

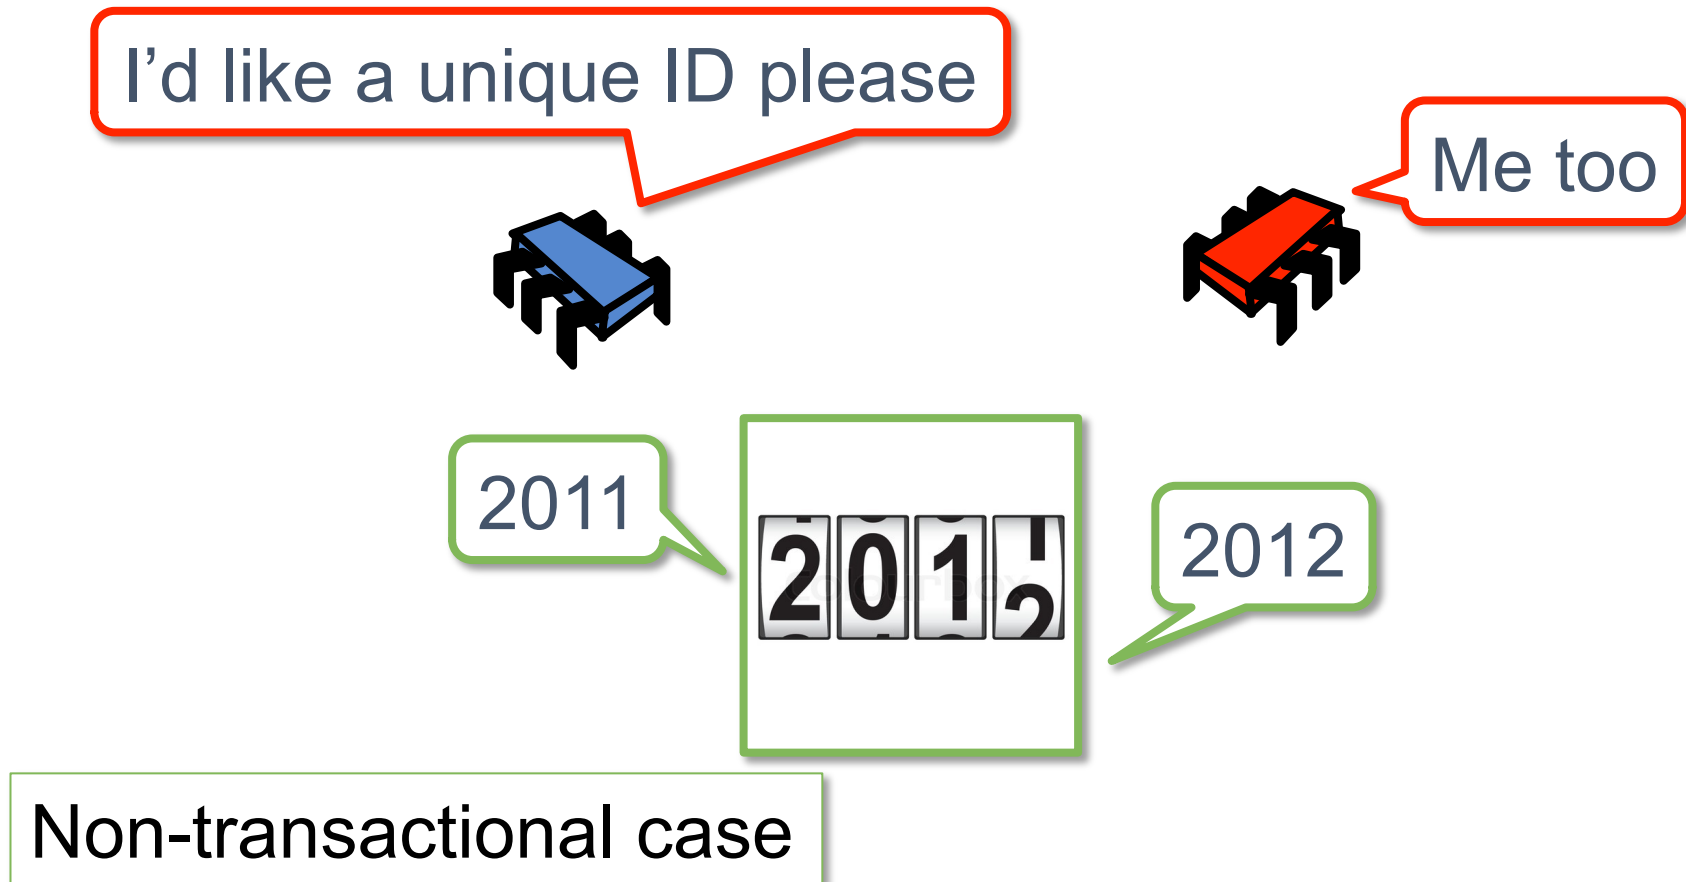
Implementation techniques for libraries of transactional concurrent data types

