Hardware-Software Integrated Diagnosis for Intermittent Hardware Faults

Majid Dadashi
Layali Rashid
Karthik Pattabiraman
Sathish Gopalakrishnan

Electrical and Computer Engineering Department
University of British Columbia

June 25, 2014

Why Intermittent Faults?

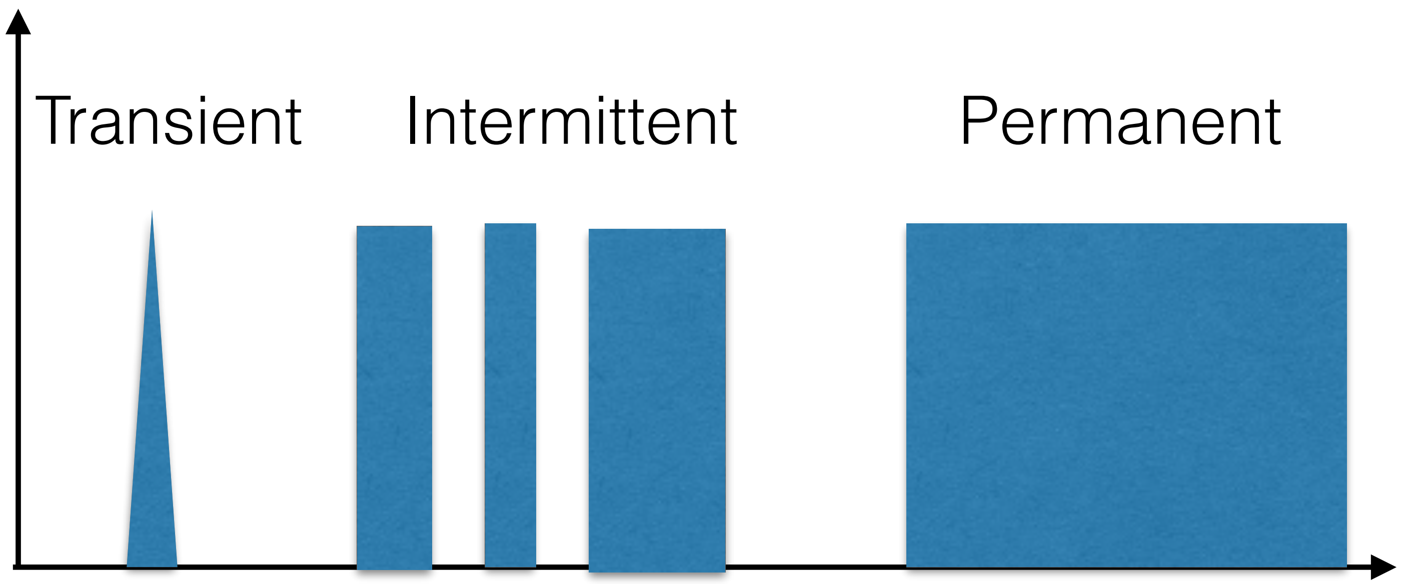

- Intermittent faults are becoming more prominent with technology scaling [Constantinescu 2008].
- One experiment has shown that intermittent faults were responsible for at least 39% of processor failures [Nightingale et al. 2011].
  - Large scale Microsoft study on **a million consumer PCs** based on Windows Error Reporting process.

Why **Online Fine-grained** Diagnosis?

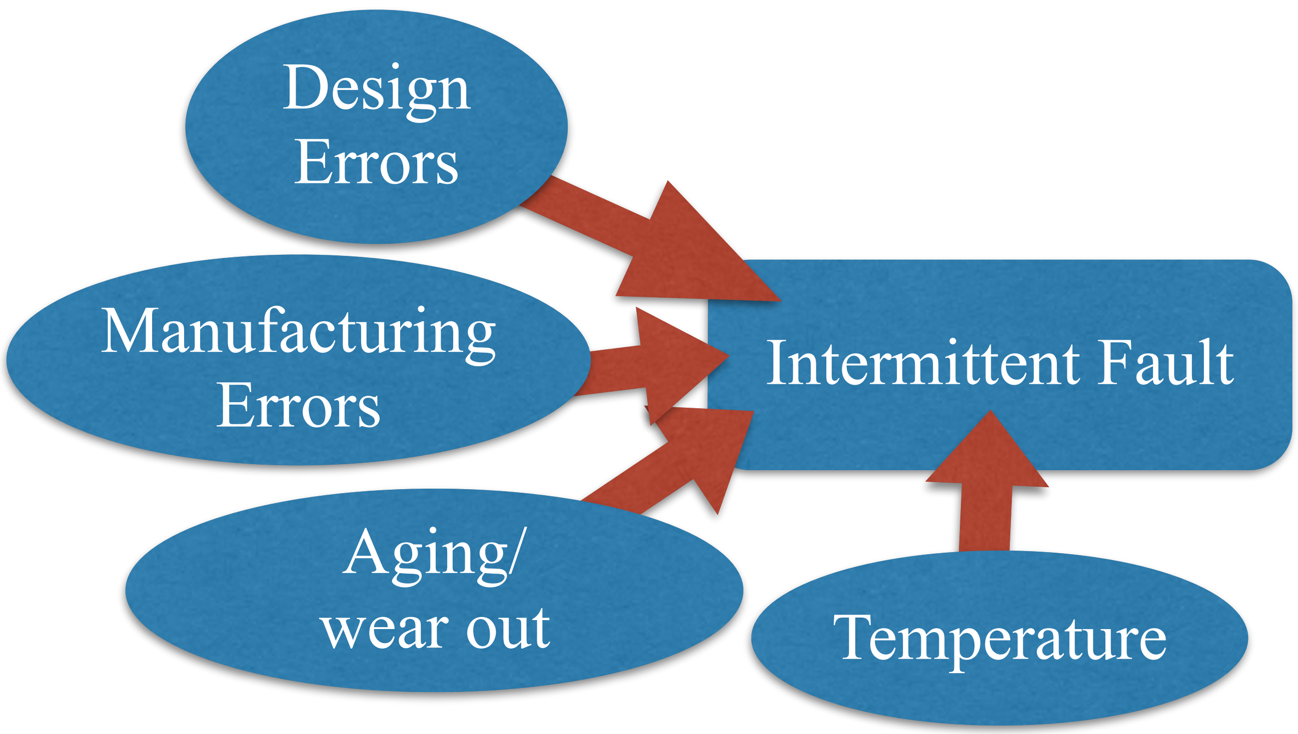

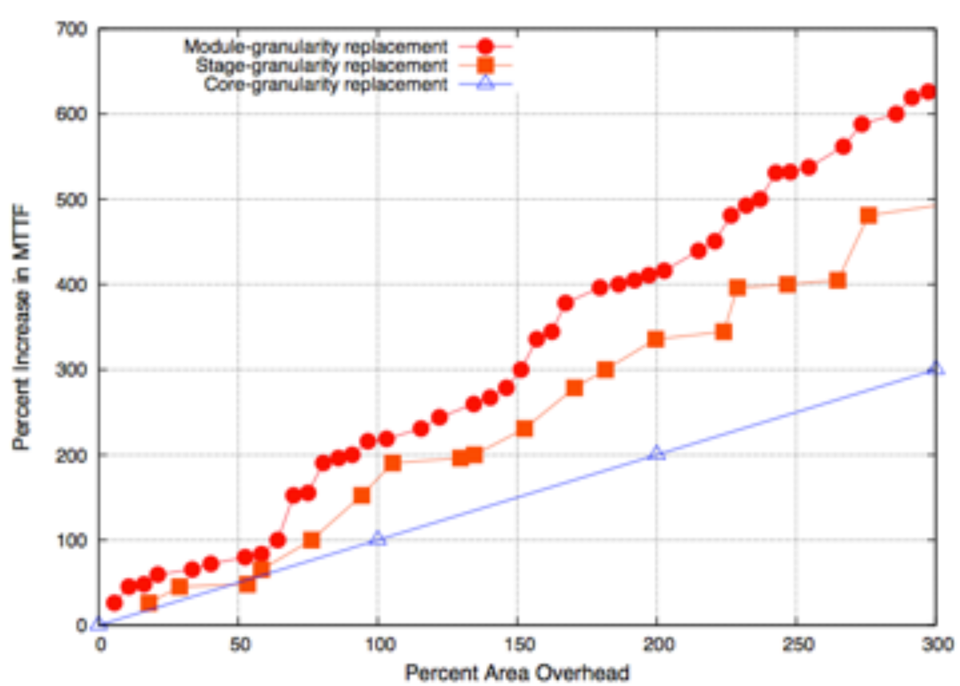

[Gupta et al. 2011]

- Intermittent faults can occur after the chip is shipped to the customer and so they need online diagnosis.
- The faulty part of chip can be disabled after diagnosis.
- The more fine-grained the diagnosis is the less slowdown will be imposed after repair.

