

# Characterizing Application Memory Error Vulnerability to Optimize Datacenter Cost via Heterogeneous-Reliability Memory

Yixin Luo, Sriram Govindan,\* Bikash Sharma,\* Mark Santaniello,\* Justin Meza, Aman Kansal,\* Jie Liu,\* Badriddine Khessib,\* Kushagra Vaid,\* Onur Mutlu  
Carnegie Mellon University,\* Microsoft Corporation

## PROBLEM

- Server memory device cost is a key component of datacenter total cost of ownership
- Traditional hardware-based techniques to reduce memory errors further increase memory cost

Technique	Added Capacity	Added Logic
Parity	1.56%	Low
SEC-DED	12.5%	Low
DEC-TED	23.4%	Low
Chipkill	12.5%	High
RAIM	40.6%	High
Mirroring	125%	Low

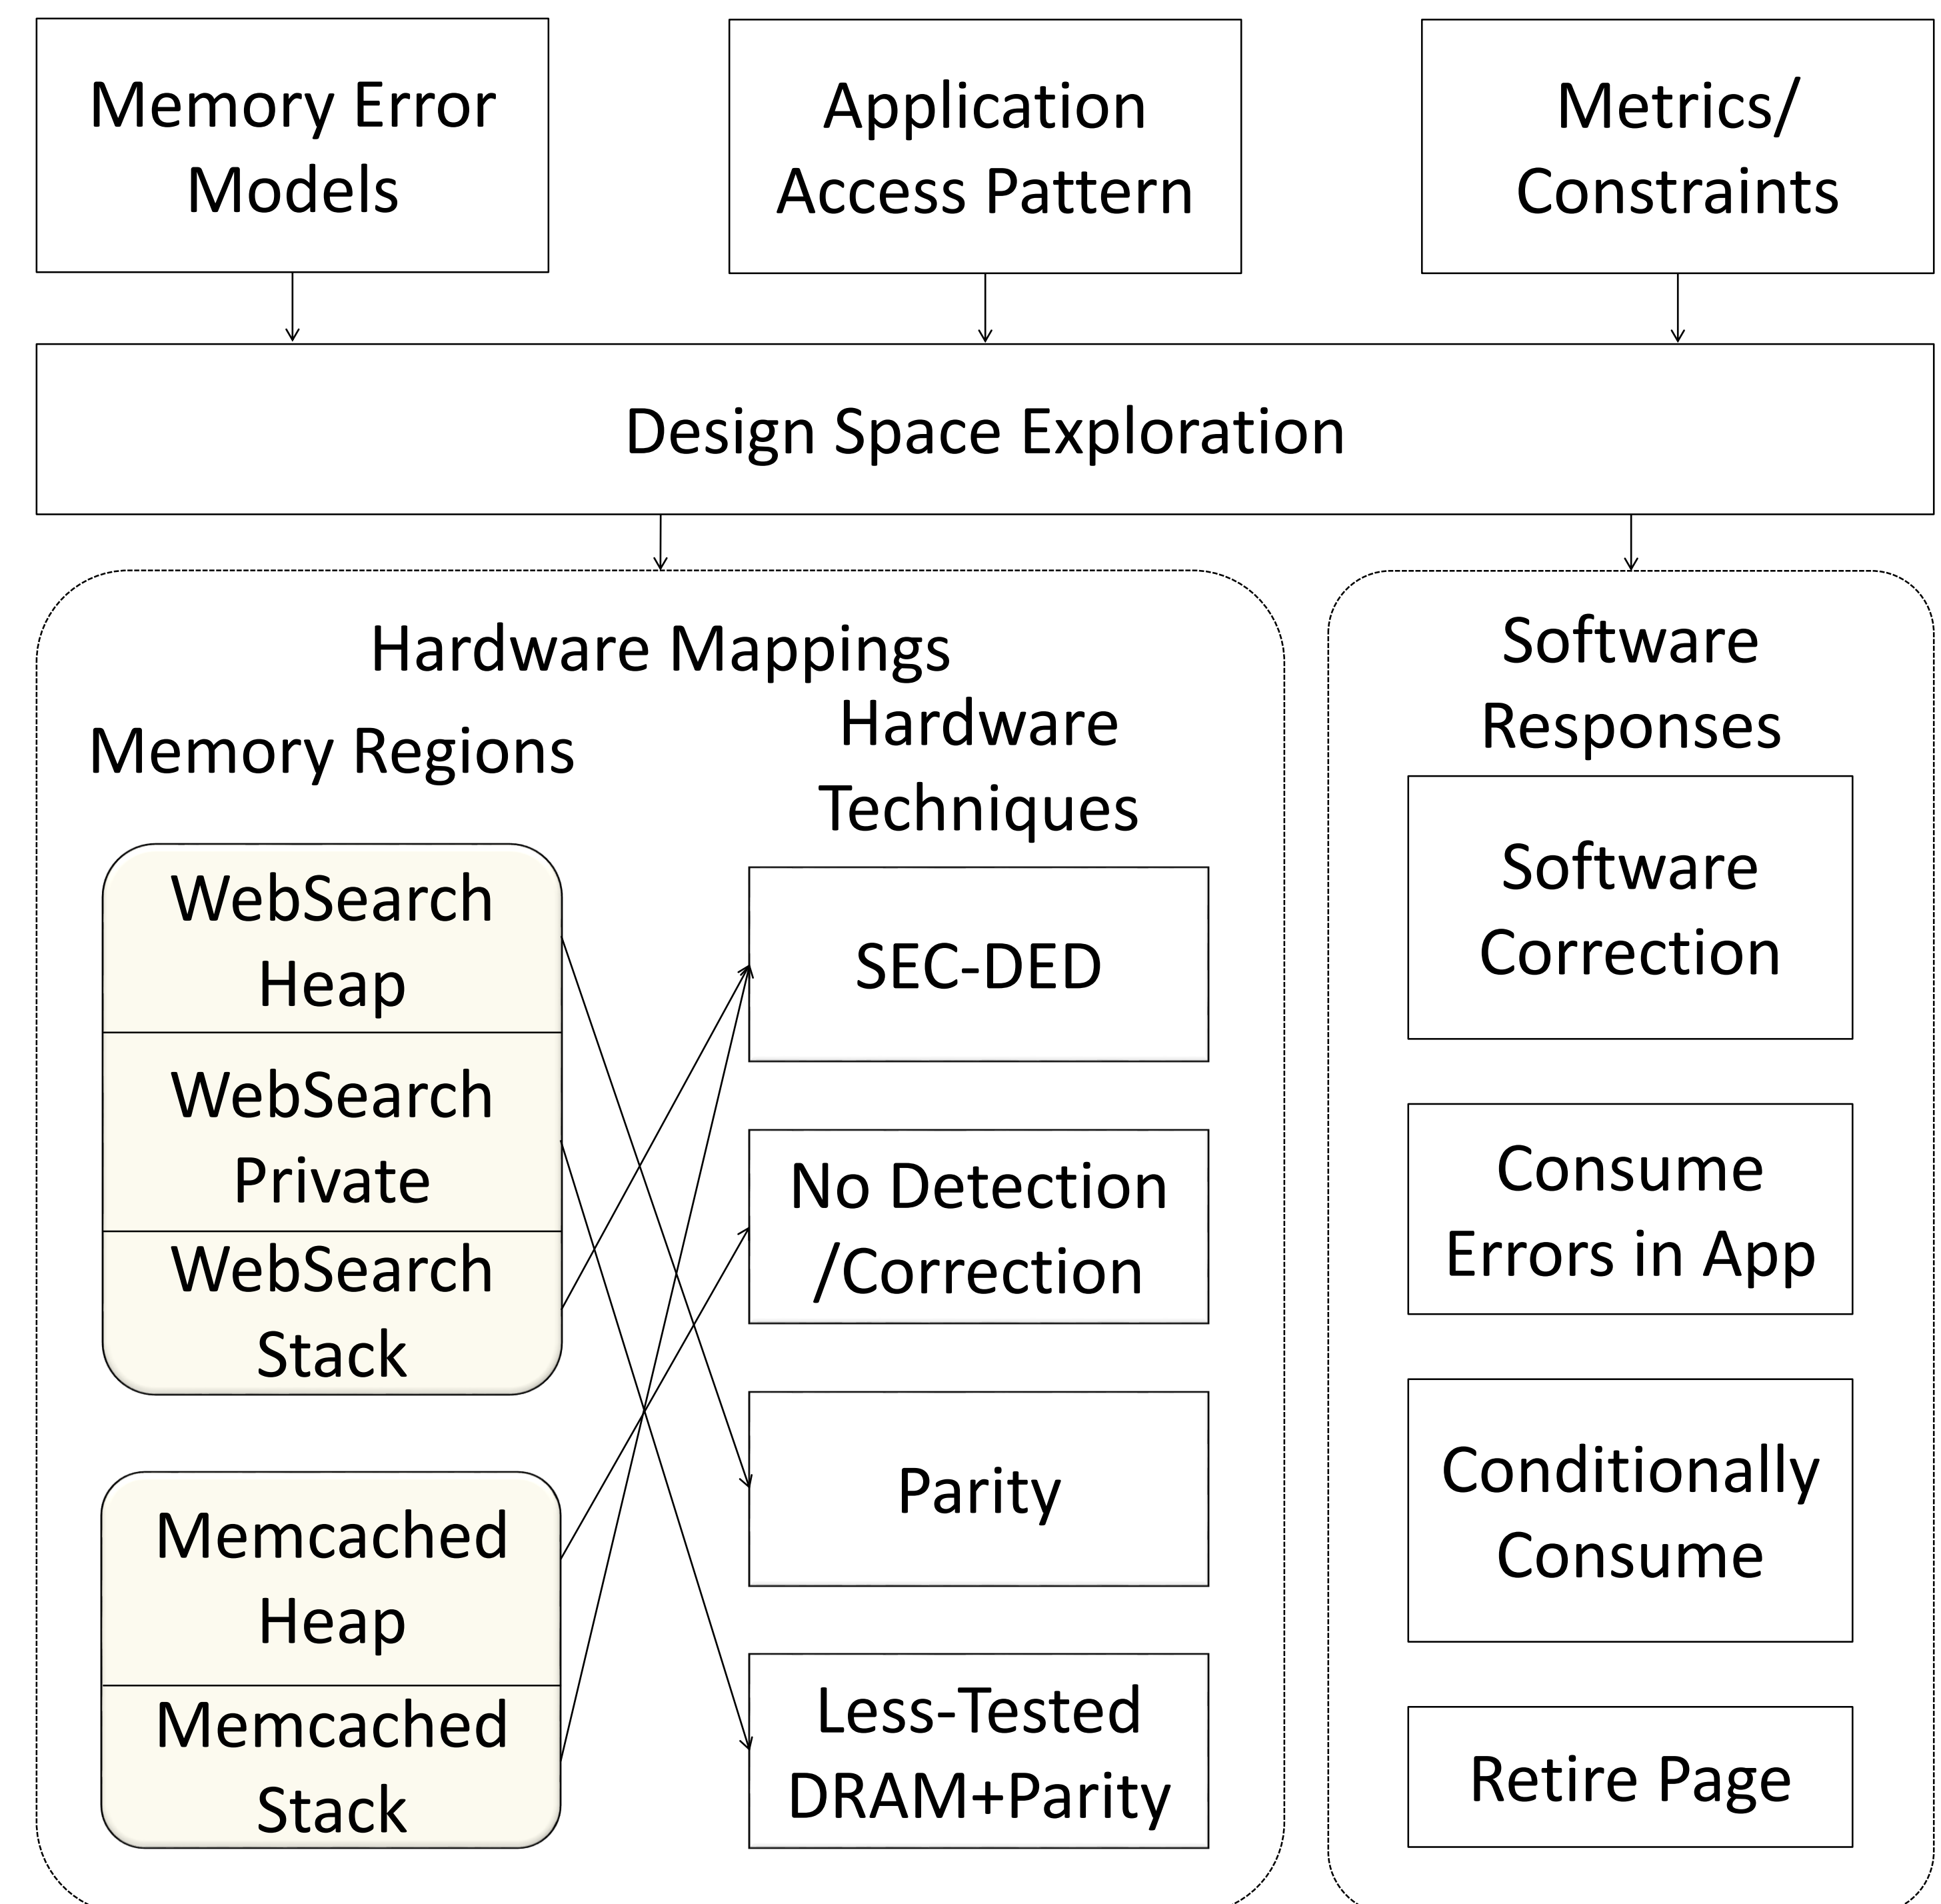