Liuba Shrira

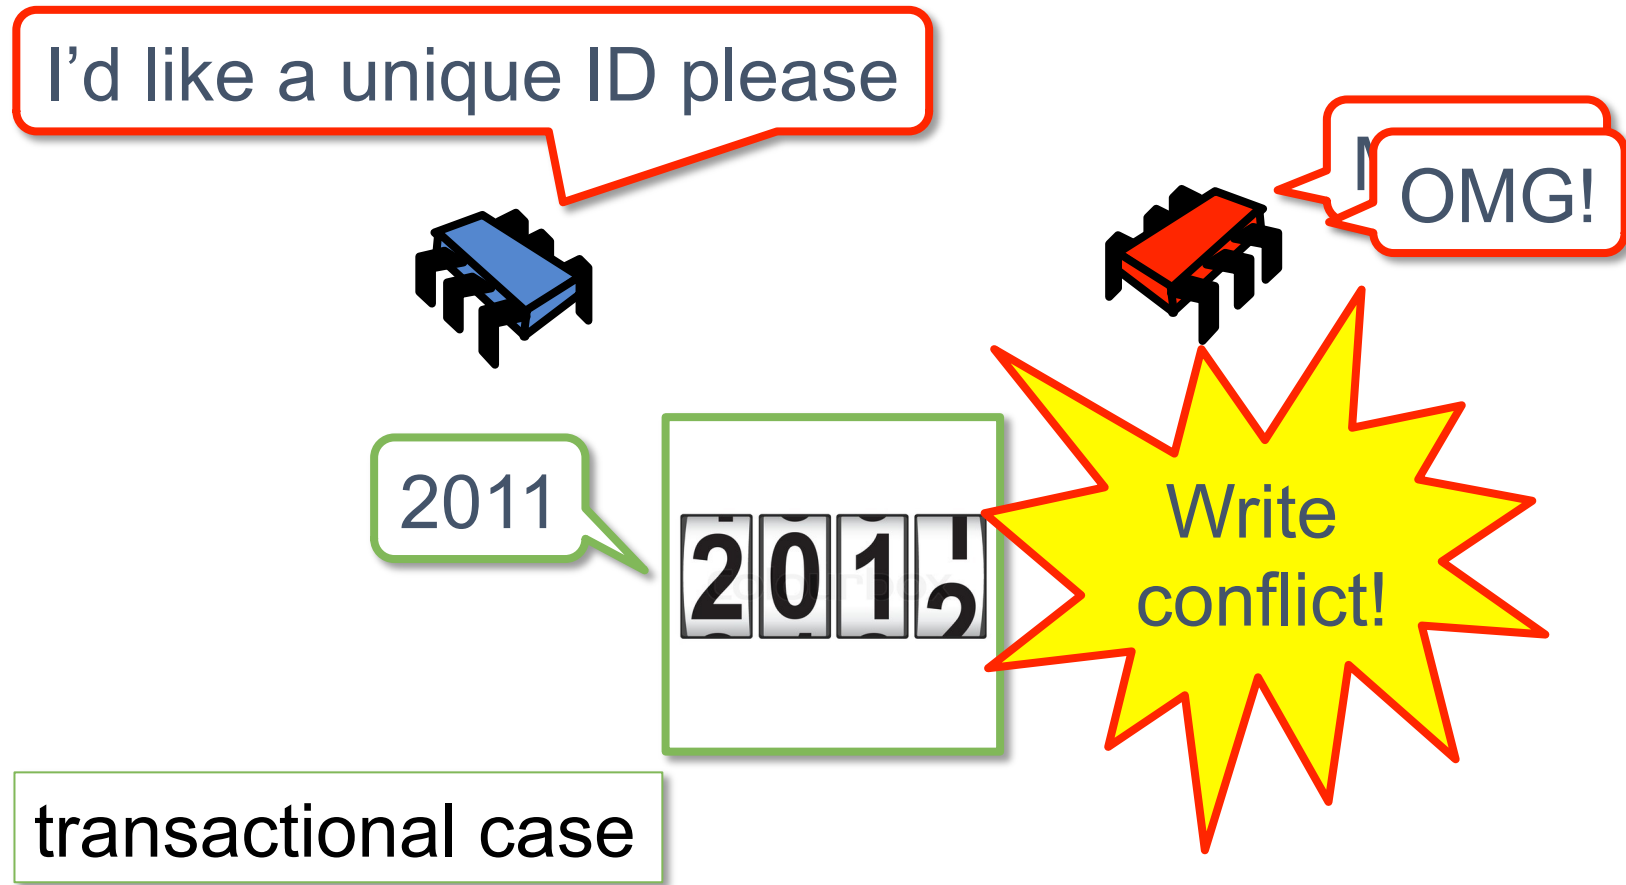
Brandeis University

Or: Type-Specific Concurrency Control and STM

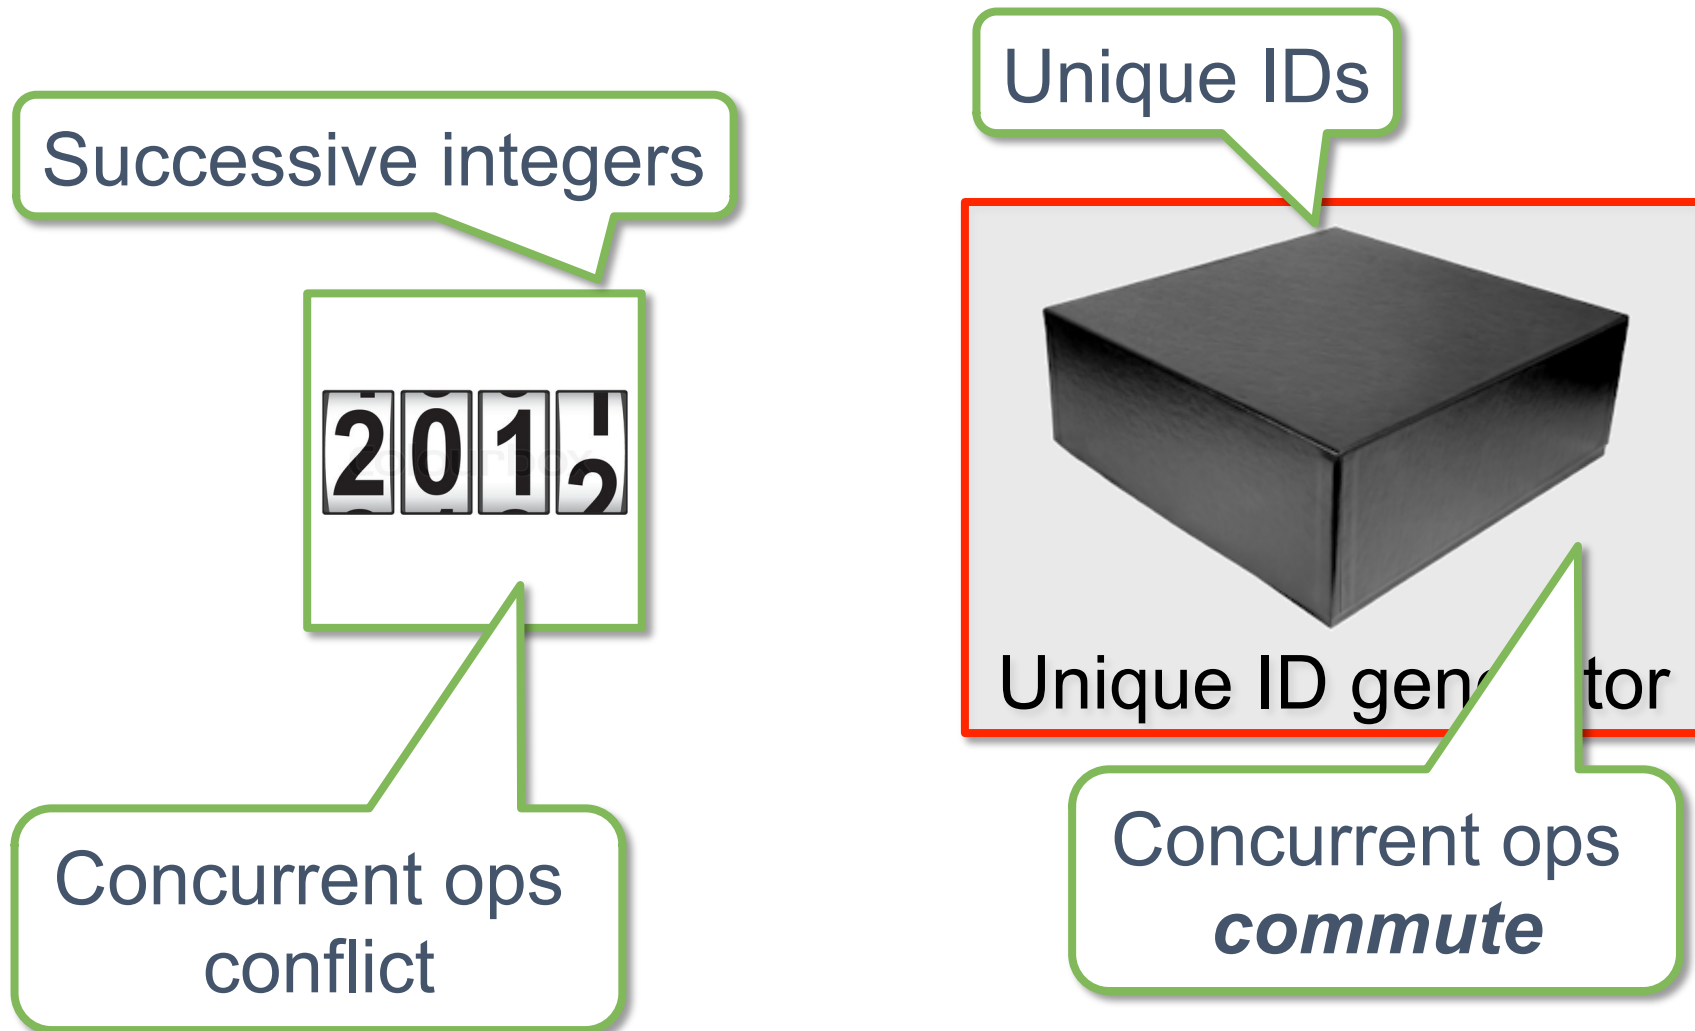
Where Modern STMs Fail




Where Modern STMs Fail



It's not the STMs problem really



Relaxed Atomicity



Early release, open-nested,
Eventual, elastic,
 ϵ -serializability, etc.

