

Scalable, Reliable, Power-Efficient Communication for Hardware Transactional Memory

*Seth H. Pugsley, Manu Awasthi, Niti Madan,
Naveen Muralimanohar, Rajeev
Balasubramonian*

UUCS-08-001

School of Computing
University of Utah
Salt Lake City, UT 84112 USA

January 25, 2008

Abstract

In a hardware transactional memory system with lazy versioning and lazy conflict detection, the process of transaction commit can emerge as a bottleneck. This is especially true for a large-scale distributed memory system where multiple transactions may attempt to commit simultaneously and co-ordination is required before allowing commits to proceed in parallel. In this paper, we propose novel algorithms to implement commit that are more scalable (in terms of delay and energy) and are free of deadlocks/livelocks. We show that these algorithms have similarities with the token cache coherence concept and leverage these similarities to extend the algorithms to handle message loss and starvation scenarios. The proposed algorithms improve upon the state-of-the-art by yielding up to a 7X reduction in commit delay and up to a 48X reduction in network messages. These translate into overall performance improvements of up to 66% (for synthetic workloads with average transaction length of 200 cycles), 35% (for average transaction length of 1000 cycles), 8% (for average transaction length of 4000 cycles), and 41% (for a collection of SPLASH-2 programs).