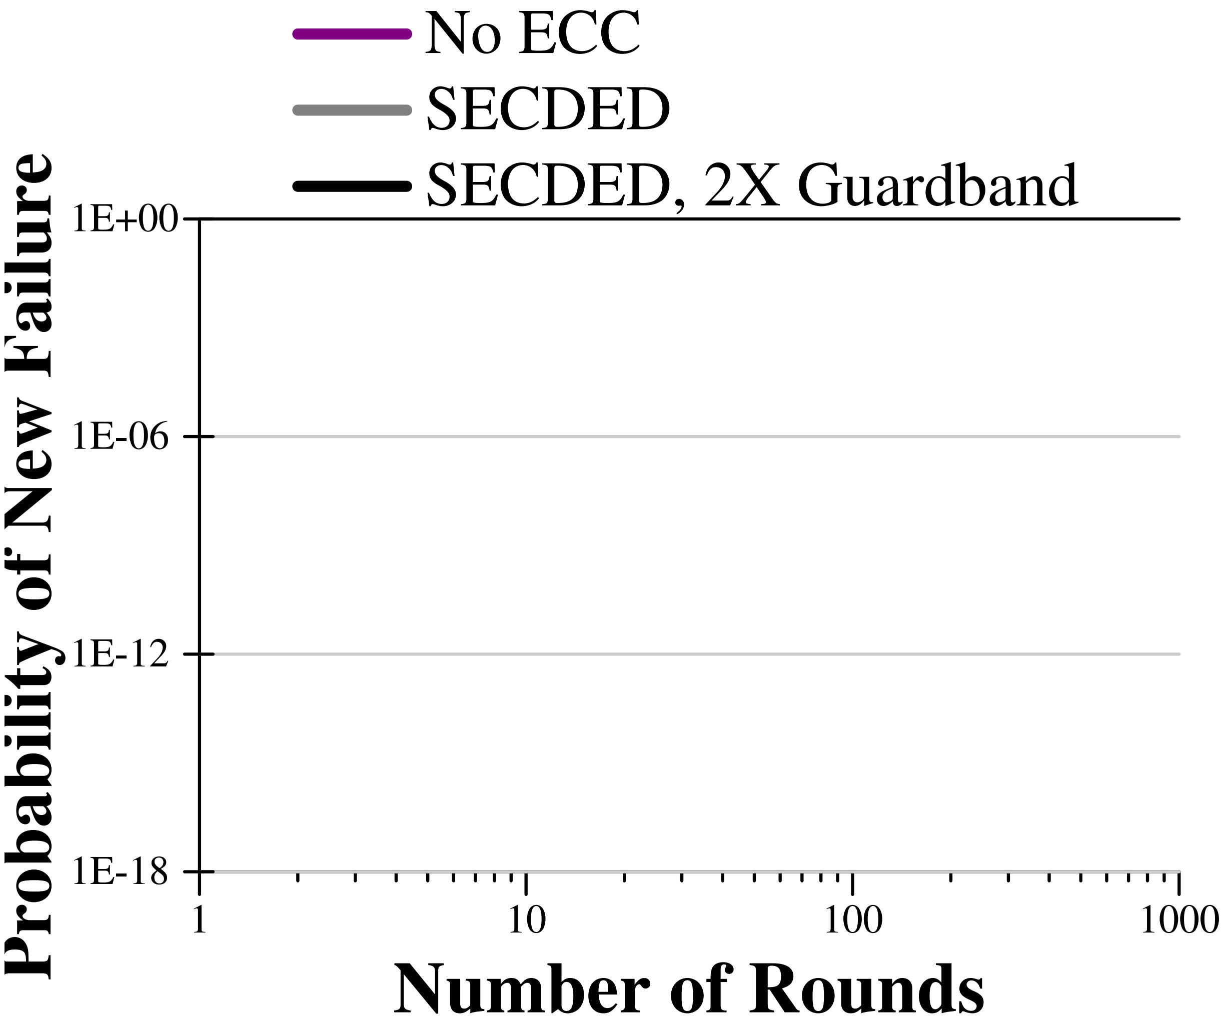

Effectiveness of ECC

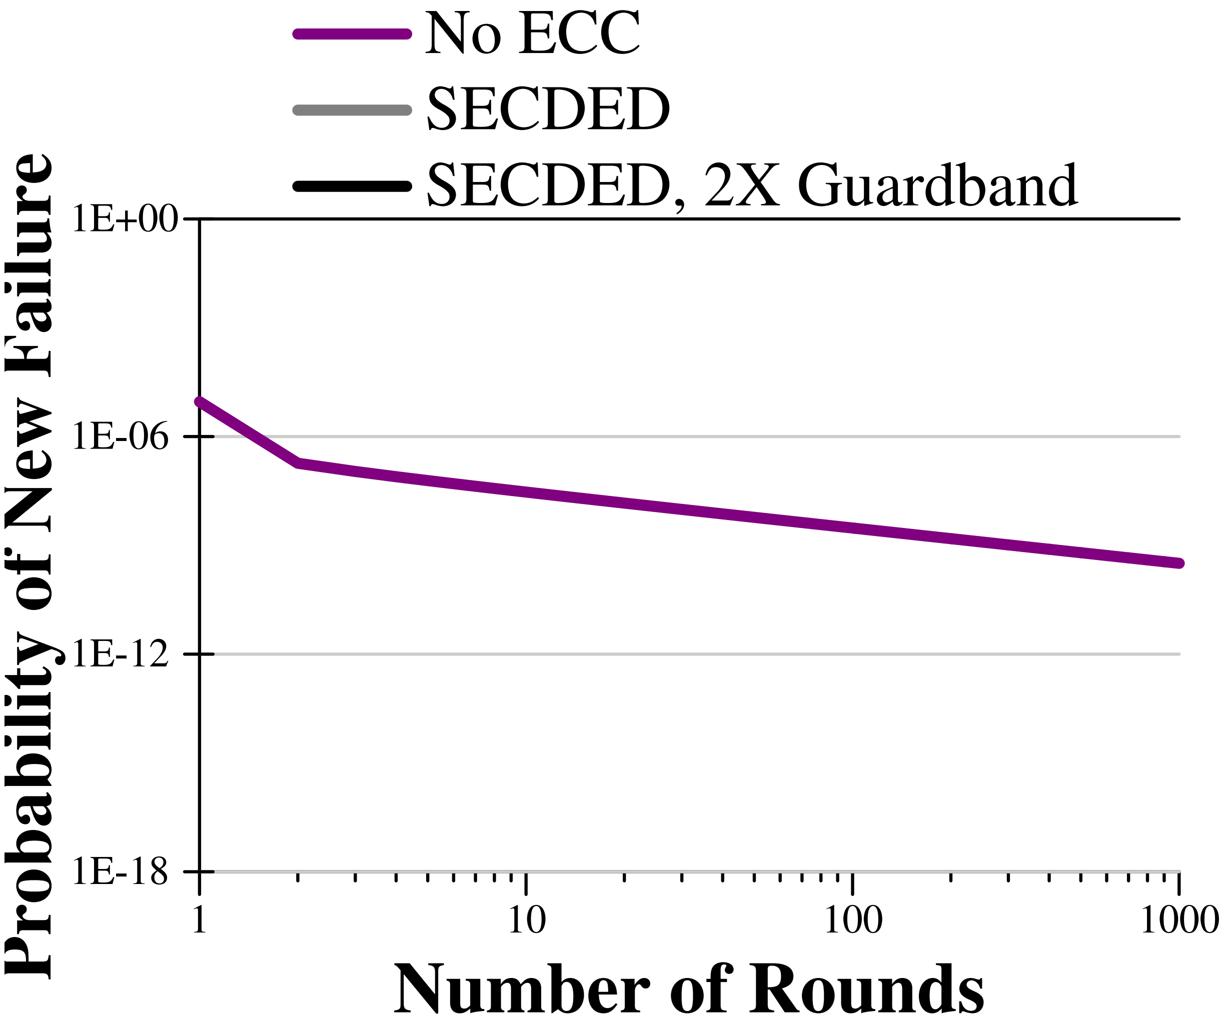

Effectiveness of ECC

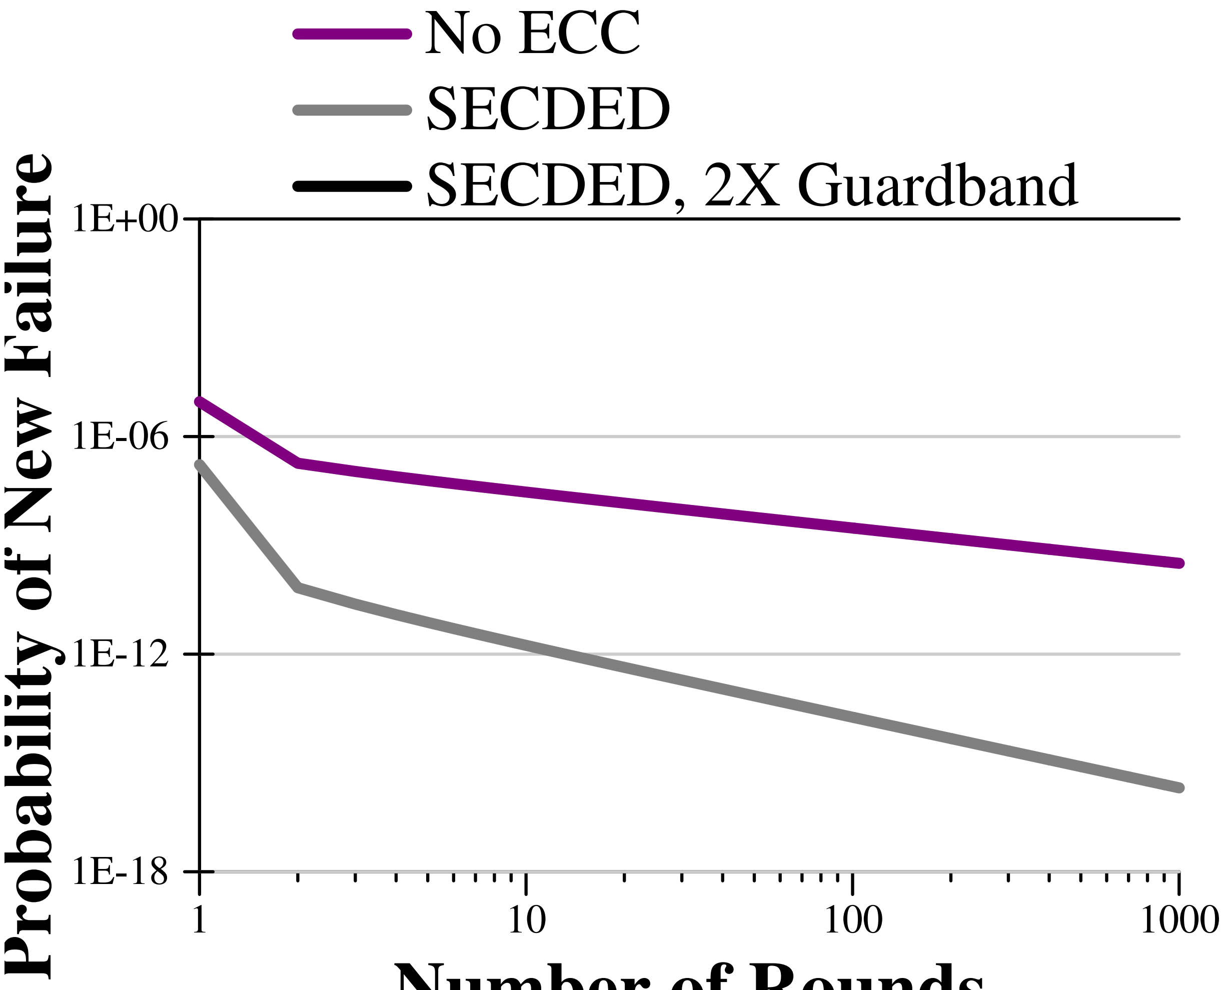

Effectiveness of ECC

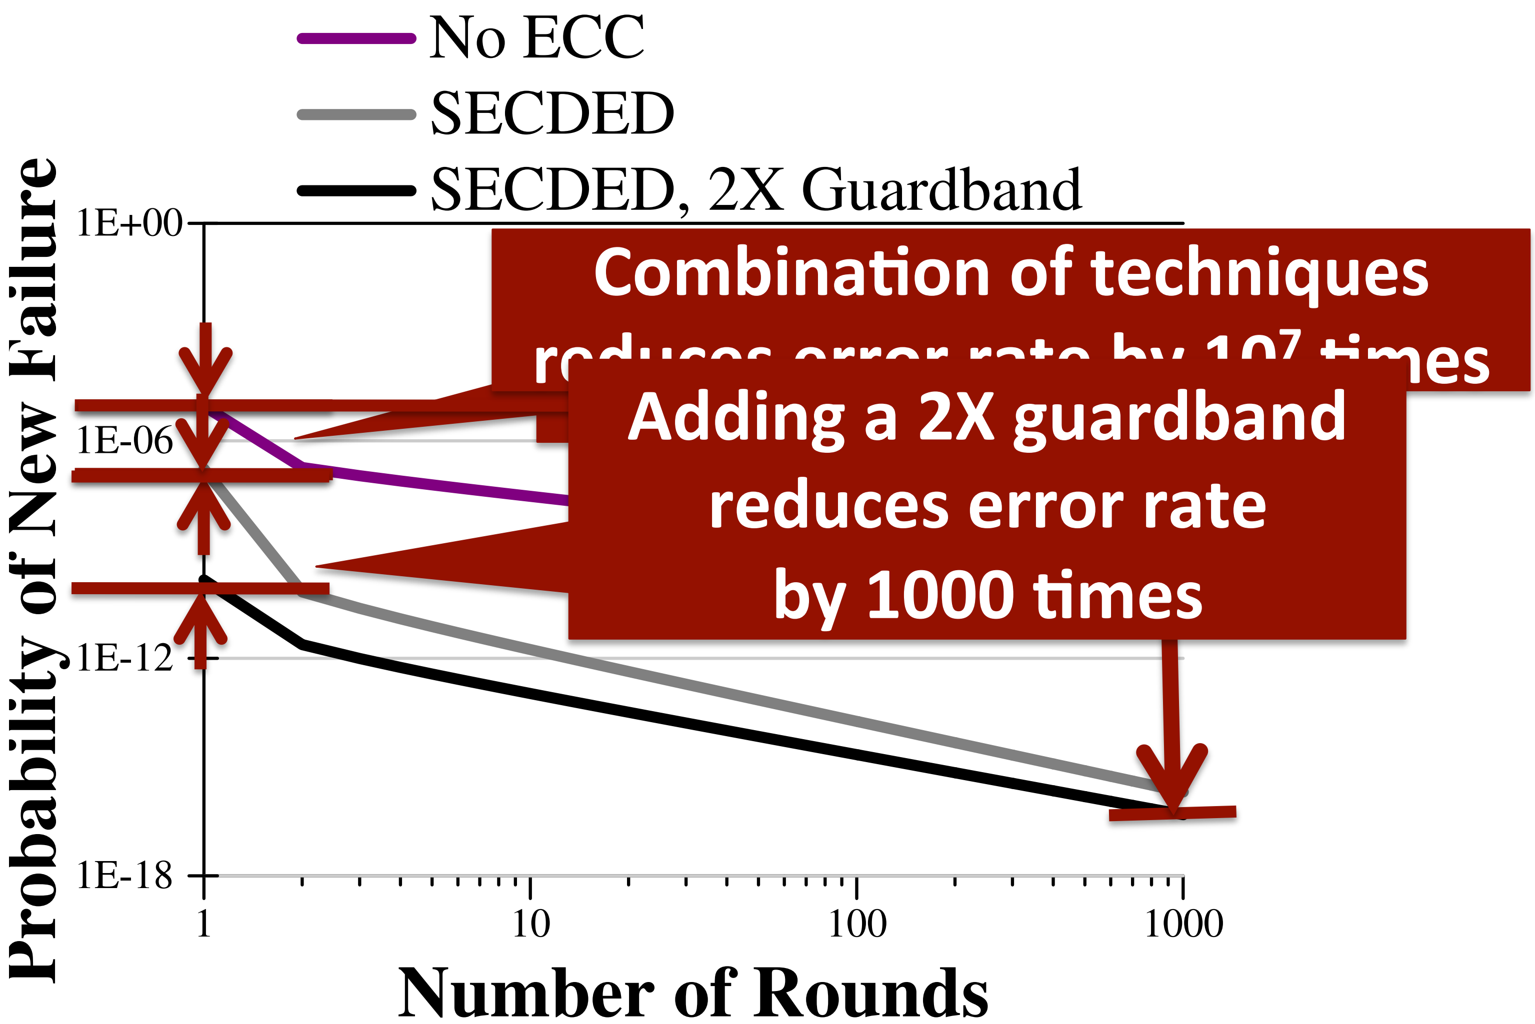

Combination of techniques reduces error rate by 10^7 times

Adding a 2X guardband reduces error rate by 1000 times

A combination of mitigation techniques is much more effective

Outline

- DRAM Scaling Problem
- Online Profiling as a Solution
- **Efficacy of System-Level Detection and Mitigation**
  - Simple Techniques
  - **Recently Proposed Techniques**
- Towards an Online Profiling System
- Conclusion

Efficacy of Recent Techniques

1. Bit Repair Techniques

In the paper

2. Variable-Strength ECC
3. **Higher-Strength ECC**

Higher Strength ECC (Hi-ECC)