Hardware/Software Co-Design

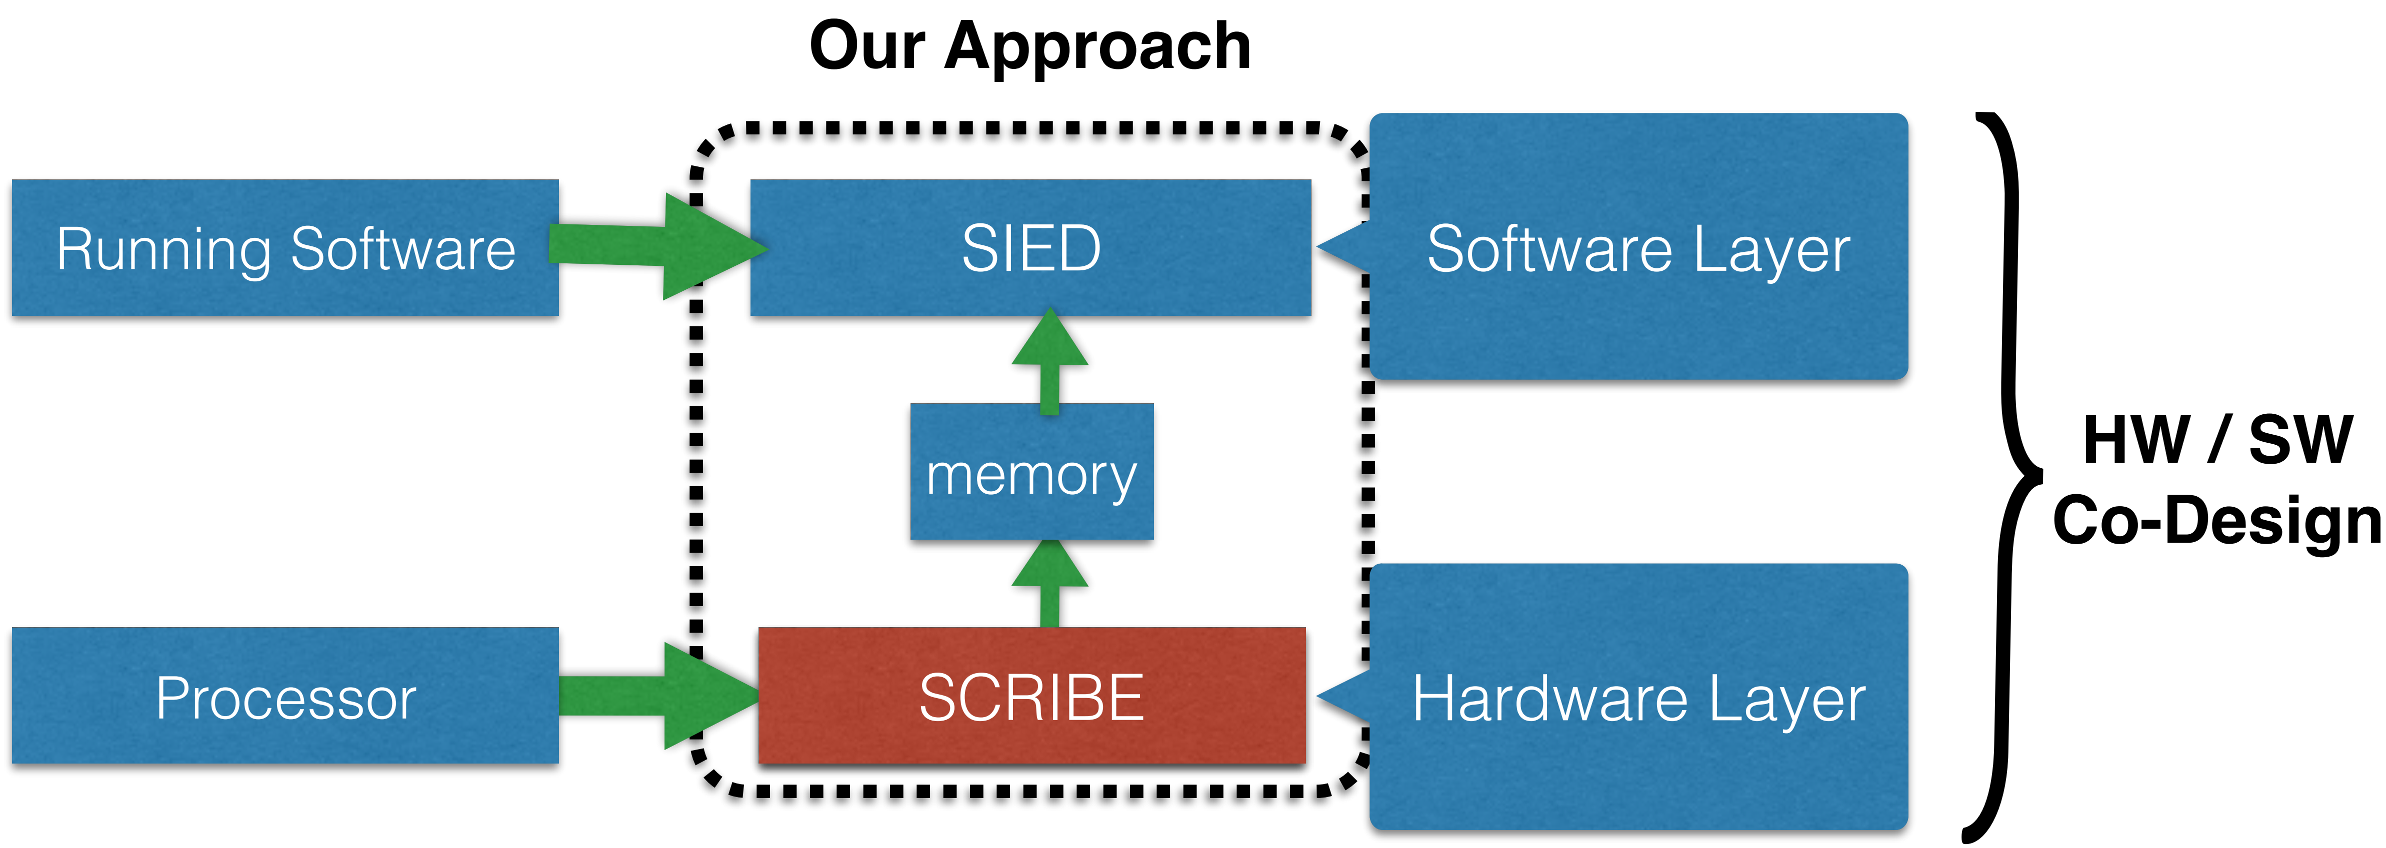

SCRIBE:
Providing Visibility through RUI Log

- Hardware layer collects the resource usage information as the instruction moves through the pipeline.
- RUI entry of an instruction is stored in a buffer to be sent to memory when the instruction is committed.

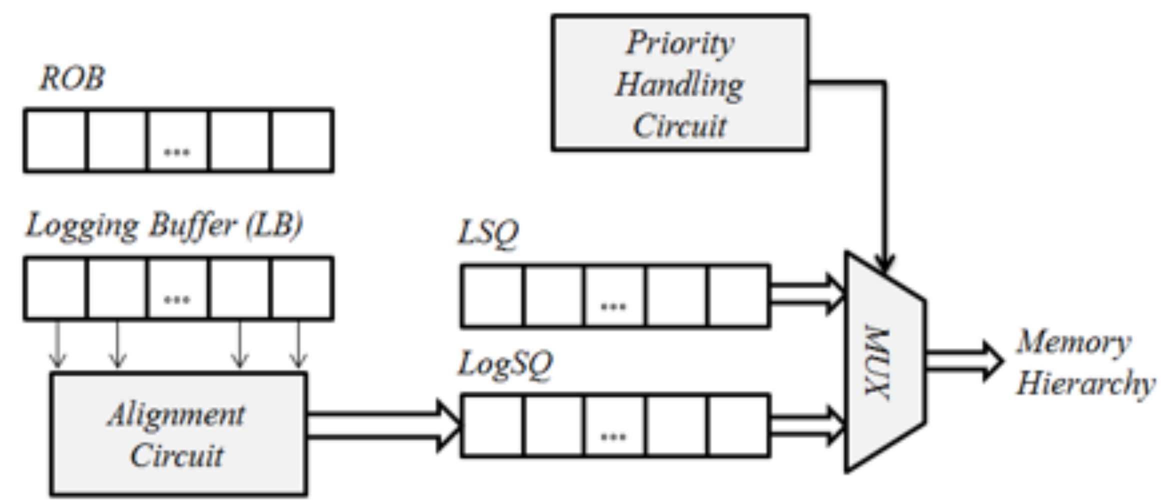

Example of an RUI Entry:

IFQ	ROB	RS	FU	LSQ
000100	0000111	0011000	0001	111111

SCRIBE: Logging RUI to Memory

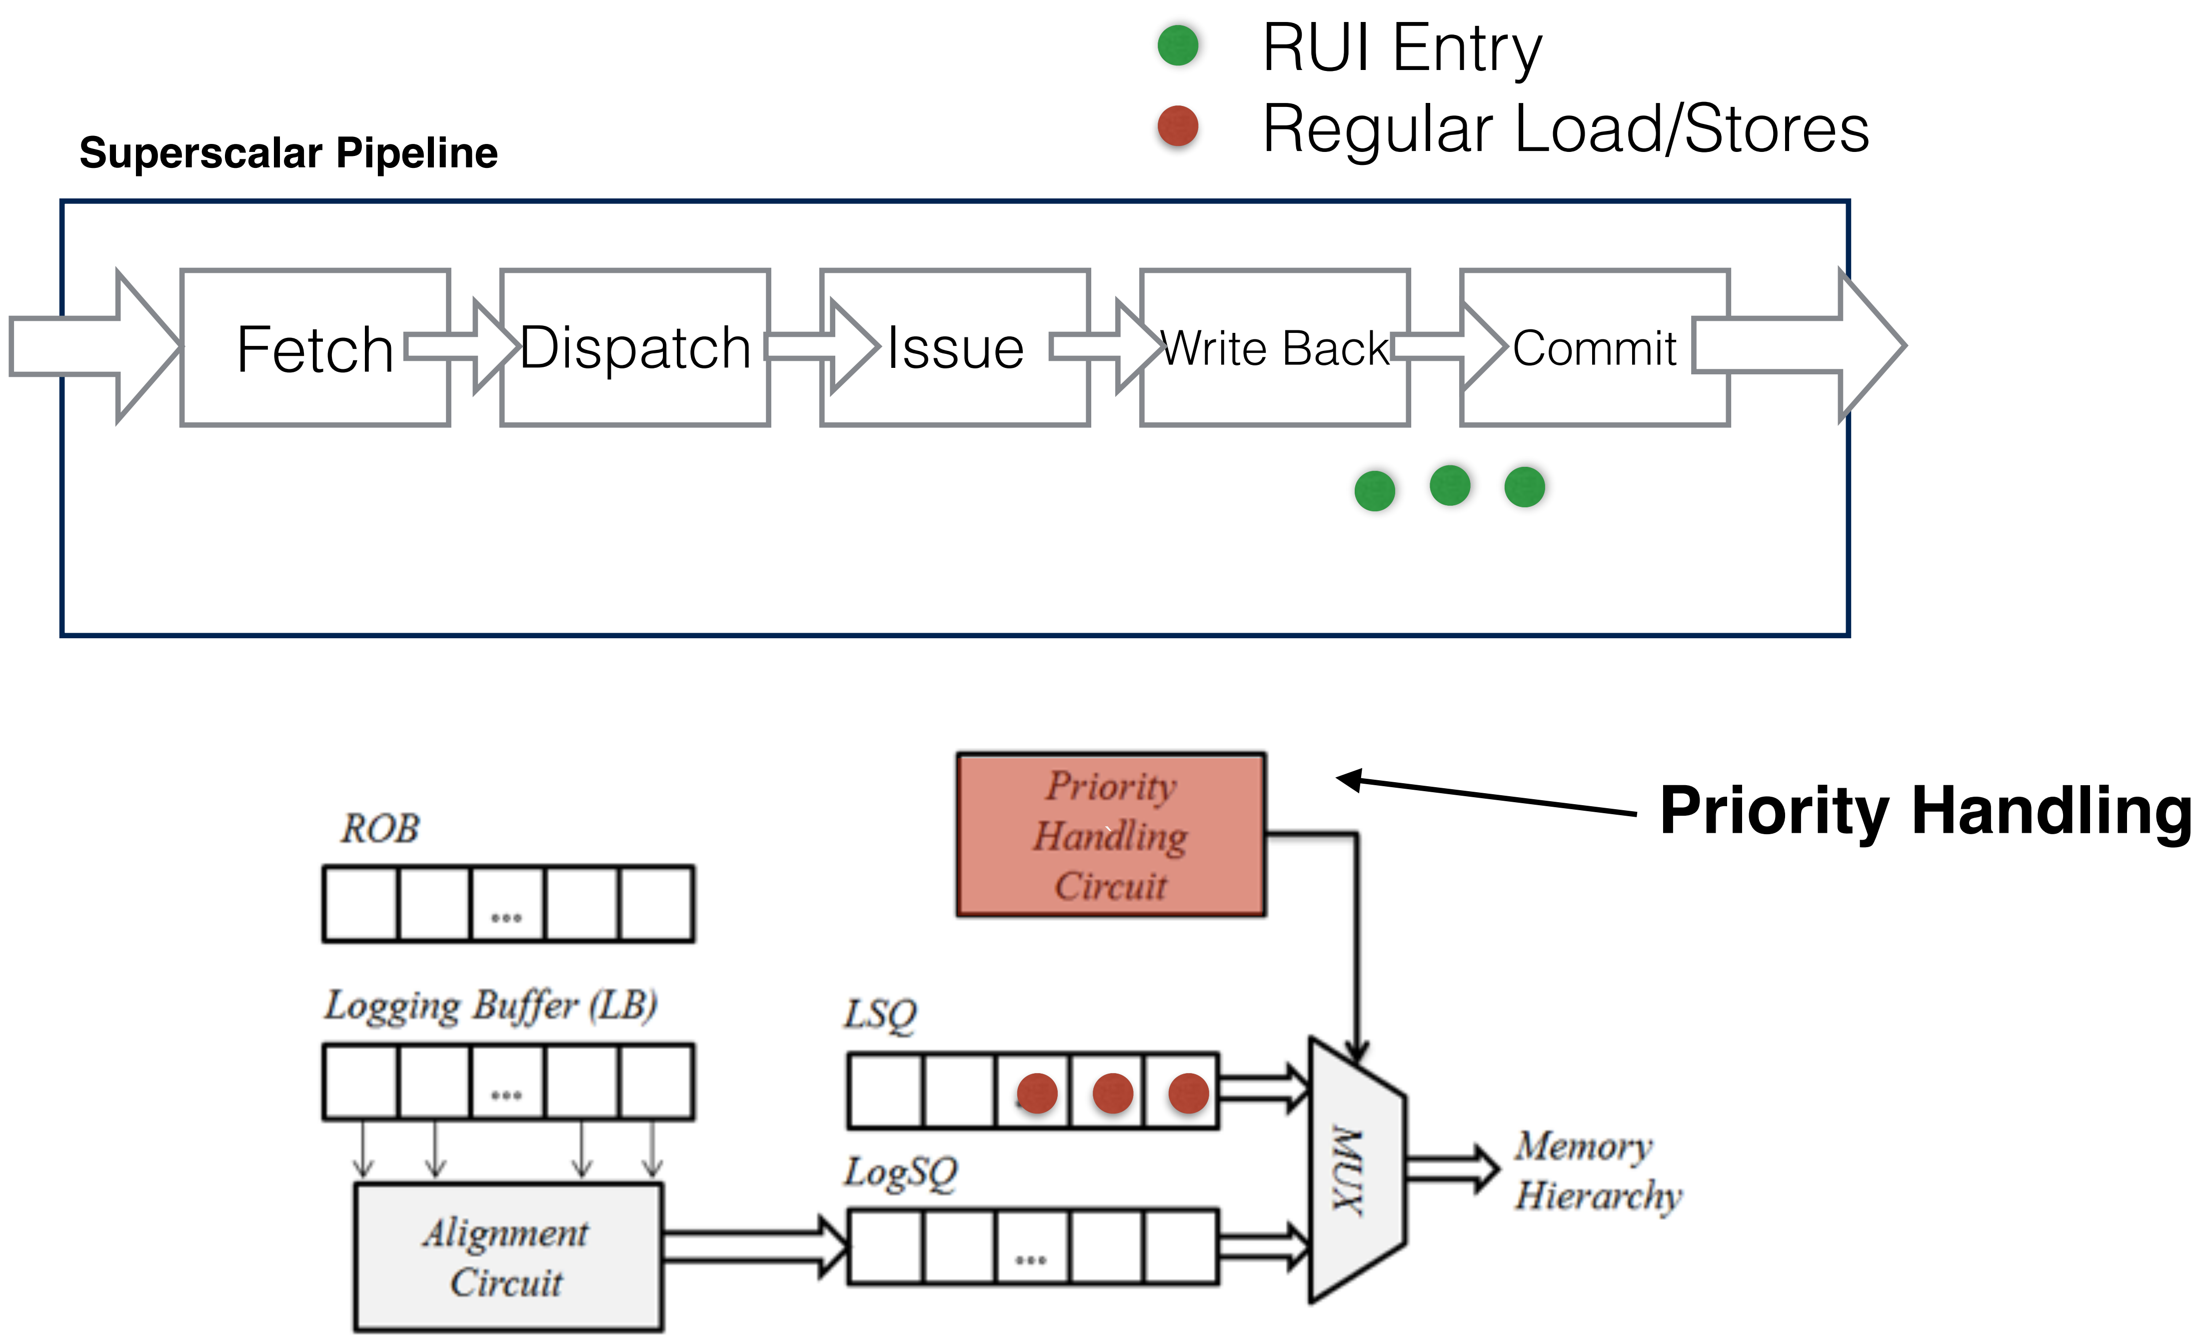

Failure and Diagnosis Scenario

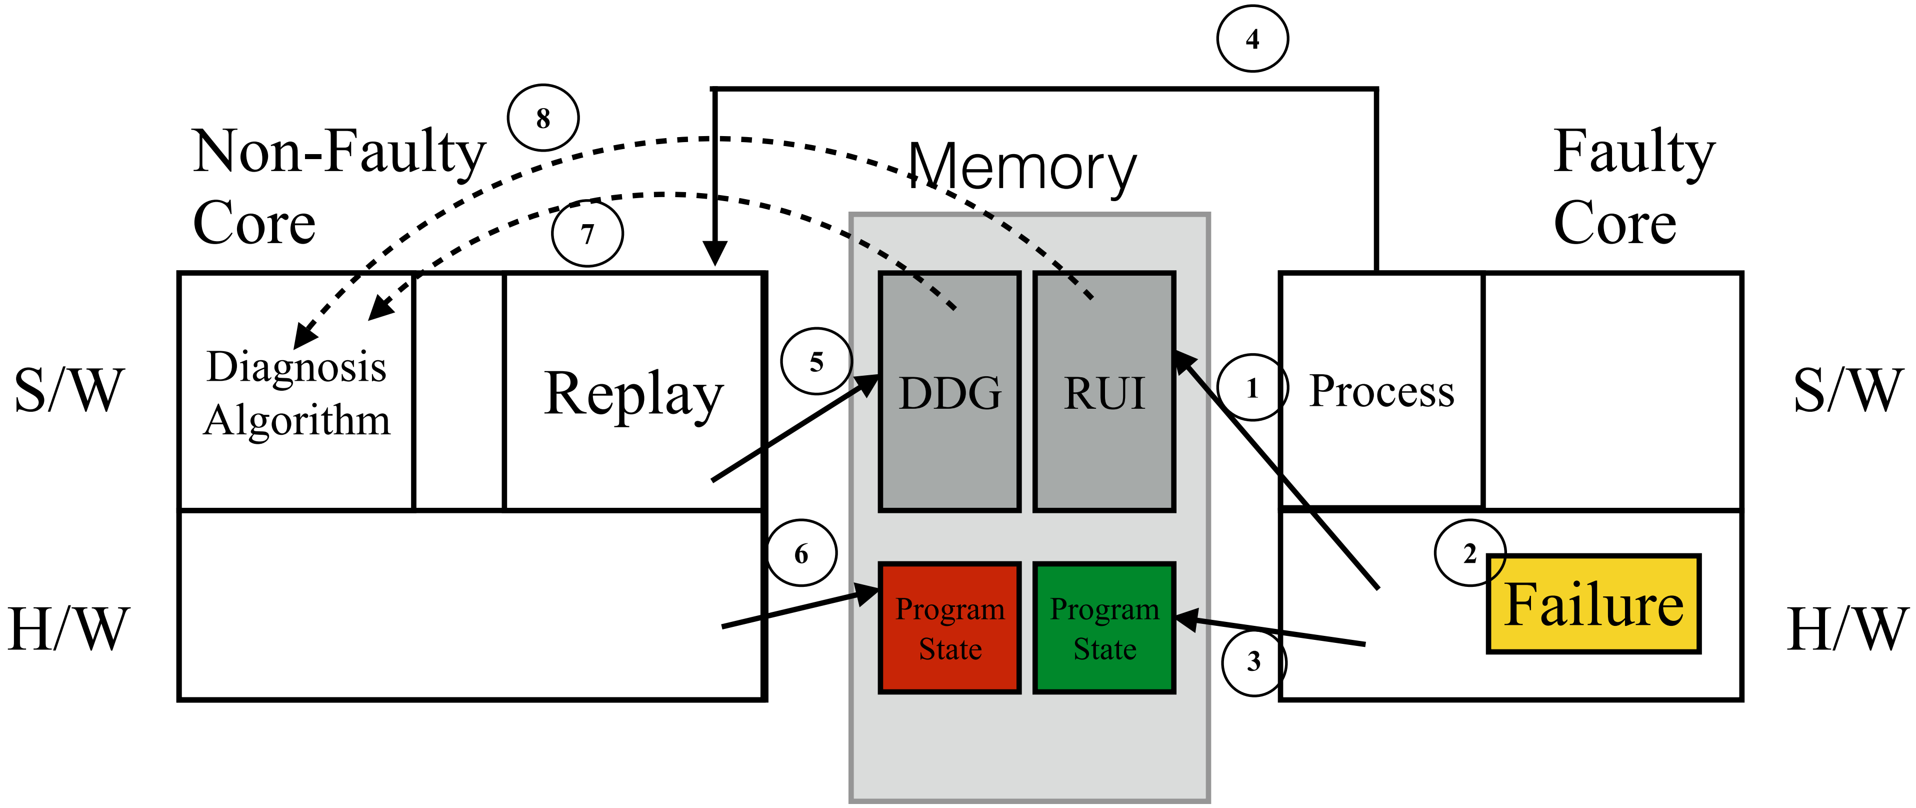

1. Gather RUI and log to memory (SCRIBE)
2. Failure due to intermittent fault
3. Log Program's register and memory state (core dump)
4. Deterministic replay on another core (SIED)
5. Construct replayed program's DDG (SIED)
6. Log replayed program's register and memory state (SIED)
7. Construct augmented DDG and backtrack using analysis heuristics (SIED)

Hardware/Software Co-Design:

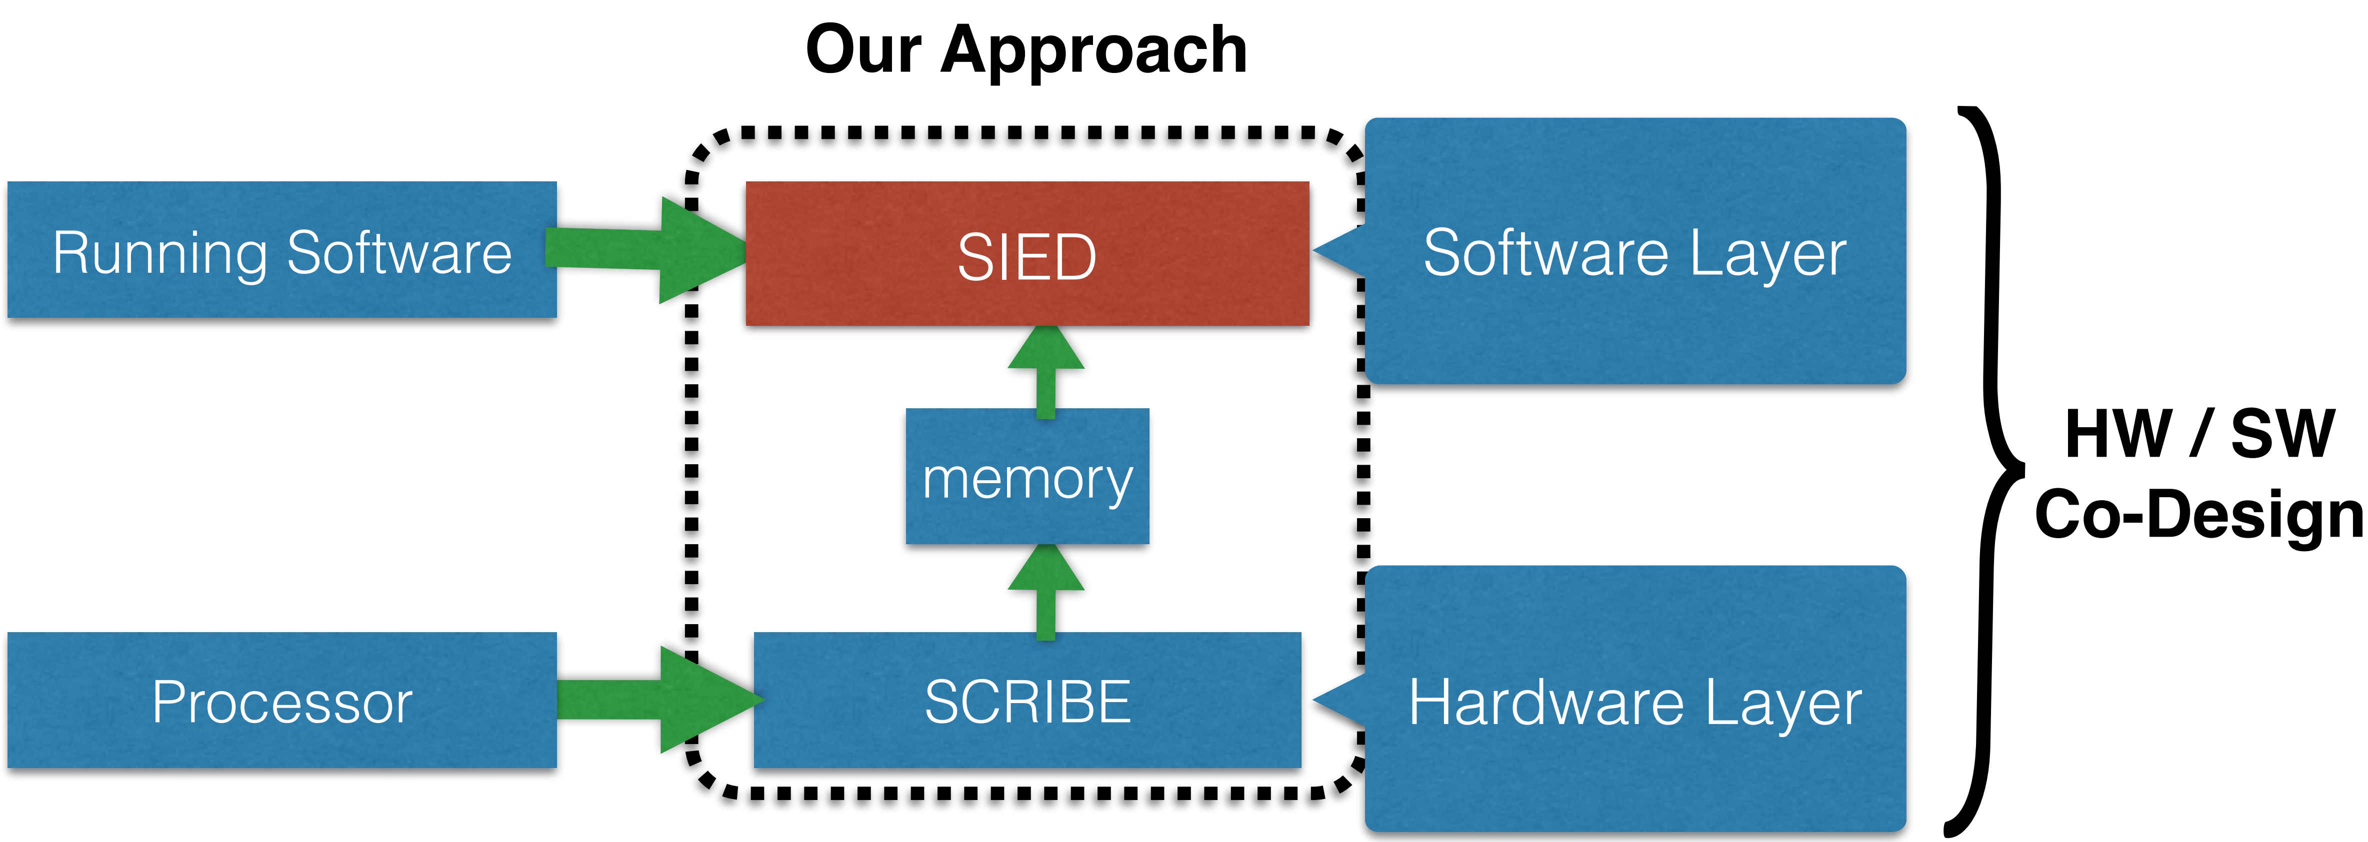

SIED: Software Layer

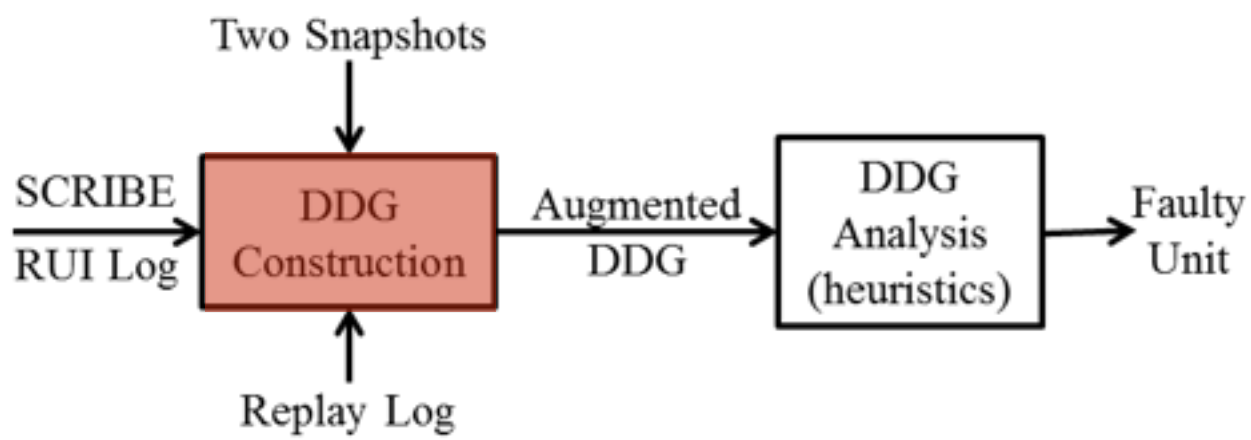

Augmented Dynamic Dependence Graph

- I1 : r3 <- r1 + r2
- I2 : r4 <- r2 + $2
- I3 : r5 <- r3 * mem[X]

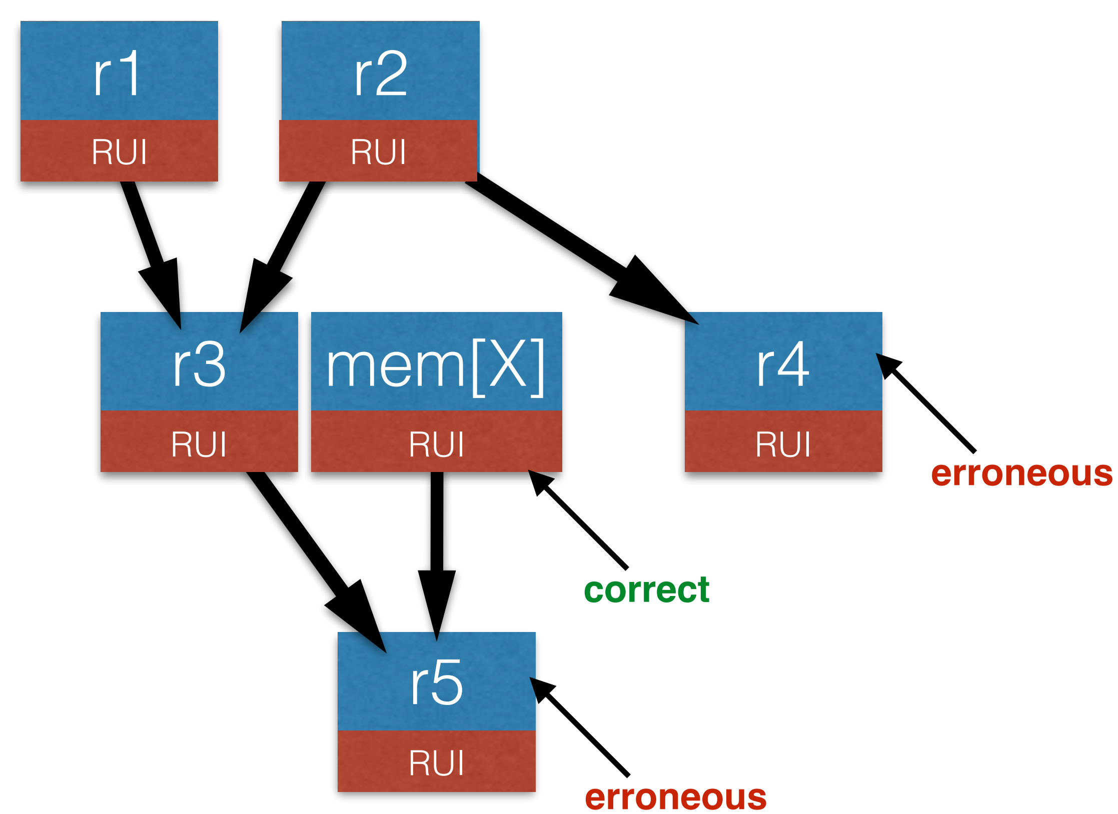

SIED: Example of DDG Analysis

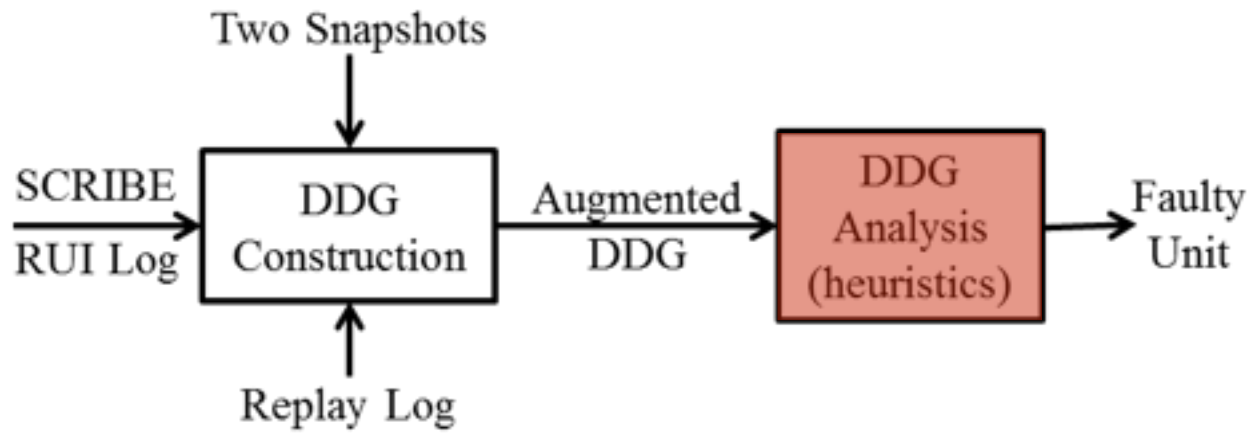

Simplified Example:

Resource Counters

Resource	Counter	
Adder 1	2	++
Adder 2	1	++
ROB[3]	1	++
ROB[4]	1	++
ROB[5]	1	++
IFQ[2]	1	++
IFQ[3]	1	++
IFQ[5]	1	++
. . .		

The resources with the highest counter values are most likely to be faulty.

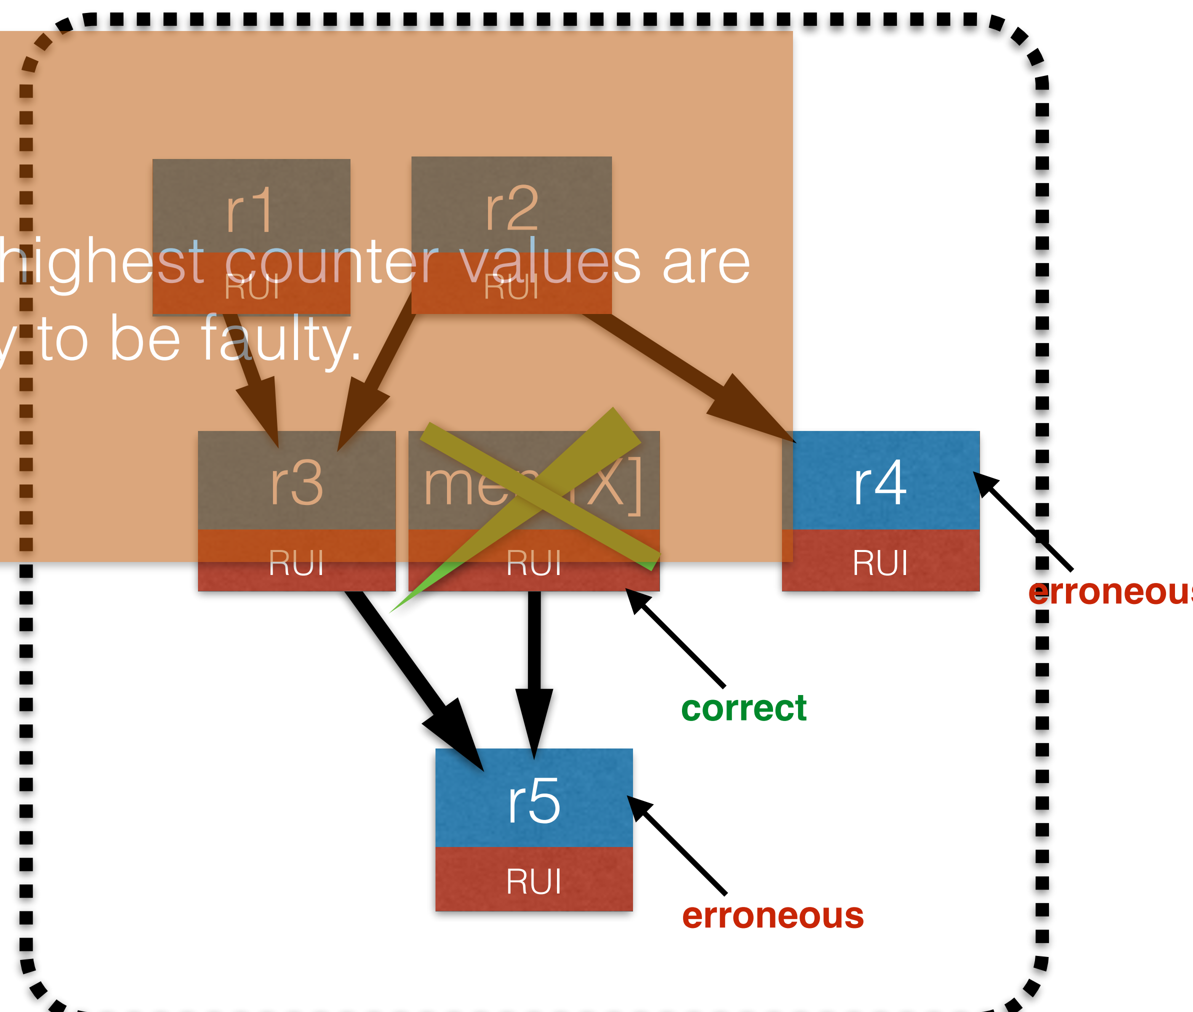

Experimental Setup

Simulation Environment	*sim-mase* from SimpleScalar package
Power Evaluation	*Wattch*
Benchmarks	SPEC 2006
Fault Model	Stuck at 1/0 - Random duration - Random starting cycle - in ROB, FU, RS, IFQ, LSQ
Simulated Processors	Narrow, Medium → Embedded/Desktop Processors Wide → High End Server Processors

