PUNO: Predictive Unicast and Notification towards Efficient Transactional Memory Execution

• Lihang Zhao, Lizhong Chen and Jeff Draper • Ming Hsieh Department of Electrical Engineering, University of Southern California • {lihangzh, lizhongc}@usc.edu, draper@isi.edu

HTM for Energy-Efficient HPC

Source: http://nerd-alert.net/

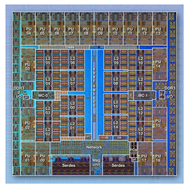

Source: R. A. Haring, et al., The IBM Blue Gene/Q Compute Chip, IEEE Micro, V.32 I.2, 2012

- Hardware Transactional Memory (HTM) is supported in four of the Top 10 supercomputers as of June, 2013.
- Main merits of HTMs: improved programmability.
- Transaction (tx): a code block that is executed atomically and in isolation with regard to other code blocks.

```
1  lock (timestamp) ;
2  lock (counter) ;
3  *t = timestamp ;
4  *r = counter++ ;
5  unlock (timestamp) ;
6  unlock (counter) ;
```

```
1  TX_BEGIN {
2    *t = timestamp ;
3    *r = counter++ ;
4  } ;
```

- Previous research on HTM largely overlooks the energy efficiency of HTM.
- Marching into exascale computing requires energy-efficient HTM designs.

Energy and Performance Pitfalls in HTM

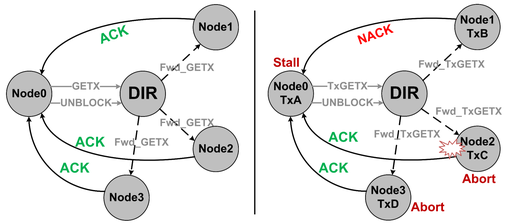

- Typical HTM designs piggyback onto the coherence protocol to detect data access conflicts among transactions.
- Eager conflict detection (CD) scheme can detect conflicts "in place" to reduce wasted transaction computation.
- False aborting: due to mismatch between the protocol and eager CD.
- False aborting incurs energy waste due to unnecessary transaction aborts and excessive on-chip communication.

41% of TX write requests cause false aborting!

Percentages of TX GETX requests causing False Abort

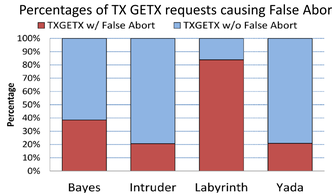

Number of TXs being falsely aborted

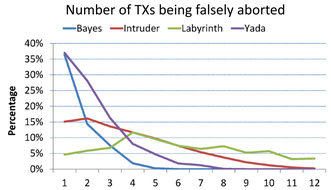

PUNO: Predictive Unicast and Notification

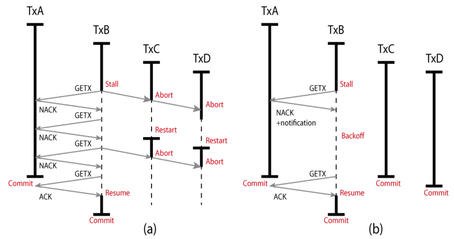

- The goal of PUNO: mitigate false aborting to improve energy efficiency.
- Two essential observations:
  - A multi-readers-single-writer conflict can be handled by one reader transaction whose priority is higher than the writer.
  - The writer transaction must be stalled until the readers with higher priority have finished executing.
- The basic idea of PUNO:

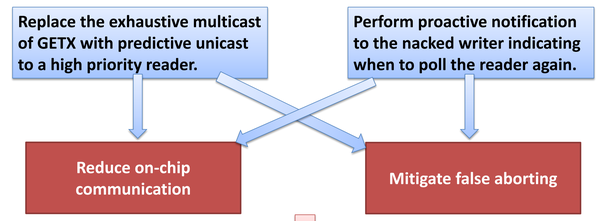

Implementation & Operation

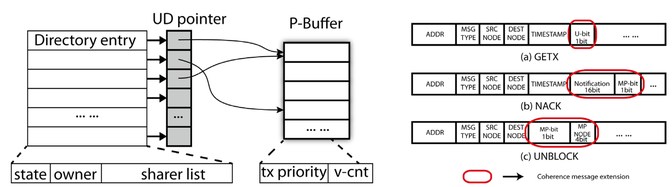

- Each coherence controller is augmented to track the priority of the active transaction on individual nodes.
- Directory entry has an extra pointer to record the sharer transaction that has the highest priority among all the sharers.
- Low hardware overhead.

