PUNO: Predictive Unicast and Notification for Energy Efficient HTM

Lihang Zhao and Jeffrey Draper, Information Sciences Institute, USC

HTM for Energy Efficient HPC

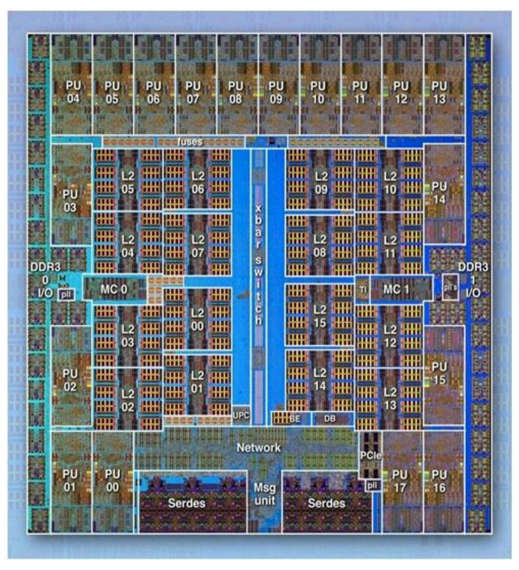

```
lock (timestamp) ;
lock (counter) ;
*t = timestamp ;
*r = counter++ ;
unlock (timestamp) ;
unlock (counter) ;
```

```
TX_BEGIN {
*t = timestamp ;
*r = counter++ ;
}
```

- Hardware Transactional Memory (HTM) is supported in four of the Top 10 supercomputers as of June, 2013.
- Main merits of HTMs: improve programmability of many-core processors.
- Transaction (TX): a code block that is executed atomically and in isolation with regard to other code blocks.
- Typical HTM designs piggyback onto the coherence protocol to detect data access conflicts among transactions.
- Previous research on HTM largely overlooks the energy efficiency issue.

Exascale computing calls for energy efficient HTM designs.

Energy Pitfall in HTM

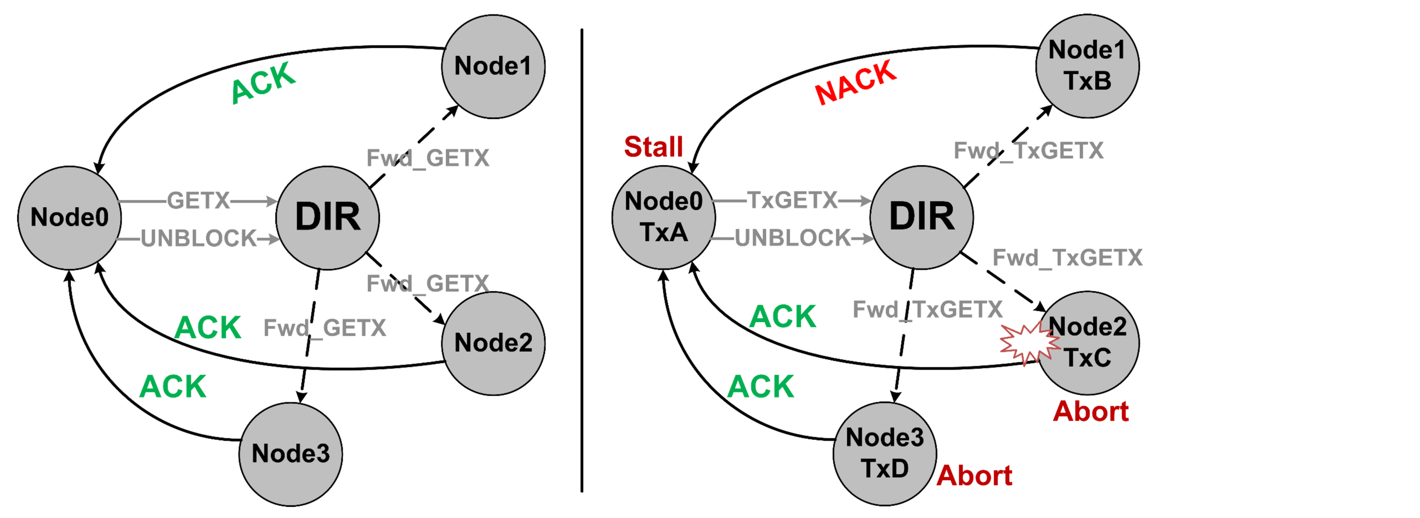

- False aborting: due to mismatch between protocol and eager CD.
- False aborting = energy waste due to excessive on-chip messaging and unnecessary transaction aborts.

41% of TX write requests cause false aborting!