Experimental Methodology

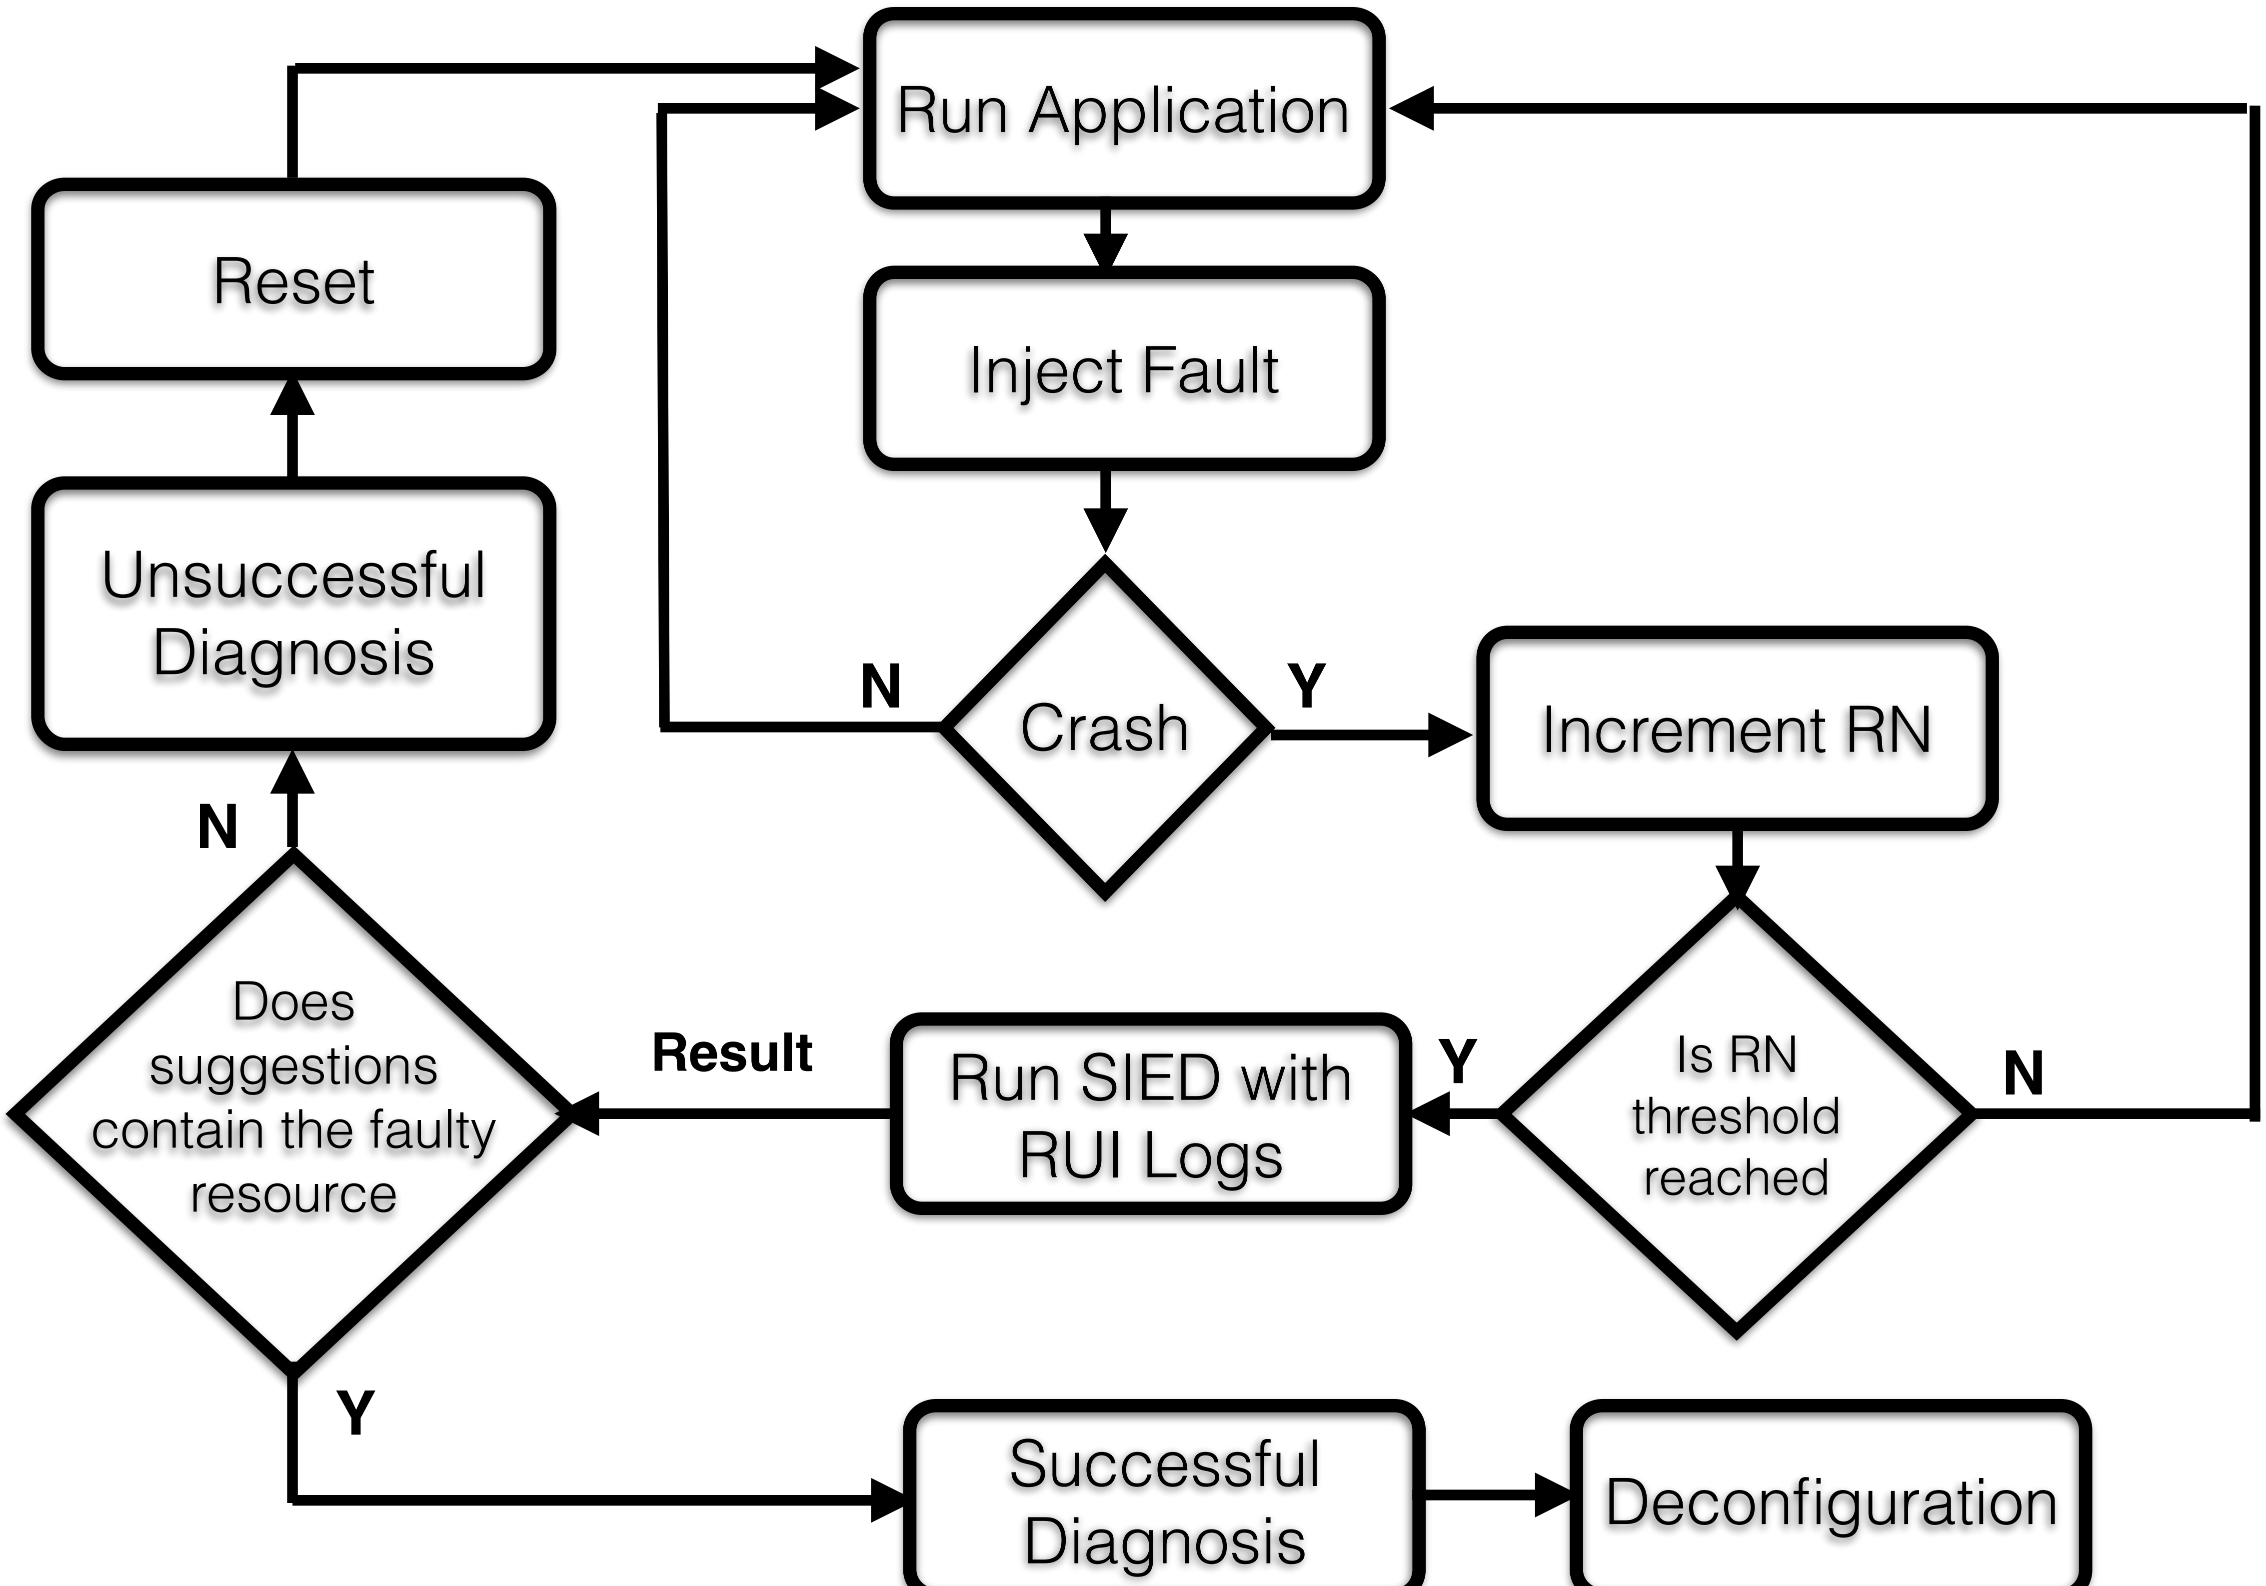

Diagnosis Accuracy

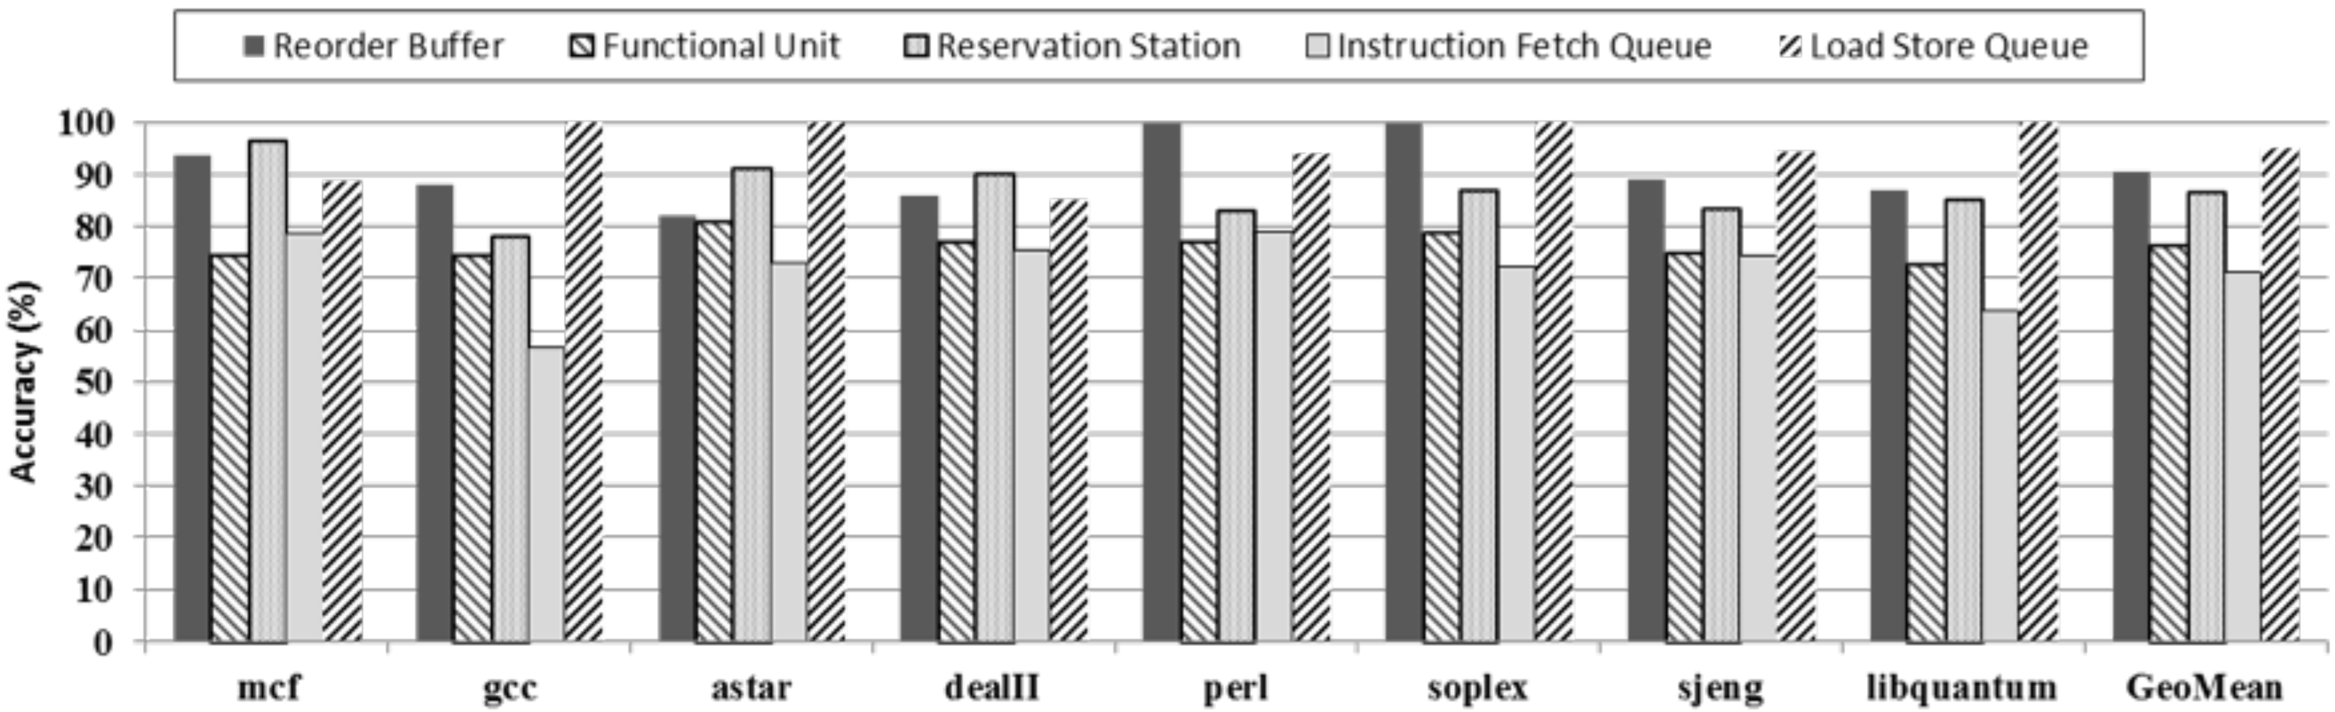

- Accuracy Definitions:
  - Probability of the faulty resource being among the resources suggested by SIED

$$accuracy = \frac{\#\ of\ Successful\ Diagnoses}{Total\ \#\ Diagnoses} \times 100$$