No testing, use strong ECC
But amortize cost of ECC over larger data chunk

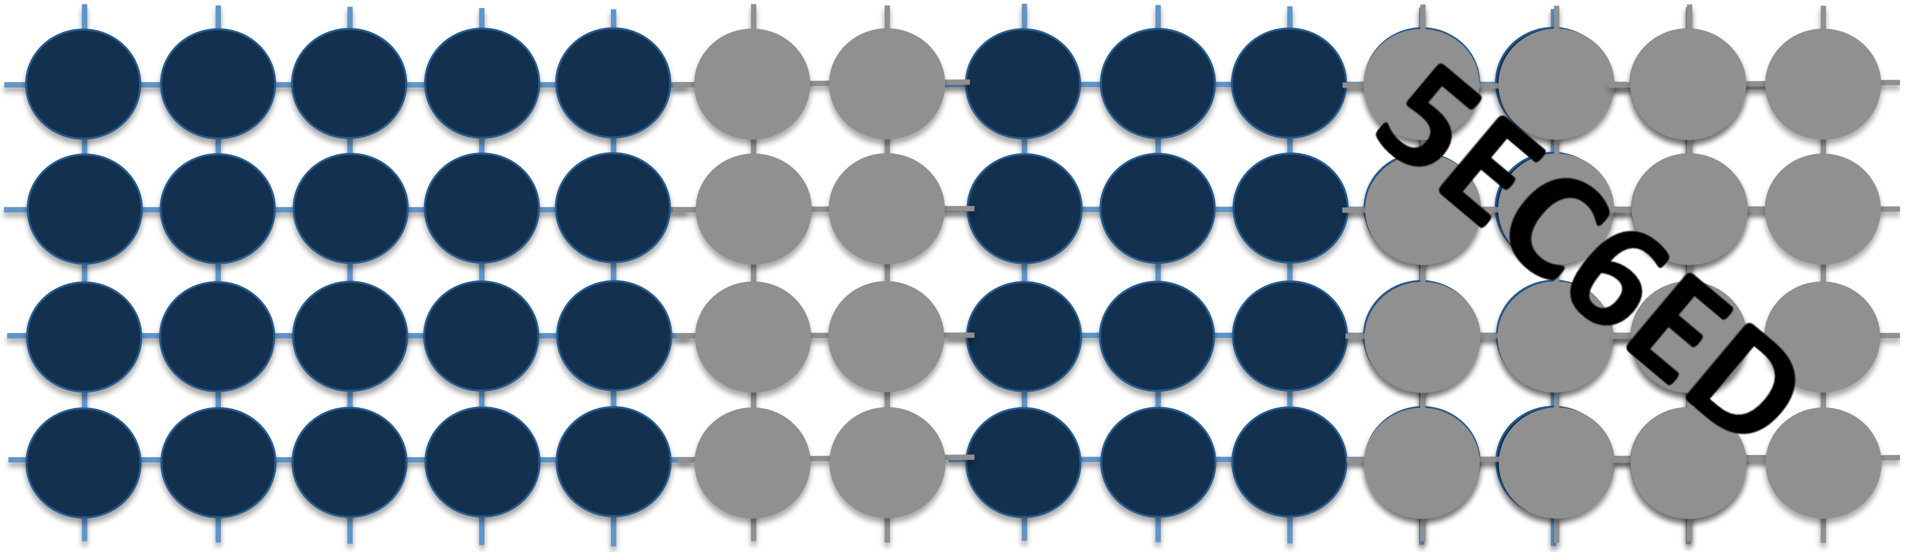

Can potentially tolerate errors at the cost of higher strength ECC

Efficacy of Hi-ECC

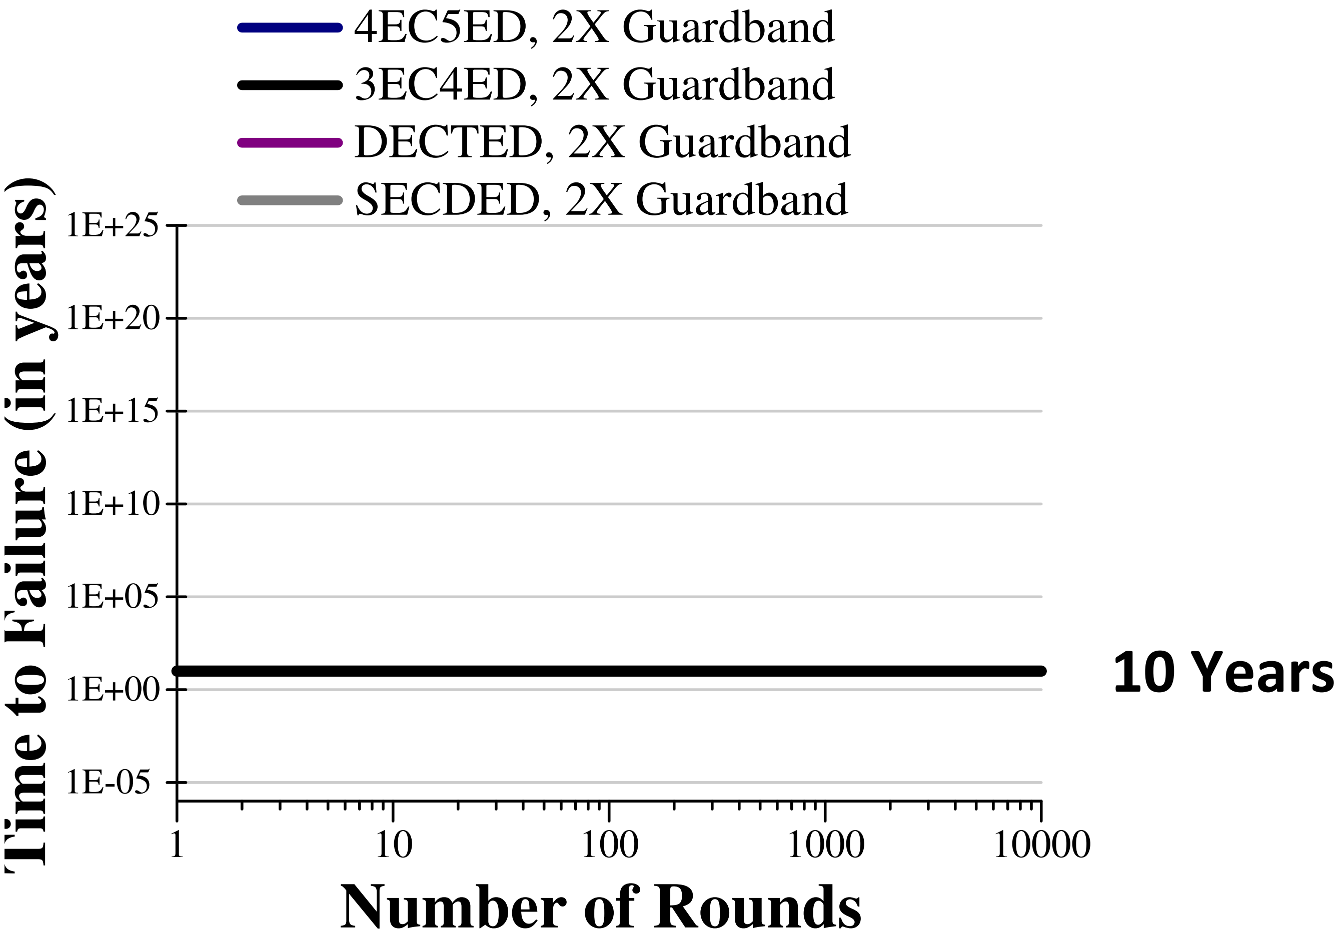

Efficacy of Hi-ECC

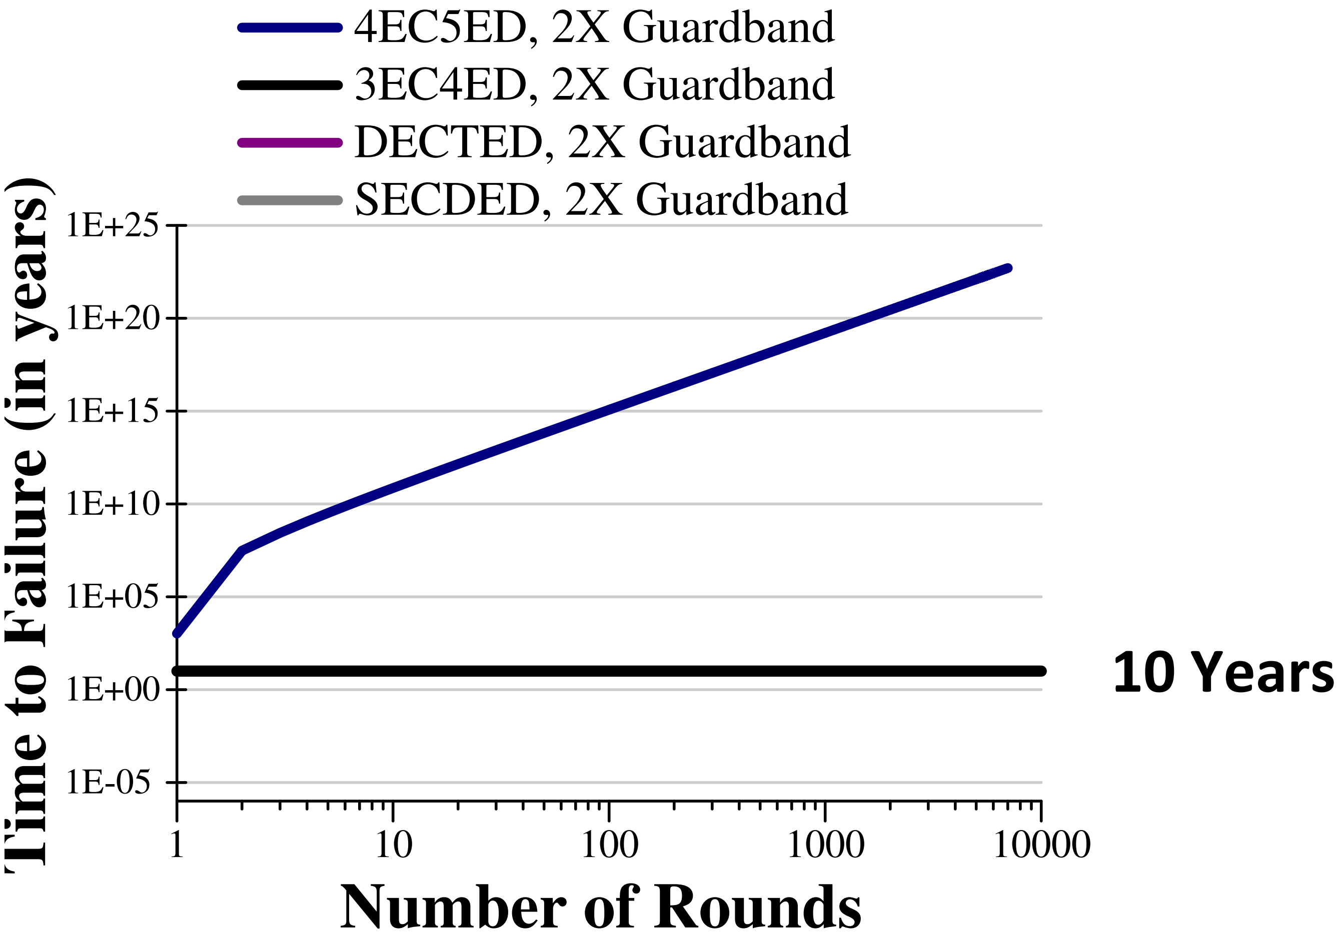

Efficacy of Hi-ECC

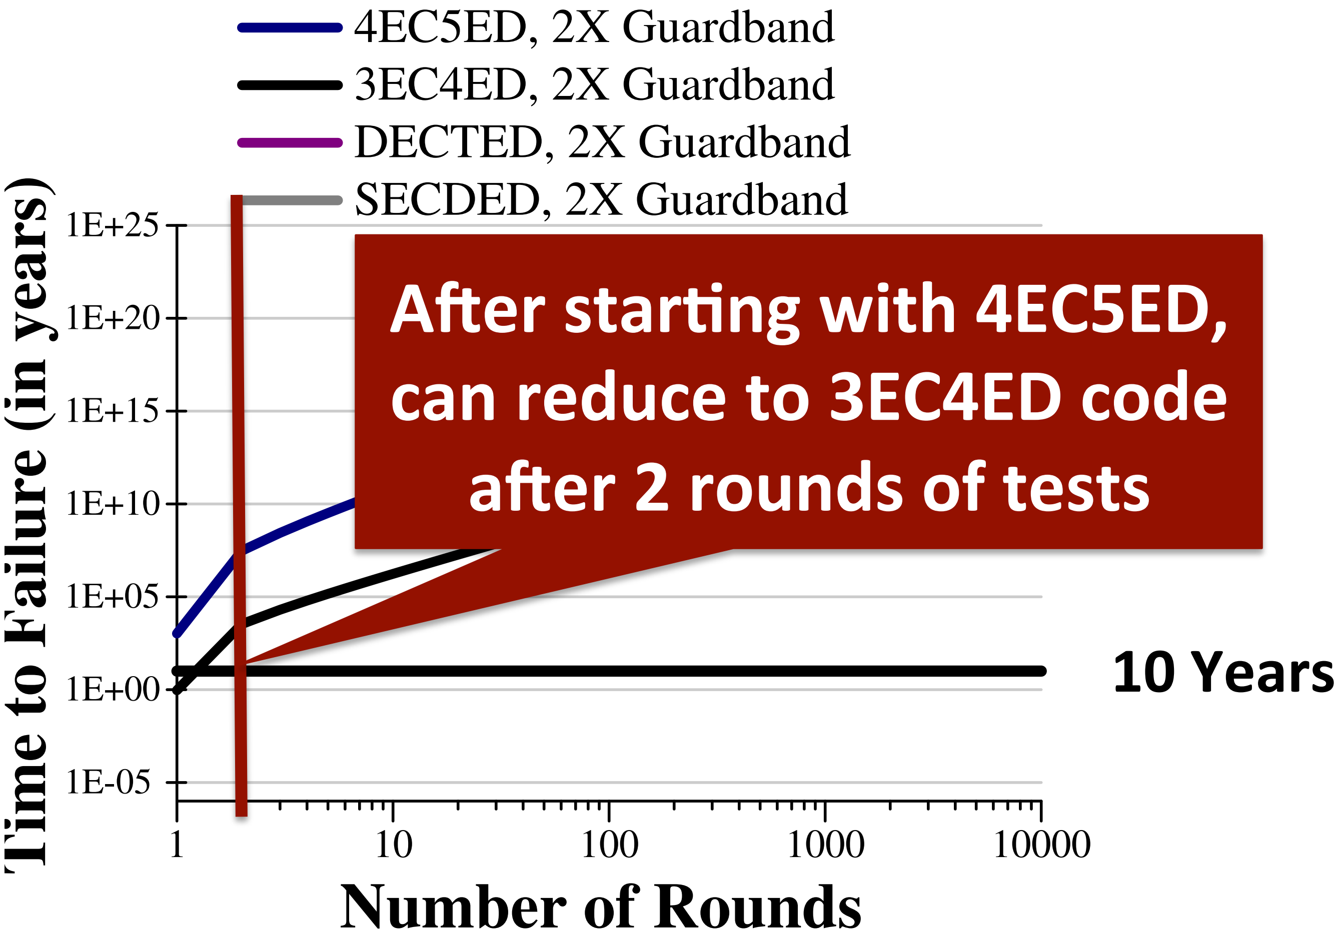

After starting with 4EC5ED, can reduce to 3EC4ED code after 2 rounds of tests

Efficacy of Hi-ECC

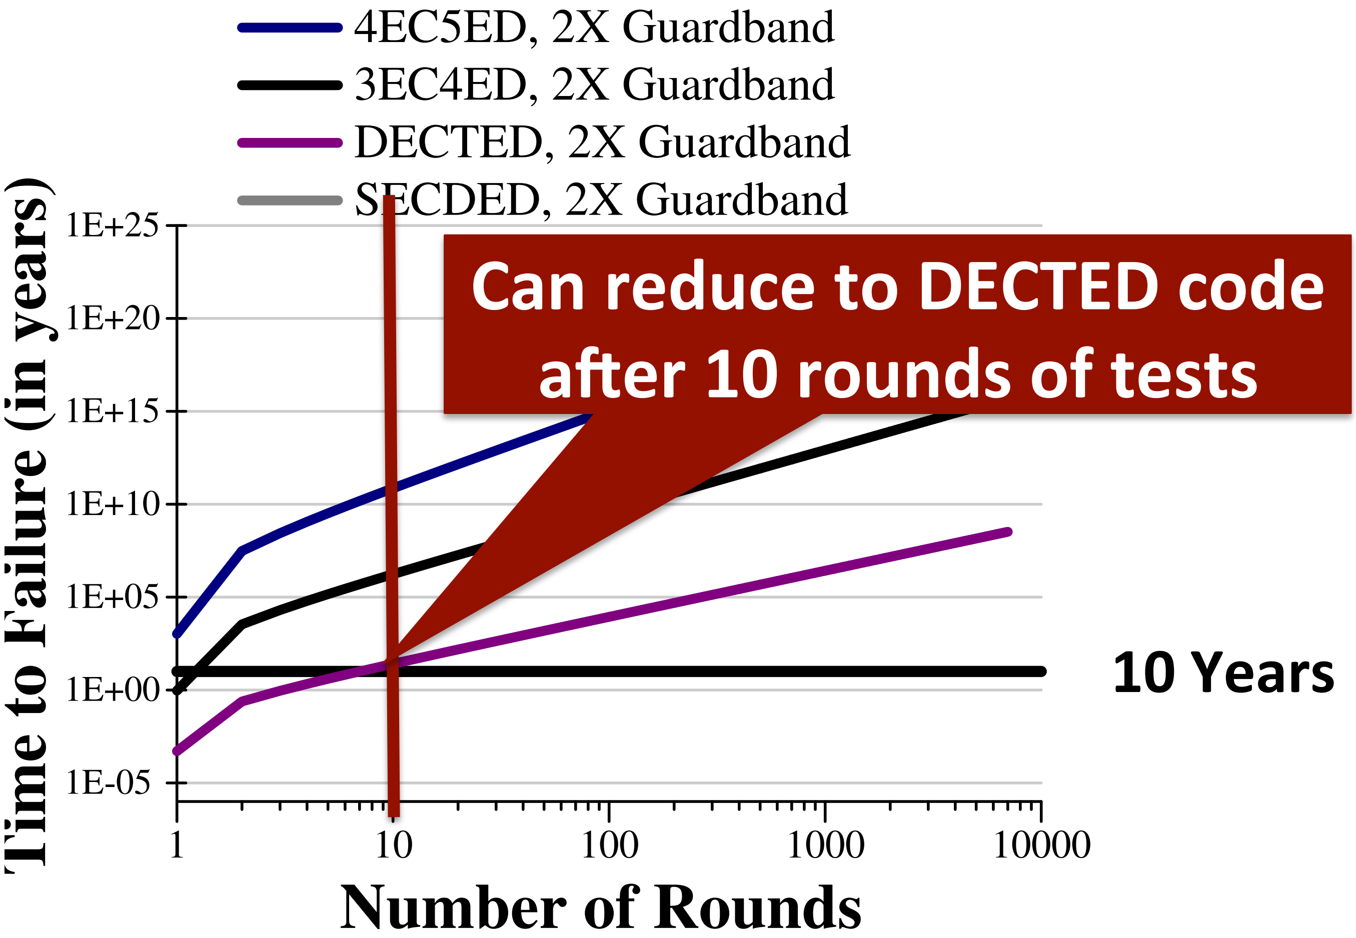

Can reduce to DECTED code after 10 rounds of tests

Efficacy of Hi-ECC

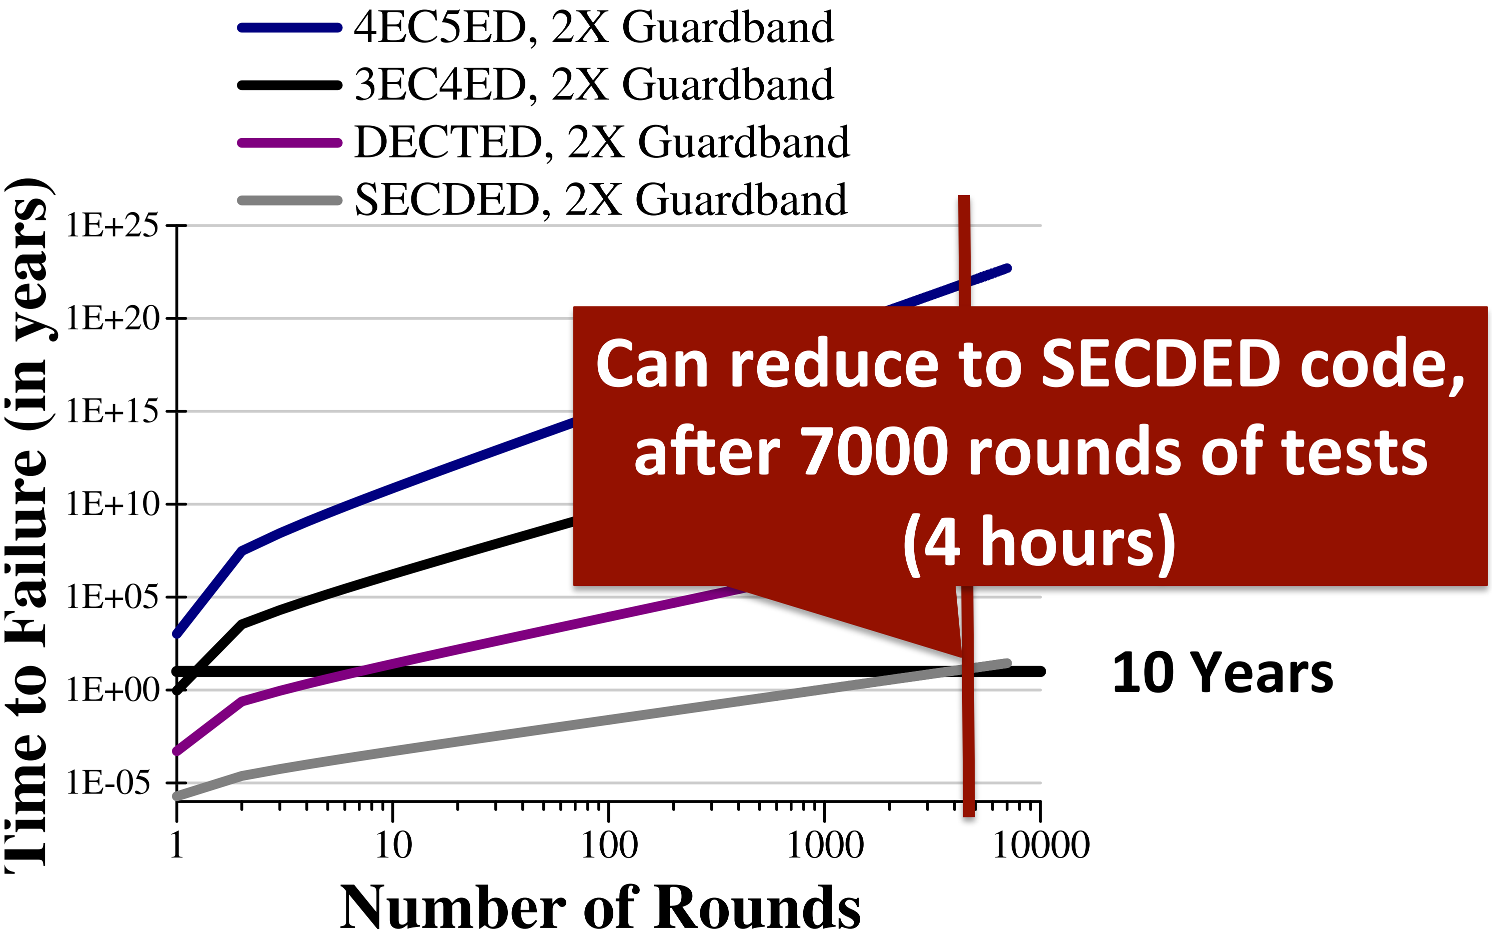

Can reduce to SECDED code, after 7000 rounds of tests (4 hours)

Testing can help to reduce the ECC strength

Outline

- DRAM Scaling Problem
- Online Profiling as a Solution
- Efficacy of System-Level Detection and Mitigation
    - Simple Techniques
    - Recently Proposed Techniques
- **Towards an Online Profiling System**
- Conclusion

Towards an Online Profiling System

Key Observations:

- Testing alone cannot detect all possible failures
- Combination of ECC and other mitigation techniques is much more effective
  - But degrades performance
- Testing can help to reduce the ECC strength
  - Even when starting with a higher strength ECC

Towards an Online Profiling System