## KEY OBSERVATIONS & FINDINGS

There exists a **diverse spectrum** of tolerance to memory errors in new data-intensive applications, and traditional **one-size-fits-all** memory reliability techniques are **inefficient** in terms of cost

1. Error tolerance varies across applications
2. Error tolerance varies within an application
3. Quick-to-crash vs. periodically incorrect behavior
4. Some memory regions are safer than others
5. More severe errors mainly decreases correctness
6. Data recoverability varies across memory regions

## OUR PROPOSAL

We propose a **hardware/software cooperative heterogeneous-reliability memory system design** that provisions the right amount of memory reliability for different applications




## CASE STUDY RESULTS FOR WEBSEARCH

Configuration	Mapping			Metrics				
	Private (36 GB)	Heap (9 GB)	Stack (60 MB)	Memory cost savings (%)	Server HW cost savings (%)	Crashes/ server/ month	Single server availability	# incorrect/ million queries
Typical Server	ECC	ECC	ECC	0.0	0.0	0	100.00%	0
Consumer PC	NoECC	NoECC	NoECC	11.1	3.3	19	99.55%	33
Detect&Recover	Par+R	NoECC	NoECC	9.7	2.9	3	99.93%	9
Less-Tested (L)	NoECC	NoECC	NoECC	27.1 (16.4-37.8)	8.1 (4.9-11.3)	96	97.78%	163
Detect&Recover/L	ECC	Par+R	NoECC	15.5 (3.1-27.9)	4.7 (0.9-8.4)	4	99.90%	12

ECC = SEC-DED memory; NoECC = no detection/correction; Par+R = parity memory and recovery from disk;

Typical one-size-fits-all baselines : Proposed heterogeneous reliability memory systems



PRINCETON UNIVERSITY

UC Berkeley.

UNIVERSITY of WASHINGTON