

# Language constructs for transactional memory

**Tim Harris** 

Disclaimer: these are my personal opinions



# Untangling "atomic" from TM

# Hiding TM from programmers

# **Current performance**



#### Example: double-ended queue



- Support push/pop on both ends
- Allow concurrency where possible
- Avoid deadlock



#### Implementing this: TM

```
Class Q {
  QElem leftSentinel;
  QElem rightSentinel;
  void pushLeft(int item) {
    QElem e = new QElem(item);
    do {
      StartTx();
      TxWrite(&e.right, TxRead(&this.leftSentinel.right));
      TxWrite(&e.left, this.leftSentinel);
      TxWrite(&TxRead(&this.leftSentinel.right).left, e);
      TxWrite(&this.leftSentinel.right, e);
    } while (!CommitTx());
  }
}
```

Broadly based on word-based STM from "Concurrent programming without locks" Keir Fraser & Tim Harris, ACM TOCS

#### Implementing this: atomic blocks

100

X 30 WD

Microsoft<sup>®</sup>

Research





#### **Design questions**



#### "Atomic blocks are transactions"

consider here the questions of external interaction through storage systems or databases, nor do we address distributed systems

Even in this setting, concurrent programming is extremely difficult. The dominant programming technique is based on locks, an syproch that is simple and direct, but that simply does not scale with program size and complexity. To ensure convertness, programmens must identify which operations conflict; to ensure liveness, they must avoid introducing deadlock; to mause good performance, they must balance the granularity at which looking is performed against the costs of fine-grain looking. Pethags the most fundamental objection, threads, in that look-haved programs do not compour. correct fragments may fail when combined

For courple, consider a hash table with thread-safe insert and fallete operations. Now suppose that we want to delete one item A from table 11, and insert it into table 12, but the informadiate state (in which neither table contains the item) must not be visible to other threads. Unless the irrelementer of the bash table articipates this read. there is simply no way to satisfy this requirement. Even if she does, all she can do is expose methods such as LockTable and Unio chTable - but as well as breaking the hash-table abstraction, they invite lock-induced deaflock, depending on the order in which the client takes the locks, or race conditions if the client Segets. Yet more complexity is required if the client wants to await the presence of A in vi, but this blocking behaviour must not lock the table (else & caused be inserted). In sheet, operations that are individually correct (insert, deletic) cannot be composed into larger consult operations.

The same phenomenon shows up trying to compose alternative blocking operations. Suppose a procedure p4 waits for one of two input pipes to have dats, using an internal call to the Unix sola et procedure; and suppose another procedure p2 does the same thing, on two different pipes. In Unix there is no way to perform a select between p1 and p2, a furdamental loss of compositionality. In-stead, Unix programmers loam awkword programming techniques to gather up all the file descriptors that must be waited for, perform a single top-level solact, and then dispatch back to the connect handler. Again, two individually-correct abstractions, p4 and p2, carnet be composed into a larger one; instead, they must be ripped spart and awkwardly morged, in direct conflict with the geals of distinction.

Rather than focing locks, a more promising and radical alterna-Eve is to base concessency control on atomic memory insurations, she known as transactional memory. We will show that instanctional memory offices a solution to the tension between concurrency and abstraction. For coample, with memory instauctions we can manipulate the hash table thus:

atmic { vr+dalata(ti,A); immert(t2,A,v) }

and to wait for either p4 or p2 we can say

atmic { pi 'arElse' pl }

These simple constructions require no knowledge of the implementation of insert, delete, pl, or p2, and they continue to work censerily if these operations may block, as we shall see.

2.1 Transctional memory

The idea of tratauctions is not now: they have been a fundamental mechanism in database design for many yours, and three has been much recent work on transactional momorias [11, 10, 9, 6, 31].

The key idea is that a block of code, including motiod culls, can be enclosed by an atomic block, with the guarantee that it rans atomically with respect to every other atomic block. Transactional memory can be implemented using optimizativ synchronization. Itstead of taking looks, an accast c block sure without looking, accamulating a flacad-local transaction log that recents every memory read and write it makes. When the block completes, it first validates

its log, to check that it has seen a consistent view of memory, and then conveits its charges to memory. If validation fails, because memory and by the method was altered by another thread during the block's execution, then the block is re-executed from senatch. Transactional memory eliminates, by construction, many of

the low-level difficulties that plague lock-based programming [6]. There are no look-induced deadlooks (because there are no looks): there is no priority invention; and there is no painful teration bebeen granularity and concurrency. However little program has been main on building transactions of the first first compose well. We should the particular problem W-ICV-Findly, sizes a transaction may be re-eff. W-ICV-

essential that it do nothing introocable. For coast

might launch a soco it might also laurch was do-schoduled af thed heldfoot both bogs for a guarante perform memory of modifications to its intercouble input/out Secondly, blocks support synchronisa and three that do re

on the storage blo from a buffer might Ites get() { atomic (a\_ites

The thread waits unit ecuting the Mock. B We cannot call get ( form an intervening in a second atomic l loss the outer block is a desseter for abst the two items) has to mentation, liferent OWNER WORKED.