Initially Protect DRAM with Strong ECC ①

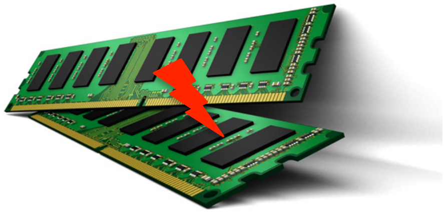

Periodically Test Parts of DRAM ②

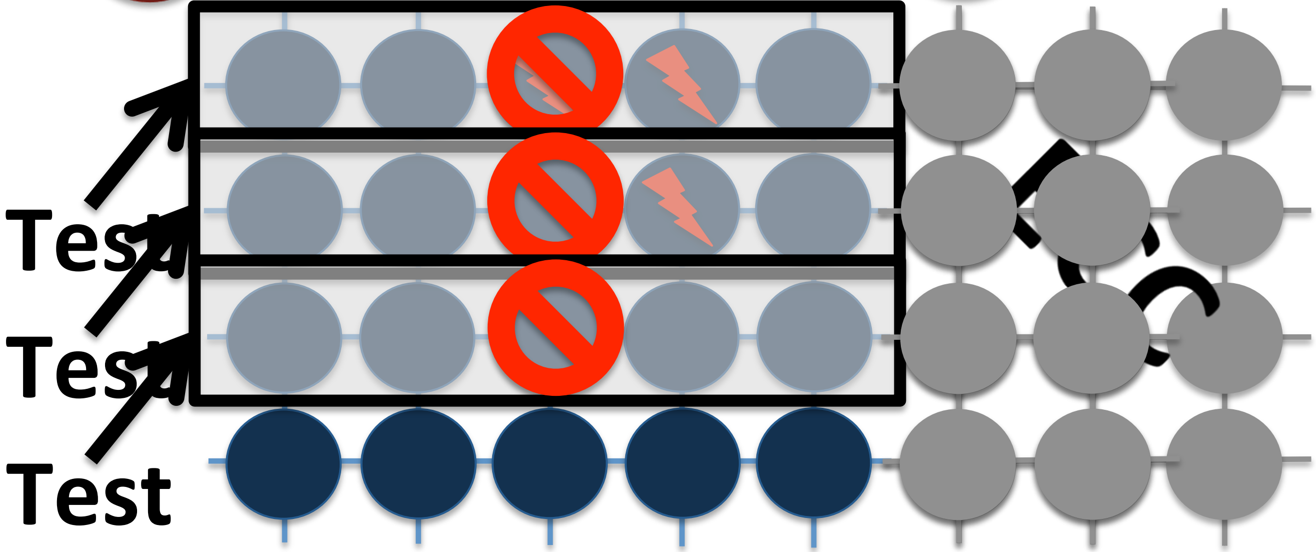

Mitigate errors and reduce ECC ③

Run tests periodically after a short interval at smaller regions of memory

Outline

- DRAM Scaling Problem
- Online Profiling as a Solution
- Efficacy of System-Level Detection and Mitigation
  - Simple Techniques
  - Recently Proposed Techniques
- Towards an Online Profiling System
- **Conclusion**

Conclusion

- ***We analyze*** **the efficacy of testing, guardbanding, ECC, and recent techniques at system-level**
  - **Using experimental data from *real DRAMs***
- ***Key Conclusions***
  - **Testing alone cannot guarantee reliable operation**
  - **A combination of techniques is more effective**
  - **Testing+ECC-based techniques block memory for significant time → Performance degradation**
- ***We propose*** **Online profiling that runs at background without disrupting current programs**
  - **Run periodically at smaller regions of memory**

Thank you

Full data set for 96 chips is available at

http://www.ece.cmu.edu/~safari/tools/dram-sigmetrics2014-fulldata.html

The Efficacy of Error Mitigation Techniques for DRAM Retention Failures

Samira Khan, Donghyuk Lee, Yoongu Kim,
Alaa R. Alameldeen, Chris Wilkerson, and
Onur Mutlu

