FaultSim: A Fast, Configurable Memory-Reliability Simulator for Conventional and 3D-Stacked Systems

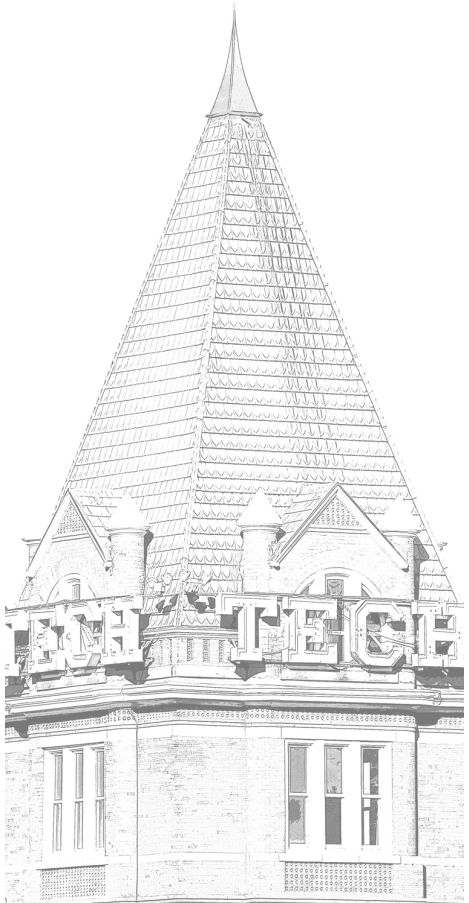

19th Jan 2016

Prague, Czech Republic

Prashant Nair - Georgia Tech

David Roberts - AMD Research

Moinuddin Qureshi- Georgia Tech

Georgia Institute of Technology

AMD

In both computers and humans, memory is perhaps the most important attribute, and the one most likely to fail

Tse-Yu Yeh (Apple)

IMPORTANCE OF MEMORY RELIBAILITY

Memory system: one of the main cause of system failure

Memory reliability seen as major challenge for Exascale

High availability servers employ Chipkill or RAID

Increasing set of challenges for future memory systems:

- Weak bits from technology scaling
- Large granularity failures as common as bit failures
- New failure modes in DRAM (e.g. TSV faults in 3D)
- New failure modes from NVRAM (disturb, endurance)

Memory reliability continues to be an important concern

HOW TO EVALUATE MEMORY RELIABILITY?

Fast and accurate simulators vital to compare effectiveness of different solutions

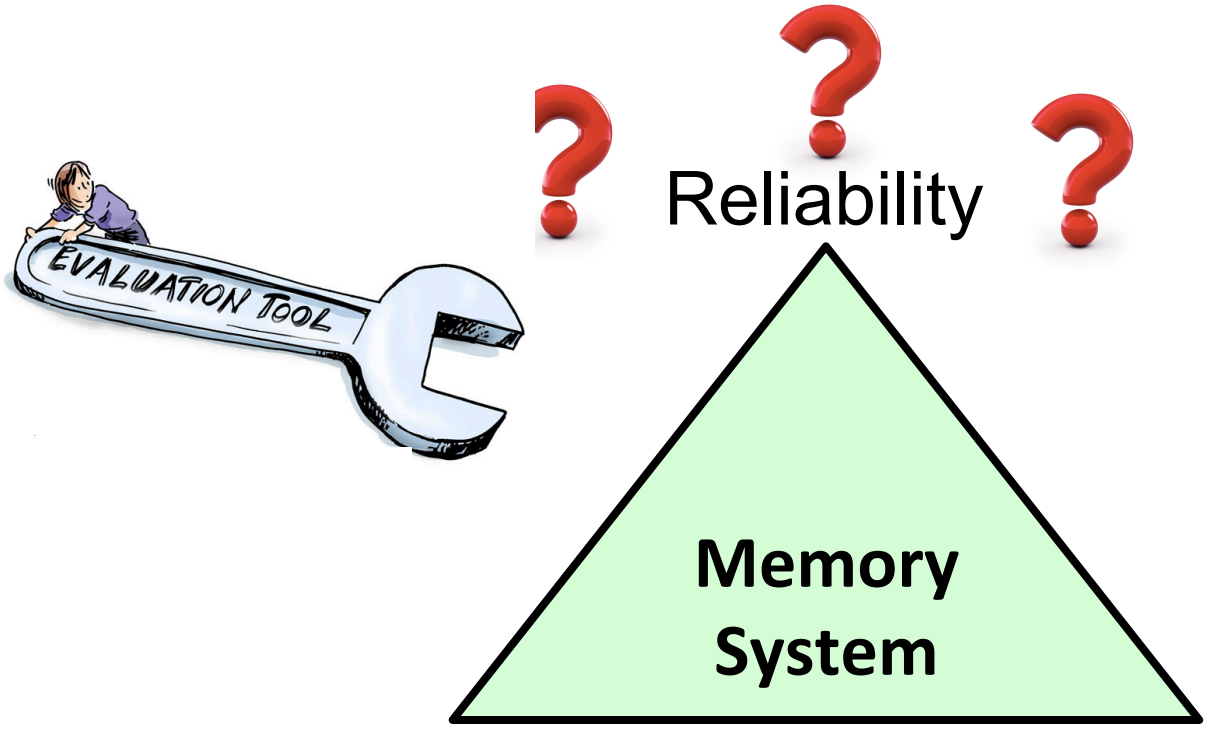

Goal: Accurately evaluate memory reliability across different systems & solutions, in less than one minute