Thirdly, no perciconstitued by the sal tion 7.2 on Congan by presenting transtive language Concu

2.2 Cencurrent E Concurrent Huskell lary, functional large abstractions for our volves side effects a necessary to be able breakthrough came i

Here is the key in when performed, ma . For courple, the ! patchar :: Char patchar :: 10 Char

> That is, yetChar takes a Char and deliven an 10 action that, when performed, prints the character on the standard cutput, while getChar is an action that, when performed, reads a character from the conside and delivers it as the result of the action. A complete

Research

Microsoft<sup>®</sup>

100

40

20

state ( if tool) the load, similar(); it on lock-induced courses there is no priority inversion; and there is no painful tension between granularity and concurrency. However little progress has been made on building transactional abstractions that compose well. We identify three particular problems.

Firstly, since a transaction may be re-run automatically, it is essential that it do nothing irrevocable. For example the transaction

atomic { if (n>k) then launch\_missiles(); S2 }

might launch a second salvo of missiles if it were re-executed. It might also launch the missiles inadvertently if, say, the thread was de-scheduled after reading n but before reading k, and another thread modified both before the thread

PPoPP105,



#### "Atomic blocks are locks"

100

40

20

**Consequences** explored methodically by Menon et al (Transact '08, SPAA '08)

> icant performance overhead. As an alternative, many in the community have informally proposed that a single global lock semantics [16, 9], where transaction semantics are mapped to those of regions protected by a single global lock, provide an intuitive and efficiently implementable model for programmers.

> In this paper, we explore the implementation and performance implications of single global lock semantics in a weakly atomic STM from the perspective of Java, and we discuss why even recent STM implementations fall short of these semantics. We describe a new weakly atomic Java STM implementation that provides single global lock semantics while permitting concurrent execution, but we show that this comes at a significant performance cost. We also



rentation and performance ntics in a weakly atomic ve discuss why even recent semantics. We describe a tation that provides single concurrent execution, but performance cost. We also five semantics that loosen widing strong guarantees to previous ones, including

training alternative to lock-for managing concurrent decade, TM research has tomatically extract concur-and provide performance fine-grain locks. M has complicated multiges. In order to write cor-sort be able to reason about Outride of TM, these has enal memory models have act of allowable to havior to the ade are intrinsically sodel is a fundamental part I to its type safety and se-

ed to this mix, there is no argaage memory model. In where all data accesses are as are obligated to handle

situations when the same data is accused both transactionally and non-transactionally. Because of this, nearly all published STM systems provide weaker guarantees for transactions than locks. Consider the publication example in Figure 1. Thread 2 reads data early into the private location top. However, that value is only used if ready is set. Otherwise, top is dead and rever accused again. If the transactions are seplaced with locks, this program will rencornectly: val would either he 0 or 1 depending upon which critical section executed first. Nevertheless, most STMs would produce the value 42 in The ad 2 given the interleaving in Figure 1. As Thread 1 writes data outside a transaction, a weakly atomic STM would not detect any conflict and, thus, would not invalidate Thread 2's read operation. It should be noted that, even under locks, this program has a

data ram: Thread 1 and Thread 2 may access data simultaneously. However, from the programmer's perspective, the value is never used in this case, and the race should be benign. Java's memory model [17] (assuming lock semantics) specifically disallows any execution that produces 42 in The ad 2. In fact, this example also has implications for correctly synchronized programs. Standard compiler seordering can inadvertently introduce a data race. Consider a consectly synchronized variant of this program in Fig-ure 2 where Thread 2 only accesses data inside the conditional Compiler optimizations such as speculative redundancy elimination or instruction scheduling may still introduce a data race by hoisting the access (s.g., to line 2 in Figure 1) if profitable. Such optimizations assume that introduced races are beingn.