① Bit Repair Techniques

Test DRAM module at boot up ①

Mitigate failures by repairing the bits ②

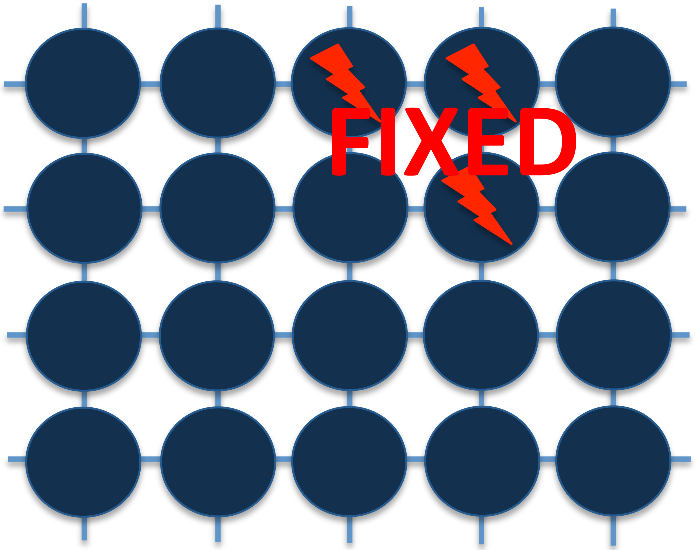

These techniques are vulnerable to new intermittent failures

ArchShield ISCA'13, RAIDR ISCA'12, SECRET ICCD'12, RAPID HPCA'06

Efficacy of Bit Repair Techniques

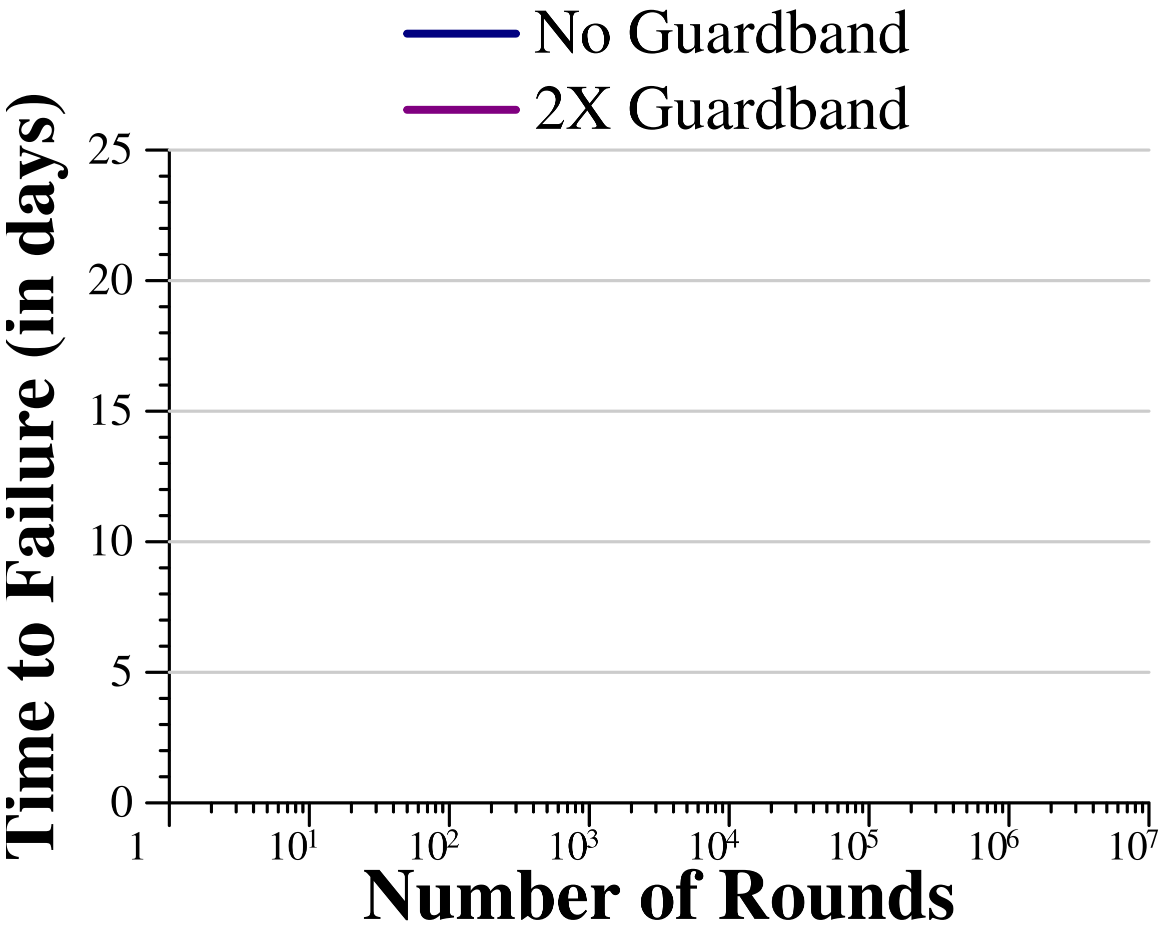

Efficacy of Bit Repair Techniques

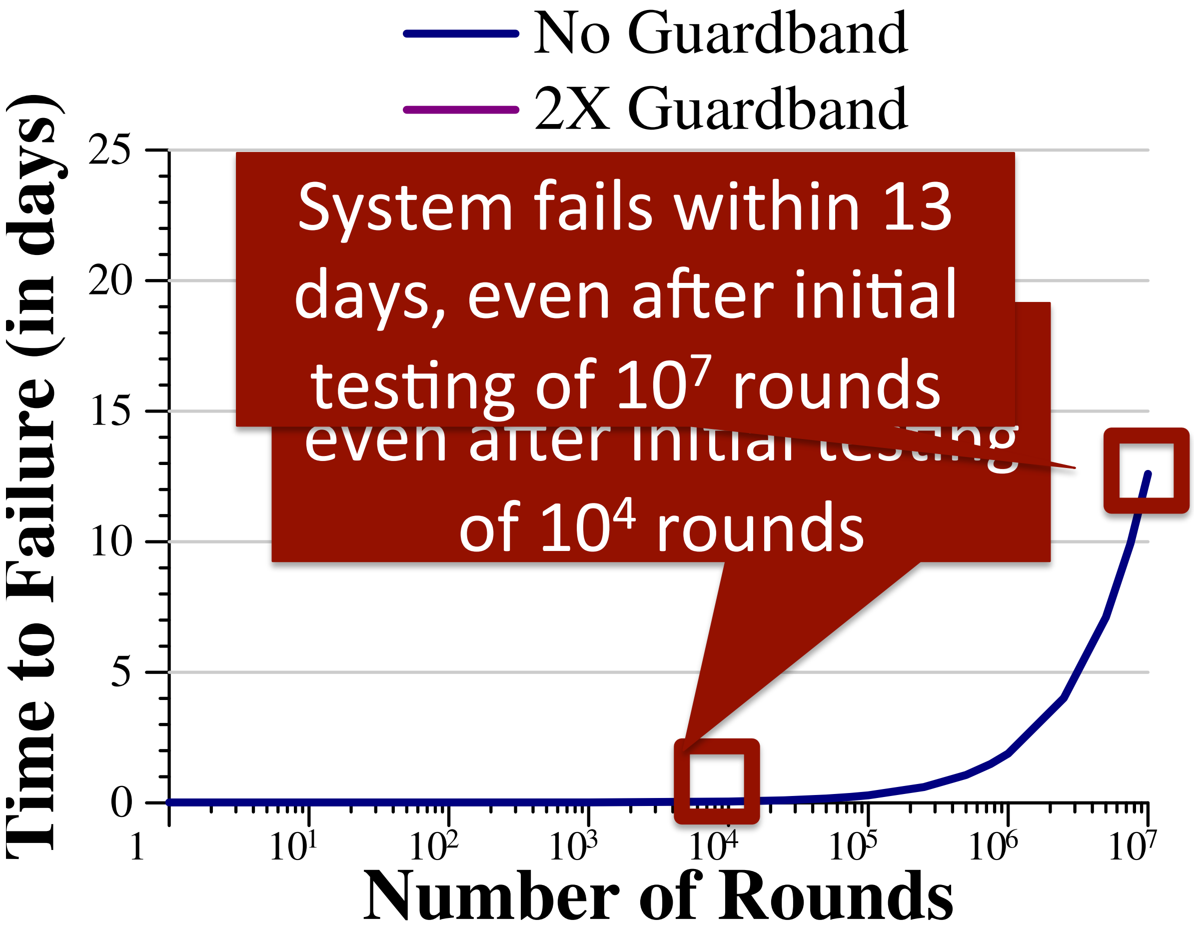

System fails within 13 days, even after initial testing of 10^7 rounds

even after initial testing of 10^4 rounds

Efficacy of Bit Repair Techniques

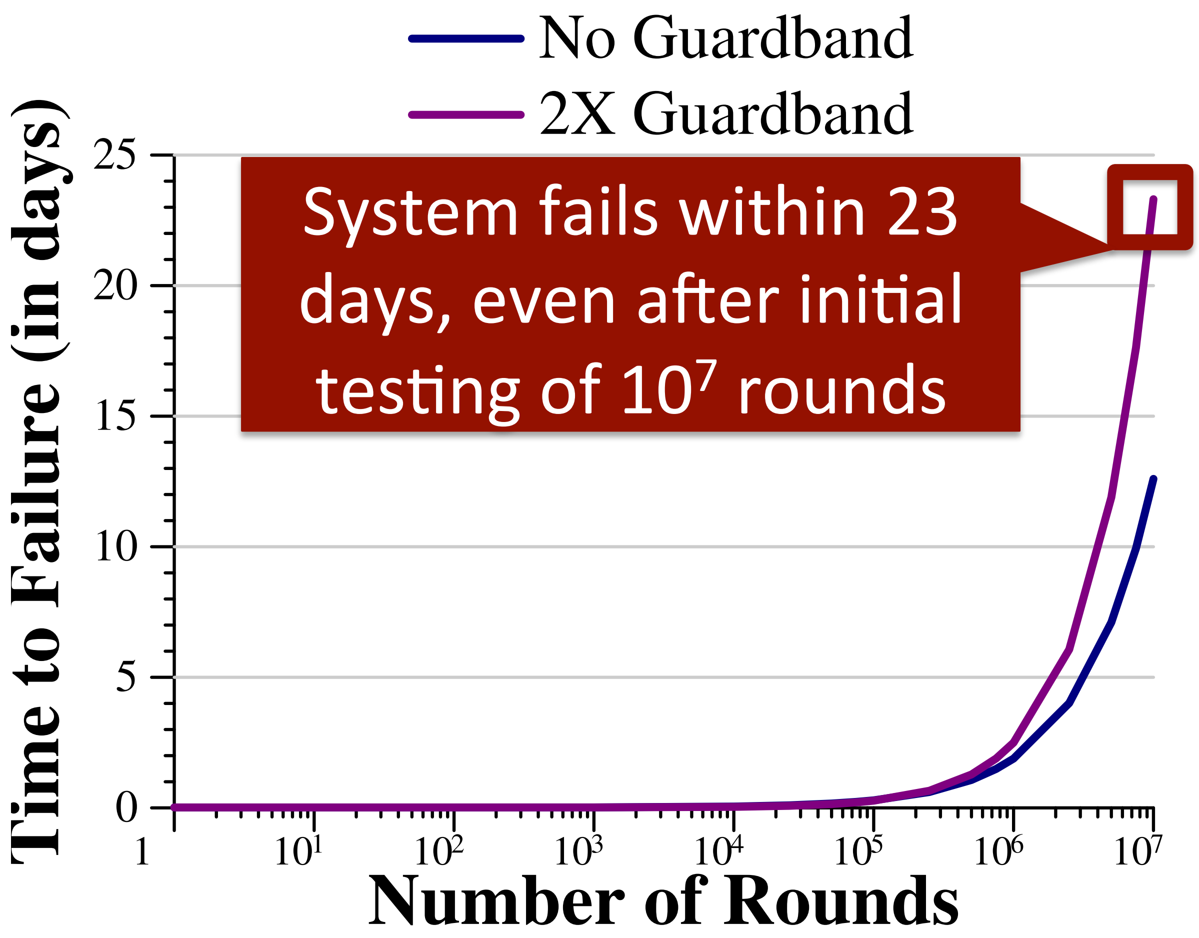

System fails within 23 days, even after initial testing of 10^7 rounds

Even longer tests are not sufficient to guarantee reliable operation

❷ Variable-Strength ECC (VS-ECC)