Average Accuracy:

84%

Diagnosis Accuracy

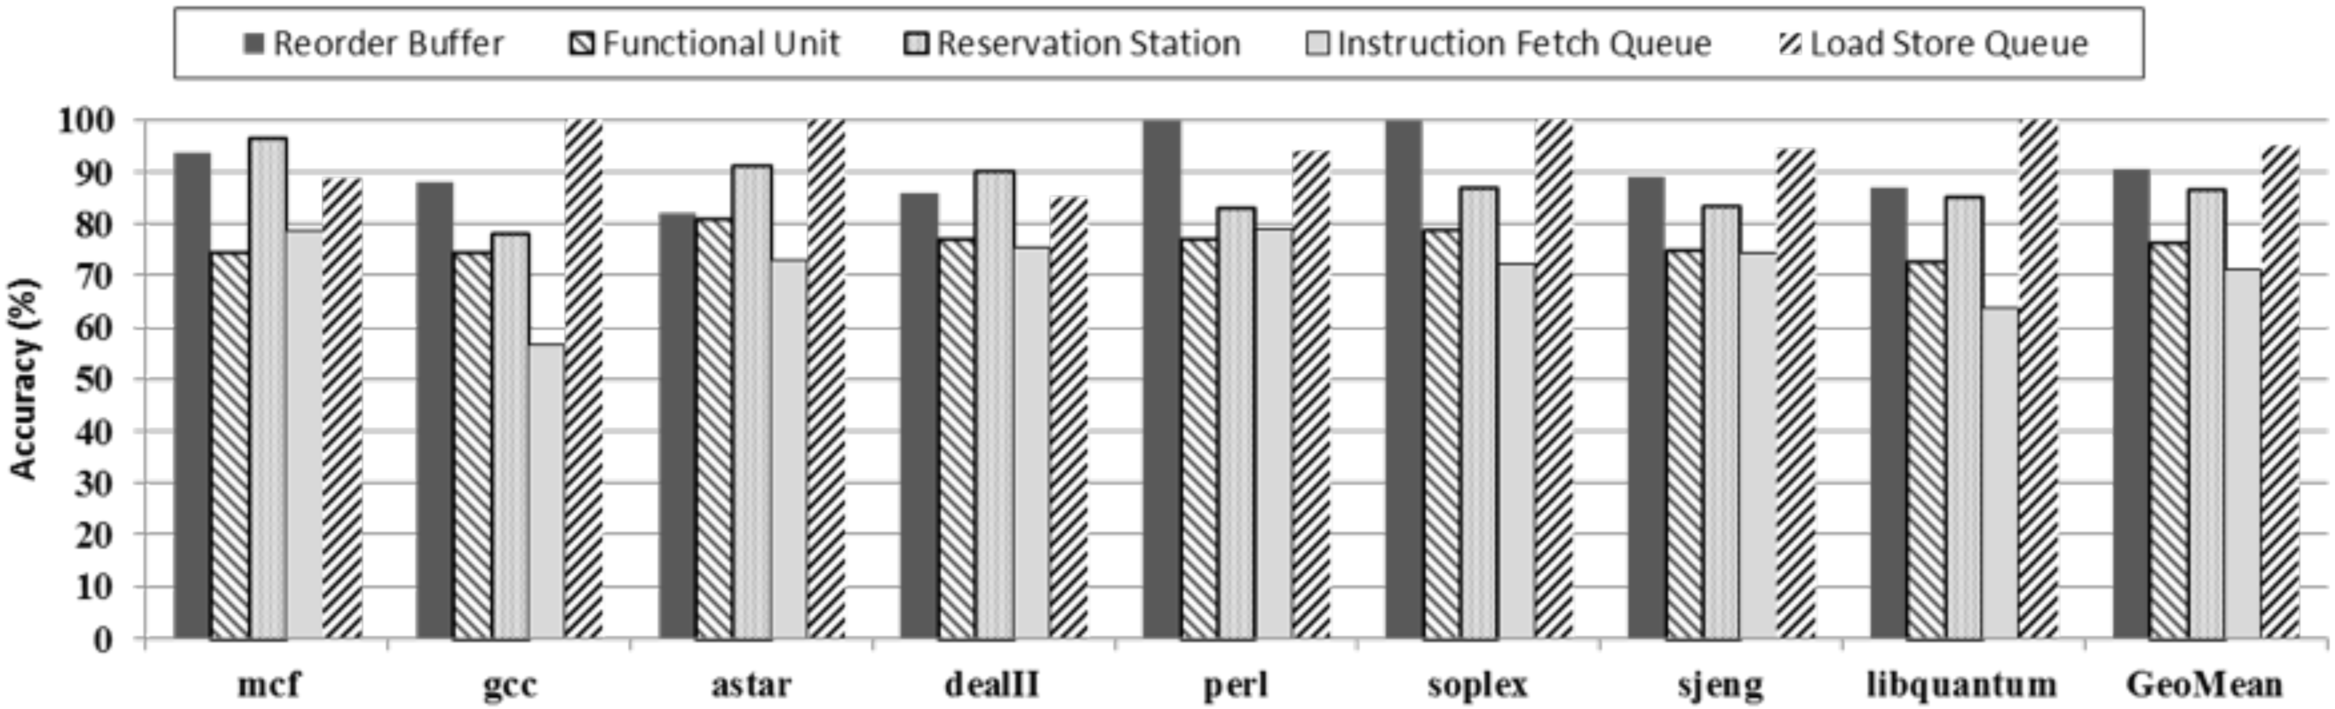

- Inaccuracy Reason:
  - SIED only knows about correctness of *final* data
  - Backtracking:
    - Resources used in correct instructions in the backward slice of erroneous instructions will also be counted.