Popular in 80s & 90s ...

DB and distributed

Mostly forgotten ...

Except for snapshot isolation.

Type-Specific Concurrency Control



Also from 80s ...

Exploit

Commutativity ...

Non-determinism

For example

Escrow ...

Exo-leasing ...

TM raises different questions

Heart of the Problem

Confusion between *thread-level* and *transaction-level* synchronization.

Needless *entanglement* kills concurrency

Relaxed consistency models are all about *more entanglement*

Heart of the Problem

Confusion between *thread-level* and *transaction-level* synchronization.

Pure Read-Write Synchronization will **never** be scalable.

Relaxed consistency models are all about *more* efficient.

It may not even be **necessary**.

50 Shades of Synchronization

Short-lived, fine-grained

Atomic instruction (CAS)

Hardware Transaction

Critical Sections

Long-lived, coarse-grained

Software transaction

Transactional Boosting

Method for transforming.....

linearizable

highly concurrent

black-box

objects

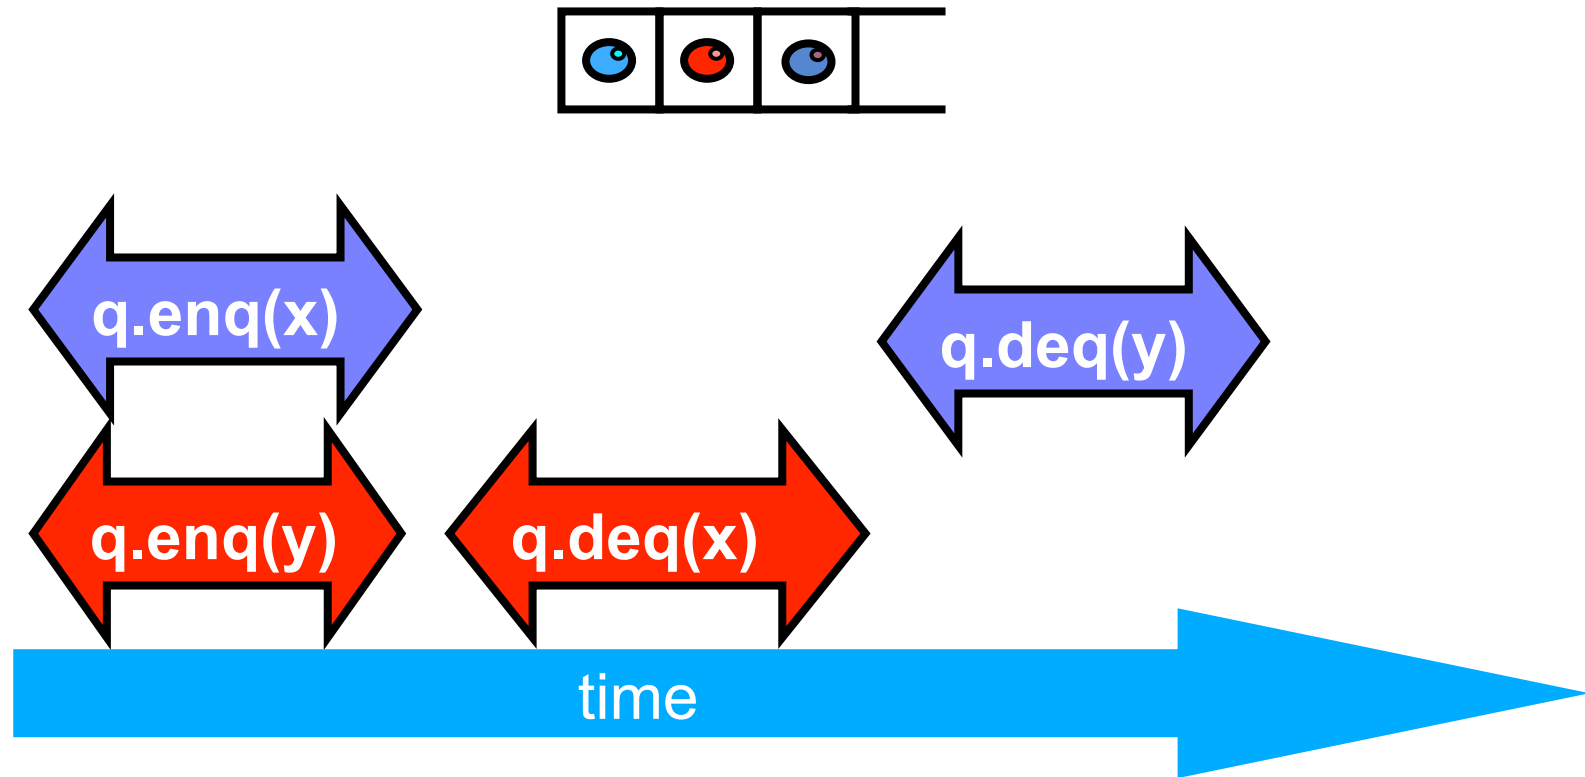
Into ...

