Unbounded Transactional Memory

C. Scott Ananian, Krste Asanović, Bradley C. Kuszmaul, Charles E. Leiserson, Sean Lie

Computer Science and Artificial Intelligence Laboratory Massachusetts Institute of Technology {cananian,krste,bradley,cel}@mit.edu, sean@slie.ca

Thanks to Marty Deneroff (then at SGI)

This research supported in part by a DARPA HPCS grant with SGI, DARPA/AFRL Contract F33615-00-C-1692, NSF Grants ACI-0324974 and CNS-0305606, NSF Career Grant CCR00093354, and the Singapore-MIT Alliance

Transactional Memory (definition)



- A transaction is a sequence of memory loads and stores that either commits or aborts
- If a transaction commits, all the loads and stores appear to have executed atomically
- . If a transaction aborts, none of its stores take effect
- Transaction operations aren't visible until they commit or abort
- Simplified version of traditional ACID database transactions (no durability, for example)
- For this talk, we assume no I/O within transactions





Infrequent, Small, Mostly-Serial?

To date, xactions assumed to be:

- . Small
 - BBN Pluribus (~1975): 16 clockcycle bus-locked "transaction"
 - Knight; Herlihy & Moss: transactions which fit in cache
- Infrequent



- Software Transactional Memory (Shavit & Touitou; Harris & Fraser; Herlihy et al)
- . Mostly-serial
 - Transactional Coherence & Consistency (Hammond, Wong, et al)

aul/Leiserson/Lie[.] Unboun



- Experiment to discover xaction properties of large real-world app.
 – First complete OS investigated!
- User-Mode Linux 2.4.19
 - instrumented every load and store, all locks
 - locks→xactions; locks not held over I/O!
 - run 2-way SMP (two processes; two processors)
- . Two workloads
 - Parallel make of Linux kernel ('make linux')
 dbench running three clients
- Run program to get a trace; run trace through Transactional Memory simulator
 - 1MB 4-way set-associative 64-byte-line cache
 - Paper also has simulation runs for SpecJVM98







Transactional Programming

- . Locks: the devil we know
- . Complex sync techniques: library-only
 - Nonblocking synchronization
 - Bounded transactions
 - Compilers don't expose memory references
 - (Indirect dispatch, optimizations, constants)
 - Not portable! Changing cache-size breaks apps.
- Unbounded Transactions:
 - Can be thought about at high-level
 - Match programmer's intuition about atomicity
 - Allow black box code to be composed safely
 - Promise future excitement!
 - . Fault-tolerance / exception-handling
 - . Speculation / search



Two new instructions

- . XBEGIN pc
 - Begin a new transaction. Entry point to an *abort handler* specified by pc.
 - If transaction must fail, roll back processor and memory state to what it was when XBEGIN was executed, and jump to pc.
 Think of this as a mispredicted branch.
- XEND
 - End the current transaction. If XEND completes, the xaction is committed and appeared atomic.
- Nested transactions are subsumed into outer transaction.

n/Asanović/Kuszmaul/Leiserson/Lie: Unbounded Transactional Memory, HPCA















<u>Original</u>	<u>Rename Table</u>	Saved set
ADD P2, P1, P1 ST 1000, P2	R1→P1,	{ P1, }
Accode XEND XBEGIN L2 ADD R1, R1, R1	R1→P2,	{ P2, }
ST 2000, R1 XEND		

































Is this good enough? • Problems solved: • Xactions as large as physical memory • Scalable overflow and commit • Easy to implement! • Low overhead • May speed up Linux! • Open Problems... • Is "physical memory" large enough? • What about duration? • Time-slice interrupts!





Caching in UTM



nal Memory HPC

- Most log entries don't need to be created
- Transaction state stored in cache/overflow DRAM and monitored using cachecoherence, as in LTM
- Only create transaction log when thread is descheduled, or run out of physical mem.
- . Can discard all log entries when xaction commits or aborts
 - Commit block left in X state in cache
 Abort use old value in main memory
- In-cache representation need not match xstate representation

/Asanović/Kuszmaul/Leiserson/Lie: Unbounded Tra

Performance/Limits of UTM

. Limits:

- More-complicated implementation
 - . Best way to create xstate from LTM state?
- Performance impact of swapping.
 - . When should we abort rather than swap?

. Benefits:

- Unlimited footprint
- Unlimited duration
- Migration and paging possible
- Performance may be as fast as LTM in the common case



Open questions

- . I/O interface?
- Transaction ordering?
 - Sequential threads provide inherent ordering

n/Lie⁻ Unbounded Trar

- Programming model?
- Conflict resolution strategies