TYPES OF MEMORY FAILURES

DRAM devices can encounter faults during operation

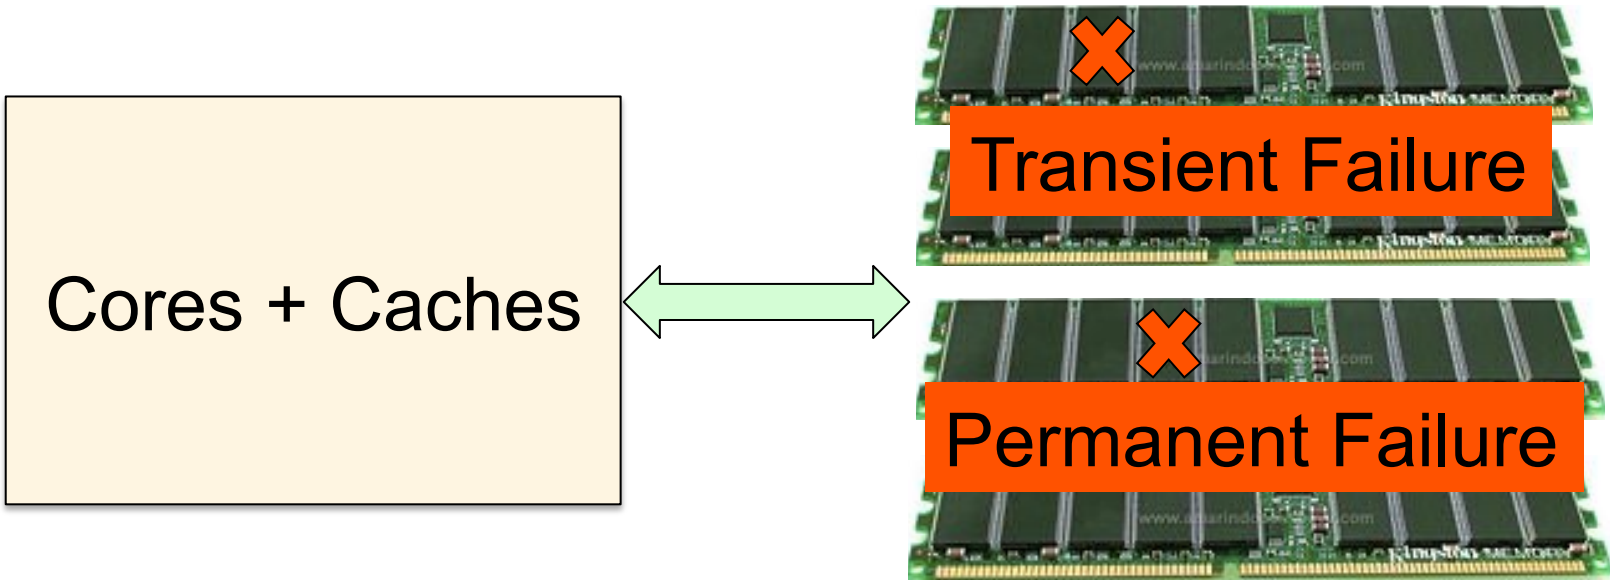

Memory reliability evaluations must account for both transient failures as well as permanent failures

GRANULARITY OF MEMORY FAILURES

Failures occur at small and large granularities:

- Bit, Word, Column, Row, Bank, Multi-Bank

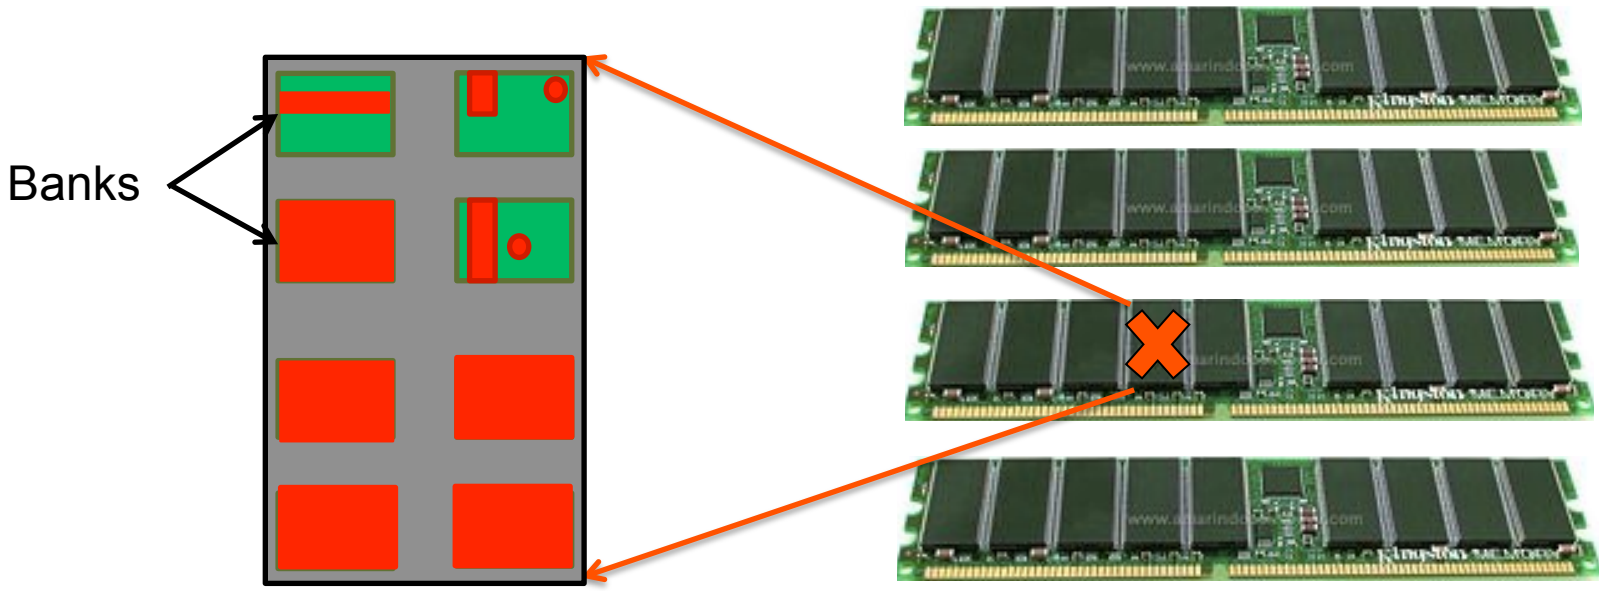

Single DRAM Die (Top View)

Memory reliability simulator should capture interaction of failures at different granularities

REAL WORLD FAILURE RATE [SRIDHARAN+ SC'13]

Failure Mode	Transient Fault Rate (FIT)	Permanent Fault Rate (FIT)
Bit	14.2	18.6
Word	1.4	0.3
Column	1.4	5.6
Row	0.2	8.2
Bank	0.8	10
Total	18	42.7

✔ *SECDED*

= 24.1

✔ *CHIPKILL*

1. Permanent faults >2x as likely as transient faults
2. Large granularity faults as common as bit faults

COMPLEXITY IN FAULT INTERACTIONS WITH ECC

Several techniques: SECDED, Chipkill, Sparing often used in combination with periodic Scrubbing

