

Coarse-Grained Transactions

Eric Koskinen
University of Cambridge

Matthew Parkinson
University of Cambridge

Maurice Herlihy
Brown University

Abstract

Traditional transactional memory systems suffer from overly conservative conflict detection, yielding so-called false conflicts, because they are based on fine-grained, low-level read/write conflicts. In response, the recent trend has been toward integrating various abstract data-type libraries using ad-hoc methods of high-level conflict detection. These proposals have led to improved performance but a lack of a unified theory has led to confusion in the literature.

We clarify these recent proposals by defining a generalization of transactional memory in which a transaction consists of coarse-grained (abstract data-type) operations rather than simple memory read/write operations. We provide semantics for both pessimistic (e.g. transactional boosting) and optimistic (e.g. traditional TMs and recent alternatives) execution. We show that both are included in the standard atomic semantics, yet find that the choice imposes different requirements on the coarse-grained operations: pessimistic requires operations be left-movers, optimistic requires right-movers. Finally, we discuss how the semantics applies to numerous TM implementation details discussed widely in the literature.

Categories and Subject Descriptors D.1.3 [Programming Techniques]: Concurrent Programming—Parallel Programming; D.3.1 [Programming Languages]: Formal Definitions and Theory—Semantics; D.3.2 [Programming Languages]: Language Classifications—Concurrent, distributed, and parallel languages; D.3.3 [Programming Languages]: Language Constructs and Features—Concurrent programming structures; F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages—Operational semantics

General Terms Languages, Theory

Keywords Coarse-grained transactions, abstract data-types, transactional memory, transactional boosting, commutativity, movers

1. Introduction

Transactional Memory (TM) comes in many varieties: hardware [13], software [7, 10, 28], and mixtures of both [6, 24]. Common to all these approaches is the need to define and detect synchronization conflicts: circumstances under which two transactions cannot proceed in parallel. Existing TM systems detect conflicts at a *fine-grained* level: each transaction maintains a read set, encompassing

the locations it read, and a write set, encompassing the locations it wrote. Two transactions are deemed to conflict if one's write set intersects the other's read or write set. Defining a synchronization conflict to be read/write conflict is appealing because it is usually easy to classify operations, whether in hardware or software, as reads or writes.

Nevertheless, there is an increasing realization that read/write conflicts support insufficient levels of concurrency for objects subject to contention. Conflict detection on the basis of read/write sets is conservative: two high-level operations may not logically conflict even though they share memory locations. For example, if two threads insert distinct items in a shared collection, then, logically, neither one should be affected by the other's actions. But if the collection type is implemented, say, as a linked list, then their read and write sets may inadvertently intersect in a harmless way, producing an unnecessary *false conflict*.

In response we have suggested replacing this low-level notion of read/write conflict with a high-level notion of conflicts between methods of abstract data-types, leveraging sequential method semantics [11]. Moreover, others have begun developing automated techniques to discover high-level semantics [2, 8, 27]. As we have shown, the resulting TMs often offer performance improvements of at least an order of magnitude.

At present, despite decades of previous work in the database community, there is no formal account for these high-level approaches in the context of programming languages. Consequently, the literature tends to be ad-hoc and numerous questions are open: What execution models are possible? How do these high-level approaches relate to traditional read/write-based transactional memory? What, precisely, are the requirements imposed on abstract data-types? Would a unified theory uncover other opportunities for improvement in the state-of-the-art?

In this paper we unify these ad-hoc techniques into a theory we call *coarse-grained transactions*, generalizing transactional memory with a semantics that answers these open questions. A coarse-grained transaction consists of operations on abstract data-types rather than simple memory read/write operations. We present both optimistic and pessimistic execution semantics of coarse-grained transactions, and show that both abide the standard atomic semantics. Rather than using a general notion of commutativity, we precisely characterize the requirements of operations in terms of left-movers and right-movers [22]. We find that the pessimistic execution requires that outstanding operations be left-movers with respect to each other, whereas the optimistic execution requires that operations be right-movers with respect to committed operations.

We then discuss how our semantics can be used to model a variety of implementation details in the literature. Our semantics encompasses both fine-grained operations (e.g., memory read/write) and coarse-grained operations (e.g., transactional boosting [11]). Our pessimistic execution semantics model, again, transactional boosting [11], whereas the optimistic execution semantics model STM implementations with lazy update strategies [7, 28] as well

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

POPL'10, January 17–23, 2010, Madrid, Spain.
Copyright © 2010 ACM 978-1-60558-479-9/10/01...\$10.00

as recent work which resembles coarse-grained transactions [3]. Finally we discuss our treatment of a variety of other issues in the literature including (closed) nested transactions [25], checkpoints [17] and deadlock detection [18].

Example

```

T1: atomic {
  x := 3;
  z++;
}
||
T2: atomic {
  z++;
  y := 4;
}

```

The transactions T1 and T2 in the example above execute a series of statements which read or write values from locations on the shared heap (x, y, z). If executed concurrently the transactions will conflict because both write to z , and one will be aborted.

However, if we assume incrementing z is a single instruction, notice any execution is serializable. This observation stems from the fact that increments to z *commute*: applying them in either order yields the same final state. Therefore, an interleaved execution:

```

T1: | (x:=3)          (z++) (cmt)
T2: |          (z++) (y:=4) (cmt)

```

is equivalent to a serial execution which preserves the commit order:

```

T1: | (x:=3) (z++) (cmt)
T2: |          (z++) (y:=4) (cmt)

```

This serial order can be obtained by reasoning about moverness [22] of the operations performed by transactions. A read(a) operation can move to the left or right of a write($b:=e$) operation (from distinct transactions) when a and b are distinct and a is not free in e . For example, to construct the serial history above, we moved the operation T2:($y:=4$) to the right of T1:($z++$). Moreover, incrementing z can be applied in either order and reach the same final state.

Coarse-Grained Transactions We saw above that two operations which each increment a variable commute, i.e. can move to the left or to the right of each other. This illustrated how the imprecise conflict detection in traditional STM implementations is unable to detect the serializability of some executions at the *fine-grained* level of memory options.

Now consider a higher level of abstraction, where transactions perform *coarse-grained* operations. In the following example two transactions perform operations on a data structure:

```

global SkipList Set;
T1: atomic {
  Set.add(5);
  Set.remove(3);
}
||
T2: atomic {
  if (Set.contains(6))
  Set.remove(7);
}

```

This illustrates the changing focus in the transactional memory community. The old read/write operations are replaced with arbitrary, coarse-grained operations. A transaction consists of operations (represented as *object methods*) on regions of the heap (represented as *object state*) which we will assume to be disjoint. The shared state above is a *Set* data structure, in this case implemented as a (linearizable) skip-list. Each transaction then proceeds by performing *Set* method calls: insert(x), remove(y), contains(z).

Now at this coarse-grained level, how can operations be synchronized? Because high-level operations may share memory locations, using a fine-grained synchronization technique (detecting read/write conflicts) impairs performance [11]. The solution is to instead base synchronization on how one method moves with respect to another. In the example above—assuming *SkipList* is linearizable—any of the methods can move to the left or right of any of the other methods provided their arguments are distinct. In

other cases some operations may only move in one direction. For example, a semaphore increment operation can move to the left of a decrement operation but not vice-versa.

Recently, ad-hoc techniques have appeared in the literature which appear to be exploiting left-movers and right-movers to some degree, yet they lack a formal footing [3, 11, 20]. In this paper we show how these and future proposals can be characterized in a semantics based on left-movers and right-movers.