highly concurrent

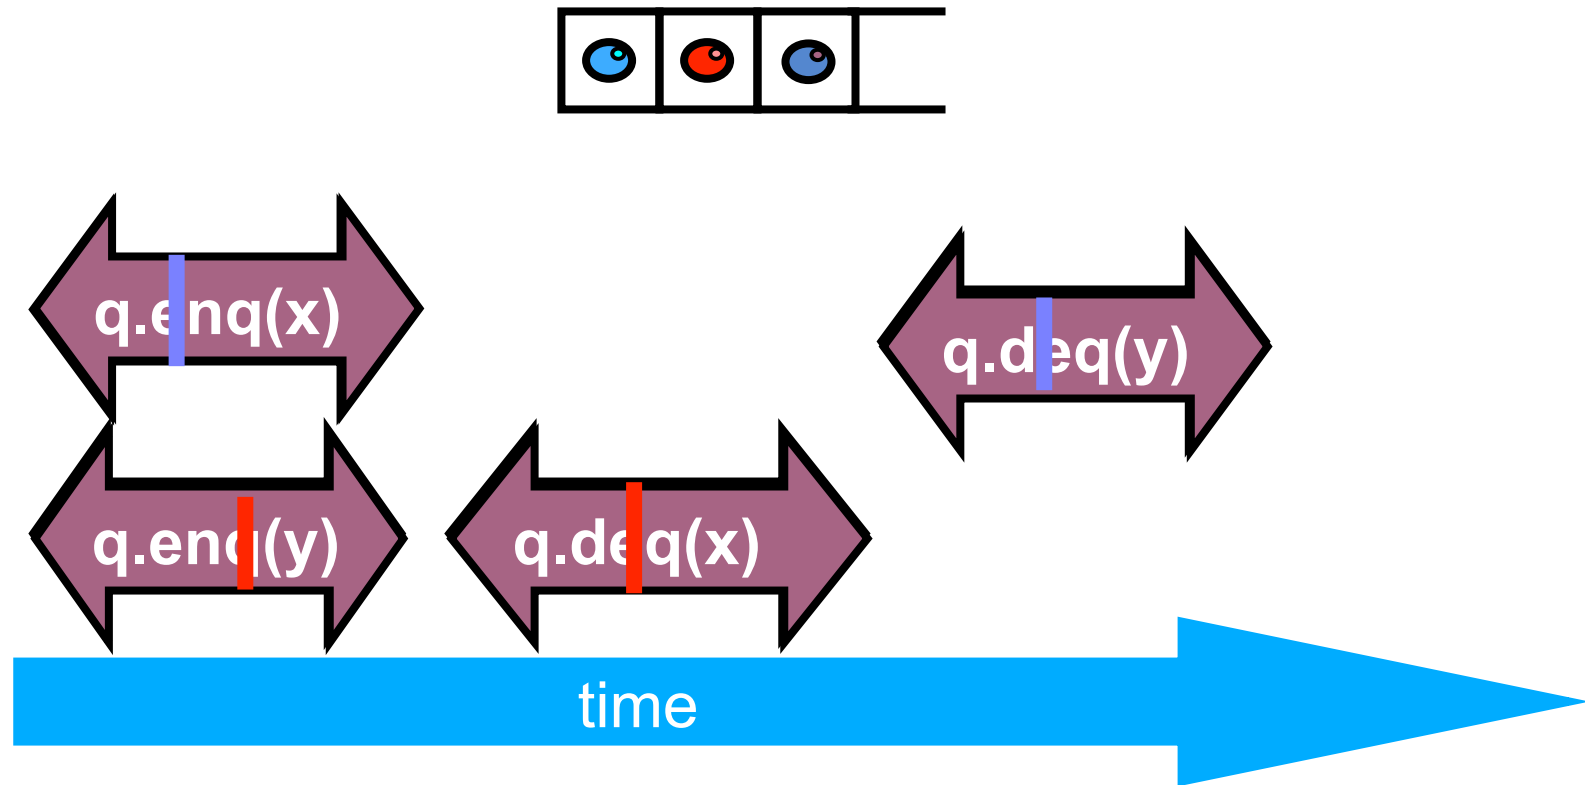
objects

transactional


Concurrent Objects




Linearizability




Linearizable Objects



Transactional Boosting



Disentangled Run-Time




Library:
abstract locks,
Inverse logs

Your favorite
fine-grained
algorithms

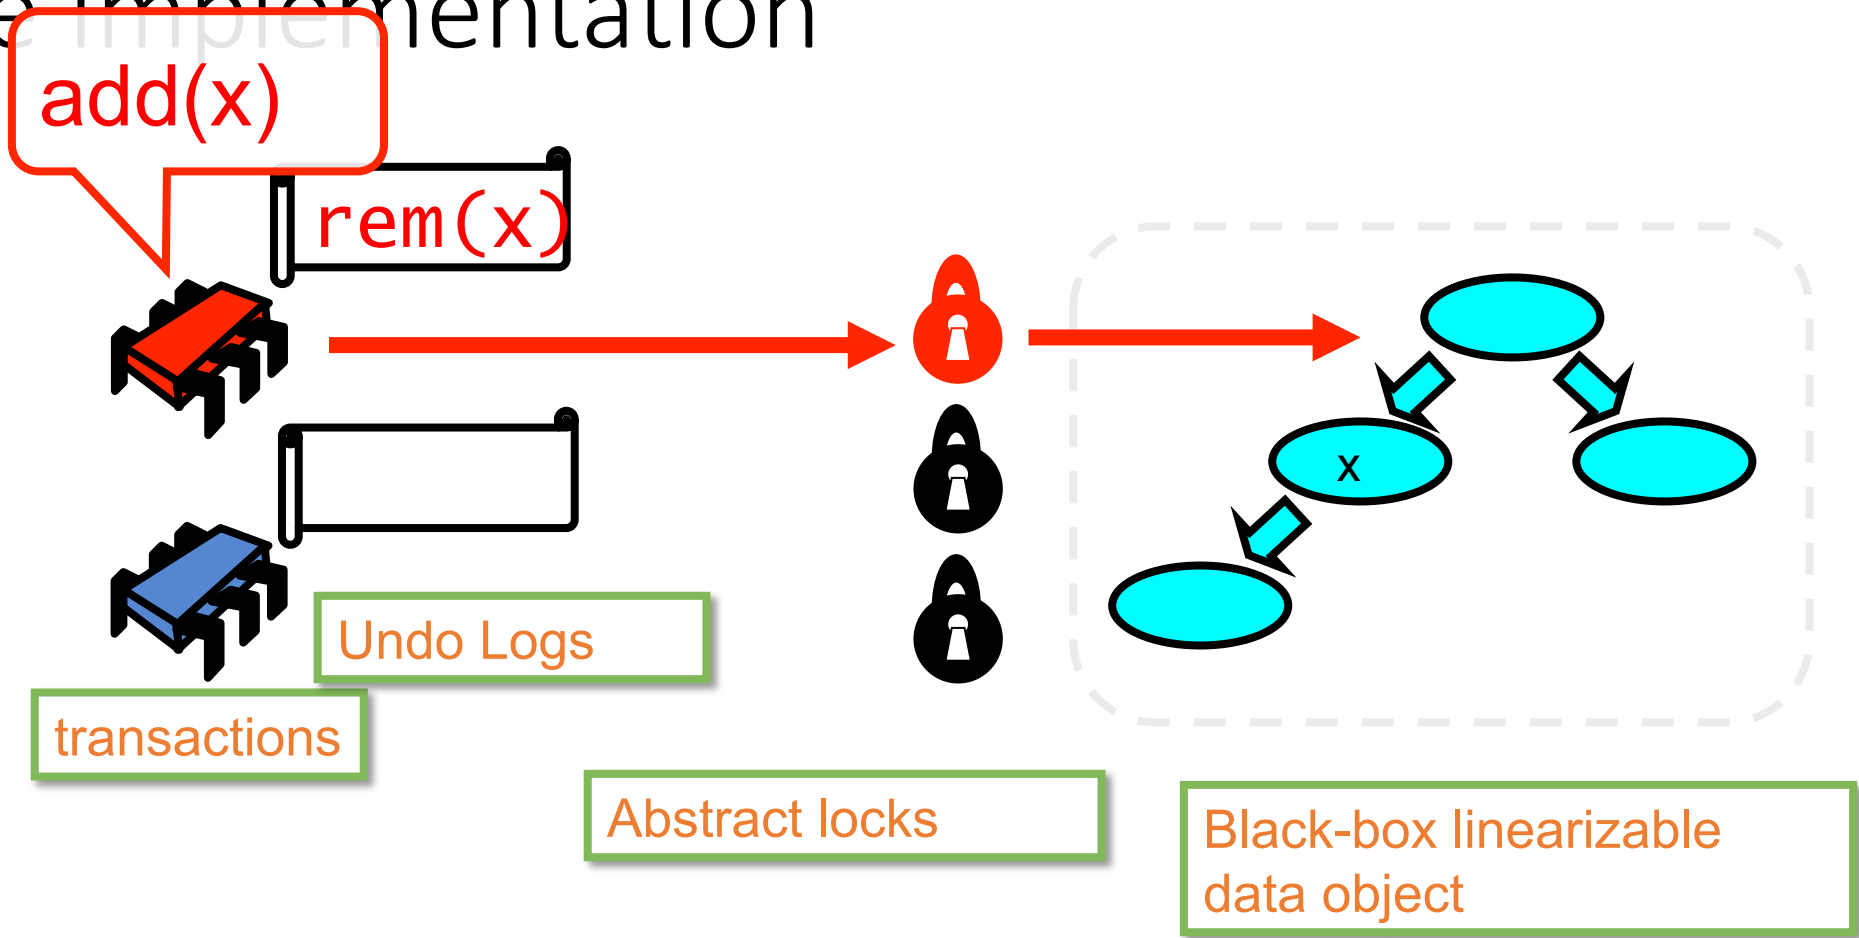
HW transac



Disentangled Reasoning



One Implementation



Lets look at some code

- Example 1: Transactional **Set**
 - implemented by boosting **ConcurrentSkipList** object, using **LockKey** for synchronization