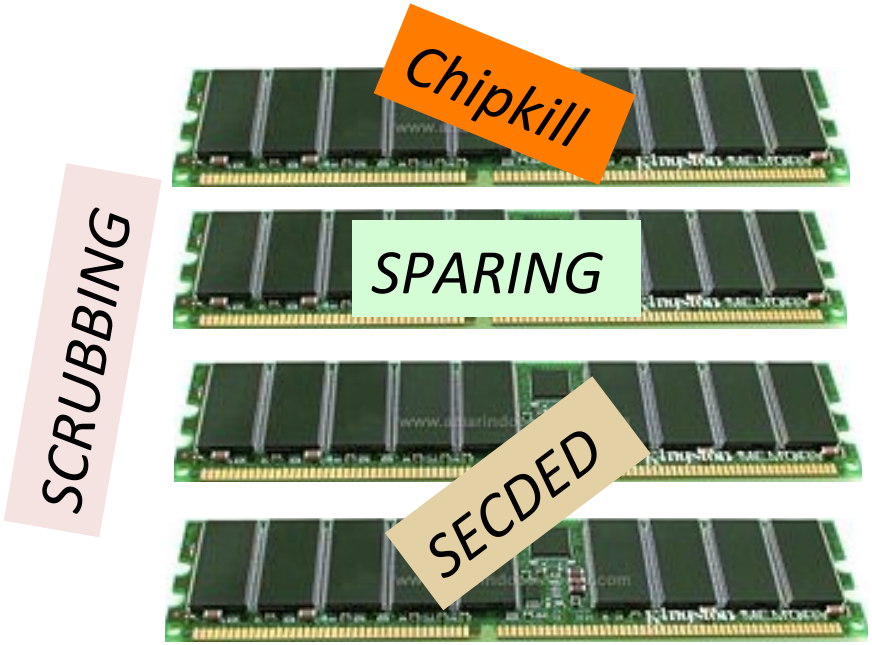

Complex interactions of techniques with fault modes and granularities ➔ How to evaluate effectiveness?

ANALYTICAL MODELS FOR MEMORY RELIABILITY

- Complex, Cumbersome, Changes with Fault Models
- A PRDC paper* has nearly 3 page model for Chipkill