Test DRAM module at boot up ①

Protect failed lines with strong ECC ②

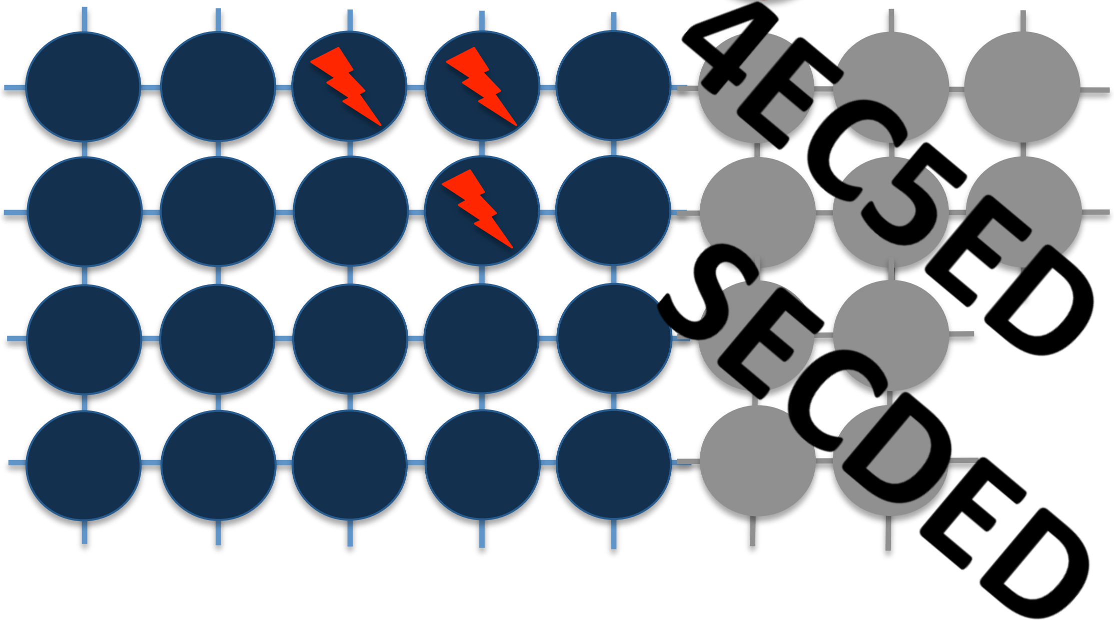

Will fail as soon as there are two bit errors in SECDED lines

Efficacy of VS-ECC

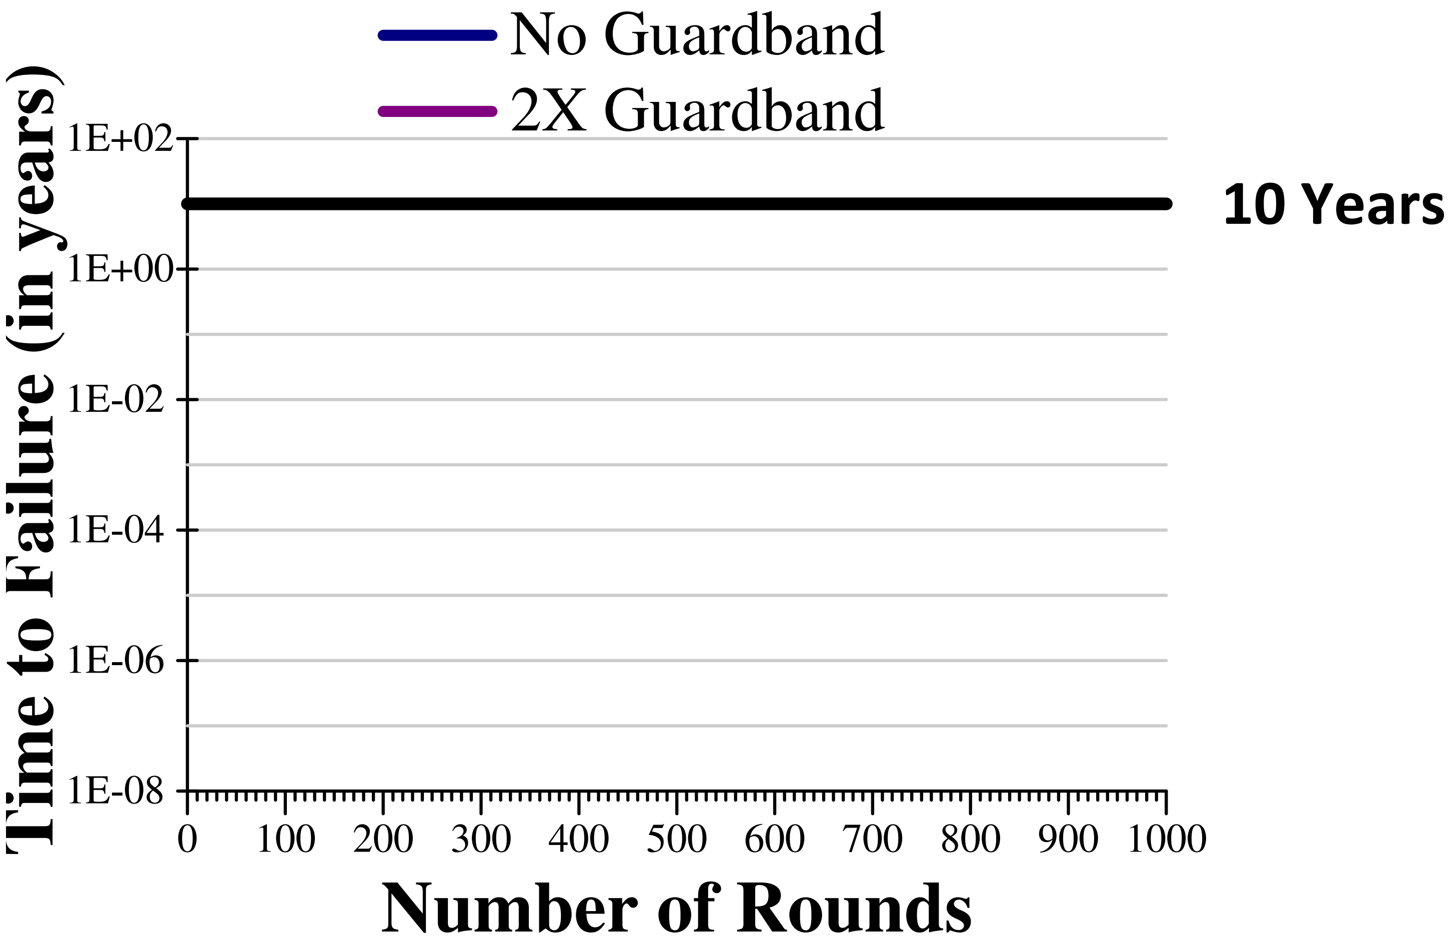

Efficacy of VS-ECC

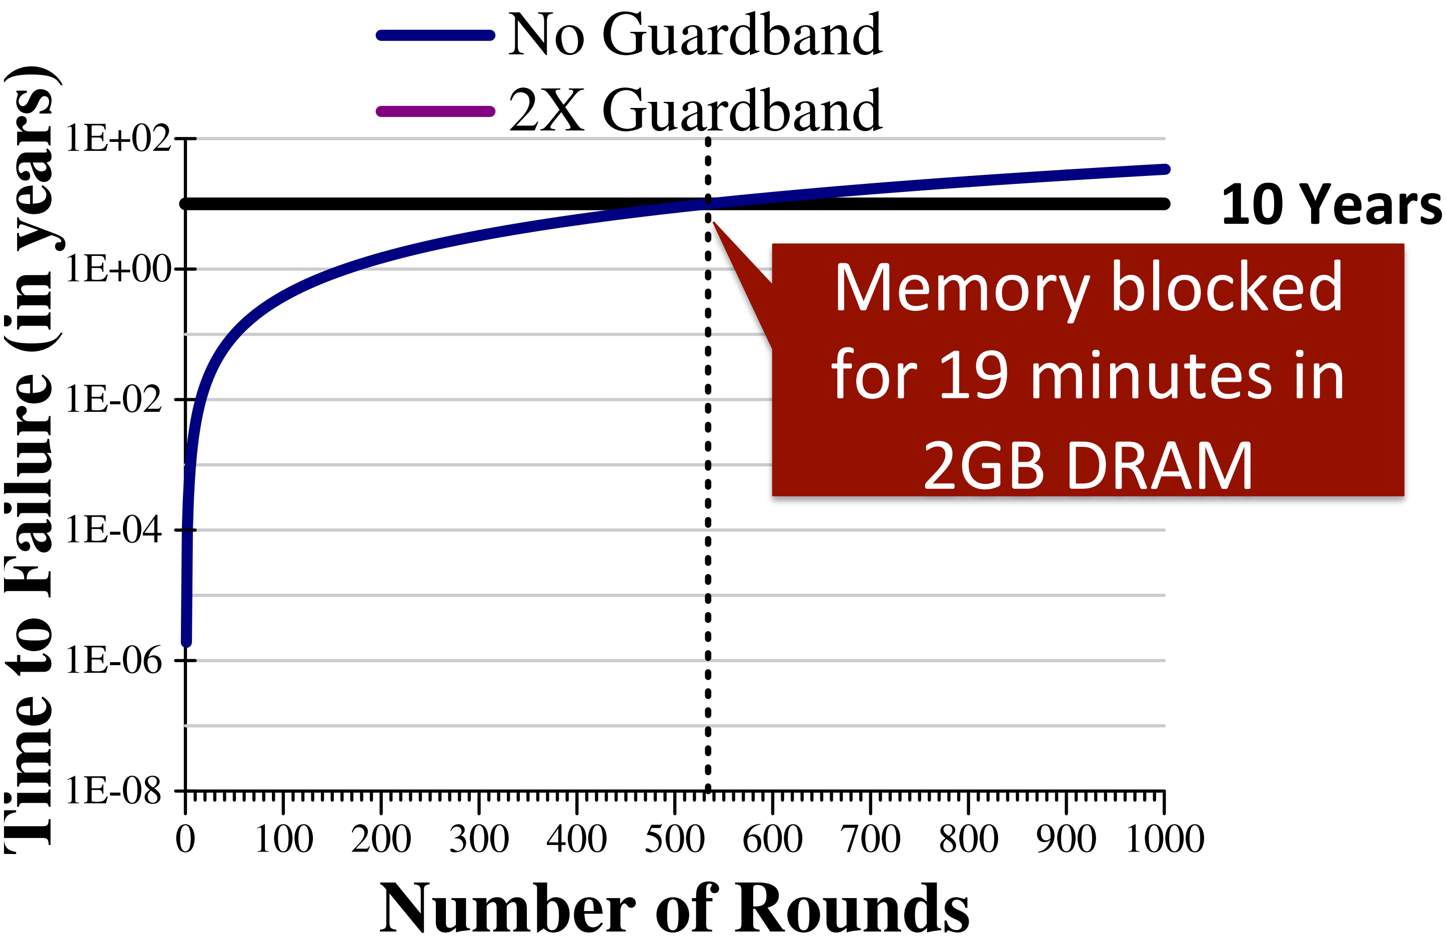

Memory blocked for 19 minutes in 2GB DRAM

Efficacy of VS-ECC

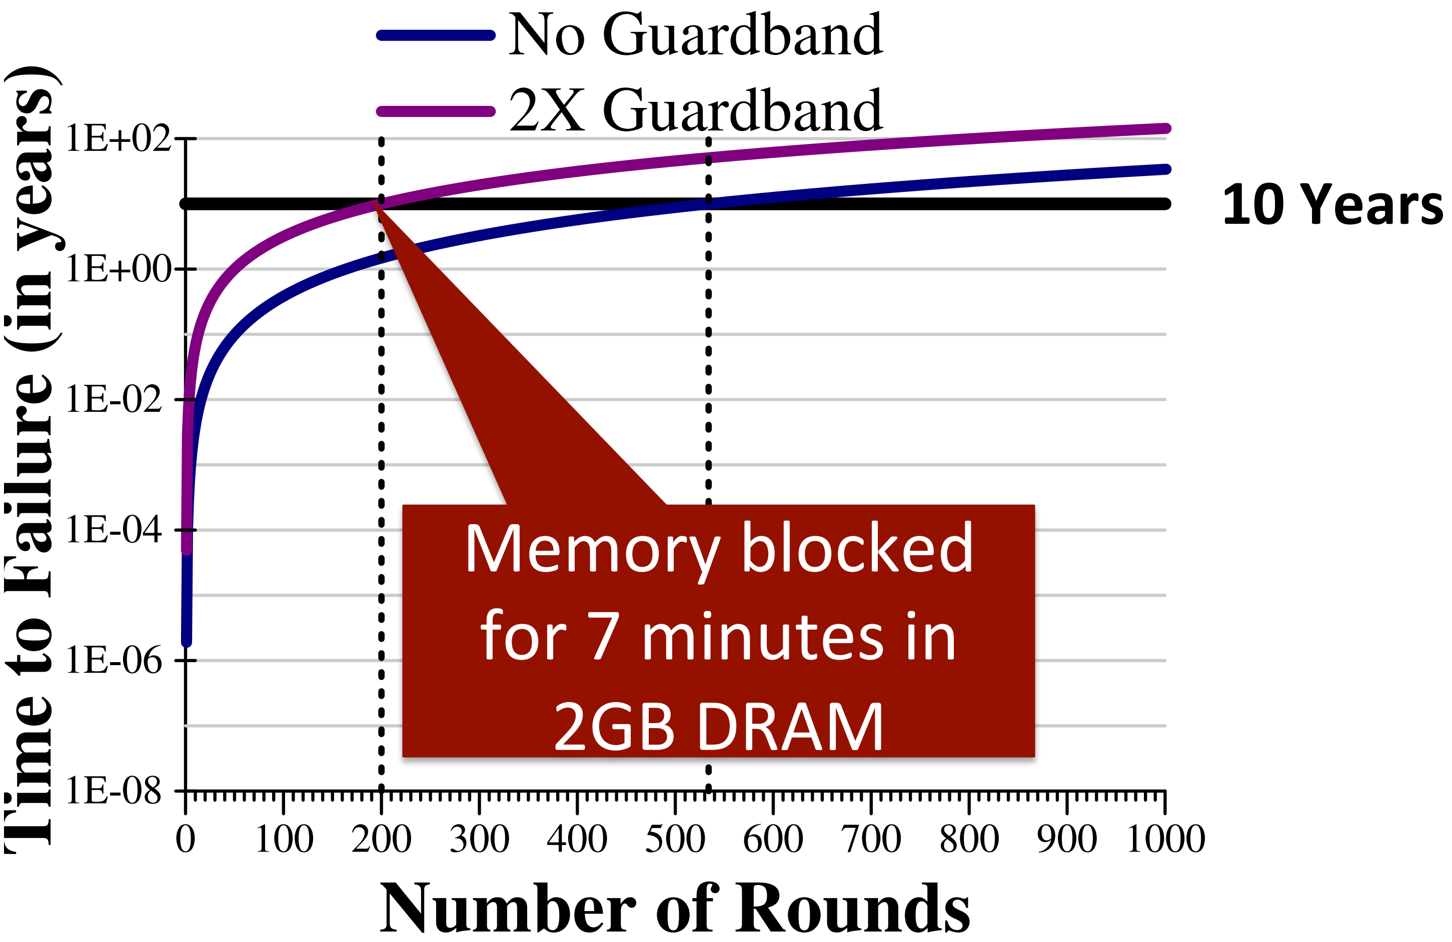

Memory blocked for 7 minutes in 2GB DRAM

With higher capacity DRAM, memory will be blocked for an unacceptable amount of time

Challenges and Opportunities

Challenges:

- Performance Overhead
- Mitigation Overhead

Opportunities:

- Enable Failure-aware Optimizations

Testing

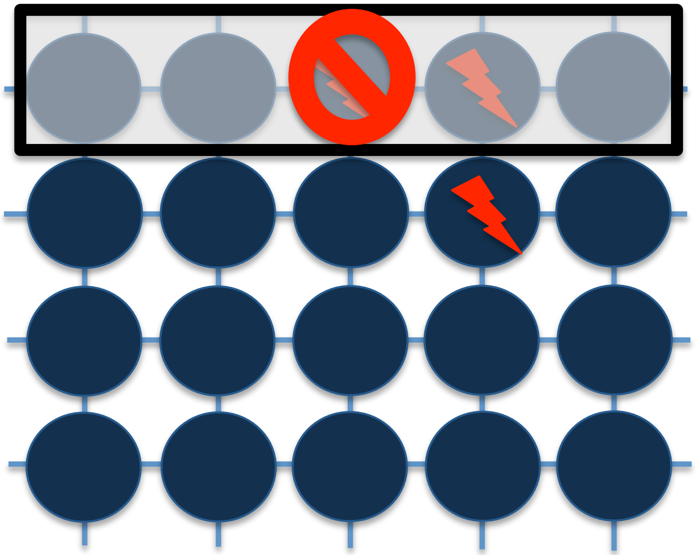

Reduction in Error Rate in all Modules

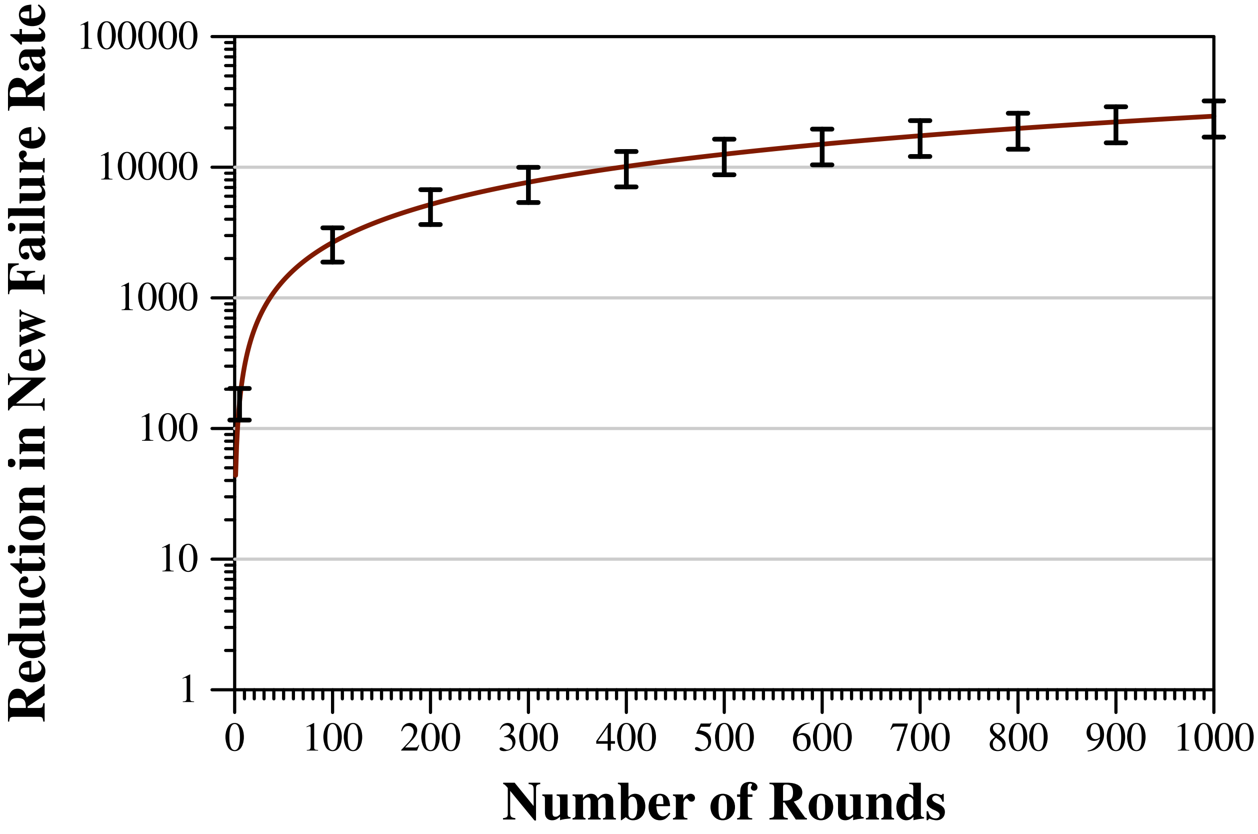

Difference in Modules

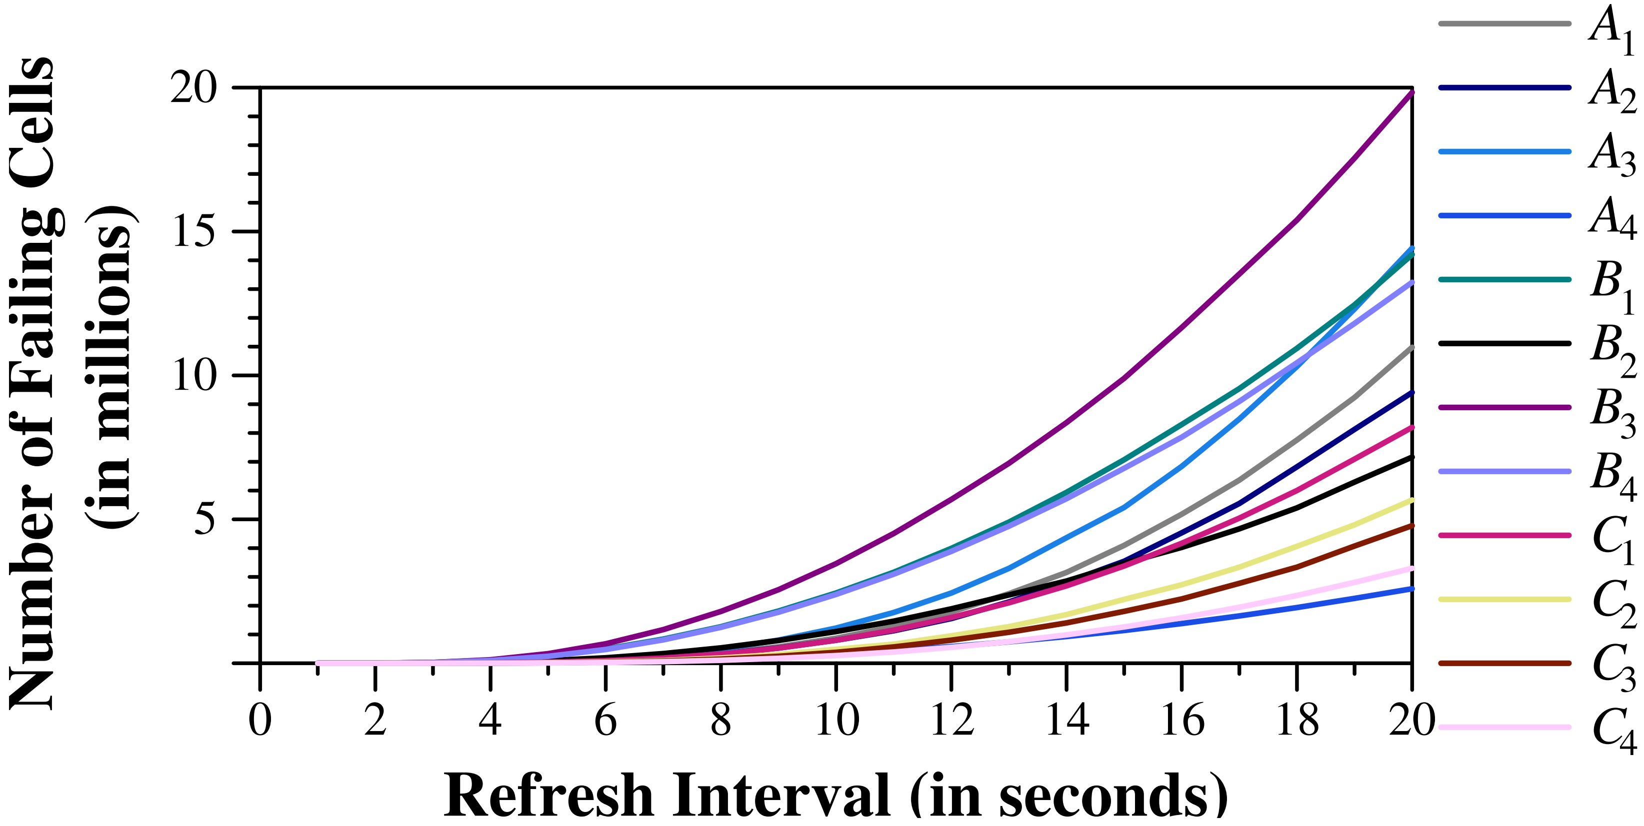

Tested DRAM Modules

Manufacturer	Module Name	Assembly Date (Year-Week)	Number of Chips
A	A1	2013-18	8
	A2	2012-26	8
	A3	2013-18	8
	A4	2014-08	8
B	B1	2012-37	8
	B2	2012-37	8
	B3	2012-41	8
	B4	2012-20	8
C	C1	2012-29	8
	C2	2012-29	8
	C3	2013-22	8
	C4	2012-29	8

Time to Test

Operation	Time (2GB)	Time (64GB)
Write/Read a Row	667.5 ns	667.5 ns
Write/Read 2GB Module	174.98 ms	5.59 s
1 round , 1 pattern	413:96 ms	11.24 s
1 round, 5 patterns	2.06 s	56.22 s
1000 rounds, 5 patterns	34 m	15.6 hours