Set Specification	
Method	Inverse
$\text{add}(x)/\text{false}$	$\text{noop}()$
$\text{add}(x)/\text{true}$	$\text{remove}(x)/\text{true}$
$\text{remove}(x)/\text{false}$	$\text{noop}()$
$\text{remove}(x)/\text{true}$	$\text{add}(x)/\text{true}$
$\text{contains}(x)/_$	$\text{noop}()$
Commutativity	
$\text{insert}(x)/_ \Leftrightarrow \text{insert}(y)/_, x \neq y$	
$\text{remove}(x)/_ \Leftrightarrow \text{remove}(y)/_, x \neq y$	
$\text{insert}(x)/_ \Leftrightarrow \text{remove}(y)/_, x \neq y$	
$\text{add}(x)/\text{false} \Leftrightarrow \text{remove}(x)/\text{false} \Leftrightarrow \text{contains}(x)/_$	

```

1 public class SkipListKey {
2     ConcurrentSkipListSet<Integer> list ;
3     LockKey lock;
4     ...
5     public boolean add(final int v) {
6         lock.lock(v);
7         boolean result = list.add(v);
8         if ( result ) {
9             Thread.onAbort(new Runnable() {
10                 public void run() { list.remove(v);}}
11             );
12         }
13         return result ;
14     }
15     ...
16 }

```

```

17 public class LockKey {
18     ConcurrentHashMap<Integer,Lock> map;
19     public LockKey() {
20         map = new ConcurrentHashMap<Integer,Lock>();
21     }
22     public void lock(int key) {
23         Lock lock = map.get(key);
24         if (lock == null) {
25             Lock newLock = new ReentrantLock();
26             Lock oldLock = map.putIfAbsent(key, newLock);
27             lock = (oldLock == null) ? newLock : oldLock;
28         }
29         if (LockSet.add(lock)) {
30             if (!lock.tryLock(LOCK_TIMEOUT,
31                 TimeUnit.MILLISECONDS);) {
32                 lockSet.remove(lock);
33                 Thread.getTransaction().abort();
34                 throw new AbortedException();
35             }
36         }
37     }
38     ...
39 }

```

More examples:

- Transactional **Priority Queue, Pipelining, UniqueID** ...
 - implemented by boosting concurrent objects from Java concurrency packages

Transactional Boosting: A Methodology for Highly-Concurrent Transactional Objects

Maurice Herlihy Eric Koskinen
Computer Science Department, Brown University
{mph,ejk}@cs.brown.edu

Abstract

We describe a methodology for transforming a large class of highly-concurrent linearizable objects into highly-concurrent transactional objects. As long as the linearizable implementation satisfies certain regularity properties (informally, that every method has an inverse), we define a simple wrapper for the linearizable implementation that guarantees that concurrent transactions without inherent conflicts can synchronize at the same granularity as the original linearizable implementation.

Categories and Subject Descriptors D.1.3 [Programming Techniques]: Concurrent Programming – Parallel Programming; D.3.3 [Programming Languages]: Language Constructs and Features
– Frameworks; Concurrent programming structures; E.1 [Data Structures]: Distributed data structures; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs

Synchronizing via read/write conflicts has one substantial advantage: it can be done automatically without programmer participation. It also has a substantial disadvantage: it can severely and unnecessarily restrict concurrency for certain shared objects. If these objects are subject to high levels of contention (that is, they are “hot-spots”), then the performance of the system as a whole may suffer.

Here is a simple example. Consider a mutable set of integers that provides $\text{add}(x)$, $\text{remove}(x)$ and $\text{contains}(x)$ methods with the obvious meanings. Suppose we implement the set as a sorted linked list in the usual way. Each list node has two fields, an integer value and a node reference next. List nodes are sorted by value and values are not duplicated. Integer x is in the set if and only if a list node has value field x . The $\text{add}(x)$ method inserts a new list node until it encounters the largest value less than x . If x is already in the set, it does nothing. Consider a node n in the list. If n is the largest value less than x , then n is the node to which $\text{add}(x)$ will add a new node.

Performance of boosting




Figure 8. Throughput for boosted STAMP benchmarks normalized against the throughput of conventional TL2, which maintains shadow copies.




Figure 9. Throughput for transactionally-boosted skip lists using simple (left) and key-based (right) two-phase locks.

What's the Catch?

Concurrent calls must ***commute***

Different orders yield same state, values

(Actually, all about ***left/right movers***)

Methods must have ***inverses***

Immediately after, restores state

What's the Catch?

Concurrent calls must

Different order of calls must result in same state values

(Abstract states must be left/right

Same functionality, conflict detection
& recovery, implicit in all TMs

Abstract (algebraic) states,
not raw, seething bits!

Boosting

- Reuse code, improve performance
- But inverses

Transactional Boosting: A Methodology for Highly-Concurrent Transactional Objects

Maurice Herlihy Eric Koskinen
Computer Science Department, Brown University
{mph,ejk}@cs.brown.edu

Abstract

We describe a methodology for transforming a large class of highly-concurrent linearizable objects into highly-concurrent transactional objects. As long as the linearizable implementation satisfies certain regularity properties (informally, that every method has an inverse), we define a simple wrapper for the linearizable implementation that guarantees that concurrent transactions without inherent conflicts can synchronize at the same granularity as the original linearizable implementation.

Categories and Subject Descriptors: D.1.3 [Programming Techniques]: Concurrent Programming – Parallel Programming; D.3.3 [Programming Languages]: Language Constructs and Features – Frameworks; Concurrent programming structures; E.1 [Data Structures]: Distributed data structures; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs

Synchronizing via read/write conflicts has one substantial advantage: it can be done automatically without programmer participation. It also has a substantial disadvantage: it can severely and unnecessarily restrict concurrency for certain shared objects. If these objects are subject to high levels of contention (that is, they are “hot-spots”), then the performance of the system as a whole may suffer.

Here is a simple example. Consider a mutable set of integers that provides $\text{add}(x)$, $\text{remove}(x)$ and $\text{contains}(x)$ methods with the obvious meanings. Suppose we implement the set as a sorted linked list in the usual way. Each list node has two fields, an integer value and a node reference next. List nodes are sorted by value and values are not duplicated. Integer x is in the set if and only if a list node has value field x . The $\text{add}(x)$ method scans the list until it encounters the largest value less than x . If x is not in the set, it creates a new node with value x and next pointer to the first node with value greater than x .

And is there ever enough performance?

How to improve performance?

Recall, we want

- Good performance when synchronization is required
- Scalability
- E.g., for in-memory key-value store

Up next: how to improve the performance of a transactional data structure?