An appealing solution is to provide strong atomicity [5, 22] 1, where non-transactional memory accesses are analogous to single instruction transactions and prevented from violating the isolation of transactions. In this model, transactions are strictly more sentric-

<sup>1</sup> The term strong technics is also used in the literature



#### Abstractions vs implementations



### Defining "atomic" without saying "TM"

100

Microsoft<sup>®</sup>

Research

- "Strong semantics"
  - Simple interleaved execution of threads

V 20

- If a thread starts an atomic block then only it can take steps
- Blocking operations (e.g. "retry", "orElse", "blockUntil") can be incorporated – see refs below
- This means:
  - Atomic blocks are atomic wrt normal memory accesses
  - Do not need to model conflict detection / resolution
  - Can choose whether or not to retain the effects of an atomic block that raises an exception

"Composable memory transactions", Tim Harris, Simon Marlow, Simon Peyton Jones, Maurice Herlihy. PPoPP '05 "Semantics of Transactional Memory and Automatic Mutual Exclusion", Martín Abadi, Andrew Birrell, Tim Harris, Michael Isard. POPL '08



#### Example: a privatization idiom

 $x_shared = true; x = 0;$ 



#### Example: a privatization idiom

Microsoft<sup>®</sup>

Research

9

Etection

× 30 40 100

 $x_shared = true; x = 0;$ 







× 30 40 100

Microsoft<sup>®</sup>

Research

 $x_shared = false; x = 100;$ 





Microsoft<sup>®</sup>

Research

9

× 30 40 100

 $x_shared = false; x = 101;$ 





#### Example: a privatization idiom

 $x_shared = true; x = 0;$ 





Microsoft<sup>®</sup>

Research

9

× 30 40 100

 $x_shared = true; x = 0;$ 





# Etection 2 Example: a privatization idiom

Microsoft<sup>®</sup>

Research

9

× 30 40 100

 $x_shared = false; x = 0;$ 



#### Example: a privatization idiom

Microsoft<sup>®</sup>

Research

9

Etection 2

× 30 40 100

 $x_shared = false; x = 1;$ 





#### **Strong semantics**

- We've not talked about "inconsistent reads", "roll backs", "in-place vs lazy updates", "weak atomicity", "strong atomicity", ...
- We've not ruled out anything (e.g. I/O)
- We've not considered program
   transformations
- Is this a pipe-dream? Can we implement it?



# Untangling "atomic" from TM

# Hiding TM from programmers

## **Current performance**



#### Example: a privatization idiom

 $x_shared = true; x = 0;$ 









× 30 40 100

Microsoft<sup>®</sup>

Research

 $x_shared = false; x = 100;$ 





× 30 40 100

Microsoft<sup>®</sup>

Research

 $x_shared = false; x = 101;$ 





# Not valid under ics execution emantics estrong emantics Example: a privatization idiom

× 30 40 100

Microsoft<sup>®</sup>

Research

 $x_shared = false; x = 0;$ 





100

40

20

**Strong semantics** 

Microsoft<sup>®</sup>

Research

atomic, retry, ..... What, ideally, should these constructs do?

Programming discipline(s)

What does it mean for a program to use the constructs correctly?

Low-level semantics & actual implementations

Transactions, lock inference, optimistic concurrency, program transformations, weak memory models, ...



#### **Programming disciplines**

 Based on a program's execution under the strong semantics





#### **Static separation**

- Atomic blocks can only access local variables and designated "atomic variables"
- "atomic variables" cannot be accessed outside atomic blocks

```
Class Q {
  atomic QElem leftSentinel;
  atomic QElem rightSentinel;
  void pushLeft(int item) {
    atomic {
      QElem e = new QElem(item);
      e.right = this.leftSentinel.right;
      e.left = this.leftSentinel;
  }
}
```



#### **Delaunay triangulation**



"Delaunay Triangulation with Transactions and Barriers" Michael L. Scott, Michael F. Spear, Luke Dalessandro, and Virendra J. Marathe, IISWC 2007



### **Delaunay triangulation (2)**





#### **Dynamic separation**

- Add explicit operations to indicate whether data is accessed inside atomic blocks, or accessed outside them
- Correctly synchronized: data is always in the correct mode when it is accessed
- Robust dynamic checking is possible:
  - Either the program runs with strong semantics
  - Or it fails with an error

<sup>&</sup>quot;Implementation and use of transactional memory with dynamic separation", Martín Abadi, Andrew Birrell, Tim Harris, Johnson Hsieh, Michael Isard, CC '09 (to appear)



#### Violation freedom (VF)

- Allow data's access mode to change implicitly
- To be correctly synchronized:
  - Conflicting data accesses must not be attempted concurrently inside & outside atomic blocks
- Reminiscent of rules for programs to be free from data races

# This etamoletee Example: a privatization idiom

× 30 40 100

 $x_shared = true; x = 0;$ 

atomic { x\_shared = false; } X++;

Microsoft<sup>®</sup>

Research



#### **Programming with violations**



C# version of Labyrinth, derived from "STAMP: Stanford Transactional Applications for Multi-Processing" Chi Cao Minh, JaeWoong Chung, Christos Kozyrakis, Kunle Olukotun , IISWC '08