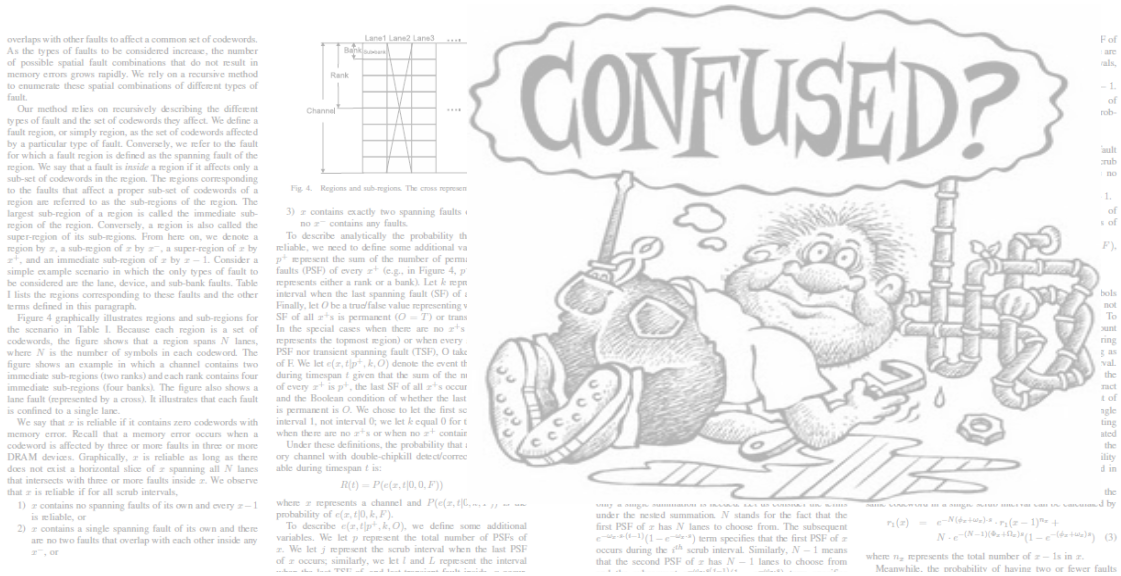

Small change in ECC → Massive changes in the model

Use empirical evaluation instead of analytical models

*Jian et. al. PRDC 2013

OVERVIEW

- WHY FAULTSIM?
- FAULTSIM: WHAT AND HOW?
- FAULTSIM: LESS THAN 1 MINUTE
- FAULTSIM: APPLIED TO 3D MEMORY
- SUMMARY

FAULTSIM: A MONTE-CARLO FAULT SIMULATOR

FaultSim is written in C++. Configuration at command line or file

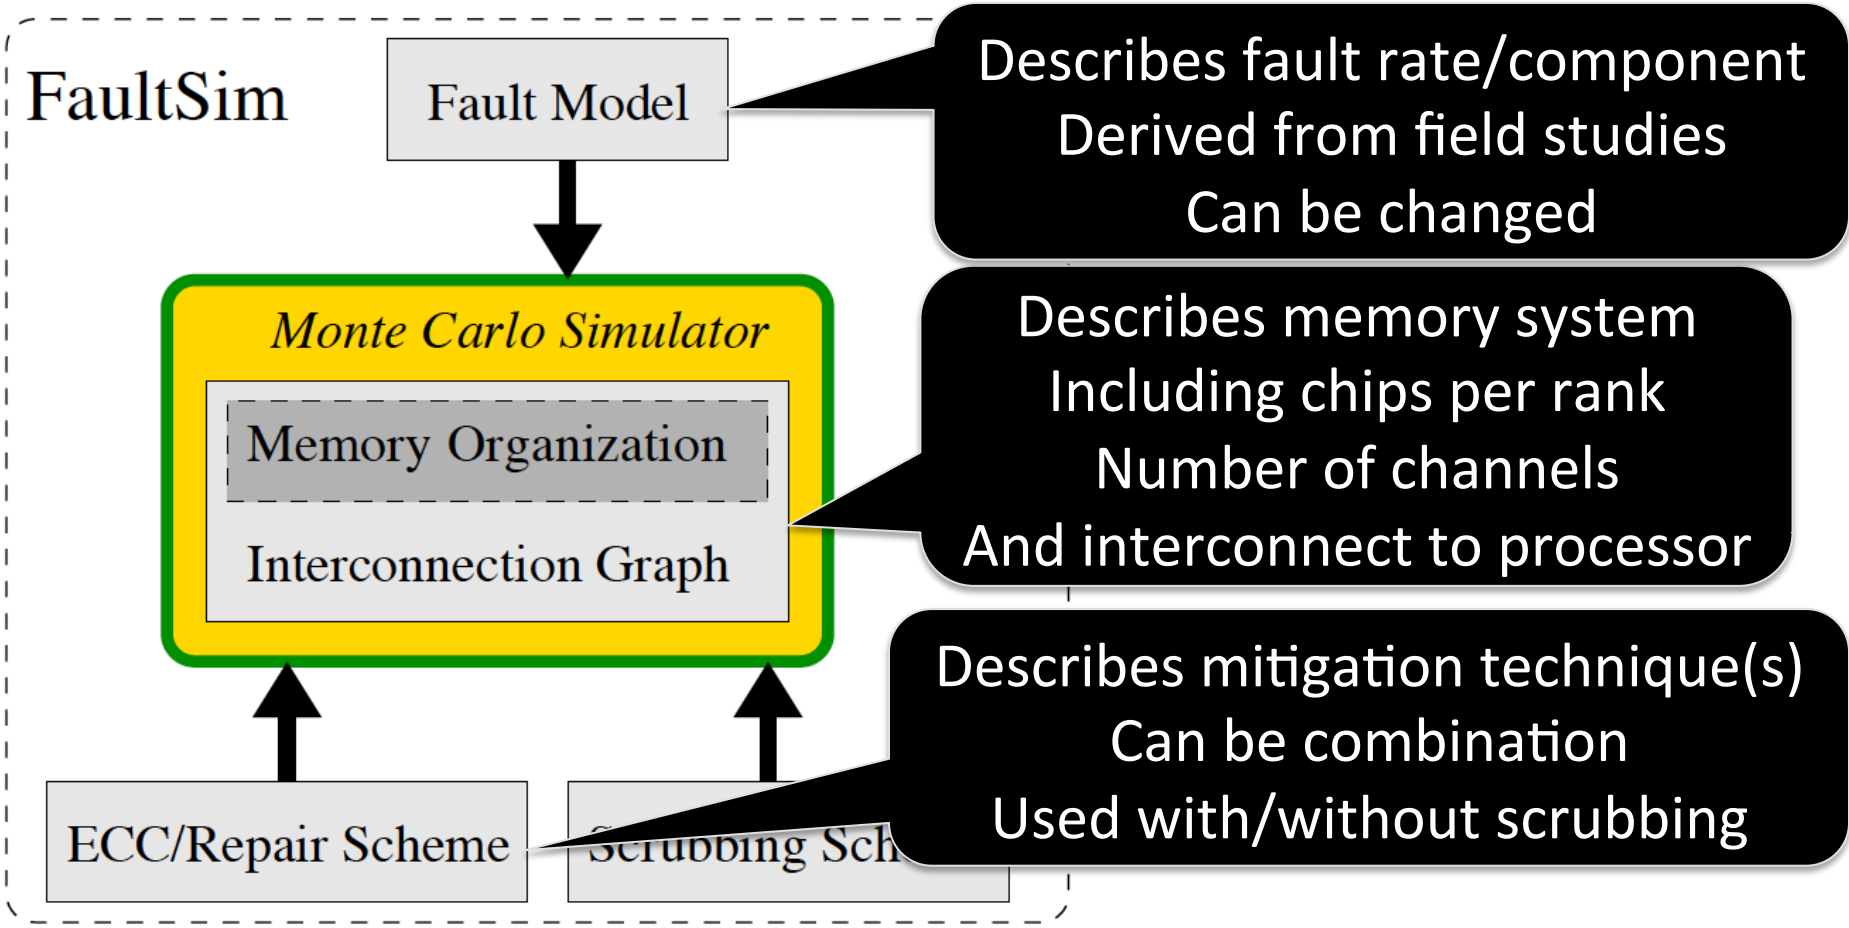

FAULTSIM OPERATION

Divide system lifetime (7 years) into smaller interval (3 hours)
Check for uncorrectable failures in every interval

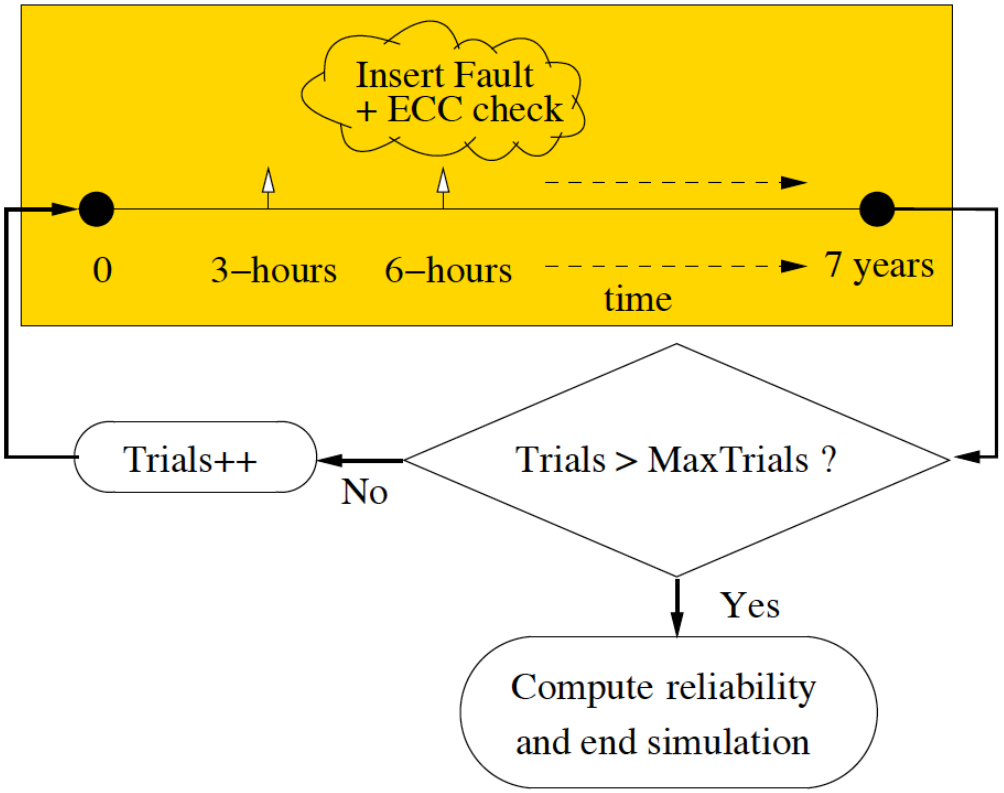

20,000 Intervals

~ 1 million trials

FaultSim performs 20K*1million interval simulations per chip for each fault type ➔ days of simulation time