- MassTree: a high performance data structure
- Silo: high performance transactions over MassTree using a different approach
- STO: a general framework and methodology for building libraries of customized high performance transactional objects

Cache Craftiness for Fast Multicore Key-Value Storage

Yandong Mao, Eddie Kohler[†], Robert Morris
MIT CSAIL, [†]Harvard University

Abstract
We present MassTree, a fast key-value database designed for SMP machines. MassTree keeps all data in memory. Its main data structure is a trie-like concatenation of B⁺-trees, each of which handles a fixed-length slice of a variable-length key. This structure effectively handles arbitrary-length keys possibly descending the tree and prefetching each tree node. Lookups use optimistic concurrency control, a read-copy-update-like technique, and do not write shared data structures; updates lock only affected nodes. Logging and checkpoints; updates appear elsewhere, MassTree's design discusses design

sufficiently fast single servers. A common route to high performance is to use different specialized storage systems for different workloads [4]. This paper presents MassTree, a storage system specialized for key-value data in which all data fits in memory, but must persist across server restarts. Within these constraints, MassTree aims to provide a flexible storage model over those keys; clients can

Speedy Transactions in Multicore In-Memory Databases

Stephen Tu, Wenting Zheng, Eddie Kohler[†], Barbara Liskov, and Samuel Madden
MIT CSAIL and [†]Harvard University

Abstract
Silo is a new in-memory database that achieves excellent performance and scalability on modern multicore machines. Silo was designed from the ground up to use system memory and caches efficiently. For instance, it avoids all centralized contention points, including that of centralized transaction ID assignment. Silo's key contention control that provides serializability while avoiding all shared-memory writes for records that were only read. Though this might seem to complicate the enforcement of a serial order, correct logging and recovery is provided by linking periodically-updated epoches with the commit protocol. Silo provides the same guarantees as any serializable database without unnecessary scalability bottlenecks or much additional overhead. Silo achieves almost 700,000 transactions per second on standard TPC-C workloads.

ization scale with the data, allowing larger databases to support more concurrency. Silo uses a MassTree-inspired tree structure for its underlying indexes. MassTree [23] is a fast concurrent B-tree-like structure optimized for multicore performance. But MassTree only supports non-serializable, single-key transactions, whereas any real database must support serial order. Our core result, the Silo commit protocol, is a minimal-contention serializable commit protocol, provides these properties. Silo uses a variant of the (OCC) [18] as a variant of

Type-Aware Transactions for Faster Concurrent Code

Nathaniel Herman
Harvard University/Dropbox
nherman@post.harvard.edu

Jeevana Priya Inala
MIT
jinala@mit.edu

Eddie Kohler
Harvard University
kohler@seas.harvard.edu

Barbara Liskov
MIT
liskov@piano.csail.mit.edu

Yi He Huang Lillian Tsai
Harvard University
yihetu@mit.edu
lilliantai@college.harvard.edu

Liuba Shrira
Brandeis University
liuba@brandeis.edu

Abstract
It is often possible to improve a concurrent system's performance by leveraging the semantics of its datatypes. We build a new software transactional memory (STM) around this observation. A conventional STM tracks read- and write-sets of memory words, even simple operations can generate large sets. Our STM, which we call STO, tracks abstract operations on transactional datatypes in general. datatypes' implementation

However, TMs have performance issues. In hardware TM, fundamental microarchitectural limitations, such as bounds on maximum transaction size, mean that transactions can never commit. STO, tracks abstract operations on transactional datatypes in general. datatypes' implementation

MassTree

- High-performance key/value store
 - In shared primary memory
 - Cores run **put**, **get**, and **delete** requests

Cache Craftiness for Fast Multicore Key-Value Storage

Yandong Mao, Eddie Kohler[‡], Robert Morris
MIT CSAIL, [†]Harvard University

Abstract

We present Masstree, a fast key-value database designed for SMP machines. Masstree keeps all data in memory. Its main data structure is a trie-like concatenation of B⁺-trees, each of which handles a fixed-length slice of a variable-length key. This structure effectively handles arbitrary-length possibly-binary keys, including keys with long shared prefixes. B⁺-tree fanout was chosen to minimize total DRAM delay when descending the tree and prefetching each tree node. Lookups use optimistic concurrency control, a read-copy-update-like technique, and do not write shared data structures; updates lock only affected nodes. Logging and checkpointing provide consistency and durability. Though some of these ideas appear elsewhere, Masstree is the first to combine them. We discuss design variants and their consequences.

On a 16-core machine, with logging enabled and queries

sufficiently fast single servers. A common route to high performance is to use different specialized storage systems for different workloads [4].

This paper presents Masstree, a storage system specialized for key-value data in which all data fits in memory, but must persist across server restarts. Within these constraints, Masstree aims to provide a flexible storage model. It supports arbitrary, variable-length keys. It allows *range queries* over those keys: clients can traverse subsets of the database, or the whole database, in sorted order by key. It performs well on workloads with many keys that share long prefixes. (For example, consider Bigtable [12], which stores information about Web pages under permuted URI paths like “edu.harvard.seas.www/news-event” and “gather information about interesting”.)

Review: Memory Model

- Each core has a **cache**
- Hitting in the cache matters a lot for reads!
- What about a write?
 - TSO (Total Store Order)

X86-TSO

- Thread t1 modifies x and later y
 - Thread t2 sees modification to y
 - t2 reads x
-
- Implies t2 sees modification of x

MassTree structure

- Nodes and records
- Nodes
 - Cover a range of keys
 - Interior and leaf nodes
- Records
 - Store the **values**