Semantics In the remainder of this paper we give a semantics for a generalized notion of transactions, which we call *coarse-grained transactions*, where transactions perform abstract data-type operations rather than simple memory read/write operations. Despite having a single atomic semantics, there are two execution semantics and each imposes different requirements on the moverness of the transactional operations:

1. *Pessimistic Execution Semantics* (Section 3.1) – Transactions will be delayed from taking steps in the semantics if the operations they intend to perform do not commute with the outstanding operations of other active transactions.

This semantics requires that operations be *left-movers with respect to other uncommitted transactions*.

This semantics is subject to *deadlock* [11] though recovery is possible [18] via *inverse* operations. This left-mover requirement also permits a transaction to perform an inverse operation that annihilates the transaction’s most recent operation. Since all the pending operations to the right must be left movers, the most recent operation can be moved to the far right.

2. *Optimistic Execution Semantics* (Section 3.2) – The shared state is logically duplicated into transaction-local state and, thus, method call steps in the semantics are always enabled. However, when transactions attempt to commit, they may be required to abort if they have performed method calls that do not commute with other concurrent transactions. Otherwise transaction-local changes are merged into the shared state.

This semantics requires that operations be *right-movers with respect to operations of other committed transactions*. This semantics includes traditional read/write software transactional memory with lazy update strategies [7, 28] as well as recent alternatives which involve user-defined conflict resolution [3]. This semantics is also subject to *livelock* (as is STM).

We show (Section 5) that both of the above high-level semantics are equivalent to atomic semantics. We then discuss (Sections 6, 7) how the semantics applies to numerous TM implementation details discussed widely in the literature.

Limitations Coarse-grained transactions have two other requirements not addressed here. First, operations must be linearizable—a property guaranteed in traditional TM since read/write operations are typically atomic. Second, for simplicity we require that moverness take isolation into consideration. If two objects share memory one may not be able to reason about them independently. Type systems or alias analyses could potentially provide such guarantees.

2. Semantics

2.1 Preliminaries

We will use the following language:

```

c ::= c | e := e | if b then c else c
s ::= s | c | beg; t; end
t ::= t | c | o.m

```

$$\begin{array}{c}
\frac{}{\langle \perp, c; c' \rangle, T, \sigma_{\text{sh}} \xrightarrow{\perp}_A \langle \perp, c' \rangle, T, \llbracket c \rrbracket \sigma_{\text{sh}}} \\
\frac{}{\langle \perp, s; s' \rangle, T, \sigma_{\text{sh}} \xrightarrow{\perp}_A \langle \perp, s' \rangle, T, \llbracket s \rrbracket \sigma_{\text{sh}}} \\
\frac{\langle \tau, t \rangle, T, \sigma_{\text{sh}} \xrightarrow{*}_A \langle \tau, \text{skip} \rangle, T, \sigma'_{\text{sh}}}{\langle \perp, \text{beg}; t; \text{end} \rangle, T, \sigma_{\text{sh}} \xrightarrow{(\tau, \text{beg}) \cdot (\tau, o.m) \cdot (\tau, \text{end})}_A \langle \perp, \text{skip} \rangle, T, \sigma'_{\text{sh}}}
\end{array}$$

Figure 1: Abridged Atomic semantics \xrightarrow{A}

The language consists of standard commands c such as assignment and branching, statements s which can be commands or initiate transactions, and transactions t which can be commands or coarse-grained operations (represented as method calls m on shared objects o). Expressions e are standard and involve variables, constants, algebraic expressions, λ -abstraction and application. Boolean expressions b are similar. Thread-local state can be modeled as shared objects that are only touched by one thread. For simplicity we assume a parallel composition rather than an explicit treatment of `fork`.

Transactions operate on a shared store of objects, manipulated with their corresponding methods. We denote the shared store σ_{sh} , and it contains shared objects, each denoted o . The shared store σ_{sh}^0 is defined as:

$$\sigma_{\text{sh}}^0 \triangleq \{\forall o \in \text{objects}(P). o \mapsto \text{init}(o)\}$$

where `objects` is the obvious function which returns the set of all objects used by the program P , and `init` returns the initial state of the given object.

The language has the abridged atomic semantics given in Figure 1 (the full atomic semantics with all reductions and a detailed treatment of the labeling is given in our Technical Report [19]). The semantics is a labeled transition system. A configuration is a pair T, σ_{sh} representing the transactions T and the shared store of objects σ_{sh} . The initial configuration is a parallel composition of threads: $C_0^a = \{T_0, T_1, \dots, T_n\}, \sigma_{\text{sh}}^0$ where each $T_i = \langle \perp, P_i \rangle$ is a tuple of transaction identifier (initially \perp) and the program text for the i th transaction. Steps are taken when one element of T is selected and any enabled reduction is taken. The domain of labels is $\perp \cup (\mathbb{N} \times (o.m \cup \{\text{beg}, \text{cmt}, \text{abt}\}))$, where the first element is the transaction identifier, and the second element is either the operation performed or a symbolic representation of transaction begin, commit or abort.

The above semantics is atomic: all operations performed by a transaction are applied to the shared state in one step, in which the transition is labeled with each of the transaction's operations. In the remainder of this paper we present two execution semantics, one pessimistic and one optimistic, which impose distinct requirements on the shared objects. We then show that both are atomic (and, in Section 6.1, opaque [9]).

2.2 Moverness

We now formally define the notions of left- and right-moverness which is due to Lipton [22]. We will use moverness to characterize the necessary restrictions placed on the object in the pessimistic and optimistic semantics.

Definition 2.1. (*Left Mover*) $o.m$ is a left-mover with $p.n$, denoted $o.m \triangleleft p.n$, if and only if for any state σ

$$\begin{aligned}
& \{\sigma'' \mid \exists \sigma'. \sigma \xrightarrow{p.n} \sigma' \wedge \sigma' \xrightarrow{o.m} \sigma''\} \\
& \subseteq \\
& \{\sigma''' \mid \exists \sigma'. \sigma \xrightarrow{o.m} \sigma' \wedge \sigma' \xrightarrow{p.n} \sigma'''\}
\end{aligned}$$

This definition gives the conditions under which one step in the semantics is a left mover with respect to another step. Specifically, in a given execution where two methods are executed, the condition is that the set of states reached is included in the set of possible states reached had the methods occurred in the opposite order. Of course when o is a distinct object from p , then for any two methods, the two are both movers with respect to each other:

It is important to note that in Definition 2.1, the relationship is an inclusion rather than an equivalence. The left-hand side of the inclusion is typically a concrete execution where specific states are known. By contrast, the right-hand side must account for non-determinism in the sequential specifications of the objects. For example, let o be a priority queue, and the method m be the removal of the highest-priority element. If there are multiple elements in the queue with the highest priority, then the semantics of the queue may state that any of them will be returned. Thus, σ' in the right-hand side is a *set* of states, as is the entire right-hand side.

We will later show that inclusion is sufficient for proving serializability of execution traces. Whereas previous proofs of serializability were based on finding *one* equivalent sequential execution for an interleaved execution, our proof technique instead shows that a given interleaved execution is contained within a set of executions, all of whom are sequential.

Definition 2.2. (*Right Mover* \triangleright , *Both Mover* \bowtie) Defined similarly.

Definition 2.3. (*Lifted Movers*) We define left mover lifted to sets, denoted $\bar{\triangleleft}$, as follows:

$$M_1 \bar{\triangleleft} M_2 \triangleq \forall (o.m, p.n) \in (M_1, M_2). o.m \triangleleft p.n$$

Lifted versions of \triangleright, \bowtie are similar and denoted $\bar{\triangleright}, \bar{\bowtie}$.

3. Execution Semantics

In this section we present the two high-level semantics—one for pessimistic execution and one for optimistic execution. These semantics illustrate how different execution models place different requirements on the moverness of the coarse-grained operations. In the subsequent sections we will show that traces in either semantics are serializable, matching the atomic semantics given in Figure 1.