FAULTSIM: DATA STRUCTURES

- Memory chips are organized as Fault Domains
- Fault Domain (FD) consists of Fault Ranges (FR)
- Each FR uses Address (ADDR) and Mask fields

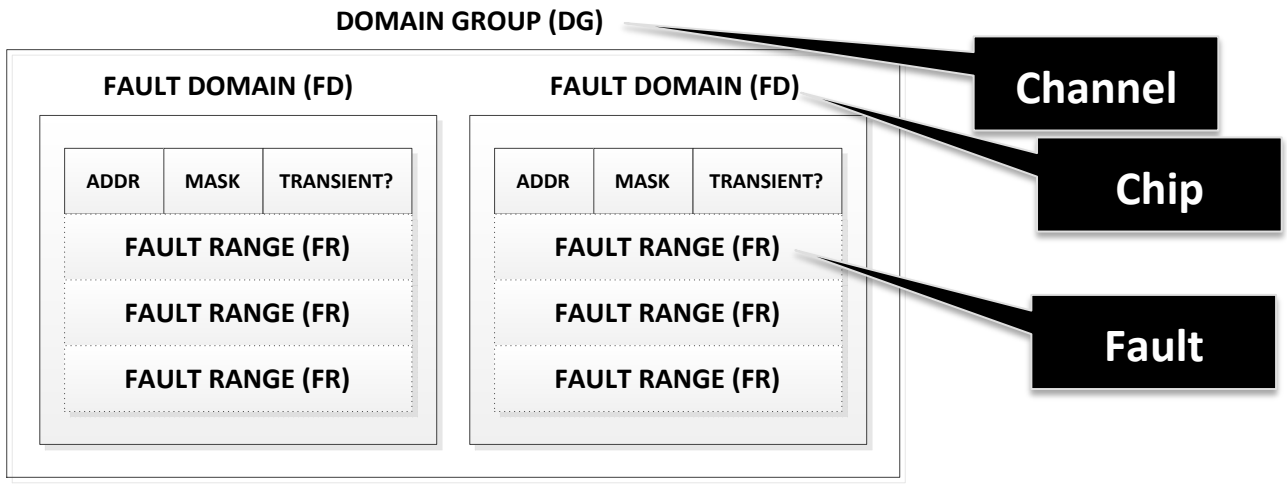

Space Efficient Representation: Large + Small granularity faults use only one type of FR data structure

FAULT REPRESENTATION: EXAMPLE

Memory with 8 rows and 8 bits per row

- Fault ranges A, B and C (A and B intersect)
- **Mask field**: fault address bit ***i*** can be 0 or 1
- **Address field**: specific address bit values where $\mathbf{Mask_i}$ == 0
- Faults intersection computed based on mask and address

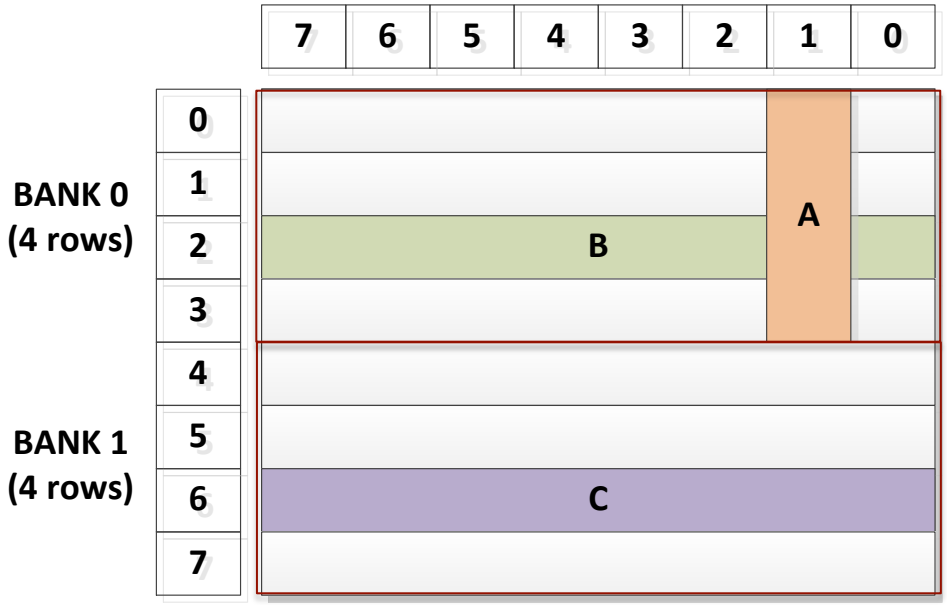

FR	Address	Mask
A	000**001**	0**11**000
B	**010**000	000**111**
C	**110**000	000**111**