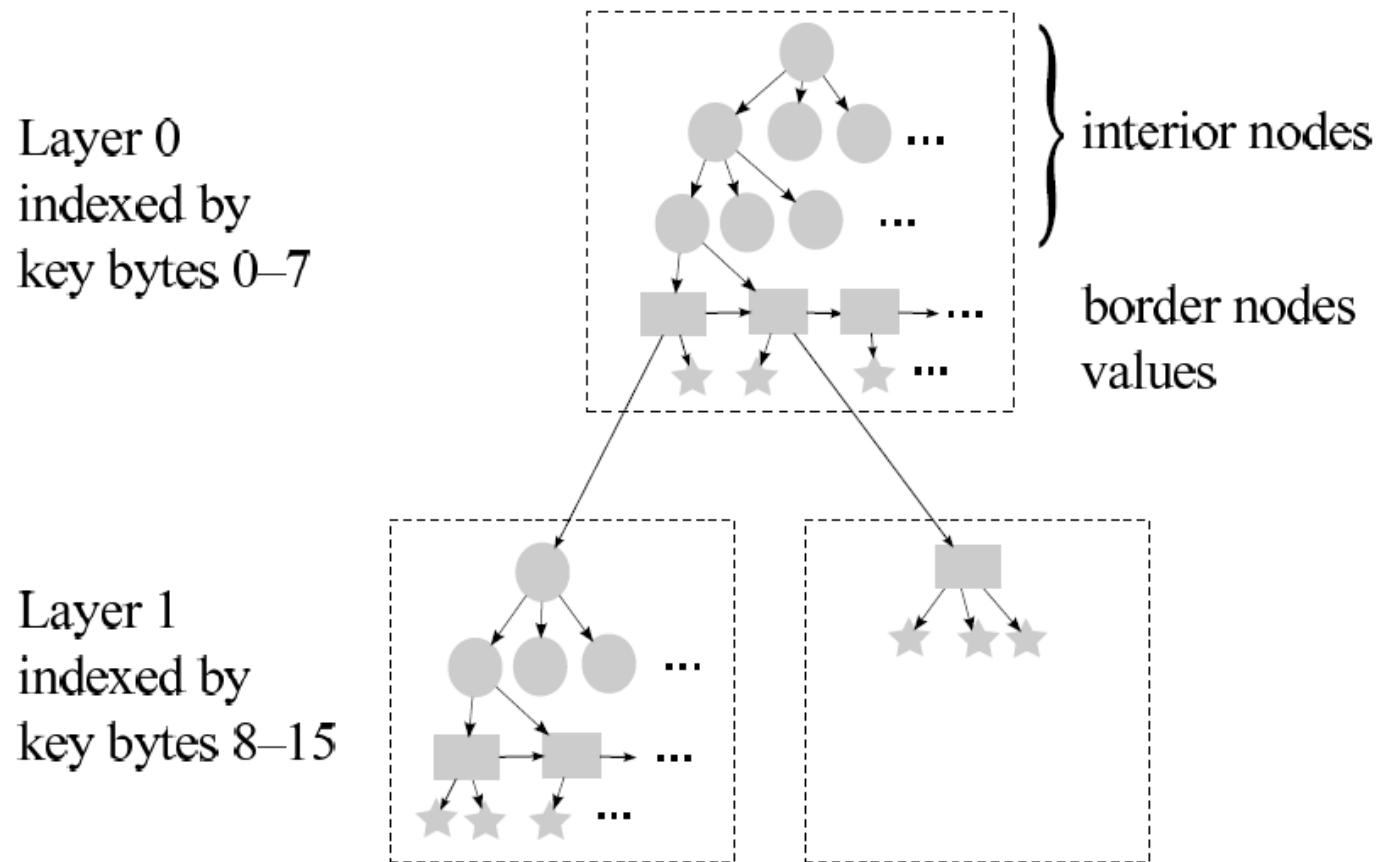


Figure 1. Masstree structure: layers of B⁺-trees form a trie.

Concurrency Control

- Reader/writer locks?

Thread-level Concurrency Control

- Base instructions
 - Compare and swap
 - On one memory word
 - Fence

Concurrency Control for multi-word

- First word of nodes and records
 - version number (v#) and lock bit

Concurrency control

- Write
 - Set lock bit (spin if necessary)
 - uses compare and swap
 - Update node or record
 - Increment v# and release lock

Concurrency control

- Write (locking)
- Read (no locking)
 - Spin if locked
 - Read contents
 - If v# has changed or lock is set, try again

Concurrency control

- Writes are pessimistic
- Reads are optimistic
- A mix!
- No writes for reads


Inserting new keys

- Into leaf node if possible
- Else split

Inserting new keys

- Into leaf node if possible
- Else split
- Split locks nodes **up the path**
 - No **deadlocks**

Interesting Issue with splitting



From MassTree to Silo

- High-performance database
- With transactions

Speedy Transactions in Multicore In-Memory Databases

Stephen Tu, Wenting Zheng, Eddie Kohler[†], Barbara Liskov, and Samuel Madden
MIT CSAIL and [†]Harvard University

Abstract

Silo is a new in-memory database that achieves excellent performance and scalability on modern multicore machines. Silo was designed from the ground up to use system memory and caches efficiently. For instance, it avoids all centralized contention points, including that of centralized transaction ID assignment. Silo's key contribution is a commit protocol based on optimistic concurrency control that provides serializability while avoiding *all* shared-memory writes for records that were only read. Though this might seem to complicate the enforcement of a serial order, correct logging and recovery is provided by linking periodically-updated *epochs* with the commit protocol. Silo provides the same guarantees as any serializable database without unnecessary scalability bottlenecks or much additional latency. Silo achieves almost 700,000 transactions per second on a standard TPC-C workload mix on a 32-core machine, as well as near-linear scalability. Considered per core, this is several times higher than previously reported.

nization scale with the data, allowing larger databases to support more concurrency.

Silo uses a MassTree-inspired tree structure for its underlying indexes. MassTree [23] is a fast concurrent B-tree-like structure optimized for multicore performance. But MassTree only supports non-serializable, single-key transactions, whereas any real database must support serial order. Our core result, the Silo commit protocol, is a minimal-contention serializable commit protocol that provides these properties.

Silo uses a variant of optimistic concurrency control (OCC) [18]. An OCC transaction tracks the records it reads and writes in thread-local storage. At commit time, after validating that no concurrent transaction's write overlapped with its read set, the transaction's write set is written to the database. This approach ensures that the database is in a consistent state.

Silo

- Database is in **primary memory**
- Runs **one-shot requests**

Silo

- Database is in **primary memory**
- Runs **one-shot requests**
- A tree for each table or index
- **Worker threads** run the requests
 - One thread per core
- Workers **share memory**

Transactions

```
begin {  
  % do stuff: run queries  
  % using insert, lookup, update, delete,  
  % and range  
}
```


Running Transactions

- MassTree operations release locks before returning
 - Hold locks longer?

Running Transactions

- OCC (Optimistic Concurrency Control)
 - Thread maintains **read-set** and **write-set**
 - Read-set contains version numbers
 - Write-set contains new state
- At end, **attempts commit**

Commit Protocol

- Phase 1: lock all objects in write-set
 - Bounded spinning

Commit Protocol

- Phase 1: lock all objects in write-set
- Phase 2: verify $v\#$'s of read-set
 - Abort if locked or changed

Commit Protocol

- Phase 1: lock all objects in write-set
- Phase 2: verify $v\#$'s of read-set
- Select Tid ($>v\#$ of r- and w-sets)
 - Without a write to shared state!

Commit Protocol

- Phase 1: lock all objects in write-set
- Phase 2: verify $v\#$'s of read-set
- Select Tid ($>v\#$ of r- and w-sets)
- Phase 3: update objects in write-set
 - Using Tid as $v\#$

Commit Protocol

- Phase 1: lock all objects in write-set
- Phase 2: verify v#'s of read-set
- Select Tid ($>v\#$ of r- and w-sets)
- Phase 3: update objects in write-set
- Release locks

Additional Issues

- Range queries
- Absent keys
- Garbage collection

Performance




Figure 4: Overhead of *MemSilo* versus *Key-Value* when running a variant of the YCSB benchmark.

Performance

- vs. Hstore
 - M. Stonebraker et al, The end of an architectural era: (it's time for a complete rewrite), VLDB '07




Figure 8: Performance of *Partitioned-Store* versus *MemSilo* as the percentage of cross-partition transactions is varied.

Silo to STO

- STO (Software Transactional Objects)

Type-Aware Transactions for Faster Concurrent Code

Nathaniel Herman
Harvard University/Dropbox
nherman@post.harvard.edu

Eddie Kohler
Harvard University
kohler@seas.harvard.edu

Jeevana Priya Inala
MIT
jinala@mit.edu

Barbara Liskov
MIT
liskov@piano.csail.mit.edu

Yihe Huang Lillian Tsai
Harvard University
yihehuang@g.harvard.edu
lilliantasai@college.harvard.edu

Liuba Shrira
Brandeis University
liuba@brandeis.edu

Abstract

It is often possible to improve a concurrent system's performance by leveraging the semantics of its datatypes. We build a new software transactional memory (STM) around this observation. A conventional STM tracks read- and write-sets of memory words; even simple operations can generate large sets. Our STM, which we call STO, tracks abstract operations on transactional datatypes instead. Parts of the transactional commit protocol are delegated to these datatypes' implementations, which can use datatype semantics, and new commit protocol features, to reduce bookkeeping, limit false conflicts, and implement efficient concurrency control.