3.1 Pessimistic Execution Semantics

Configurations in the pessimistic execution semantics are similar to the atomic semantics given in Section 2.1. However, the elements of T are of the form $\langle \tau, s, M, \sigma_\tau \rangle$, where $\tau \in \mathbb{N}_\perp$ identifies the transaction if one is currently running, s denotes the program code, M denotes the sequence of object methods that τ has applied, and σ_τ is the transaction-local state. The initial configuration is a parallel composition of threads: $C_0^p = \{T_0, T_1, \dots, T_n\}, \sigma_{\text{sh}}^0$, where each $T_i = \langle \perp, P_i, \emptyset, \sigma^0 \rangle$ and σ_{sh}^0 is as defined in the previous section.

The relevant reductions for pessimistic execution are given in Figure 2 (for the complete semantics, see our Technical Report [19]). The PCMD step involves a thread executing a statement which impacts only thread-local state and control-flow. In the PBEG rule the system makes a transition in which a thread begins to execute a transaction. The effect is that the sequence of applied operations M is initialized to the empty sequence and the new transaction tuple is added to the set of active transactions. $\tau \in \mathbb{N}$ and the predicate `fresh`(τ) indicates that τ is fresh and strictly increasing (the monotonicity requirement is unneeded here, but required for the optimistic execution semantics discussed next, which use τ as a timestamp).

$$\begin{array}{c}
\frac{}{\langle \tau, s \in \{:=, \text{if}\}; s', M, \sigma_\tau \rangle, T, \sigma_{\text{sh}} \xrightarrow{P} \langle \tau, s', M, \llbracket s \rrbracket \sigma_\tau \rangle, T, \sigma_{\text{sh}}} \text{PCMD} \\
\frac{}{\langle \perp, \text{beg}; s, M, \sigma_\tau \rangle, T, \sigma_{\text{sh}} \xrightarrow{P} \langle \text{fresh}(\tau), s, \llbracket \cdot \rrbracket, \sigma_\tau \rangle, T, \sigma_{\text{sh}}} \text{PBEG} \\
\frac{\{o.m\} \nabla \text{meths}(T)}{\langle \tau, x := o.m; s, M, \sigma_\tau \rangle, T, \sigma_{\text{sh}} \xrightarrow{P} \langle \tau, s, M :: (\text{"}x := o.m; s\text{"}, \sigma_\tau, \text{"}o.m\text{"}), \sigma_\tau[x \mapsto rv(\llbracket o \rrbracket \sigma_{\text{sh}}.m)] \rangle, T, \sigma_{\text{sh}}[o \mapsto \llbracket o \rrbracket \sigma_{\text{sh}}.m]} \text{PAPP} \\
\frac{}{\langle \tau, s, M :: (s', \sigma'_\tau, \text{"}o.m\text{"}), \sigma_\tau \rangle, T, \sigma_{\text{sh}} \xrightarrow{P} \langle \tau, s', M, \sigma'_\tau \rangle, T, \sigma_{\text{sh}}[o \mapsto \llbracket o \rrbracket \sigma_{\text{sh}}.m^{-1}]} \text{PUNDO} \\
\frac{}{\langle \tau, \text{end}; s, M, \sigma_\tau \rangle, T, \sigma_{\text{sh}} \xrightarrow{P} \langle \perp, s, \emptyset, \sigma_\tau \rangle, T, \sigma_{\text{sh}}} \text{PCMT} \\
\begin{array}{l}
\text{meths}(\emptyset) = \emptyset \\
\text{meths}(T' \cup \{\langle \tau, C, M, \sigma_\tau \rangle\}) = \text{meths}_M(M) \cup \text{meths}(T') \\
\text{meths}_M(\llbracket \cdot \rrbracket) = \emptyset \\
\text{meths}_M(M' :: (s, \sigma_\tau, \text{"}o.m\text{"})) = \text{meths}_M(M') \cup \{\text{"}o.m\text{"}\}
\end{array}
\end{array}$$

Figure 2: Pessimistic Semantics \rightarrow_P

The PAPP rule allows a transaction to apply a method m on a shared object o . The rule indicates the initiation of a method call, which is appended to the sequence of operations. The log also includes the current state and program counter, which may be used later in the PUNDO rule. The return value (given by rv) is stored in the thread-local variable x . (Note that we assume the shared object is linearizable.)

In these pessimistic execution semantics synchronization is considered throughout a transaction rather than simply at commit-time. Thus the PAPP step involves a check. The commit order of the active transactions is not determined *a priori*, so any active transaction may commit at any time. Therefore there is a restriction for the reduction: the new method call must be a left-mover with respect to all other active method calls. Left-moverness ensures that the current transaction could commit before all of the others.

Since synchronization is considered during the PAPP rules, the commit rule PCMT is trivial. We will see in the next section that these pessimistic semantics stand in contrast to the optimistic semantics in which the PAPP rule is always enabled, but a synchronization check is performed in the commit rule.

The commit rule PCMT merely involves discarding the set of active methods and terminating the transaction. All effects have already been applied to the shared state.

Finally the PUNDO rule allows a transaction to move backwards to the program location immediately preceding the most recently applied object method. Longer jumps are possible by multiple uses of the PUNDO rule. PUNDO is always enabled when M is non-empty. The PUNDO rule (as we discuss in Section 6) is useful to allow transactions to escape from deadlock [18].

There is no moverness requirement for the PUNDO operation—there is a subtle reason why it is unneeded. As we show formally in Section 5.1, the “doomed” operation that the transaction is undoing is, in effect, a right-mover since all subsequent operations from other transactions are left-movers. Thus, the doomed operation can be moved to the far right and then annihilated with the inverse operation.

3.2 Optimistic Execution Semantics

We now present the optimistic execution semantics. We then, in Section 3.3, compare it to the pessimistic counterpart. The most sig-

nificant difference between the two is that the optimistic semantics ensure synchronization at commit-time, whereas the pessimistic semantics ensure synchronization with each operation.

Before we continue, we emphasize that even though our model involves thread-local copies of the shared state, in many cases this is not a performance problem since (i) copies are unneeded when operations are read-only and (ii) copy-on-write can be used elsewhere. The existing TL2 [7] and Intel [28] implementations are evidence of this fact.

Configurations in the optimistic semantics are a triple $T, \sigma_{\text{sh}}, \ell_{\text{sh}}$ where T is a set of per-transaction state, σ_{sh} is the same as the pessimistic shared state, and ℓ_{sh} is a log of pairs: each committed transaction and the operations it has performed. Elements of T are of the form $\langle \tau, s, \sigma_\tau, \overleftarrow{\sigma}_\tau, \ell_\tau \rangle$. τ is a unique transaction identifier, s is the next code to be executed, σ_τ is the transaction-local copy of the shared objects, $\overleftarrow{\sigma}_\tau$ is a back-up copy of the state (used for aborting), and ℓ_τ is a log of the operations performed during a transaction (to be used to check moverness at commit-time). The initial configuration is a parallel composition of threads: $C_0^{\text{opt}} = \{T_0, T_1, \dots, T_n\}, \sigma_{\text{sh}}^0, \llbracket \cdot \rrbracket$, where each $T_i = \langle \perp, P_i, \sigma^0, \sigma^0, \llbracket \cdot \rrbracket \rangle$.

Labeled transitions occur when an element of T is selected and one of the reductions given in Figure 3 applies (for the complete semantics, see our Technical Report [19]). Like PCMD, the OCMD reduction merely applies changes to the local state, and leaves $\overleftarrow{\sigma}_\tau$ unchanged. The OBEG reduction represents the initiation of a transaction. A new identifier is generated, and the existing thread-local state is copied to $\overleftarrow{\sigma}_\tau$ to be used if the transaction is later aborted. A snapshot $\text{snap}(\sigma, \sigma_{\text{sh}})$ is taken of the current state of all shared objects, duplicating them from the shared store to the local store. The current s is also stored in $\overleftarrow{\sigma}_\tau$.