VALIDATION: WITH ANALYTICAL MODELS

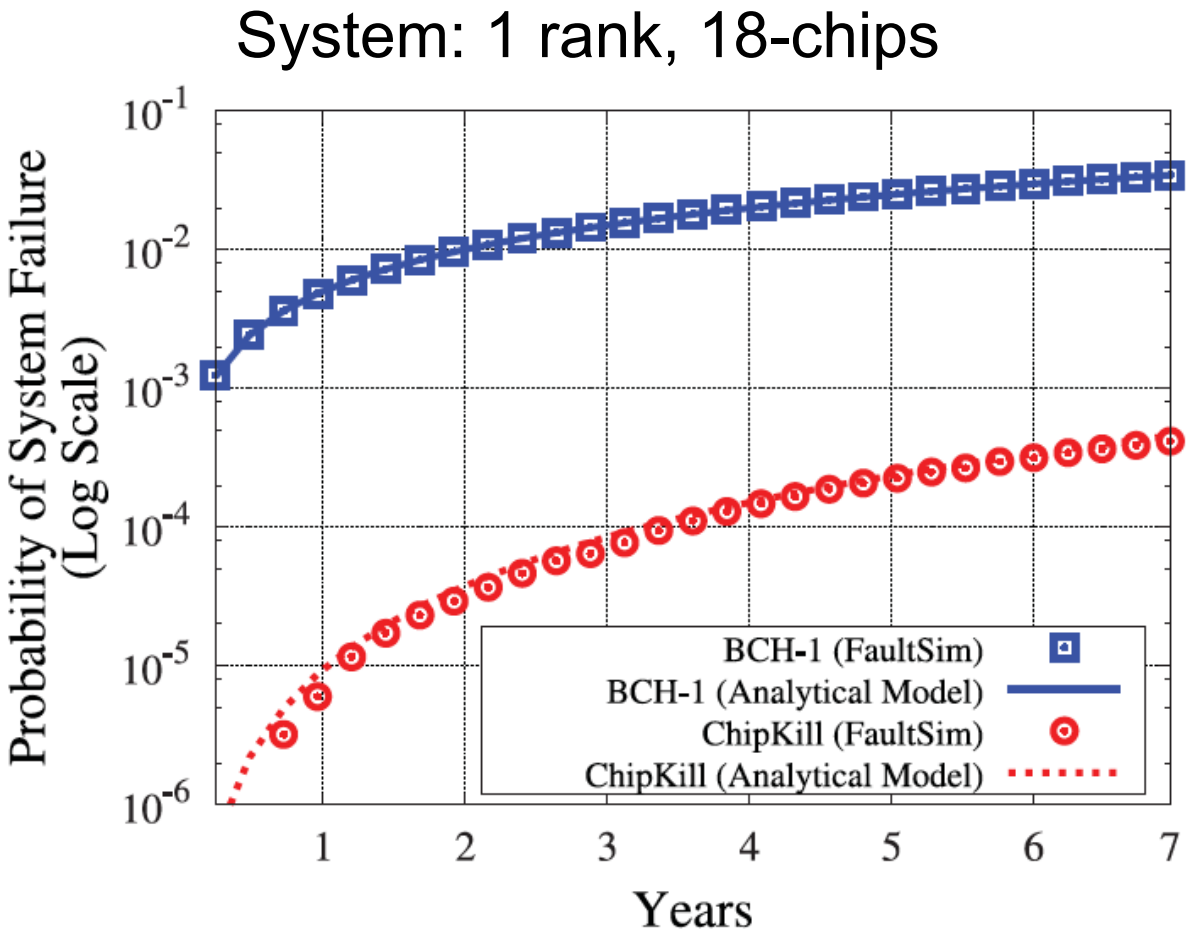

FaultSim closely follows the analytical model (within 2%)

RESULTS: SIMULATION TIME

Time for a million trials with FaultSim

REPAIR SCHEME	Simulation Time (Wall Clock)
SECDED	49.5 hours
ChipKill	49.2 hours

FaultSim still has simulation time in the order of days
How to we reduce this to less than a minute?

OVERVIEW

- WHY FAULTSIM?
- FAULTSIM: WHAT AND HOW?
- FAULTSIM: LESS THAN 1 MINUTE
- FAULTSIM: APPLIED TO 3D MEMORY
- SUMMARY

OBSERVATION: FEW FAULTS IN SYSTEM LIFETIME

FaultSim consults random number generator at-least once during each interval (20K)

System with 2 DIMMs, 9 chips each, over 7 years

Num. Faults Encountered (Total)	TRIALS
0	92.9%
1	6.7%
2	0.2%
3+	0.2%

Can we consult random number generator in proportion to faults, instead of every time interval?

INSIGHT: COMPUTE DISTANCE TO NEXT FAULT

The time between events in a process in which events occur continuously and independently at a constant average is ***exponentially distributed***

Example: Let the likelihood of a lottery ticket be a winner be 1/1000. We buy 5000 tickets. What is the likelihood of "X" winning tickets?

Naïve Method: Draw 5000 tickets, for each ticket check if it is winner

Distance Method: Compute distance to winning ticket using exponential distribution (avg=1000). Do until sum of distance > 5000.

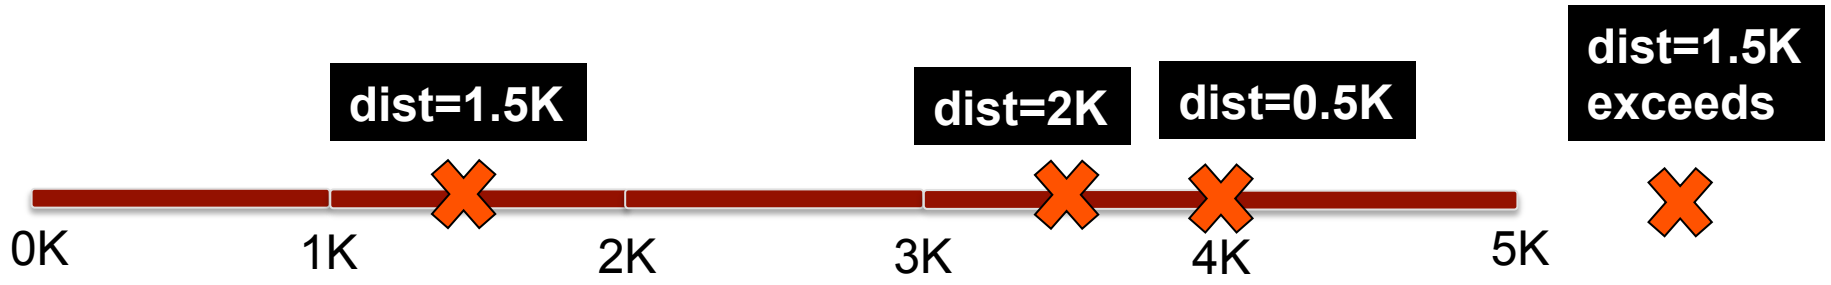

FAULTSIM: EVENT-BASED FAULT INJECTION

Event-Based Fault Injection: When is the next fault?