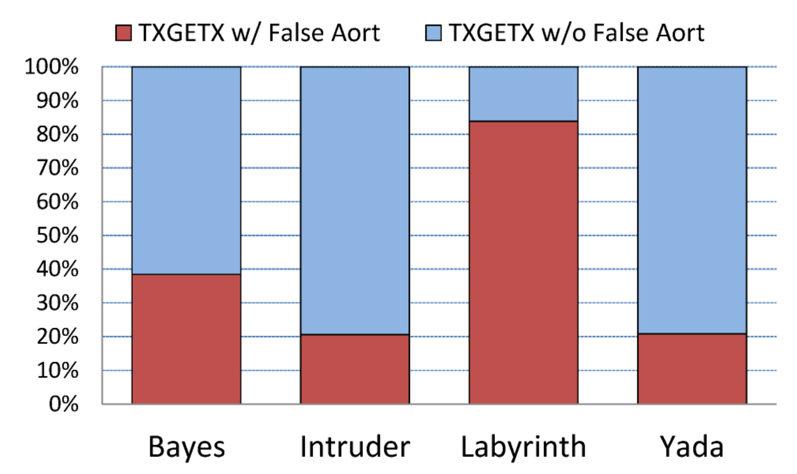

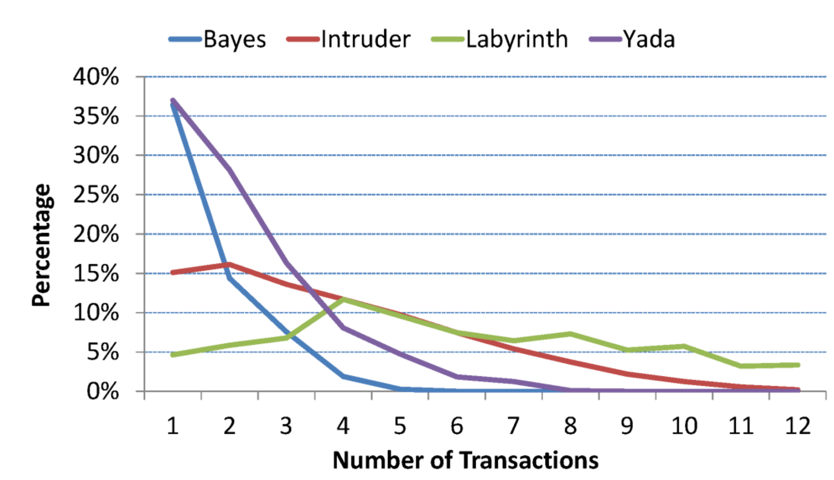

Predictive Unicast and Notification: Mitigating False Aborting

- Two essential observations:
  - The 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:

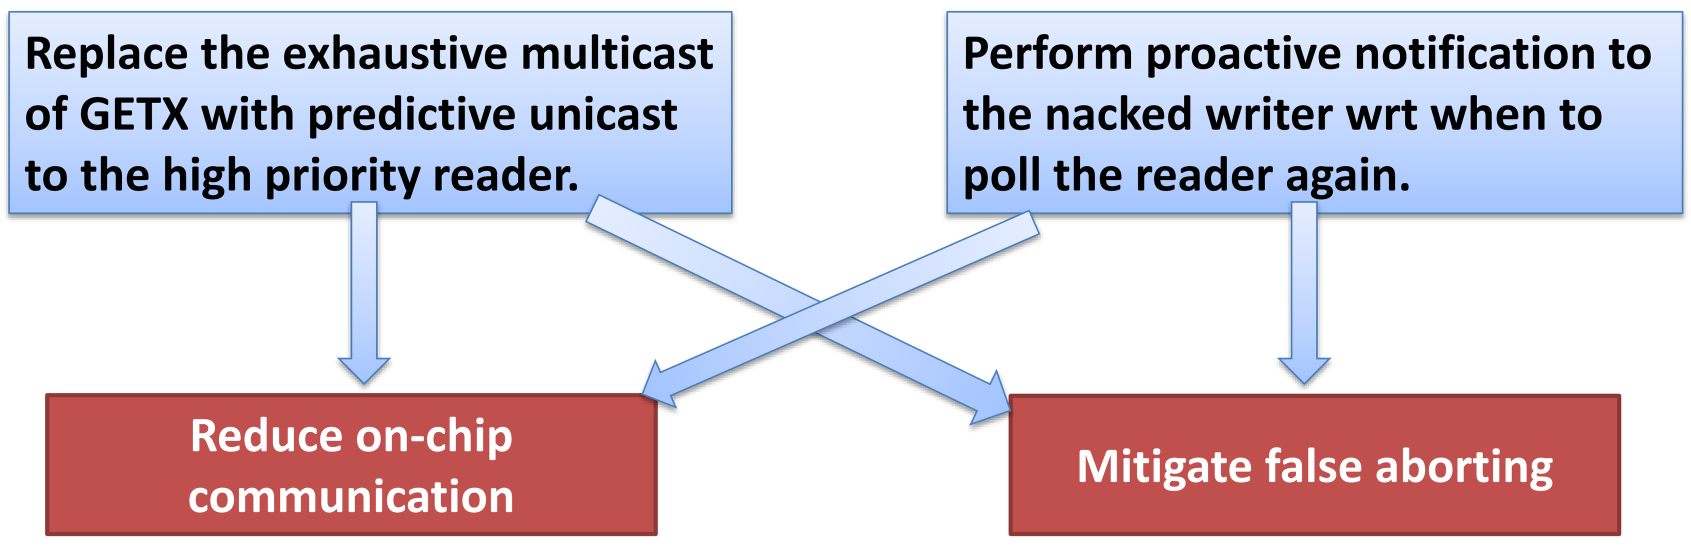

- Illustraion of predictive unicast and notification

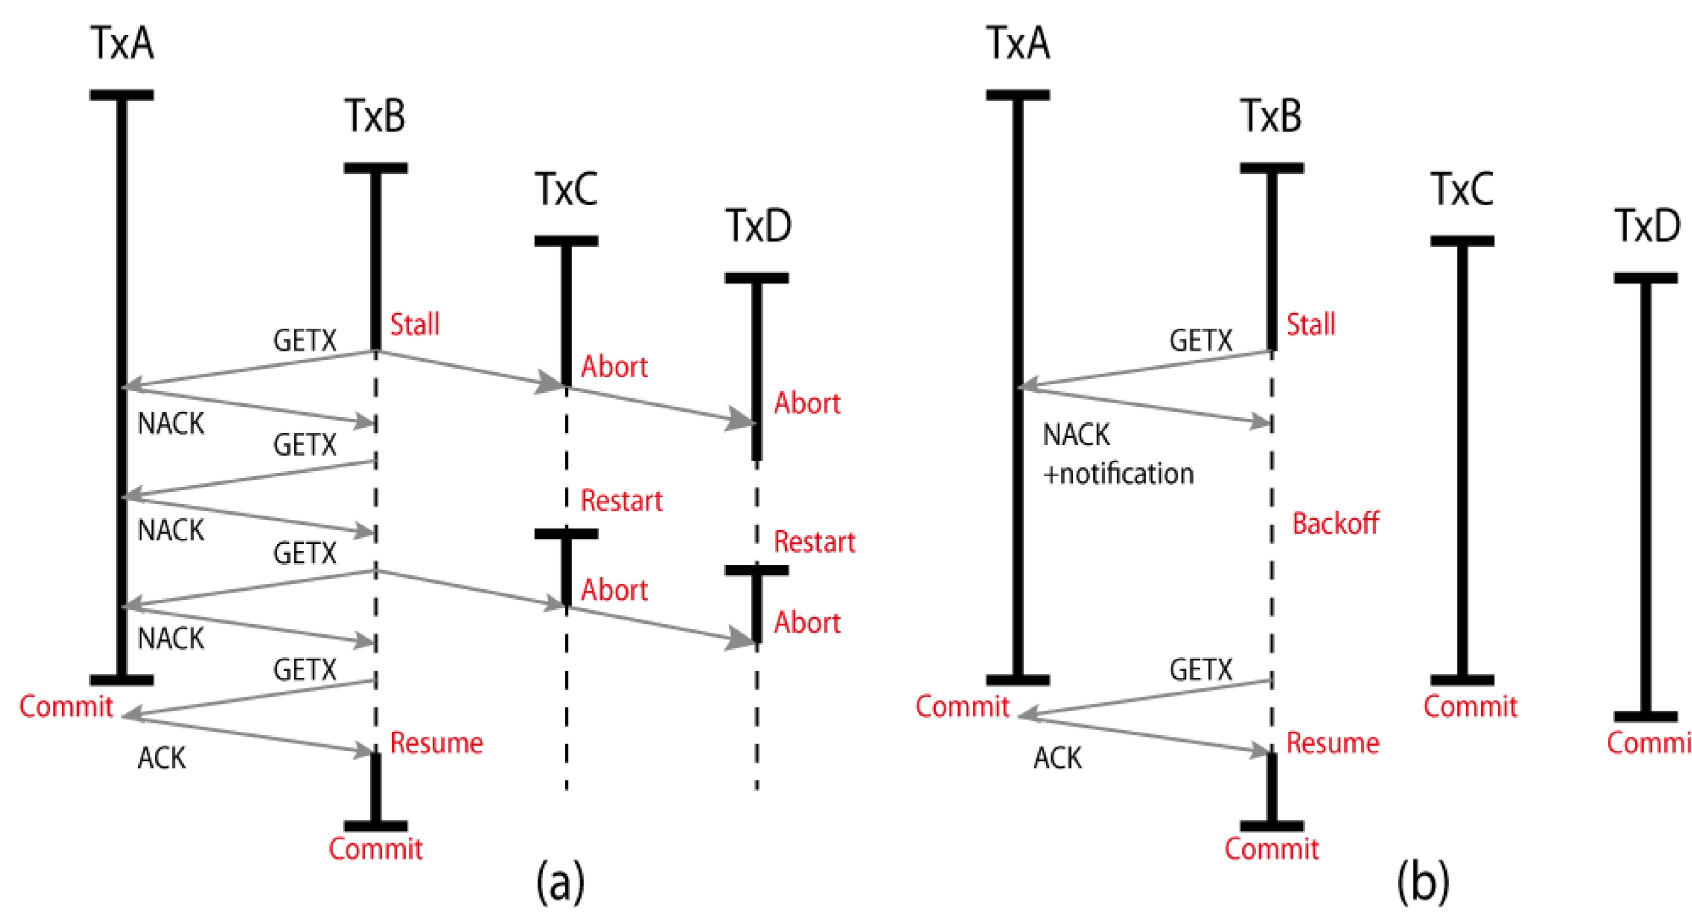

Experimental Setup and Results

- **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.

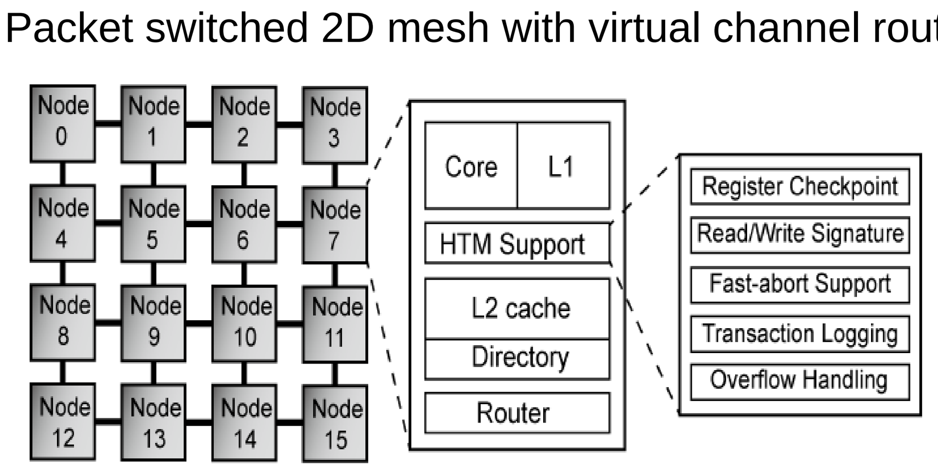

Unit	Value
Processor	16 Sun UltraSPARC III+ cores
L1 Cache	32 KB, 4-way associative, write-back, 1-cycle latency
L2 Cache	8 MB, 8-way associative, 20-cycle latency
Coherence	MESI protocol, static cache bank directory
Memory	4 GB, 4 memory controller, 200-cycle latency
Network	4X4 2D mesh, dimension-order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