In the reduction OAPP, a transaction applies a method. The operation is applied to the local store (containing copies of shared objects), and the operation is logged so it can be replayed on the shared state later at commit. $\overleftarrow{\sigma}_\tau$ is unchanged.

The OCMT reduction first identifies this set of recently committed transactions by comparing the timestamp of committed transactions with the begin timestamp of the current transaction. The reduction is only taken if all of the transaction’s methods are right-movers with respect to methods of the recently committed transactions. If the reduction is taken, the thread-local operations from ℓ_τ

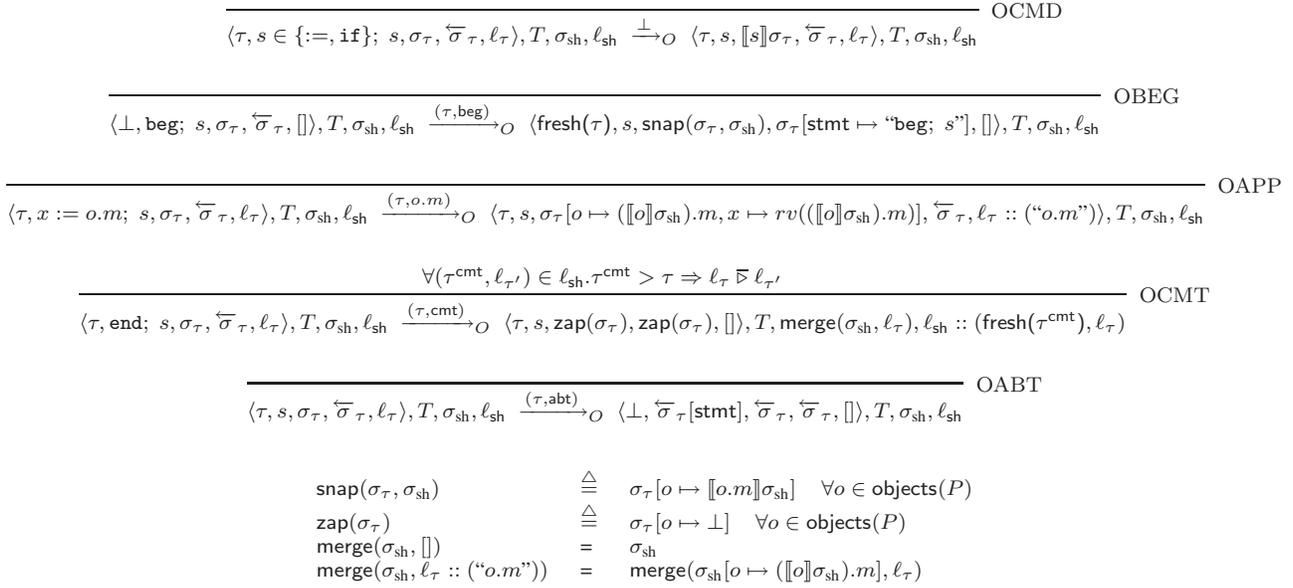


Figure 3: Optimistic Semantics \rightarrow_O

are propagated to the shared state σ_{sh} via the merge function, and a new timestamp (denoted $\text{fresh}(\tau^{\text{cm}})$) is appended to the commit log. The reduction then clears the thread local state, restoring it to $\sigma_\tau^0 \triangleq [\cdot \mapsto \perp]$, clearing the thread-local log ℓ_τ and discarding the unused $\overleftarrow{\sigma}_\tau$.

The last rule OABT corresponds to aborting a transaction. The saved thread-local store $\overleftarrow{\sigma}_\tau$ restored, code (i.e. program counter) is restored via $\overleftarrow{\sigma}_\tau$, and the log is discarded.

Remark: Partial Aborts Note that by contrast to the pessimistic version, this semantics does not allow partial aborts [17]. Adding such a facility can certainly be accomplished, but for ease of presentation, and since no existing implementations provide it, we have omitted it.

3.3 Comparison

As noted above, the most significant distinction between the two semantics is when synchronization occurs. The pessimistic semantics perform a check (left-moverness) with each coarse-grained operation, whereas the optimistic semantics perform a check (right-moverness) at commit-time. Intuitively, the pessimistic left-moverness preserves the invariant that any live transaction *can commit at any time*. By contrast, the optimistic right-moverness checks to see if a transaction’s *thread-local store is logically consistent* given that other transactions may have committed since the copy was made. There are several considerations involved in the choice between execution semantics:

- **Aborts** – Optimistic transactions can always proceed with operations immediately modifying their thread-local store, however they may need to be aborted. Pessimistic transactions are always able to commit (modulo deadlock) but may have to be delayed before applying an operation.
- **State Duplication** – The optimistic semantics require, at least logically, thread-local duplication of the shared state. Sometimes duplication can be reduced (e.g. with copy-on-write strategies) or even eliminated (e.g. when a return value is unneeded or can be otherwise calculated). Alternatively, the pessimistic semantics does not require duplication.

- **Merge** – The optimistic semantics involve a commit-time merge of operations into the shared state. In the worst case this may involve replaying the entire operation, perhaps favoring the pessimistic semantics. However, there may be instances (such as in read/write STMs) where merge can be implemented efficiently.
- **Inverses** – When transactions are aborted, the pessimistic semantics performs inverse operations (PUNDO). If these operations are expensive or difficult to define, the optimistic semantics may be preferred, which simply discards the thread-local scratch pad.
- **Objects** – In some applications, left-movers are more common than right-movers or vice-versa. In such instances the corresponding semantics (pessimistic or optimistic, respectively) is more amenable because more concurrency is possible.

Some examples of the two semantics are as follows. An example of the optimistic semantics are the TL2 [7] STM implementation and Intel STM compiler [28], which we discuss in Section 6. In each, locks are used at commit-time, ensuring the left-moverness of pending updates. Another example of the optimistic semantics, discussed in Section 7.2, is recent work [3] in which operations are more coarse-grained and conflicts are resolved by programmer-provided rules. Finally, transactional boosting [11] is an example of pessimistic coarse-grained transactions. We discuss boosting in Section 7.1. A side-by-side empirical evaluation of pessimism versus optimism is left for future work.

Remark One might wonder whether the optimistic semantics are contained within the pessimistic semantics (or vice-versa). Unfortunately this is not the case because of the fundamental differences described earlier, and leads to a more-or-less disjoint union of the two semantics. Whereas the pessimistic semantics maintain an commutativity invariant (via the PAPP reduction), the optimistic semantics check at commit. The two semantics have different kind of logs: optimistic transactions have “intention logs” (ℓ_τ) where as pessimistic transactions have “undo logs” (M). One might try to think of the intention logs as being empty for pessimistic semantics and vice-versa, but then neither is contained within the other.

4. Execution Traces

In this section we define execution traces and prove several important properties of them. These properties will be used in Section 5 to show that pessimistic and optimistic execution models are both serializable.

Recall that labels are given by the reductions in the semantics (either \rightarrow_P or \rightarrow_O). They are taken from the domain $\perp \cup (\mathbb{N} \times (o.m \cup \{\text{beg}, \text{cmt}, \text{abt}\}))$, representing transaction commencement, completion and abortion, as well as applying operations on the shared objects.

Definition 4.1. (*Execution Trace*) A trace ϵ is a series of labels:

$$\epsilon = \mathcal{L}_0 \cdot \mathcal{L}_1 \cdot \mathcal{L}_2 \cdot \dots$$

We denote an empty execution trace as \circ .