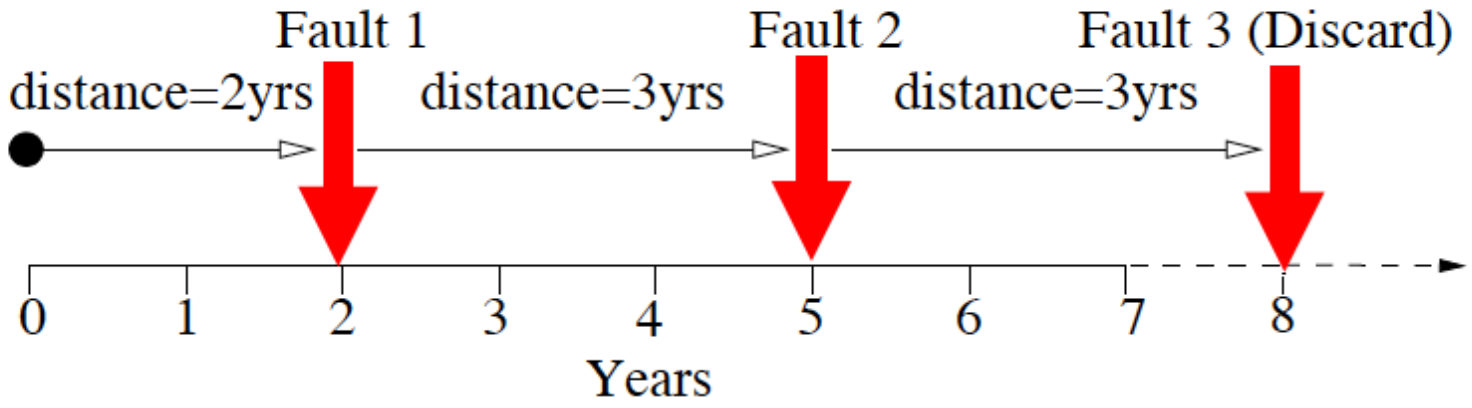

Time-Stamp of all faults computed at start of simulation.
Simulation skips from one fault to another

Calls to random number reduced from 20K to 1 (or 2)

RESULTS: SIMULATION TIME

Time for a million trials with FaultSim

SCHEME	Simulation Time (Wall Clock)
SECDED (Interval Based)	49.5 hours
SECDED (Event Based)	**34 seconds**
ChipKill(Interval Based)	49.2 hours
Chipkill (Event Based)	**33 seconds**

FaultSim ~5000x faster with Event-Based Fault Injection ➔ reliability simulation in less than one minute

OVERVIEW

- WHY FAULTSIM?
- FAULTSIM: WHAT AND HOW?
- FAULTSIM: LESS THAN 1 MINUTE
- FAULTSIM: APPLIED TO 3D MEMORY
- SUMMARY

INCORPORATE NEW TECHNOLOGIES: 3D MEMORY

Industry moving towards 3D DRAM for higher BW
New failure modes due to Through-Silicon Via (TSV)

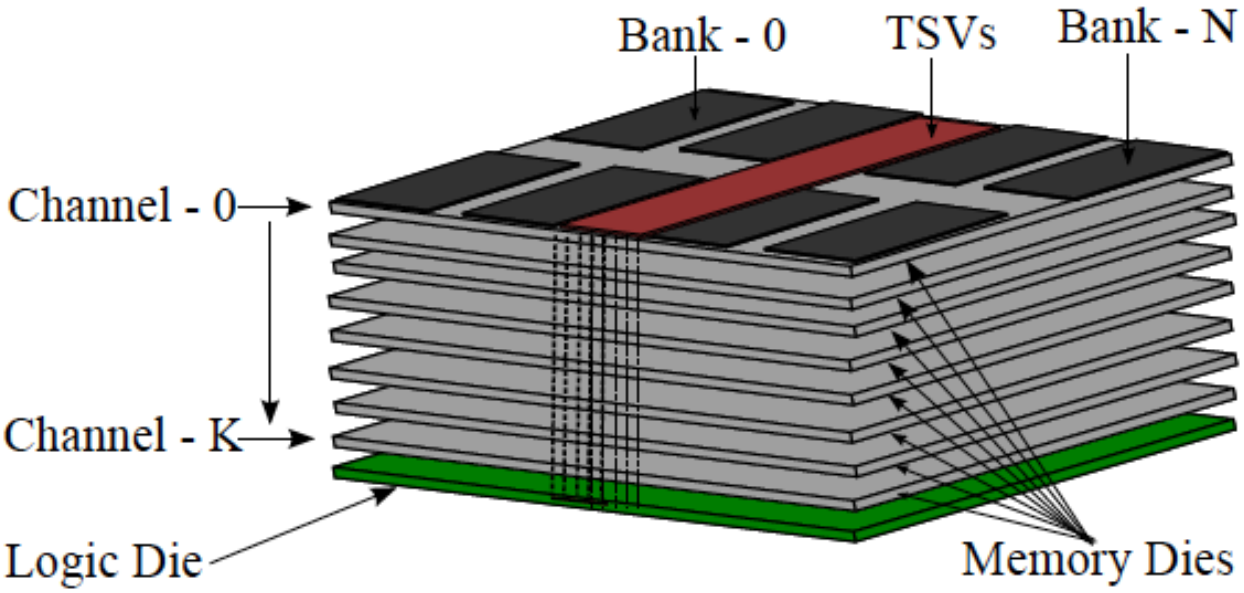

FaultSim can model new components like TSVs

CAPTURING THE EFFECT OF TSV FAULTS

- Data TSV Fault ➡ Few Columns Faulty
- Address TSV Fault ➡ 50% Memory Loss

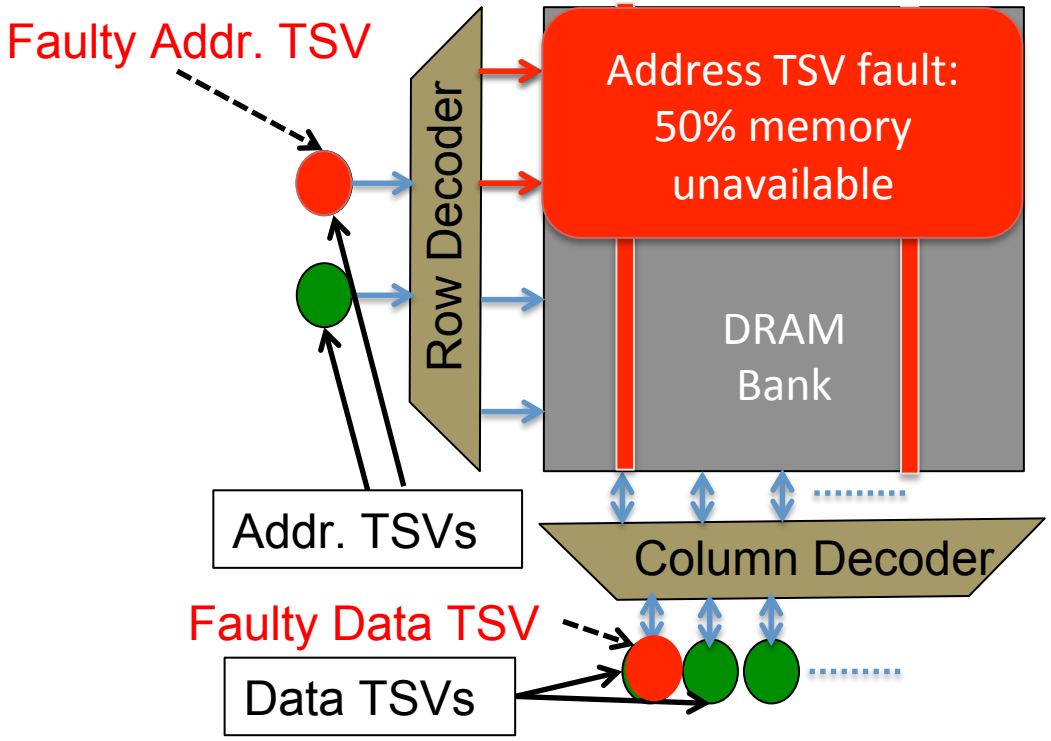

TSVs faults manifested as column/bank failure

USING FAULTSIM TO EVALUATE 3D MEMORY

["Citadel", MICRO 2014]