- **Experimental Results**

PUNO's Impact on HTM Energy Efficiency

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

Normalized Transaction Abort Count

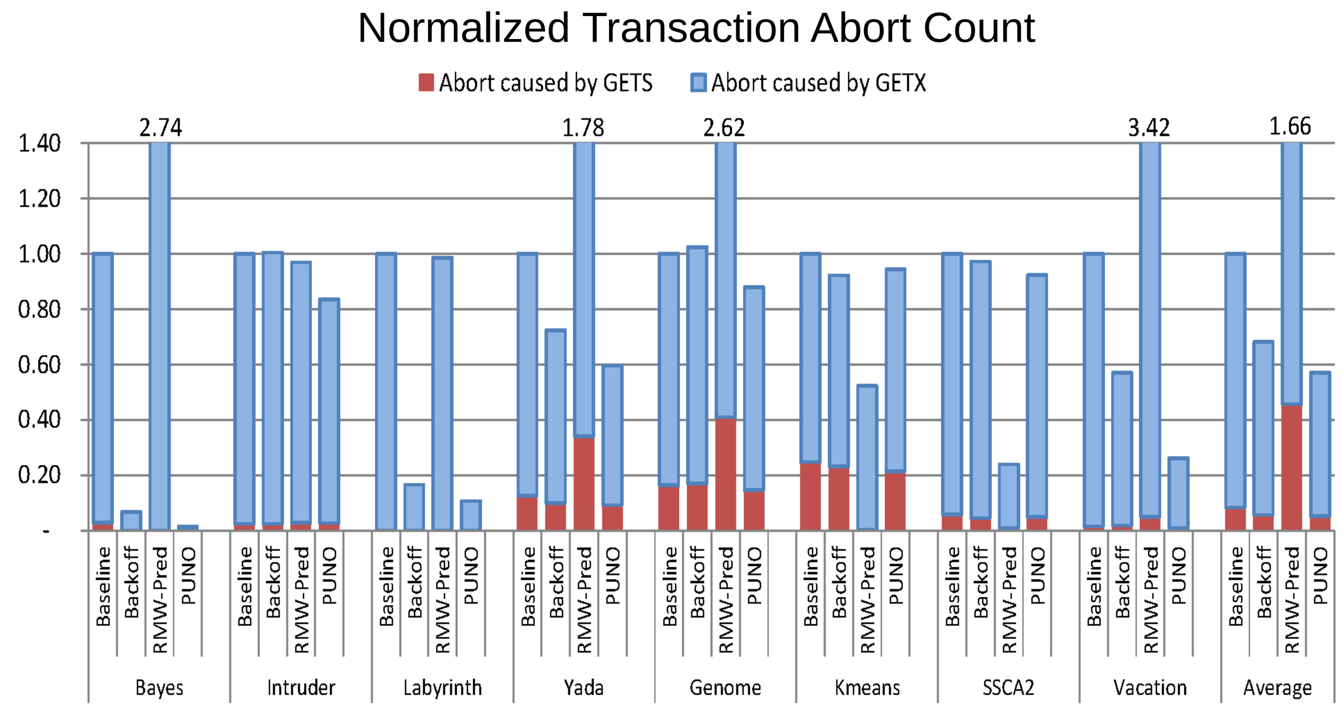

Normalized On-Chip Network Traffic

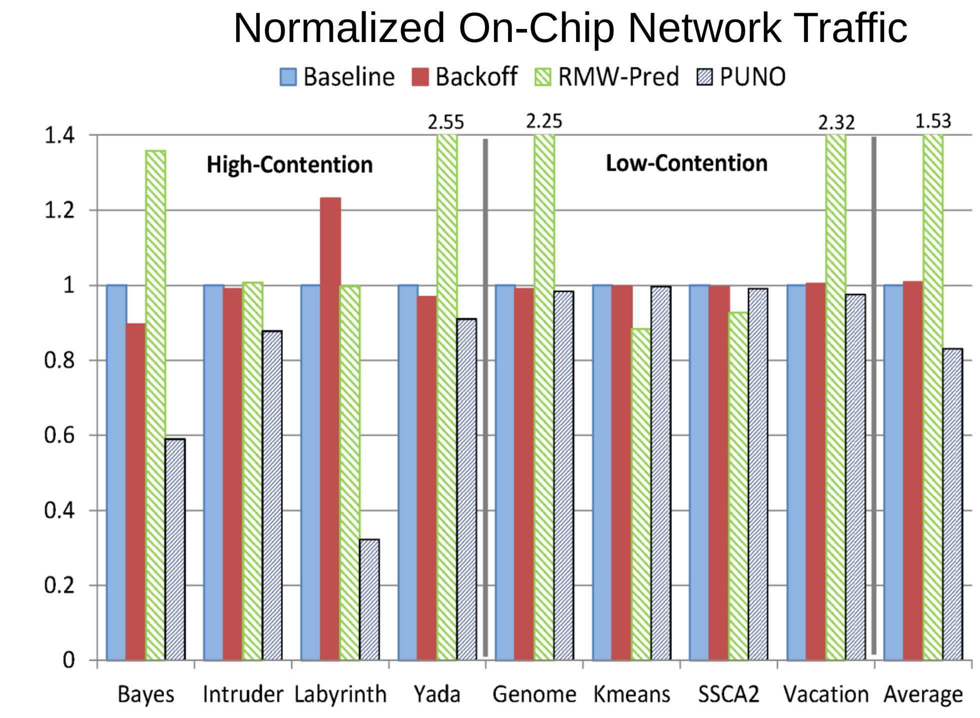

lihangzh@isi.edu