Definition 4.2. (*Finals*) fin is an inductively defined predicate, mapping a set of initial configurations and an execution trace (sequence of labels) to a set of final configurations:

$$\begin{aligned} [\text{fin } \mathbf{C} \circ] &\triangleq \mathbf{C} \\ [\text{fin } \mathbf{C} \mathcal{L} \cdot \epsilon] &\triangleq [\text{fin } \{C' \mid \forall C \in \mathbf{C}. C \xrightarrow{\mathcal{L}} C'\} \epsilon] \end{aligned}$$

For shorthand, we often use a single configuration as the first argument to fin , though we mean the usual singleton set. Observe that we can always expand subsequences of any execution:

Property 4.1. (*Containment Folding*)

$$[\text{fin } \mathbf{C} \epsilon \cdot \epsilon'] = [\text{fin } [\text{fin } \mathbf{C} \epsilon] \epsilon']$$

Proof: Proof by induction on ϵ . For details see our Technical Report [19].

Lemma 4.1. (*Lifting containment*)

$$\mathbf{C} \subseteq \mathbf{C}' \Rightarrow [\text{fin } \mathbf{C} \mathcal{L}] \subseteq [\text{fin } \mathbf{C}' \mathcal{L}]$$

Proof: Unrolling fin once and then proof by contradiction. For details see our Technical Report [19].

Definition 4.3. (*Movers lifted to labels*)

$$\mathcal{L}_1 \hat{\triangleleft} \mathcal{L}_2 \iff \forall C. [\text{fin } C \mathcal{L}_2 \cdot \mathcal{L}_1] \subseteq [\text{fin } C \mathcal{L}_1 \cdot \mathcal{L}_2]$$

Right- and both-movers $\hat{\triangleleft}$ and $\hat{\triangleleft} \hat{\triangleleft}$ are similarly lifted.

Property 4.2. $o.m \triangleleft p.n \iff (\tau, o.m) \hat{\triangleleft} (\tau', p.n)$ and similar for $\hat{\triangleleft}$ and $\hat{\triangleleft} \hat{\triangleleft}$.

Proof: Correspondence is definitional. For details see our Technical Report [19].

Definition 4.4. (*Execution containment*)

$$\epsilon \sqsubseteq \epsilon' \triangleq \forall C. [\text{fin } C \epsilon] \subseteq [\text{fin } C \epsilon']$$

Lemma 4.2.

$$\epsilon_a \sqsubseteq \epsilon_b \Rightarrow \epsilon_a \cdot \epsilon' \sqsubseteq \epsilon_b \cdot \epsilon'$$

Proof: By induction on ϵ' . For details see our Technical Report [19].

4.1 Serial Traces

Recall that execution traces are given as series of labels in the transition systems of the appropriate operational semantics. For example, here is a trace of two transactions:

$$\epsilon = (\tau, \text{beg}) \cdot (\tau, o.m) \cdot (\tau', \text{beg}) \cdot (\tau', o.m) \cdot (\tau', \text{cmt}) \cdot (\tau, \text{cmt})$$

The following atomicity predicate indicates whether a transaction occurs within an execution trace without any interleaved labels from other transactions.

Definition 4.5. (*Serial Trace*) A serial trace is given by the following predicate, inductively defined over an execution ϵ :

$$\begin{aligned} \frac{\tau, \tau \cup \{\tau\} \vdash \text{serial } \epsilon}{\perp, \tau \vdash \text{serial } (\tau, \text{beg}) \cdot \epsilon} \quad \tau \notin \tau & \quad \frac{\perp, \tau \vdash \text{serial } \epsilon}{\tau, \tau \vdash \text{serial } (\tau, \text{cmt}) \cdot \epsilon} \\ \frac{\tau, \tau \vdash \text{serial } \epsilon}{\tau, \tau \vdash \text{serial } (\tau, o.m) \cdot \epsilon} & \quad \frac{}{\perp, \tau \vdash \text{serial } \circ} \end{aligned}$$

This definition states that a trace is serial if each transaction appears in the trace without any interleaved events from other transactions. We now show that the atomic semantics given in Figure 1 are serializable:

Lemma 4.3. (*Atomic Serializability*)

$$\text{serial } \epsilon \quad \forall \epsilon = C_0^a \rightarrow_A^* C^a$$

Proof: Trivial since all executions are of the form

$$\begin{aligned} &(\tau, \text{beg}) \cdot (\tau, o.m) \cdot (\tau, p.n) \cdot \dots \cdot (\tau, \text{cmt}) \\ &(\tau', \text{beg}) \cdot (\tau', o.m) \cdot (\tau', p.n) \cdot \dots \cdot (\tau', \text{cmt}) \\ &\dots \end{aligned}$$

For details see our Technical Report [19].

In the following lemma, let \mathcal{L}_τ^i denote an operation corresponding to transaction τ .

Lemma 4.4.

$$\text{serial } \epsilon \Rightarrow \text{serial } \epsilon \cdot (\tau, \text{beg}) \cdot \mathcal{L}_\tau^0 \cdot \dots \cdot \mathcal{L}_\tau^n \cdot (\tau, \text{cmt})$$

when none of the labels in ϵ correspond to τ .

Proof: By derivation over serial . For details see our Technical Report [19].

5. Serializability

5.1 Pessimistic Execution Model

We now give our serializability result for the pessimistic semantics. The structure of our proof of serializability of executions is standard. We show that for any execution of the semantics, there exists an equivalent execution in which all transactions appear in a serial order (without interleavings of operations from other transactions).

Left-moverness Invariant

An important part of the correctness of the pessimistic semantics is the fact that an invariant is maintained about method calls of all active transactions. This is a result of the PAPP reduction. We define a predicate on a configuration which holds if all operations of uncommitted transactions are left-movers with respect to all operations of other uncommitted transactions. Note that we use the notation \mathcal{L}_τ to mean any label involving τ , i.e. $(\tau, _)$.

Definition 5.1. (*Active Left-Movers*)

$$\text{comm}_\triangleleft \triangleq \lambda \epsilon. \forall \mathcal{L}_\tau, \mathcal{L}_{\tau'} \in \epsilon. \mathcal{L}_\tau \hat{\triangleleft} \mathcal{L}_{\tau'} \text{ where } \tau \neq \tau'$$

Lemma 5.1.

$$\text{comm}_\triangleleft \epsilon \Rightarrow \forall \tau. \epsilon \sqsubseteq \epsilon_\tau \cdot \epsilon_{U \setminus \tau}$$

Proof: Double induction (i) showing that for the first τ -label $\mathcal{L}_\tau = (\tau, _)$ in ϵ , that $\text{comm}_\triangleleft \epsilon \Rightarrow \epsilon_a \mathcal{L}_\tau \epsilon_b \sqsubseteq \mathcal{L}_\tau \epsilon_a \epsilon_b$ and (ii) showing that $\text{comm}_\triangleleft \epsilon \Rightarrow \epsilon \sqsubseteq \epsilon_\tau \cdot \epsilon_{U \setminus \tau}$ by induction on the sequence of τ -labels occurring in ϵ . For details see our Technical Report [19].

