ADAPTING TO INTERMITTENT FAULTS IN MULTICORE SYSTEMS

Philip Wells, Koushik Chakraborty, and Guri Sohi {pwells, kchak, sohi}@cs.wisc.edu

THE UNIVERSITY of WISCONSIN MADISON

Introduction

Intermittent hardware faults

Circuit timing errors from:

- Manufacturing variation, in-progress wear-out
- Voltage, temp fluctuations, etc.

Occur in bursts (few cycles to a few seconds or more)

Adapting to intermittent faults

➢ Suspend use of faulting core

How do you suspend use of a core?

Eval of 3 existing techniques yields several drawbacks

➢ Propose a new solution: Use overcommitted system

Assumptions

1) Intermittent faults will occur "frequently"

- Process Variation ↑
- Wear-out ↑
- Noise Margins ↓
  - ➢ More susceptibility to intermittent faults [1,2,3]

2) *Practical* circuits cannot mask all intermittent faults (e.g. Razor [4])

- Ways faults can manifest ↑
- Complexity & OH of circuit reliability mechs. ↑

3) Suspending the use of a core helps

- Factors causing faults ↓
- Faults to correct with other mechanisms ↓
- Useful for other purposes (e.g., aid in diagnosis of permanent damage)

Methodology

Microarchitecture:
- Core: 128-entry, 4-wide, OOO @ 3GHz; 16k L1I&D, 512k L2
- Chip: 8-cores, shared L3, point-to-point network, MOSI directory

Circuit Reliability:
- Assume circuit mechs. to detect & recover from faults (e.g. [4])
- Checkpoint-based recovery: 10k cycle penalty
- Results hold for different mechs.

Simulation:
- Virtutech Simics
- SPARC V9 w/ Solaris 9
- Commercial and other workloads

OS Reconfiguration:
- Send interrupt to Solaris
- Hijack call to "unconfigure" CPU

Processor Virtualization:
- Firmware w/ HW support
- Model VCPU save, restore, migrate
- Spin Detection Buffer [6]

Existing Techniques

Pause Execution

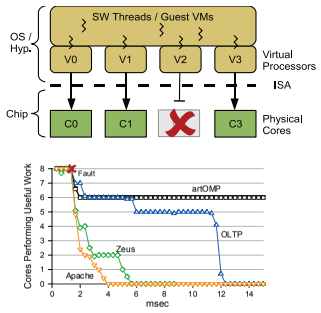

- May cause cascading livelock
- Significant performance concerns

Spare Cores

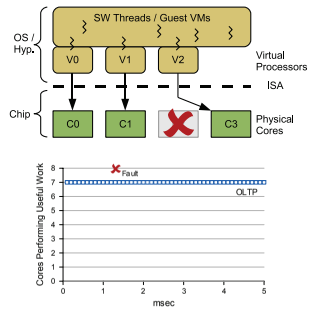

- High fault-free cost
- More concurrent faults than spares?

System SW Reconfig

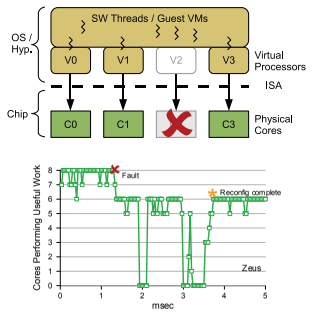

- High reconfig latency & overhead
- All cores must work during reconfig

Proposal: Overcommitted System

System SW uses more *virtual procs (VCPUs)* than available cores [6]

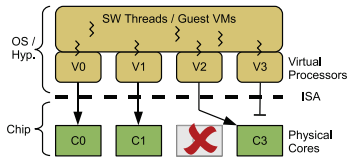

Thin hardware/firmware layer
- Manages VCPUs
- Rotates sharers for fairness
- Transparent to ISA

Modest complexity
- Maintain VCPU-core map
- Manage VCPU state
- Spin detection hardware

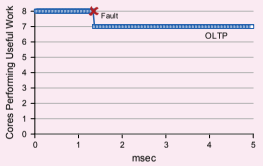

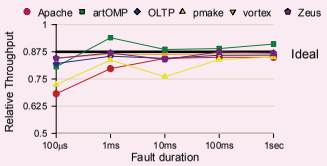

Fairness (µ-bench)
- Fair Speedup: 94% of base

Transaction Latency (µ-bench)
- Avg: 12% worse than base
- No additional outliers

Throughput during a fault

Small OH for short faults from:
- Fault recovery
- Migration of VCPU state

Little OH for faults ≥1ms

Summary

	Throughput	Latency	Fairness	Fault-free	Concurrent	Complexity	Timescales
Overcommitted	✓	✓	✓	✓	✓	M	✓
Pause Execution	✗	✗	✗	✓	✓	L	✗
Spare Cores	✓	✓	✓	✗	✗	M	✓
SW Reconfig	✗	✗	✗	✓	✗	H	✗

(Details in Tech Report [5])

1) Intermittent faults are important considerations

2) Existing techniques have negative consequences

3) Suspending core while overcommitted:
- Has high performance across timescales
- Handles concurrent failures
- Involves only modest complexity

References

[1] S. Borkar, T. Karnik, and V. De. Design and reliability challenges in nanometer technologies. In *Proc. of 41th DAC*, 2004.
[2] C. Constantinescu. Trends and challenges in VLSI circuit reliability. *IEEE Micro*, 23(4):14–19, 2003.
[3] C. Constantinescu. Intermittent faults in VLSI circuits. In *Proc. of Workshop on SELSE*, 2007.
[4] D. Ernst et al. Razor: A low-power pipeline based on circuit-level timing speculation. In *Proc. of 36th MICRO*, 2003.
[5] P. Wells, K. Chakraborty, and G. Sohi. Adapting to intermittent faults in multicore systems. University of Wisconsin-Madison Technical Report CS-TR-1606-2007, Aug. 2007
[6] P. Wells, K. Chakraborty, and G. Sohi. Hardware support for spin management in overcommitted virtual machines. In *Proc. of 15th PACT*, 2006.

Acknowledgements

This work is supported in part by National Science Foundation (NSF) grants CCF-0702313 and CNS-0551401, funds from the John P. Morgridge Chair in Computer Sciences and the University of Wisconsin Graduate School. Sohi has a significant financial interest in Sun Microsystems. The views expressed herein are not necessarily those of the NSF, Sun Microsystems or the University of Wisconsin. Jichaun Chang and Matthew Allen contributed to many discussions.