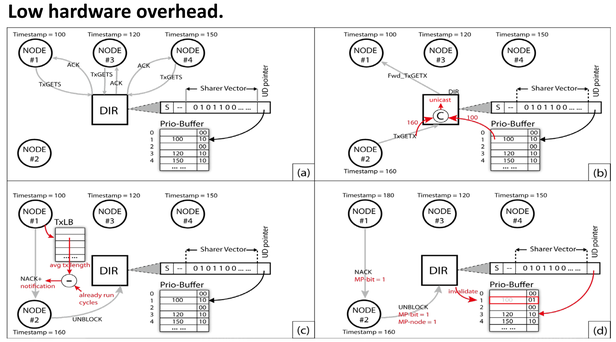

Experimental Setup and Results

- **Experimental setup**
  - Full system simulation using SIMICS+GEMS+Garnet.
  - 16-core tiled CMP with MESI cache coherence protocol.
  - Log-based HTM mode: eager VM and eager CD.
  - Packet switched 2D mesh with virtual channel router.

Unit	Value
Processor	16 Sun UltraSPARC III+ cores
L1 Cache	32 KB, 4-way associative, write-back, 1-cycle latency
L2 Cache	8 MB, 8-way associative, 20-cycle latency
Coherence	MESI protocol, static cache bank directory
Memory	4 GB, 4 memory controller, 200-cycle latency
Network	4X4 2D mesh, dimension-order routing, 4-stage router, 128-bit flit
PUNO	16-entry Priority Buffer; 32-entry Transaction Length Buffer.

- **Experimental results**

Normalized Transaction Abort Count

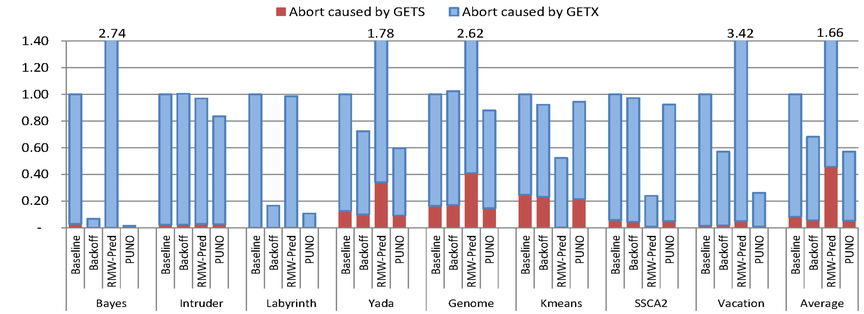

Transaction abort is reduced by 43% across a wide spectrum of TM workloads.

Impact on Energy

- Eliminate 33% of the network traffic in high contention workloads
- Avoid 43% of the transaction aborts -> less discarded computation.

Impact on Performance

- Eliminate 18% of the directory blocking cycles, more concurrency.
- Boost performance by 12% (up to 31%) in high contention workloads.

Normalized On-Chip Network Traffic

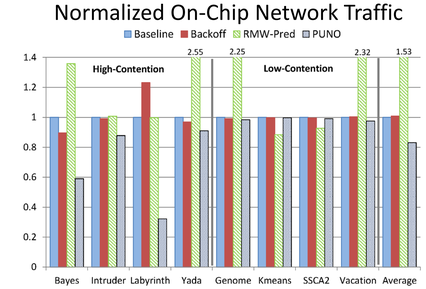

Network traffic is reduced by 18% on average.

Normalized Directory Blocking Cycles

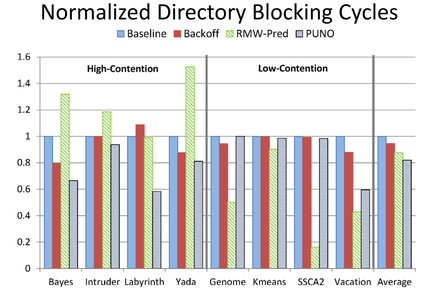

DIR blocking is reduced by 18% on average.

References

[1] P. S. Magnusson, M. Christensson, J. Eskilson, D. Forsgren, G. Hallberg, J. Hogberg, F. Larsson, A. Moestedt, and B. Werner, "SIMICS: A Full System Simulation Platform," *Computer*, vol.35, 2002

[2] M. M. K. Martin, D. J. Sorin, B. M. Beckmann, M. R. Marty, M. Xu, A. R. Alameldeen, K. E. Moore, M. D. Hill, and D. A. Wood, "Multifacet's General Execution-driven Multiprocessor Simulator (GEMS) Toolset," *SIGARCH Compute Architecture News*, vol. 33, November 2005.

[3] C. C. Minh, J. Chung, C. Kozyrakis, and K. Olukotun, "STAMP: Stanford Transactional Applications for Multi-Processing," in *Proceedings of International Symposium on Workload Characterization*, 2008.

[4] K. E. Moore, J. Bobba, M. J. Moravan, M. D. Hill, and D. A. Wood, "LogTM: Log-based Transactional Memory," in *Proceedings of the 12th International Symposium on High Performance Computer Architecture (HPCA)*, 2006

[5] M. M. K. Martin, P. J. Harper, D. J. Sorin, M. D. Hill, and D. A. Wood, "Using Destination-set Prediction to Improve the Latency/bandwidth Tradeoff in Shared-memory Multiprocessors," in Proceedings *of ISCA*, 2003