

Hardware-driven Undo+Redo Logging

Jishen Zhao

https://users.soe.ucsc.edu/~jzhao/

Computer Engineering, UC Santa Cruz

October 11, 2017





Research Overview

• Main research topics

Processing Persistent Memory In Memory Memory Network

• Research efforts

- Workload characterization
- Software/hardware cooperative mechanisms
- Software/hardware interface design
- Hardware-based accelerators

Memory and Storage Systems



Bridging software and hardware design

Typical memory and storage hierarchy:



here! Persistent memory is coming!

Not flash memory... Hardware – Nonvolatile random-access memories (NVRAMs)

Battery-backed DRAM





DDR3 Compatible MRAM



DRAM w/ Ultra-capacitor

No

NO

NO

NO

NO

Yes



Software – Persistent-memory-aware system software



here! Persistent memory is coming!

...but unlocking its full potential isn't easy



Dry it with "Logs"

And copy-on-write, checkpointing, eťc.



- Persistence
 - Used to be a property of storage systems
 - Now needs to be maintained in the memory system



Opportunity Hardware-driven undo + redo Logging



Persistence requirement in cache-memory hierarchy

Update persistent memory with a transaction





Preview of performance benefit with undo+redo logging



Benefits of undo + redo logging



Benefits of undo + redo logging



Preview of performance benefit with undo+redo logging



Inefficiency of logging by software in persistent memory

Conservative

Extra instructions

Increased



Risks with

Performance cost of increased memory traffic



What can we leverage from the hardware? Cache hierarchy maintains undo + redo information by nature



Our undo+redo logging: Take a ride given by CPU caching

Maintain the order between log and data updates by nature



How about cache flushes?

Decouple cache flushes and transaction execution



Commit the transaction

Design principles

- Undo+redo logging taking a ride given by CPU caching
- Cache flushes decoupled from transaction execution

Transaction T _A
Tx_begin
do some reads
do some computation
write A
clwb
Tx_commit

Software and hardware implementation cost

• Software support





• Hardware overhead

Major Components	Logic Type	Size
Transaction ID register	Flip-flop	1 byte per HW thread
Log head and tail registers	Flip-flop	16 bytes
Fwb cache tag bit	SRAM	1 bit per cache line

Key performance results



Processor configuration: Core i7, 22nm, 4-core, 2.5GHz, 2 threads/core

Other results: energy consumption, instruction increase, IPC, sensitivity studies, etc.



• Key points

- Rethink the way traditional software schemes are used
- Exploit opportunities in existing hardware can support data persistence by nature



Hardware-driven Undo+Redo Logging

Jishen Zhao

https://users.soe.ucsc.edu/~jzhao/

Computer Engineering, UC Santa Cruz

October 11, 2017