Maximum Accuracy:
LSQ - 95%

Minimum Accuracy:
IFQ - 71%

Deconfiguration Granularity

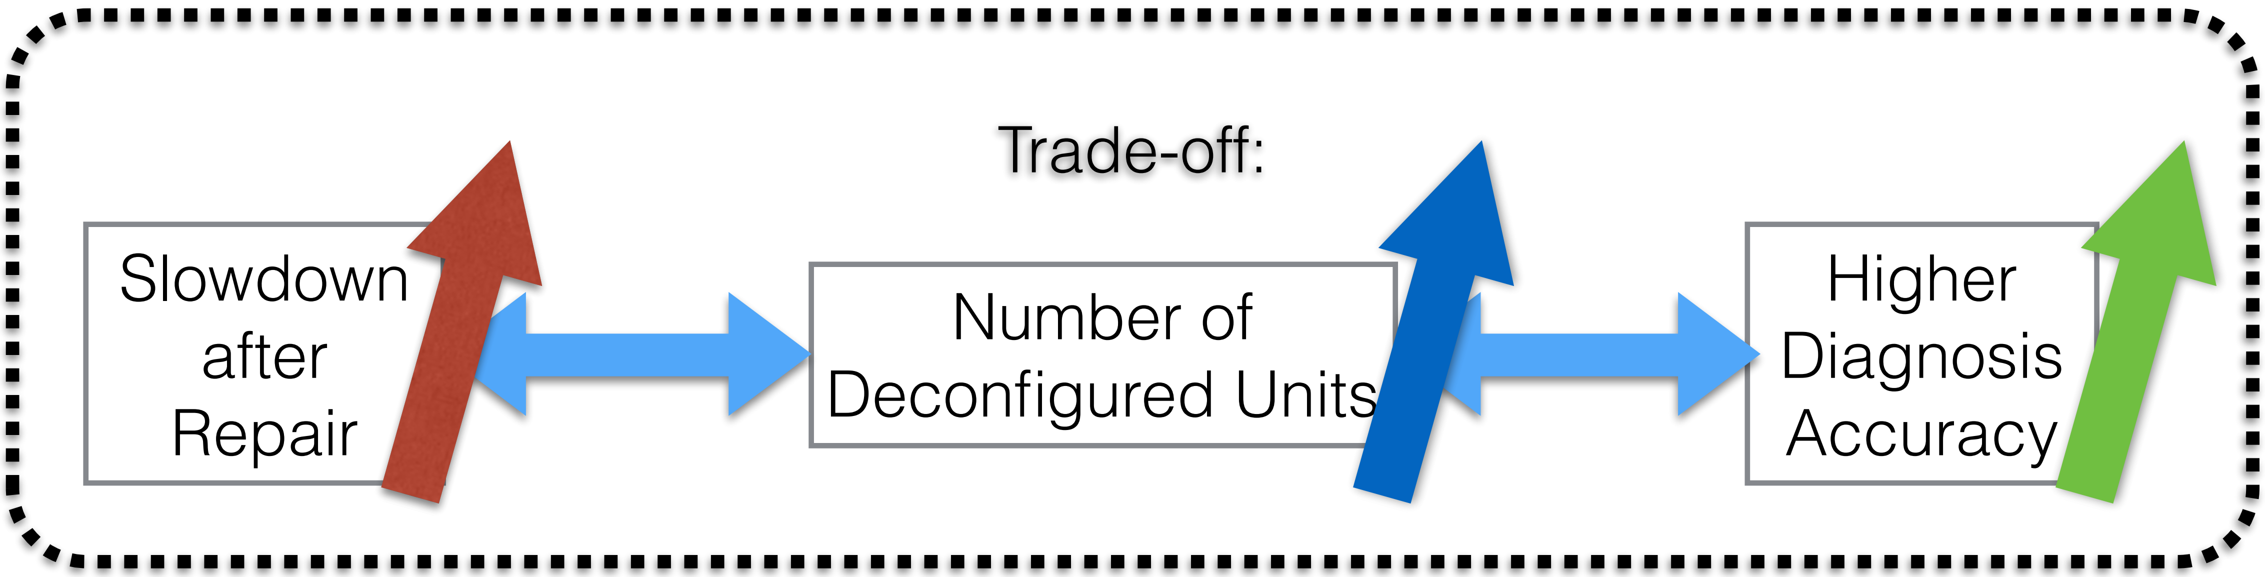

N_{deconf} = 5
Slowdown = 1.6%

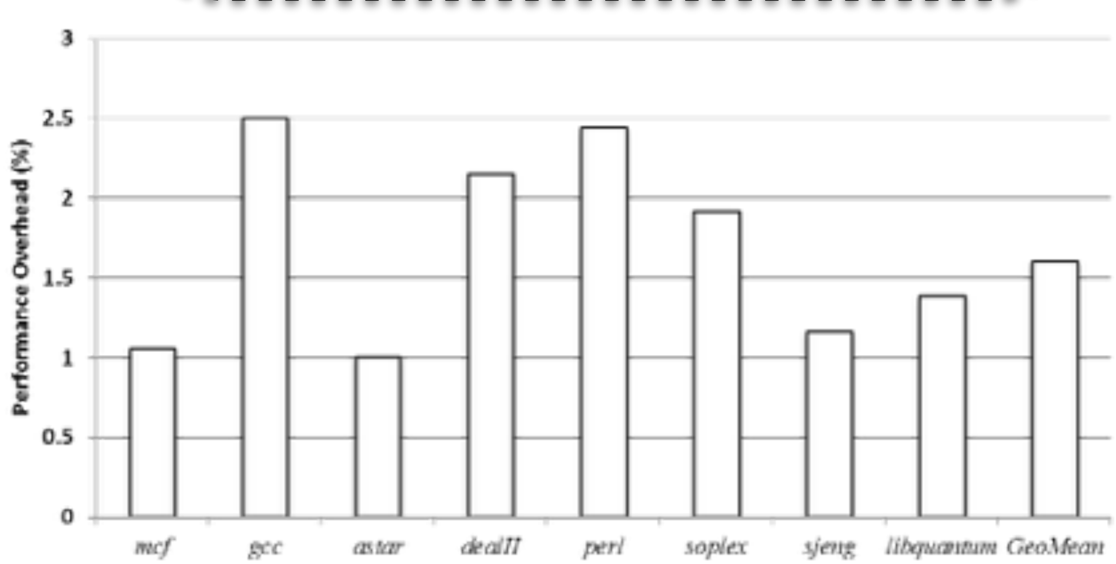

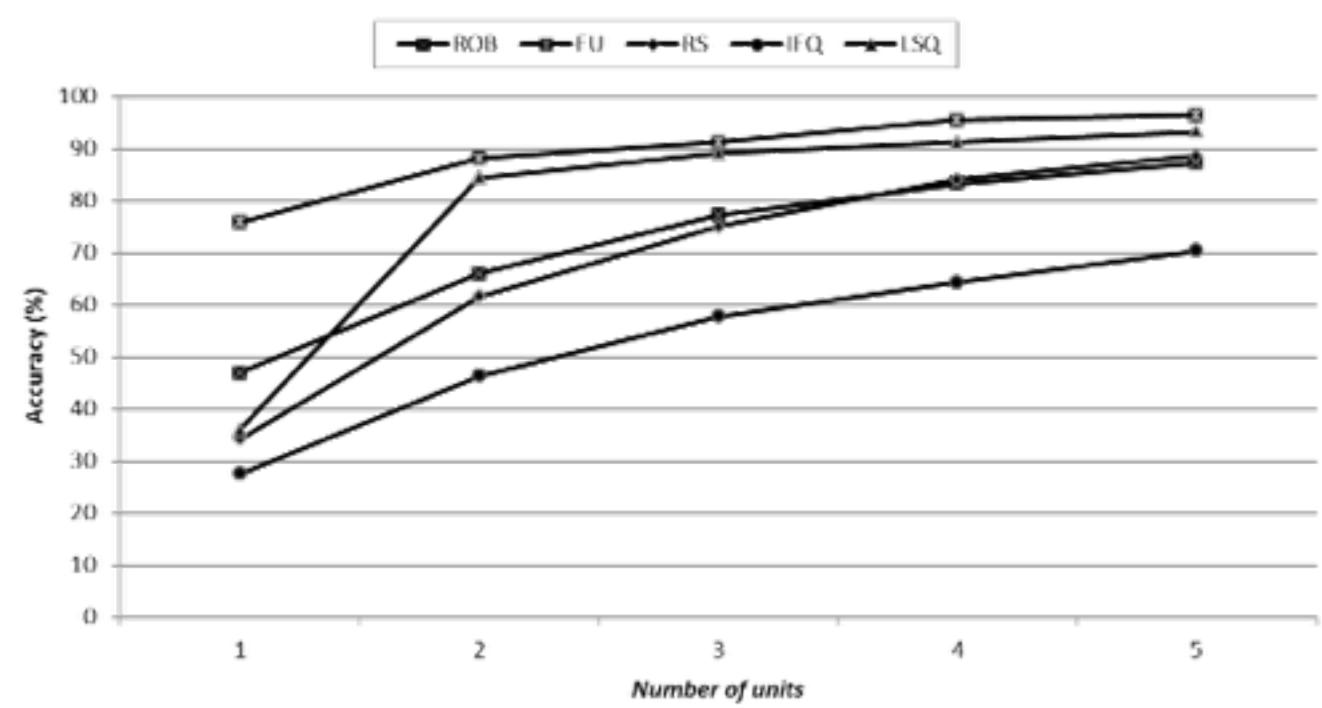

Failure Recurrence

- Intermittent faults are non-deterministic but recurrent.
- Every diagnosis of a recurrent failure provides more information.
- Resource counters are the average of the resources counters among multiple recurrences.
- We report the accuracy after the 4th recurrence.

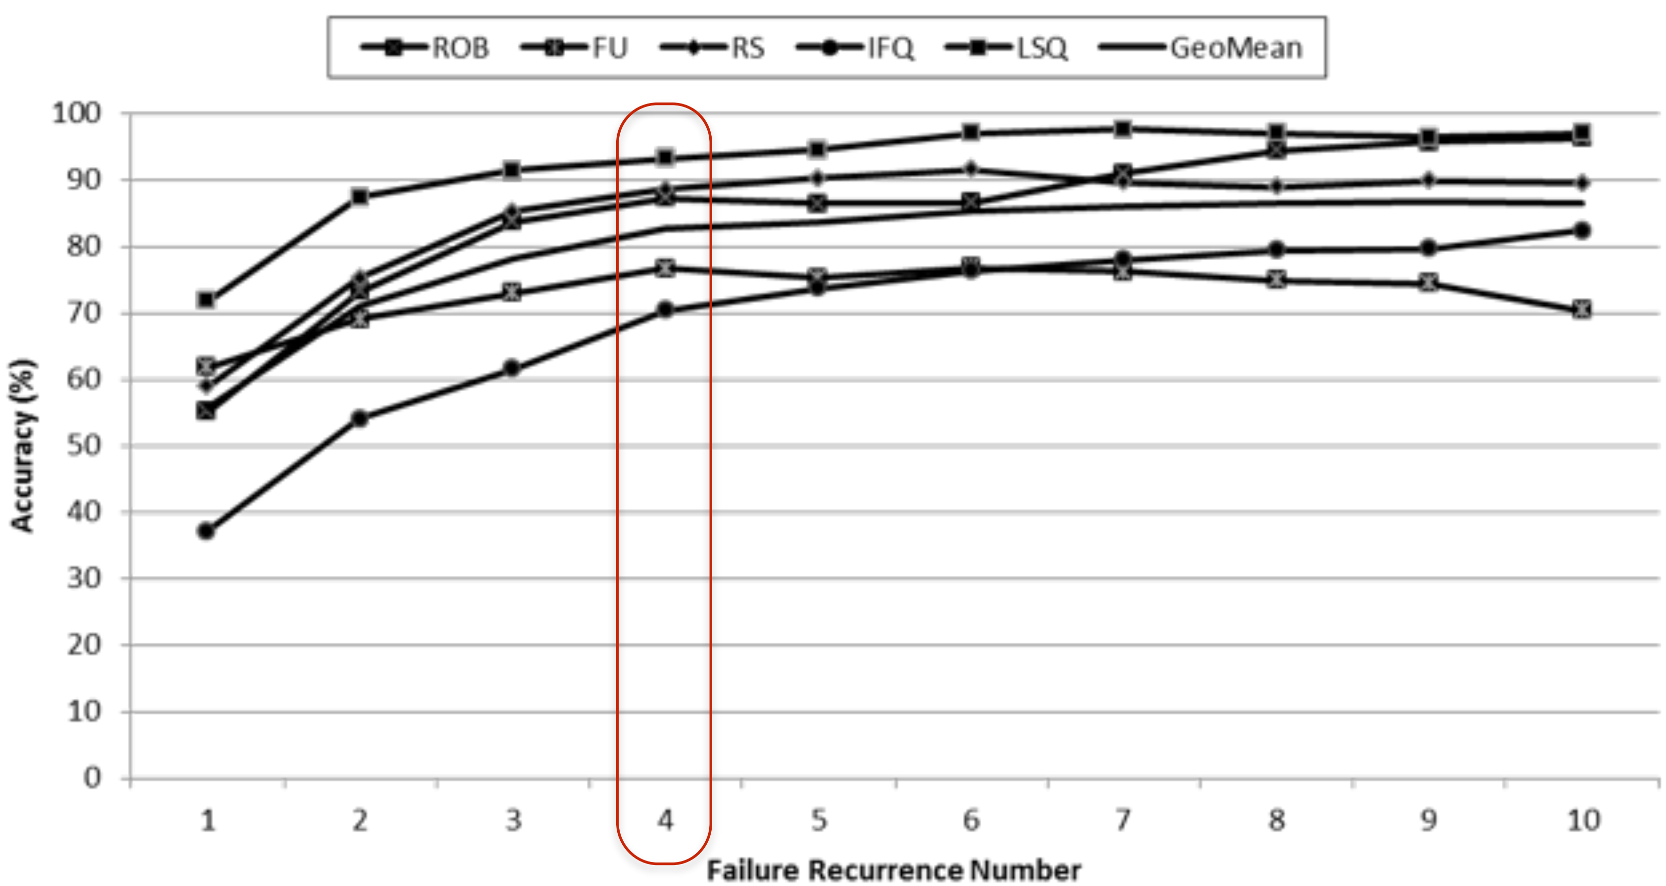

Performance and Power Overhead

- Online performance overheads:
  - Narrow: 11.53%
  - Medium: 11.88%
  - Wide: 23.21%
  - More Pipeline Stalls => Less Scribe Overhead
    - RUI entries keep being sent even when the processor is stalled.
- Largest source of power overhead:
  - D-Cache accesses

Average Performance Overhead:

14.7%

Average Power Overhead:

9%

Performance

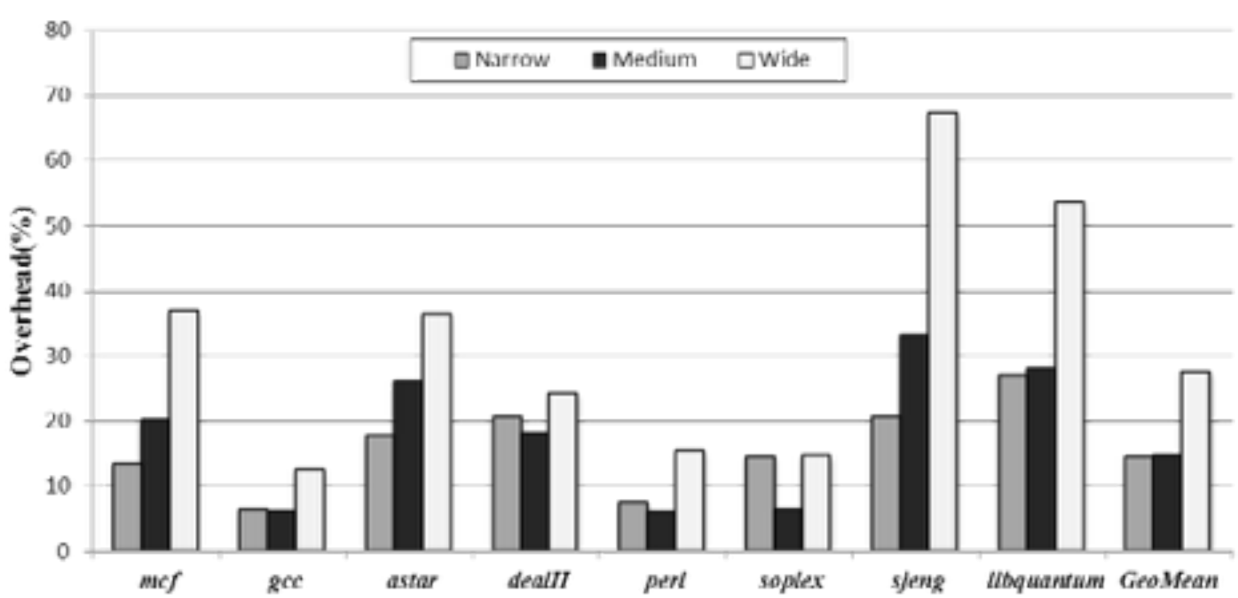

Power

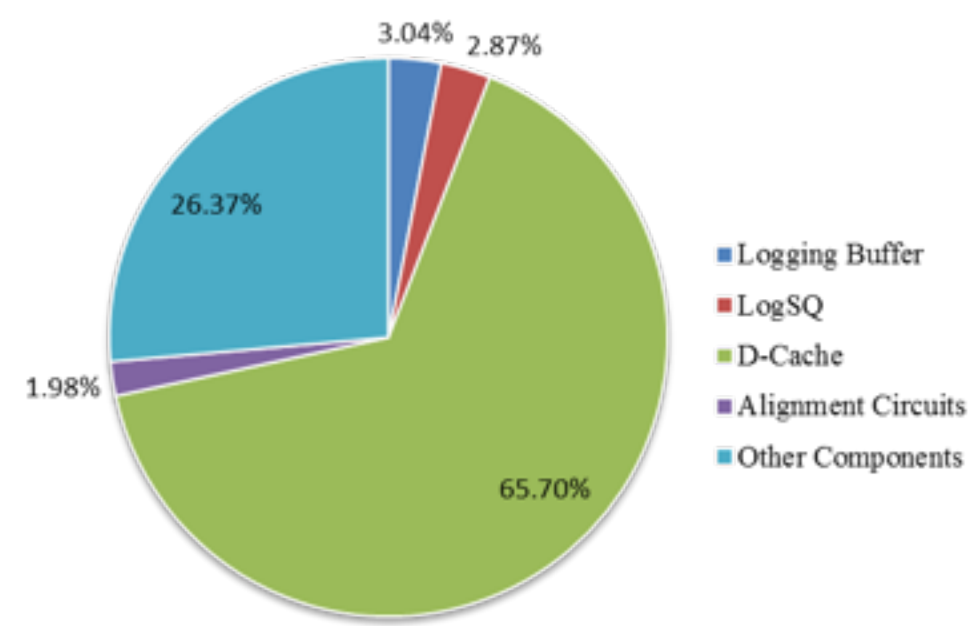

Related Work

- Diagnosis:
  - [Bower et al. 2005] - Duke University
    - Hardware only approach
    - Based on resource counters
    - Relies on DIVA fault detector
  - [Li et al. 2008] - University of Illinois at Urbana-Champaign
    - Hybrid
    - For permanent faults (relying on determinism)
  - [Carretero et al. 2011] - Intel
    - Hybrid
    - Only for Load Store Unit
    - Goal: diagnose design faults

None of the previous works has studied the **fine grained** diagnosis of **intermittent faults.**

Summary

- Introduced a Hybrid Hardware-Software technique for intermittent hardware fault diagnosis
  - SCRIBE: Provide resource usage visibility to SW layer
    - Performance Overhead : 14.7%
    - Power Overhead : 9%
  - SIED: Use the information provided by SCRIBE for diagnosis
    - Accuracy: 84%
- Diagnosis with such a fine granularity enables chip repair using deconfiguration with less than 2% slowdown.
- First framework to decouple
  - diagnosis information and
  - diagnosis algorithms
- Building block for other diagnosis algorithms

Oracle Mode

- Only correctness of final data is known.
- Option: Save output of every instruction along with its RUI.
  - Correctness of all data in the DDG will be known.

2 to 3 times the overhead of sending only RUI

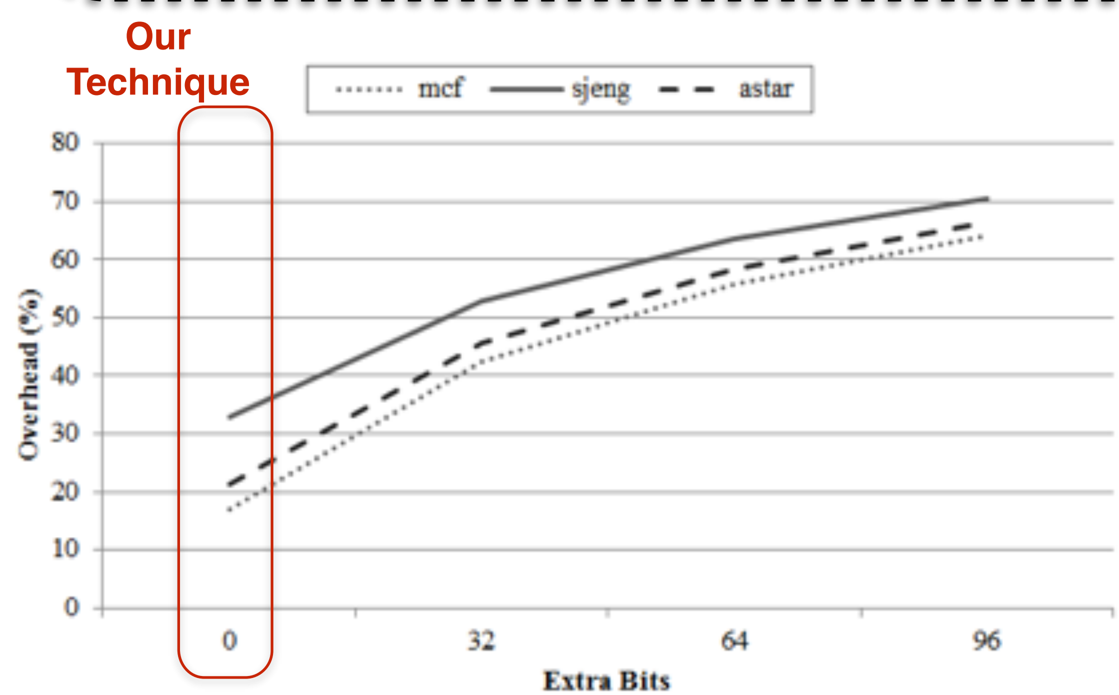

- I1 : r3 <- r1 + r2
- I2 : r4 <- r2 + $2
- I3 : r5 <- r3 * mem[X]

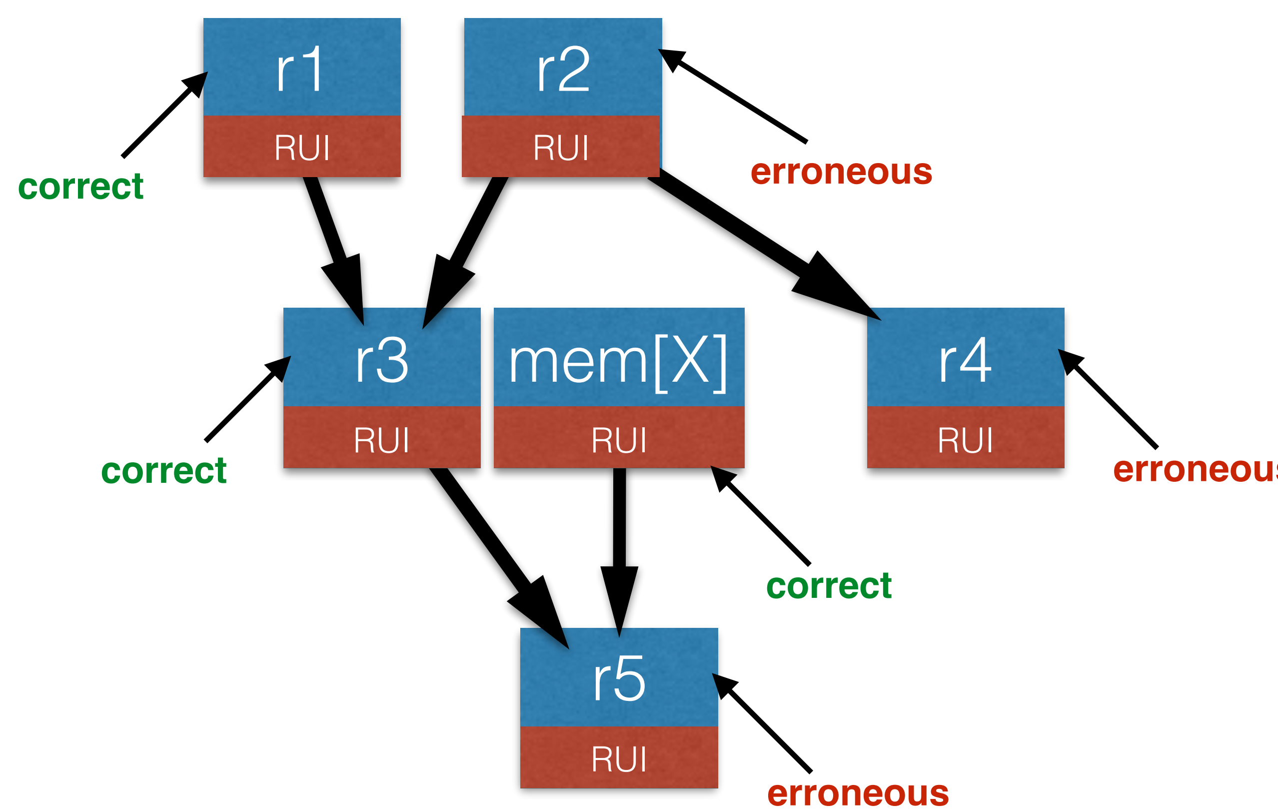

References

- [Nightingale et al. 2011] E. B. Nightingale, J. R. Douceur, and V. Orgovan, “Cycles, cells and platters: An empirical analysis of hardware failures on a million consumer PCs,” ser. EuroSys, 2011, pp. 343–356.
- [Bower et al. 2005] Fred A. Bower, Daniel J. Sorin, and Sule Ozev. 2005. A Mechanism for Online Diagnosis of Hard Faults in Microprocessors *(MICRO)*. 197–208.
- [Li et al. 2008] Man-Lap Li, P. Ramachandran, S.K. Sahoo, S.V. Adve, V.S. Adve, and Yuanyuan Zhou. 2008. Trace-based microarchitecture- level diagnosis of permanent hardware faults (DSN). 22–31.
- [Carretero et al. 2011] J. Carretero, X. Vera, J. Abella, T. Ramirez, M. Monchiero, and A. Gonzalez. 2011. Hardware/software-based diagnosis of load-store queues using expandable activity logs *(HPCA)*. 321–331.
- [Constantinescu 2008] C. Constantinescu, “Intermittent faults and effects on reliability of integrated circuits,” ser. RAMS, 2008, pp. 370–374.
- [Rashid et al. 2012] L. Rashid, K. Pattabiraman, and S. Gopalakrishnan, “Intermittent hardware errors recovery: Modeling and evaluation,” ser. QEST, 2012, pp. 220–229.
- [Gupta et al. 2011] Gupta, Shantanu, et al. "Stagenet: A reconfigurable fabric for constructing dependable cmps." Computers, IEEE Transactions on 60.1 (2011): 5-19.

Configurations

Topic	Parameter	Machine Width		
		Nar.	*Med.*	*Wide*
Pipeline Width	Fetch	2	4	8
	Decode	2	4	8
	Issue	2	4	8
	Commit	2	4	8
Array Sizes	ROB Size	64	128	256
	LSQ Size	32	32	32
Number of Functional Units	Integer Adder	2	4	8
	Integer Multiplier	1	1	1
	FP Adder	1	1	2
	FP Multiplier	1	1	1