Definition 5.2. (Active transactions) The function *active* returns the set of transaction identifiers in ϵ which have not committed:

$$\begin{aligned} \text{active}(\circ, A) &\triangleq A \\ \text{active}((\tau, \text{beg})\cdot\epsilon, A) &\triangleq \text{active}(\epsilon, A \cup \{\tau\}) \\ \text{active}((\tau, \text{cmt})\cdot\epsilon, A) &\triangleq \text{active}(\epsilon, A \setminus \{\tau\}) \\ \text{active}((\tau, \text{abt})\cdot\epsilon, A) &\triangleq \text{active}(\epsilon, A \setminus \{\tau\}) \\ \text{active}((\tau, _)\cdot\epsilon, A) &\triangleq \text{active}(\epsilon, A) \end{aligned}$$

Definition 5.3. (Live methods) The function *live*(ϵ, τ) takes an execution ϵ and set of active transactions τ and returns the set of operations performed by transactions which have not committed:

$$\begin{aligned} \text{live}(\circ, \tau) &\triangleq \emptyset \\ \text{live}((\tau, o.m)\cdot\epsilon, \tau) &\triangleq \begin{cases} \{o.m\} \cup \text{live}(\epsilon, \tau) & \text{if } \tau \in \tau \\ \text{live}(\epsilon, \tau) & \text{otherwise} \end{cases} \\ \text{live}((_, _)\cdot\epsilon, \tau) &\triangleq \text{live}(\epsilon, \tau) \end{aligned}$$

Note that the (τ, abt) is not used in the pessimistic semantics, but will appear in the optimistic semantics. The following predicate holds if both an execution trace and a configuration agree on the set of all operations of uncommitted transactions.

Definition 5.4. (Uncommitted consistency)

$$\text{consist}(\epsilon, T) \triangleq \text{meths}(T) = \text{live}(\epsilon, \text{active}(\epsilon, \emptyset))$$

Moverness in Pessimistic Semantics

We now show some moverness relationships which are specific to the pessimistic semantics and used throughout the proof.

Lemma 5.2.

$$(\tau, \text{beg}) \hat{\bowtie} (\tau', _)$$

$$(\tau, \text{cmt}) \hat{\bowtie} (\tau', _)$$

Proof: By showing that reordering the execution leads to the same final configuration *fin*. For details see our Technical Report [19].

Note that (τ, cmt) is not a right-mover with $(\tau', _)$. The proof would fail in the final case for Eqn. 2 where

$$[\text{fin } C (\tau', o.m)\cdot(\tau, \text{cmt})] = \emptyset$$

Moving an operation to the left of a commit may lead to an execution which is not valid since the antecedent of the PAPP rule may not hold. Informally, the transaction performing the operation may need to wait until the other commits.

In the following let $\epsilon_{\neg\tau}$ denote an execution which contains no labels corresponding to τ .

Corollary 5.3. (To Lemma 5.1) In the pessimistic semantics,

$$\text{comm}_{\triangleleft} \epsilon_{\neg\tau} \Rightarrow \epsilon_{\neg\tau} \cdot (\tau, \text{cmt}) \sqsubseteq (\tau, \text{cmt}) \cdot \epsilon_{\neg\tau}$$

Proof: This is an immediate consequence of Lemma 5.1, noting the moverness of (τ, cmt) given by Equation 2.

Lemma 5.4. (Implied right-moverness)

$$\text{comm}_{\triangleleft} \epsilon \Rightarrow \epsilon \cdot (\tau, o.m) \cdot \epsilon_{\neg\tau} \sqsubseteq \epsilon_{\neg\tau} \cdot (\tau, o.m)$$

Proof: Trivial induction. For details see our Technical Report [19].

Serializability

The next theorem shows that executions of the pessimistic semantics are strictly serializable. The theorem proceeds by showing inclusion in another execution of two parts: a prefix in which all com-

mitted transactions appear in a serial order (using *serial*) and a suffix of uncommitted operations in which $\text{comm}_{\triangleleft}$ is invariant. We must also maintain as invariant the consistency ($\text{consist}(\epsilon, C^P.T)$) between operations of uncommitted transactions in the execution trace and operations of uncommitted transactions as given by the configuration. Finally, note that when the PUNDO rule is taken, the corresponding doomed operation is moved to the right and then annihilated by the inverse operations.

Theorem 5.5. (Pessimistic Serializability)

For any $\epsilon = C_0^P \xrightarrow{*} C^P$,

$$\exists \epsilon_C \cdot \epsilon_U. \text{serial } \epsilon_C \wedge \text{comm}_{\triangleleft} \epsilon_U \wedge \text{consist}(\epsilon_U, C^P.T) \wedge \epsilon \sqsubseteq \epsilon_C \cdot \epsilon_U$$

Proof: We prove the theorem by induction on the execution trace. constructing an appropriate $\epsilon_C \cdot \epsilon_U$ at each step. Consider the empty execution $\epsilon = \circ$, which starts in configuration C_0^P . We can define $\epsilon_C = \epsilon_U = \circ$. *serial* ϵ_C holds by axiom. $\text{comm}_{\triangleleft} \epsilon_U$ and $\text{consist}(\epsilon_U, C^P.T)$ hold trivially. We now show that

$$\begin{aligned} \exists \epsilon_C \cdot \epsilon_U. \epsilon \sqsubseteq \epsilon_C \cdot \epsilon_U \wedge \text{serial } \epsilon_C \wedge \text{comm}_{\triangleleft} \epsilon_U \wedge \text{consist}(\epsilon_U, C^P.T) \\ \wedge \epsilon \xrightarrow{\mathcal{L}}_P \epsilon' \\ \Rightarrow \\ \exists \epsilon_C' \epsilon_U'. \epsilon' \sqsubseteq \epsilon_C' \epsilon_U' \wedge \text{serial } \epsilon_C' \wedge \text{comm}_{\triangleleft} \epsilon_U' \wedge \text{consist}(\epsilon_U', C^{P'}.T) \end{aligned}$$

We case split on \mathcal{L} .

Case $\mathcal{L} = \perp$: $\epsilon' = \epsilon$. Trivial.

Case $\mathcal{L} = (\tau, \text{beg})$: Let $\epsilon_C' = \epsilon_C$ and $\epsilon_U' = \epsilon_U \cdot (\tau, \text{beg})$. $\epsilon' \sqsubseteq \epsilon_C' \cdot \epsilon_U' \cdot (\tau, \text{beg})$ by Lemma 4.2. *serial* ϵ_C' holds because *serial* ϵ_C holds. $\text{consist}(\epsilon_U', C^{P'}.T)$ holds because no new operations have been performed. Equation 1 maintains the truth of $\text{comm}_{\triangleleft} \epsilon_U'$.

Case $\mathcal{L} = (\tau, o.m)$: Let $\epsilon_C' = \epsilon_C$ and $\epsilon_U' = \epsilon_U \cdot (\tau, o.m)$. $\epsilon' \sqsubseteq \epsilon_C' \cdot \epsilon_U' \cdot (\tau, o.m)$ by Lemma 4.2. *serial* ϵ_C' holds because *serial* ϵ_C holds. By the PAPP rule, $C^{P'}.T = C^P.T \cup \{o.m\}$ and it is easy to see that $\text{live}(\epsilon_U', \emptyset) = \text{live}(\epsilon_U, \emptyset) \cup \{o.m\}$. Thus $\text{consist}(\epsilon_U', C^{P'}.T)$ holds. $\text{comm}_{\triangleleft} \epsilon_U'$ holds because $\text{comm}_{\triangleleft} \epsilon_U$ and PAPP requires that $\{o.m\} \nabla \text{meths}(C^P.T)$. Lifting the antecedent, we have that

$$\begin{aligned} \forall \mathcal{L}' \in \epsilon_U. \mathcal{L} \hat{\bowtie} \mathcal{L}' &\implies \forall \mathcal{L}_1, \mathcal{L}_2 \in \epsilon_U'. \mathcal{L}_1 \hat{\bowtie} \mathcal{L}_2 \\ &\implies \text{comm}_{\triangleleft} \epsilon_U' \end{aligned}$$

Case $\mathcal{L} = (\tau, o.m^{-1})$: Without loss of generality, let $\epsilon_U = \epsilon_a \cdot (\tau, o.m) \epsilon_{\neg\tau}$. By Lemma 5.4 we know that $\epsilon_U \sqsubseteq \epsilon_a \epsilon_{\neg\tau} (\tau, o.m)$. Now we append $(\tau, o.m^{-1})$ to both sides, and then chose $\epsilon_U' = \epsilon_a \cdot \epsilon_{\neg\tau}$.