### Strong atomicity

- Similar to typical HTM behavior
- Trade off implementation complexity for (hopefully) scalability & straight-line speed,
- Two recent approaches:
  - "Dynamic Optimization for Efficient Strong Atomicity", Schneider et al, OOPSLA '08
  - "Transactional memory with strong atomicity using off-the-shelf memory protection hardware", Abadi et al, PPoPP '09



#### Strong atomicity =/> strong semantics

```
atomic {
	ready = true;
	data = 1;
}
```

tmp1 = ready; if (tmp1 == true) { tmp2 = data;}

- Can tmp1==true, tmp2==0?
- Under strong semantics: no
- Under plausible implementations with strong atomicity: yes

Example from "What do high-level memory models mean for transactions?" Dan Grossman, Jeremy Manson, William Pugh, MSPC '06





80 40

100

#### Open nesting, boosting, system calls

100

× 30 40

Microsoft<sup>®</sup>

Research





# Untangling "atomic" from TM

# Hiding TM from programmers

**Current performance** 



### Perf. figures depend on...

- Workload : What do the atomic blocks do? How long is spent inside them?
- Baseline implementation: Mature existing compiler, or prototype?
- Intended semantics: Support static separation? Violation freedom? Strong atomicity?
- STM implementation: In-place updates, deferred updates, eager/lazy conflict detection, visible/invisible readers?
- STM-specific optimizations: e.g. to remove or downgrade redundant TM operations
- Integration: e.g. dynamically between the GC and the STM, or inlining of STM functions during compilation
- Implementation effort: low-level perf tweaks, tuning, etc.
- Hardware: e.g. performance of CAS and memory system



#### Labyrinth



- STAMP v0.9.10
- 256x256x3 grid
- Routing 256 paths
- Almost all execution inside atomic blocks
- Atomic blocks can attempt 100K+
  updates
- C# version derived from original C
- Compiled using Bartok, whole program mode, C# -> x86 (~80% perf of original C with VS2008)
- Overhead results with Core2 Duo running Windows Vista

"STAMP: Stanford Transactional Applications for Multi-Processing" Chi Cao Minh, JaeWoong Chung, Christos Kozyrakis, Kunle Olukotun, IISWC 2008



100

40

20



STM implementation supporting static separation In-place updates Lazy conflict detection Per-object STM metadata Addition of read/write barriers before accesses Read: log per-object metadata word Update: CAS on per-object metadata word Update: log value being overwritten



















#### Scaling – Labyrinth





#### Scaling – Delaunay





#### Scaling – Genome





#### Scaling – Vacation





# Untangling "atomic" from TM

# Hiding TM from programmers

**Current performance** 



#### Abstractions vs implementations





#### **Future directions**

- Which programming discipline should we settle on
  - …in a language like C#?
  - ...in future languages?
- H/W acceleration based on mature optimized S/W implementations
- Progress guarantees, interactions with implementation techniques and performance
- What asymptotic bounds on STM performance can we give when supporting different programming disciplines?
- How do we define correctness of an STM interface, as opposed to the whole language implementation?



# Acnkowledgements

100

40

20

- Most of the work described in these slides has been collaborative; I'd like to thank colleagues at MSR Cambridge, MSR Redmond, MSR Mountain View, the Microsoft Parallel Computing Platform Group, the University of Cambridge Computer Lab, and the MSR-BSC joint research centre.
- Material is drawn from the following publications:
  - Martín Abadi, Andrew Birrell, Tim Harris, Johnson Hsieh, Michael Isard. "Implementation and use of transactional memory with dynamic separation". *Compiler Construction, March 2009*
  - Martín Abadi, Tim Harris, Mojtaba Mehrara. "Transactional memory with strong atomicity using offthe-shelf memory protection hardware". *PPoPP, February 2009*
  - Martín Abadi, Tim Harris, Katherine Moore. "A model of dynamic separation for transactional memory". CONCUR, August 2008
  - Martín Abadi, Andrew Birrell, Tim Harris, Michael Isard. "Semantics of Transactional Memory and Automatic Mutual Exclusion". POPL, January 2008
  - Keir Fraser, Tim Harris. "Concurrent programming without locks". ACM TOCS, May 2007
  - Tim Harris, Mark Plesko, Avraham Shinnar, David Tarditi. "Optimizing Memory Transactions" *PLDI*, June 2006
  - Tim Harris, Simon Marlow, Simon Peyton Jones, Maurice Herlihy. "Composable memory transactions" *PPoPP*, June 2005
  - Tim Harris, Keir Fraser. "Language Support for Lightweight Transactions". OOPSLA, October 2003
- The material on performance is current at Jan 2009, and reflects a slightly later more optimized implementation than that described in the PPoPP 2009 paper