However, TMs have performance issues. In hardware TM, fundamental microarchitectural limitations, such as bounds on maximum transaction size, mean some valid transactions can never commit [51]. Implementations will gradually improve, but general-purpose transactions will be backed up by software. Unfortunately, software TM performance severely lags that of purpose-built concurrent code [7], since STM implementations have high costs for bookkeeping and concurrency control. Transaction must track all objects accessed during. Transaction validation, either by locking them during or by

STO

- Silo trees are an **highly concurrent data structures**
 - **Specification** determines potential concurrency
 - **Implementation is hidden**
 - Including concurrency control

A vision for concurrent code

- Apps run **transactions**

A vision for concurrent application code, like boosting

- Apps run **transactions**
- Using **transaction-aware datatypes**
 - E.g., sets, maps, arrays, boxes, queues

Transactions

```
begin {  
  % do stuff: run queries  
  % using insert, lookup, update, delete,  
  % and range  
}
```

Back to our vision for concurrent code

- Apps run **transactions**
- Using fast **transaction-aware datatypes**
 - Designed by experts
 - Require sophistication to implement
 - But so are concurrent datatypes in Java

STO

- Think Silo broken into two parts:
 - STO platform
 - Transaction-aware datatypes

STO Platform

- Runs transactions
 - Transaction { ... }
- Provides **transaction state**
 - Read- and write-sets
- Runs commit protocol using **callbacks**

Transaction-aware datatypes

- Provide ops for user code
 - E.g., lookup, update, insert, delete, range
- Record reads and writes via platform
- Provide callbacks
 - lock, unlock, check, install, cleanup

Transaction-aware datatypes

- Provide ops for user code
- Record reads and writes via platform
- Provide callbacks
 - lock, unlock, check, install, cleanup
 - cleanup for abort, after-commit
 - E.g., deleting a key

Transaction-aware types

- Maps
- Hash tables
- Counters
 - `void incr()` vs. `int incr()`
 - Uses **check** and **install**

Designing fast STO's data types:

- Specification
- Some common tricks
 - Inserted elements: direct updates
 - Absent elements: extra version numbers
 - Read-my-writes: adjustments
- Correctness

Specification

Inserted elements and repeated lookup

- Hybrid strategy
 - T1: insert “poisoned” element
 - T2: abort on observing a “poisoned” element
 - T1: no need to validate insertion at commit

Absent elements

- T1: `get(K)` : K is absent
 - How to validate at commit?
- Extra version numbers
 - For hash table: on bucket of absent key
 - BTree : on parent node of absent key

Read-my-writes

- T1: scan a range A..Z; insert a key C
 - how to validate range ?


Correctness

- Version numbers on all shared state
- Exclusive locks
- Check must fail if segment locked or version number changed
- Modifications invisible to other transactions before install

Performance

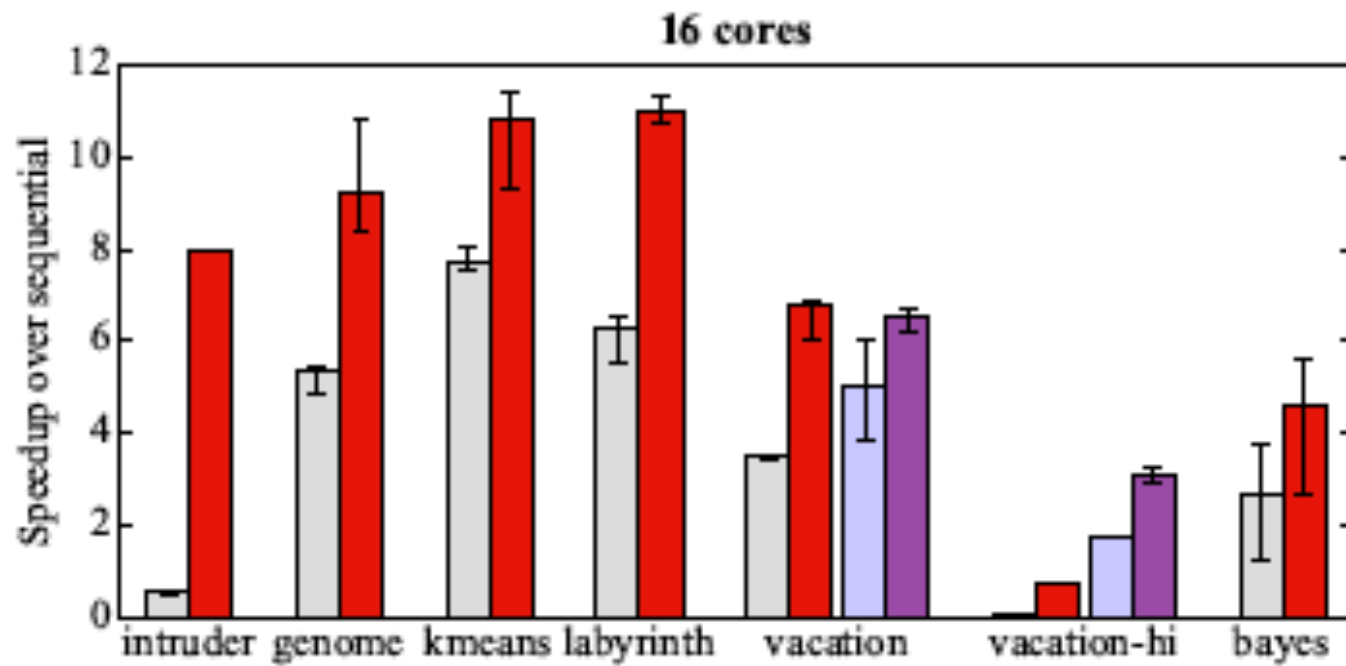
Implementation

- Silo: 7000 lines of code
- STO-Silo: 3000 lines of code
 - Uses hash tables and trees



Performance

- vs. TL2 (grey)
- And boosting (lilac)



Optimism vs Pessimism?

Effects of pessimism and boosting on a hash table micro benchmark.

Numbers are speedup at 16 threads relative to single-threaded STO

STO	12.91x
Boosting	7.02x
Boosting in STO	6.67x
Pessimistic STO	9.82x

More examples of powerful optimizations

Type-Aware Transactions for Faster Concurrent Code

Nathaniel Herman
Harvard University/Dropbox
nherman@post.harvard.edu

Eddie Kohler
Harvard University
kohler@seas.harvard.edu

Jeevana Priya Inala
MIT
jinala@mit.edu

Barbara Liskov
MIT
liskov@piano.csail.mit.edu

Yihe Huang Lillian Tsai
Harvard University
yihehuang@g.harvard.edu
lilliant sai@college.harvard.edu

Liuba Shrira
Brandeis University
liuba@brandeis.edu

Abstract

It is often possible to improve a concurrent system's performance by leveraging the semantics of its datatypes. We build on software transactional memory (STM) around this idea. A conventional STM tracks read- and write-operations on memory words; even simple operations can generate memory words. Our STM, which we call STO, tracks abstract operations on transactional datatypes instead. Parts of the conventional commit protocol are delegated to abstract implementations, which can use datatypes to avoid conflicts. Our protocol features a new abstraction for

However, TMs have performance issues. In hardware TMs, fundamental microarchitectural limitations place bounds on maximum transaction size. Software transactions can never commit size. Hardware transactions gradually improve, but software transactions can be backed up by hardware transactions. Hardware transactions perform

STO: last word for exploiting ADT in TM?

- Needs more work
 - More datatypes
 - Methodology
 - Programming language integration
 - Distribution

Summary:

Implementing a Library of Transactional Data types:

- Distinction between short Thread level vs coarse grain Transaction-level coordination is key
- Can re-use data structure code or co-design and customize:
 - Boosting: a black box approach, first ADT/STM (code re-use, restrictions)
 - STO: high-performance pessimistic/optimistic approach (co-design and customize)
- (Thanks to M.Herlihy and B. Liskov for help with slides!)

Questions