Case $\mathcal{L} = (\tau, \text{cmt})$: We first show the following inclusion:

$$\begin{aligned} \epsilon' &= \epsilon \cdot (\tau, \text{cmt}) \\ &\sqsubseteq \epsilon_C \cdot \epsilon_U \cdot (\tau, \text{cmt}) && \text{I.H. } \epsilon \sqsubseteq \epsilon_C \cdot \epsilon_U \\ &\sqsubseteq \epsilon_C \cdot \epsilon_{\tau} \cdot \epsilon_{U \setminus \tau} \cdot (\tau, \text{cmt}) && \text{Lemma 5.1} \\ &\sqsubseteq \epsilon_C \cdot \epsilon_{\tau} \cdot (\tau, \text{cmt}) \cdot \epsilon_{U \setminus \tau} && \text{Corollary 5.3} \end{aligned}$$

Now let $\epsilon_C' = \epsilon_C \cdot \epsilon_{\tau} \cdot (\tau, \text{cmt})$. *serial* ϵ_C' holds by Lemma 4.4. Letting $\epsilon_U' = \epsilon_{U \setminus \tau}$, observe that ϵ_U' is a subsequence of ϵ_U , so $\text{comm}_{\triangleleft} \epsilon_U'$ holds.

5.2 Optimistic Execution Model

We now prove serializability of the optimistic model. Unfortunately we cannot reuse the proof of the pessimistic model. The optimistic model is fundamentally different: shared memory is not modified in-place. Instead transactions modify local copies of memory and update the shared state at commit.

The optimistic execution model performs synchronization at *commit-time* rather than synchronizing with each operation. The effect is two-fold: (i) If a transaction aborts then recovery is achieved by simply discarding the transaction-local copy rather than invoking inverse operations. (ii) When a transaction τ commits it must

check that its operations can be moved to the right of operations corresponding to any transaction τ' which has committed after τ began. This is why, for example, in TL2 [7] at commit-time transactions acquire a lock for each modified memory location.

Moverness in Optimistic Semantics

Lemma 5.6. *In the optimistic semantics,*

$$(\tau, beg) \hat{\boxtimes} (\tau', beg) \quad (3)$$

$$(\tau, beg) \hat{\boxtimes} (\tau', o.m) \quad (4)$$

$$(\tau, o.m) \hat{\boxtimes} (\tau', p.n) \quad (5)$$

$$(\tau, cmt) \hat{\boxtimes} (\tau', o.m) \quad (6)$$

$$(\tau, cmt) \hat{\boxtimes} (\tau', beg) \quad (7)$$

$$(\tau, cmt) \cdot (\tau', cmt) \quad (8)$$

$$(\tau, beg) \hat{\boxtimes} (\tau', abt) \quad (9)$$

$$(\tau, abt) \hat{\boxtimes} (\tau', abt) \quad (10)$$

Proof: For details see our Technical Report [19].

The proof of the above Lemma involves showing that reordering the execution leads to the same final configuration fin . Eqn. 3 holds because swapping the order of which transaction begins first has no effect on the eventual commit order. Eqns. 4 and 5 hold because $o.m$ operations are applied to transaction-local copies of the shared state and does not impact concurrent transactions. Eqn. 7 holds because moving (τ, cmt) to the left of (τ', beg) reduces the burden on τ' when it ultimately tries to commit— τ' begins with a newer copy of the shared state and can ignore the effects of τ when scanning the commit log. Eqn. 8 holds because if a commit is moved to the left it may prevent the other transaction from committing since the commit log will be longer. Finally, Eqns. 10 and 9 hold because the OABT rule discards transaction-local changes.

Aborting a Transaction

Lemma 5.7. *In the optimistic semantics,*

$$(\tau, beg) \cdot \epsilon_{\tau} \cdot (\tau, abt) \sqsubseteq \circ$$

where ϵ_{τ} is an execution consisting of labels of the form $(\tau, o.m)$.

Proof: First we show that this holds when $\epsilon_{\tau} = \circ$, which requires us to show that $[\text{fin } C(\tau, beg) \cdot (\tau, abt)] \sqsubseteq [\text{fin } C \circ]$. Or rather that $[\text{fin } C(\tau, beg) \cdot (\tau, abt)] = C$. From the optimistic semantics it is easy to see that the OABT reduction restores the state which held before the OBEG reduction fired.

The lemma can now be proven by a rather trivial induction on ϵ_{τ} . In the base case when $\epsilon_{\tau} = (\tau, o.m)$, the claim above still holds because OAPP1 and OAPP2 only perform changes to transaction-local σ_{τ} and ℓ_{τ} , which are discarded by the OABT reduction. This of course holds for any larger ϵ_{τ} .

Lemma 5.8. *In the optimistic semantics,*

$$\epsilon \cdot (\tau, beg) \cdot \epsilon_U \cdot (\tau, abt) \sqsubseteq \epsilon \cdot \epsilon_U \setminus \tau$$

Proof: Using induction and the moverness give in Eqns. 4 and 6, it is easy to show that

$$\epsilon \cdot (\tau, o.m) \cdot \epsilon_{\neg\tau} \sqsubseteq \epsilon \cdot \epsilon_{\neg\tau} \cdot (\tau, o.m)$$

where $\epsilon_{\neg\tau}$ consists of labels of the form $(\tau', o.m)$ or (τ', beg) . We then show that the set of all n τ -labels in ϵ , that

$$\epsilon \cdot \mathcal{L}_{\tau}^n \cdot \epsilon_{\neg\tau} \cdot \mathcal{L}_{\tau}^1 \cdot \epsilon_{\neg\tau} \sqsubseteq \epsilon \cdot \epsilon_{\neg\tau} \cdot \mathcal{L}_{\tau}^1 \cdot \epsilon_{\neg\tau} \cdot \mathcal{L}_{\tau}^n \cdot \mathcal{L}_{\tau}^1$$

For details see our Technical Report [19].

Serializability

We prove the theorem again by showing that executions are contained within an alternate serial execution. However, there are some differences. Since transactions operate on thread-local state, the $(\tau, o.m)$ labels appear but do not correspond to modifications to the shared state. We are therefore free to discard these labels (so long as a commit label has not appeared). Thus when a transaction takes the OCMT step it can move all of its operations to the right of other committed transactions because of the antecedent and it need not consider other uncommitted transactions (because we temporarily abort them).

The proof depends on the following lemmata. We introduce the notation $\epsilon_{\{\tau, \tau', \dots\}}$ to refer to an execution consisting of operations performed by τ, τ' , or any other transactions. The notations $\epsilon_{\{\tau, \dots\}}$, $\epsilon_{\{\tau, \tau'\}}$ and $\epsilon_{\{\tau', \dots\}}$ are similar.

Lemma 5.9. *(Right of committed) If $\ell_{\tau} \bar{\boxtimes} \ell_{\tau'}$, then*

$$\begin{aligned} & \epsilon \cdot \epsilon_{\{\tau, \tau', \dots\}} \cdot (\tau', cmt) \cdot (\tau, cmt) \cdot \epsilon' \\ & \sqsubseteq \\ & \epsilon \cdot \epsilon_{\{\tau', \dots\}} \cdot (\tau', cmt) \cdot \epsilon_{\{\tau\}} \cdot (\tau, cmt) \cdot \epsilon' \end{aligned}$$

Proof: First, by induction we can move all τ operations to the right within $\epsilon_{\{\tau, \tau'\}}$ because of Eqn. 5. Then each $(\tau, o.m)$ can move to the right of (τ', cmt) because $\ell_{\tau} \bar{\boxtimes} \ell_{\tau'}$. For details see our Technical Report [19].