Temperature Controlled Environment

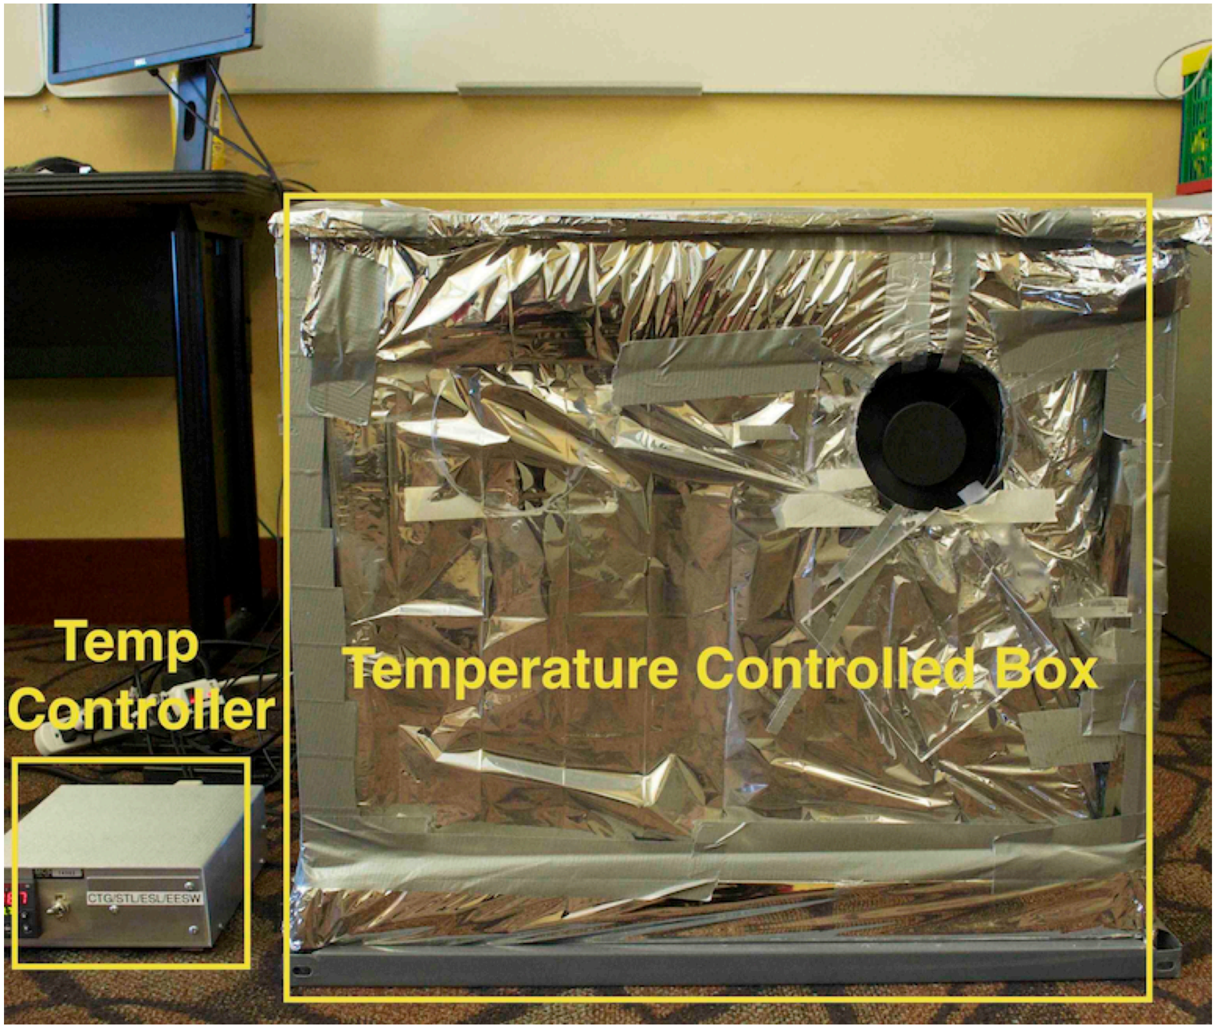

Dependence of Retention Time on Temperature

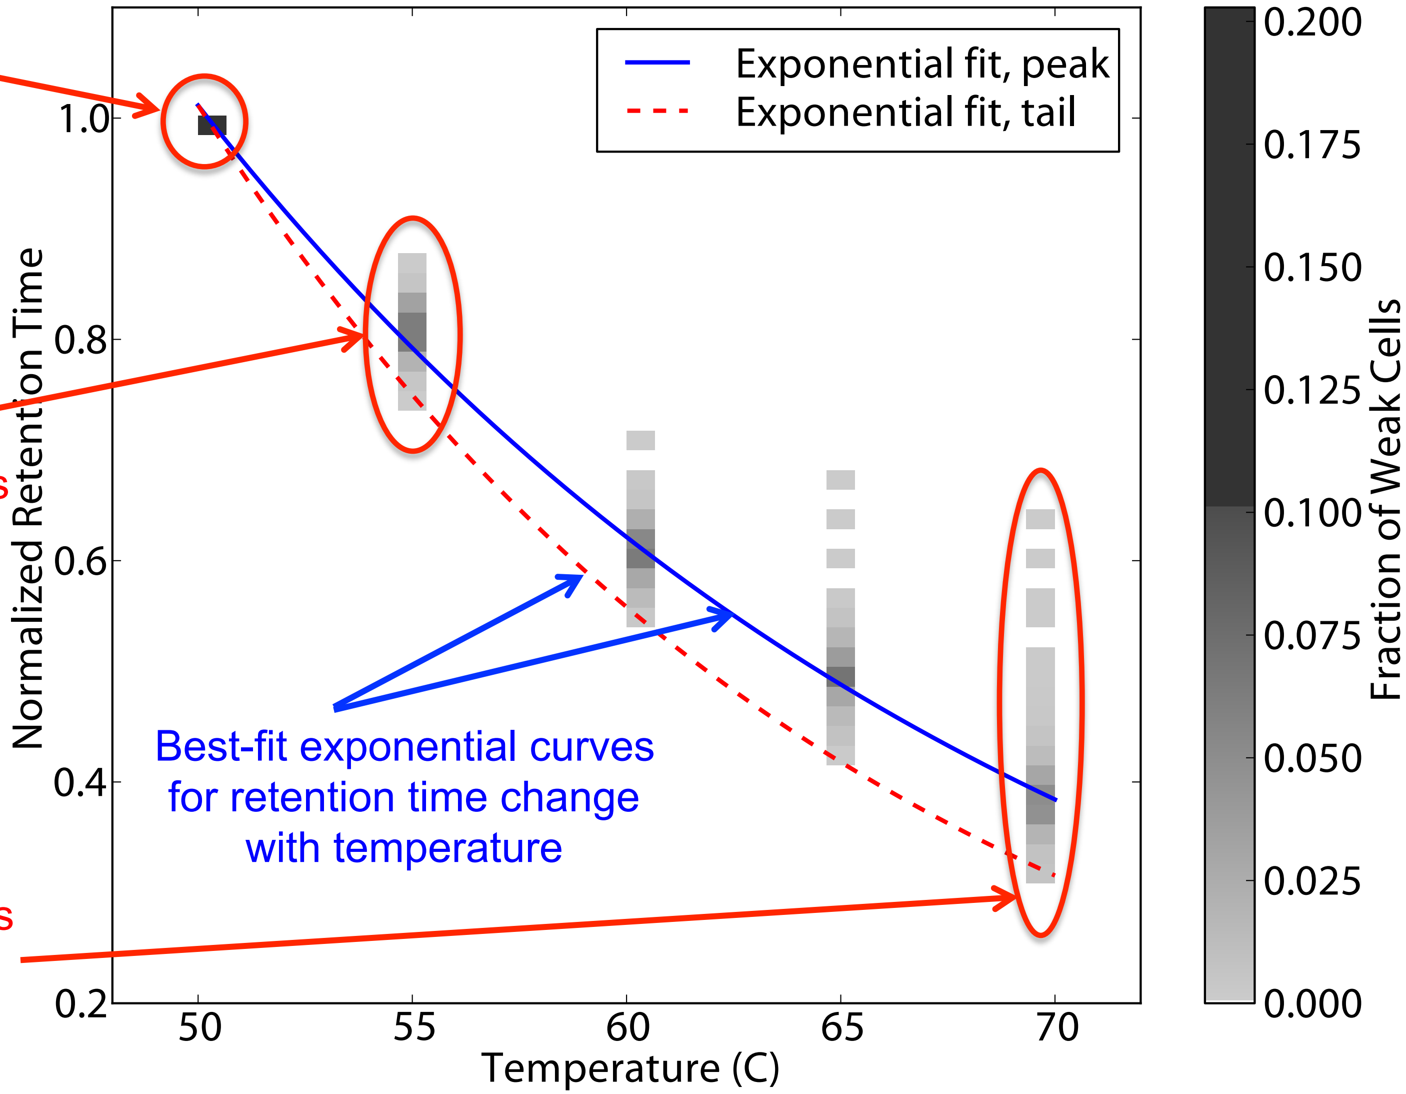

Fraction of cells that exhibited retention time failure at any *tWAIT* for any data pattern at 50°C

Normalized retention times of the same cells at 55°C

Normalized retention times of the same cells At 70°C

Best-fit exponential curves for retention time change with temperature

Slide Courtesy Onur Mutlu ISCA'13

Dependence of Retention Time on Temperature

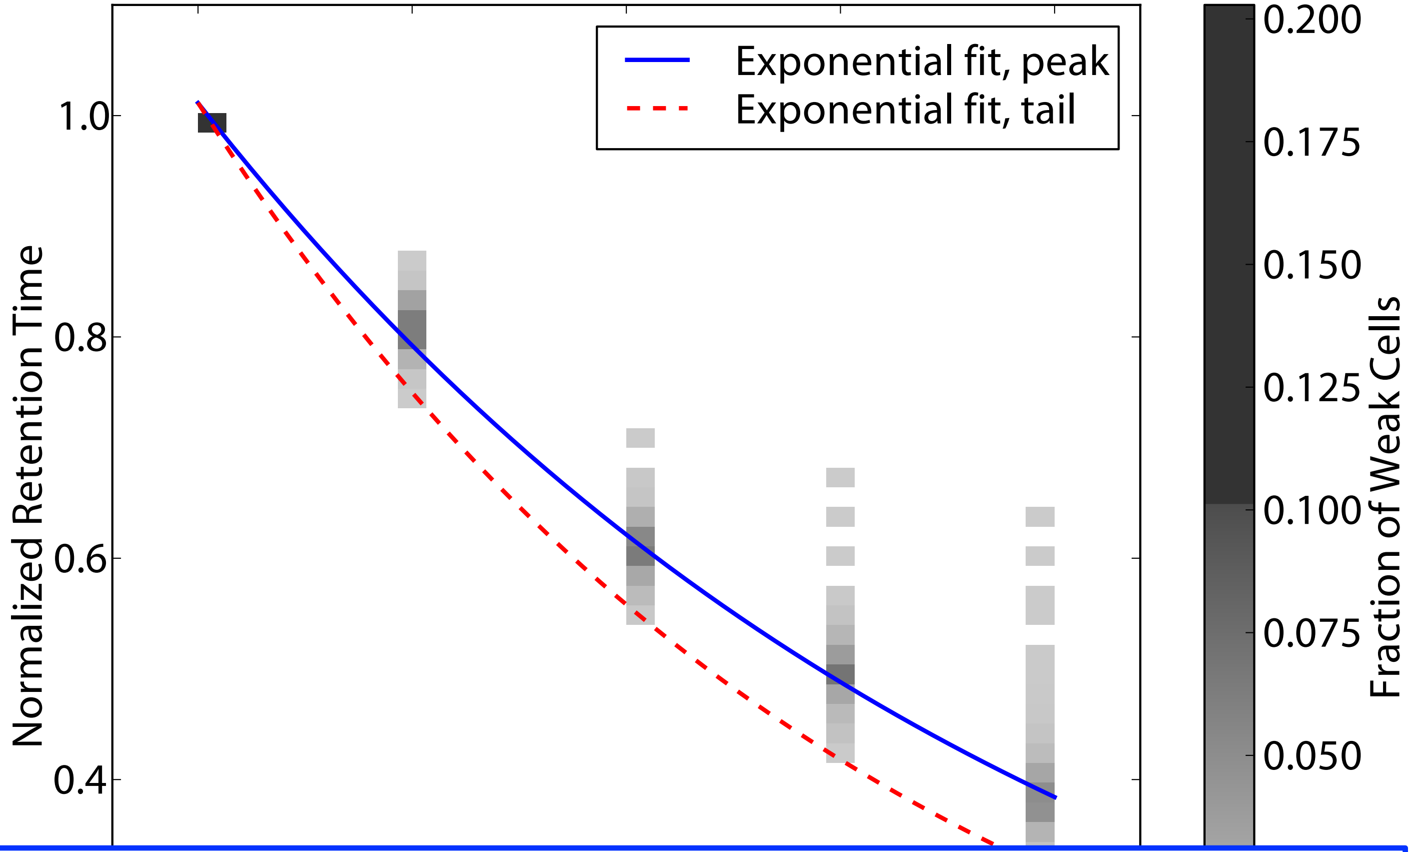

Relationship between retention time and temperature is consistently bounded (predictable) within a device

Every 10^{O}C temperature increase
→ 46.5% reduction in retention time in the worst case

Effect of Temperature

- Worst fit curve for retention time at different temperature corresponds to e^-0.0625T , where T is the temperature [ISCA'13]
- A 10 C increase in temperature results in a reduction of 1 – e^-0.0625*10 = 46.5%
- 1 second → 82 ms at 45 C
- 20 seconds → 1640 ms at 85 C

Characteristics not Dependent on Refresh Interval

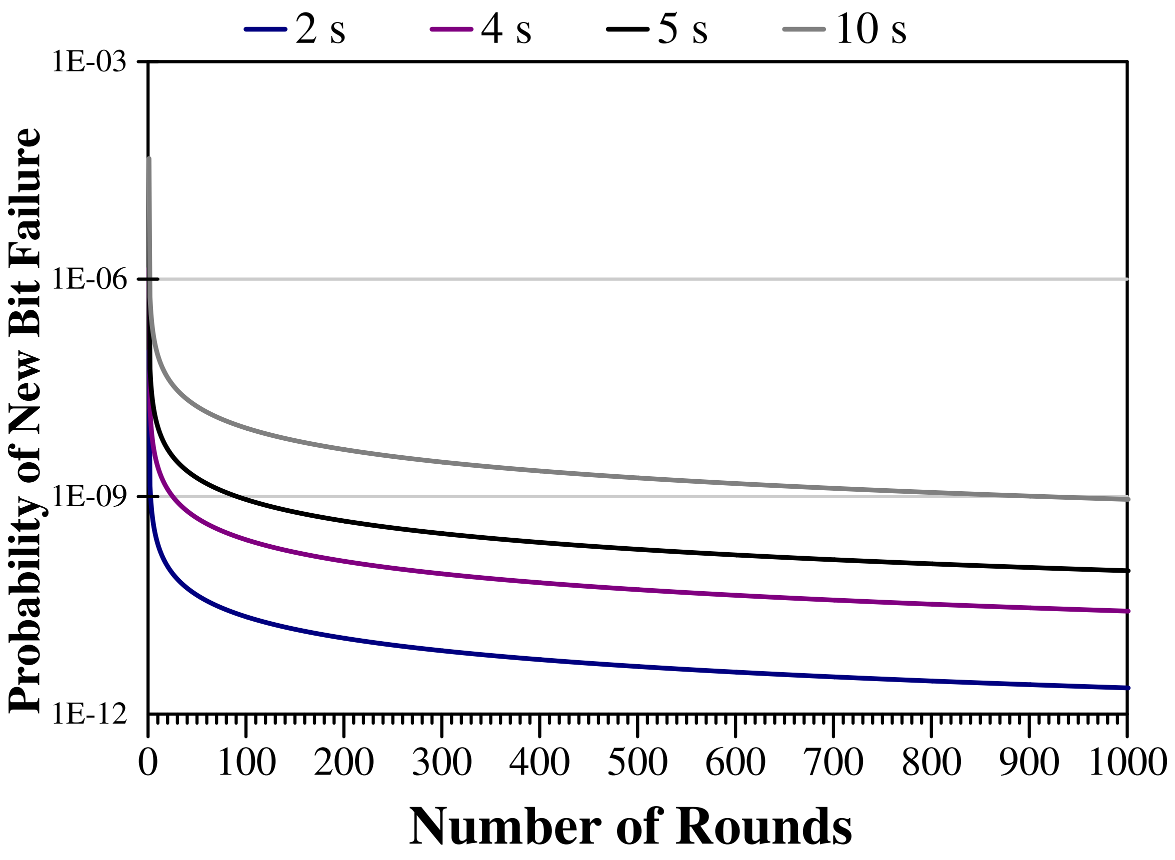

Expected Number of Multi-Bit Failures