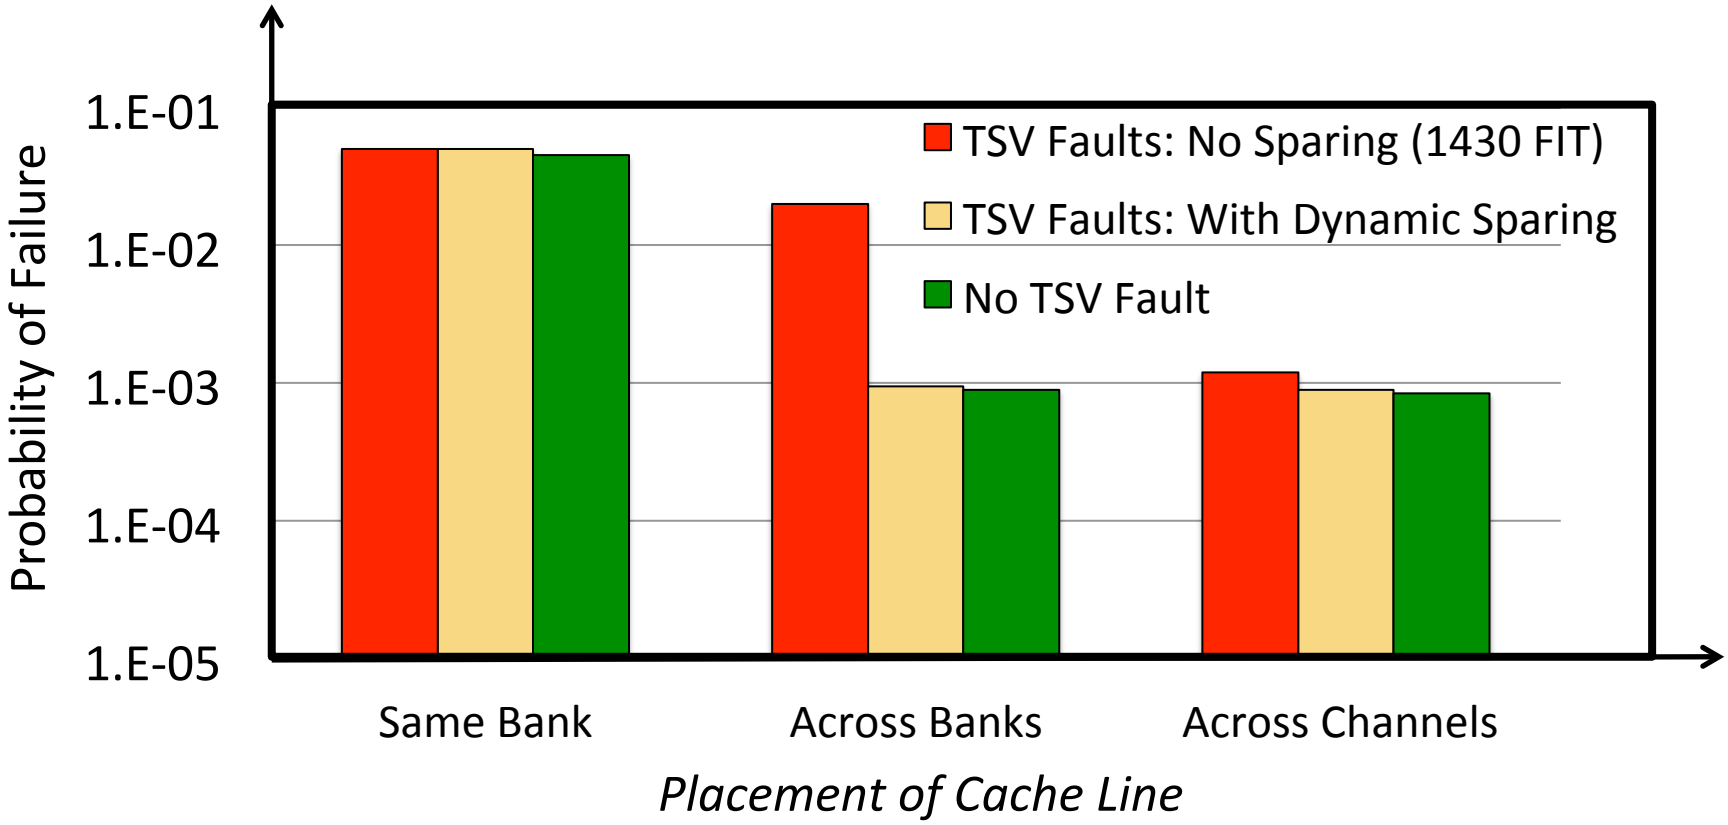

FaultSim used to evaluate TSV sparing in 3D memory

OVERVIEW

- WHY FAULTSIM?
- FAULTSIM: WHAT AND HOW?
- FAULTSIM: LESS THAN 1 MINUTE
- FAULTSIM: APPLIED TO 3D MEMORY
- SUMMARY

SUMMARY

- Memory-Reliability is becoming increasing important and there is a need for evaluation tools
- We introduce FaultSim → An efficient and fast memory reliability simulator
- FaultSim uses event based simulation, efficient representation and quickly computable functions
- FaultSim enables evaluating memory-reliability within 2% of the analytical model
- FaultSim is ~ 5000x faster than interval-based Monte-Carlo simulator

OBTAINING AND RUNNING FAULTSIM

Clone it from github

$git clone https://github.com/Prashant-GTech/FaultSim-A-Memory-Reliability-Simulator

Running FaultSim

`./faultsim --help` for a list of command line parameters

```
./faultsim --configfile configs/DIMM_none.ini --outfile
out.txt
```