Agenda

Designing Transactional Memory Systems

Part III: Lock-based STMs

Pascal Felber University of Neuchatel Pascal.Felber@unine.ch



Based on joint work with Christof Fetzer & Torvald Riegel with slides borrowed from several other people

- Part I: Introduction
- Part II: Obstruction-free STMs
- Part III: Lock-based STMs
 - WSTM: a lock-based STM design
 - TINYSTM: a time-based lock-based STM design

Transactional Memory: Part III - P. Felber

Why not obstruction-free? [Ennals]

- OF prevents a long-running transaction blocking others
 - Not true: neither OF nor LF guarantee that!
- OF prevents the system locking up if a thread is switched part-way through a transaction
 - But this is rare and temporary event...
- OF prevents the system locking up if a thread fails
 - But failure would break the original, non-STM program as well...

Why lock-based? [Ennals]

- More efficient implementation
 - No extra indirections for data accesses
 - Better cache locality
 - Simpler (and often faster) algorithm
- Typical implementation
 - Revocable two phase locking for writes (blocking)
 - Preventive abort to avoid deadlocks
 - Optimistic concurrency control for reads
 - Lazy conflict detection (if data has changed)



ΙŪΛ

WSTM [Harris & Fraser, 2003]

- Word-based STM
 - Granularity of conflict detection is (roughly) memory location
- Basic principle
 - For every word, there exists a version number
 - Transactions don't update memory or version numbers (i.e., no synchronization) until commit
 - They update (and consult) thread-local log
 - Modified words are only "locked" during commit

Transactional Memory: Part III - P. Felber

• Parallel reads don't cause aborts

WSTM: data structures



WSTM: load and store





3/17/2008

STM design choices

- We have already seen a number of designs
- There are many more design choices:
 - Obstruction-free vs. lock-based (blocking)
 - Object-based vs. word-based
 - Visible vs. invisible reads
 - Encounter-time vs. commit-time locking
 - Write-through vs. write-back
 - Etc.

STM design choices

- The "right" design depends on the workload
 - E.g., ratio of update to read-only transactions, number of locations read or written, contention on shared memory locations, etc.

There is no "one-size-fits-all" STM

Object-based vs. word-based

Transactional Memory: Part III - P. Felber

Object-based

- Granularity of conflict detection is object
- Store metadata in object (or proxy)
- Need language support
- Good with small objects (cloning cost)
- Coarse-grained "lock"

Word-based

- Granularity of conflict detection is memory location
- Separate metadata
- Load/store API
- Good for low-level, unmanaged languages
- Fine-grained "lock"

Visible vs. invisible reads

Transactional Memory: Part III - P. Felber

Visible reads

- Let writers detect conflicts with readers
- Need to maintain shared reader lists
- Pessimistic design
- Abort early may save useless processing
- Better progress upon high contention

Invisible reads

- Writers do not detect read-write conflicts
- Incremental validation costly (use time-based)
- Optimistic design
- Much faster when there is no conflict
- No (costly) reader list
- May fall back to visible





Implementation notes

- Implementing concurrent algorithms efficiently is challenging
 - Need to understand how MP systems and multicore processors work
 - Need to use the cheapest primitives that are sufficient (for safety)
 - Need to understand the memory models (when defined!)

If this scares you, leave it to "experts" and use STM!

Transactional Memory: Part III - P. Felber

This is not wrong!

- We are assuming sequentially consistent behavior
 - But time is a relative dimension!
- Computers don't care about your intuition regarding time across multiple threads
 - Compilers, processors, caches can reorder instructions
- Preventing reordering must be explicitly requested
 - Using synchronization operations (lock/unlock)
 - Using memory barriers

3/17/2008

Memory models



Memory barriers

- A processor can execute hundreds of instructions during a memory access
 - Why delay on every memory write?
 - Instead, keep value in register or cache



- Memory barrier instruction (expensive)
 - Flush unwritten caches
 - Bring caches up to date
 - Added by compiler (synchronization, volatile)

3/17/2008

Transactional Memory: Part III - P. Felber

Based on documentation by Doug Lea

• LD1 L/L LD2

LD1's data loaded before data accessed by LD2 and all subsequent loads are loaded

• ST1 <mark>S/S</mark> ST2

ST1's data visible to other processors before data associated with ST2

• LD1 L/S ST2

LD1's data loaded before data associated with ST2 and all subsequent store instructions are flushed

• ST1 <mark>S/L</mark> LD2

Based on documentation by Doug Lea

ST1's data visible to other processors before data accessed by LD2 and all subsequent loads are loaded

Transactional Memory: Part III - P. Felber

Based on documentation by Doug Lea

Memory barriers example (Java)





// Write value and release lock
*w->addr = w->val;
S/S;
*w->lock = tx->timestamp;
S/L;





iinü

TANGER: transactional C compiler

- Problem: explicit load/store instructions are cumbersome and error-prone
- TANGER: open source transactional compiler
 - Application code with transaction boundaries transformed into STM transactional code
 - Uses LLVM's open-source compiler framework
 - Intermediate representation for a load/store architecture (no stack)
 - Supports aggressive optimizations at compile-, link-, or run-time

Transactional Memory: Part III - P. Felber

Easily extensible by "compiler passes"

What TANGER does

- Application code uses minimal API
 - TX begin/commit
 - Language syntax unchanged
- Transform code to use word-based STM API
 - Create transactional version of functions
 - Within TX begin/commit, redirect:
 - (Heap) memory accesses to STM load/store
 - Calls to transactional versions of functions
 - Dynamic memory management (malloc, free)

Transactional Memory: Part III - P. Felber









3/17/2008

TANGER: performance



Compiler: overhead [Intel]



Compiler: scalability [Intel]



Conclusion (Part III)

Transactional Memory: Part III - P. Felber

- Lock-based STM designs are very efficient
 - No progress guarantee, but sufficient from a pragmatic perspective
- Word-based designs good for OS, low-level languages, and compiler integration
 - Fine conflict detection granularity
 - Can be used to build object-based STMs
- The design space is very large
 - No "one-size-fits-all" STM
 - The right design depends on the workload!

3/17/2008