Lemma 5.10. *(Move ops right)*

$$\begin{aligned} & \epsilon \cdot (\tau, beg) \cdot \epsilon_{\{\tau_1, \dots, \tau_n\}} \cdot (\tau_1, cmt) \cdot \epsilon_{\{\tau_2, \dots, \tau_n\}} \cdot (\tau_2, cmt) \cdot \dots \cdot (\tau_n, cmt) \cdot \epsilon_{\tau} \\ & \sqsubseteq \\ & \epsilon \cdot (\tau, beg) \cdot \epsilon_{\{\tau, \tau_1, \dots, \tau_n\}} \cdot (\tau_1, cmt) \cdot \epsilon_{\{\tau, \tau_2, \dots, \tau_n\}} \cdot (\tau_2, cmt) \cdot \dots \cdot (\tau_n, cmt) \end{aligned}$$

Proof: Induction on n , using Lemma 5.9. For details see our Technical Report [19].

Lemma 5.11. *(Move beg right) Let ϵ_{τ}^{oc} be an execution consist of operations and commits of transactions other than τ . If $\ell_{\tau} \bar{\boxtimes} \ell_{\tau'}$ for all $\tau' \in \bar{\tau}$ then*

$$\epsilon \cdot (\tau, beg) \cdot \epsilon_{\tau}^{oc} \cdot \epsilon_{\{\tau\}} \cdot (\tau, cmt) \sqsubseteq \epsilon \cdot \epsilon_{\tau}^{oc} \cdot (\tau, beg) \cdot \epsilon_{\{\tau\}} \cdot (\tau, cmt)$$

Proof: Induction and Eqn. 7. For details see our Technical Report [19].

Definition 5.5. *(Abort Factory)*

$$\epsilon_{\text{abt}}(\epsilon, A) \triangleq (\tau_1, abt) \cdot \dots \cdot (\tau_n, abt) \quad \forall \tau_i \in \text{active}(\epsilon, \emptyset)$$

Theorem 5.12. *(Optimistic Serializability)*

For any $\epsilon = C_0^{\text{opt}} \xrightarrow{*} C^{\text{opt}}$,

$$\exists \epsilon_S. \text{serial } \epsilon_S \in \epsilon_{\text{abt}}(\epsilon, \emptyset) \wedge \epsilon \cdot \epsilon_{\text{abt}}(\epsilon, \emptyset) \sqsubseteq \epsilon_S \in \epsilon_{\text{abt}}(\epsilon, \emptyset)$$

Proof: By induction on the execution trace, where all uncommitted transactions have been removed (i.e. aborted). Specifically, at each step of the proof we construct an appropriate serial ϵ_S for which we can show ϵ is included so long as both are appended by $\epsilon_{\text{abt}}(\epsilon_U, \emptyset)$.

Base case: Consider the empty execution $\epsilon = \circ$, which starts in configuration C_0^{opt} . We can define $\epsilon_S = \circ$. $\text{serial } \epsilon_S$ holds by axiom, and \sqsubseteq holds trivially. We now show that

$$\begin{aligned} & \exists \epsilon_S. \text{serial } \epsilon_S \cdot \epsilon_{\text{abt}}(\epsilon, \emptyset) \wedge \epsilon \cdot \epsilon_{\text{abt}}(\epsilon, \emptyset) \sqsubseteq \epsilon_S \cdot \epsilon_{\text{abt}}(\epsilon, \emptyset) \wedge \epsilon \xrightarrow{\mathcal{L}} \epsilon' \\ & \Rightarrow \\ & \exists \epsilon_S'. \text{serial } \epsilon_S' \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset) \wedge \epsilon' \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset) \sqsubseteq \epsilon_S' \cdot \epsilon_{\text{abt}}(\epsilon_S', \emptyset) \end{aligned}$$

Now case split on \mathcal{L} .

Case $\mathcal{L} = \perp$: $\epsilon' = \epsilon$. Trivial.

Case $\mathcal{L} = (\tau, beg)$: Let $\epsilon_S' = \epsilon_S \cdot (\tau, beg)$.

The predicate $\text{serial } \epsilon_S' \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset)$ is equivalent (by Lemma 5.8) to $\text{serial } \epsilon_S \cdot \epsilon_{\text{abt}}(\epsilon, \emptyset)$ which holds by induction. Again using

Lemma 5.8 and the inductive hypothesis, the other obligation holds:

$$\epsilon' \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset) \sqsubseteq \epsilon_{S'} \epsilon_{\text{abt}}(\epsilon_{S'}, \emptyset)$$

Case $\mathcal{L} = (\tau, \text{cmt})$: First we show

$$\begin{aligned} & \epsilon(\tau, \text{cmt}) \cdot \epsilon_{\text{abt}}(\epsilon^1, \emptyset) \\ & \sqsubseteq \epsilon_{S'}(\tau, \text{cmt}) \cdot \epsilon_{\text{abt}}(\epsilon^1, \emptyset) \\ & \sqsubseteq \epsilon_{\{\tau', U\}}(\tau, \text{beg}) \cdot \epsilon_{\{\tau, \tau', U\}}^1(\tau, \text{cmt}) \cdot \epsilon_{\text{abt}}(\epsilon^1, \emptyset) \quad (\text{Rewrite}) \\ & \sqsubseteq \epsilon_{\{\tau'\}}(\tau, \text{beg}) \cdot \epsilon_{\{\tau, \tau'\}}^1(\tau, \text{cmt}) \quad (\text{Lm. 5.8}) \\ & \sqsubseteq \epsilon_{\{\tau'\}}(\tau, \text{beg}) \cdot \epsilon_{\{\tau'\}}^1(\tau, \text{cmt}) \end{aligned}$$

The last step is accomplished by induction with Lemmas 5.10 and 5.11. Now let $\epsilon' = \epsilon_{\{\tau'\}}(\tau, \text{beg}) \cdot \epsilon_{\{\tau'\}}^1(\tau, \text{cmt})$. By the inductive hypothesis, $\text{serial } \epsilon_{\{\tau'\}}(\tau, \text{beg}) \cdot \epsilon_{\{\tau'\}}^1(\tau, \text{cmt})$ holds and it is easy to show that $\text{serial } (\tau, \text{beg}) \cdot \epsilon_{\{\tau'\}}^1(\tau, \text{cmt})$ holds. Thus $\text{serial } \epsilon'$ holds.

Case $\mathcal{L} = (\tau, \text{abt})$: Let $\epsilon_{S'} = \epsilon_S(\tau, \text{abt})$. Predicate $\text{serial } \epsilon_{S'} \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset)$ holds because $\text{serial } \epsilon_S \epsilon_{\text{abt}}(\epsilon, \emptyset)$ holds and appending an abort τ can be removed by Lemma 5.8 (and ϵ_{abt} now has a lesser obligation). So it is also trivial to show that

$$\epsilon' \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset) \sqsubseteq \epsilon_{S'} \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset)$$

Case $\mathcal{L} = (\tau, o.m)$: Let $\epsilon_{S'} = \epsilon_S(\tau, o.m)$. Predicate $\text{serial } \epsilon_{S'} \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset)$ holds because $\text{serial } \epsilon_S \epsilon_{\text{abt}}(\epsilon, \emptyset)$ holds and appending a single operation adds an event that will be removed by ϵ_{abt} . Using similar reasoning and we can show that

$$\epsilon' \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset) \sqsubseteq \epsilon_{S'} \cdot \epsilon_{\text{abt}}(\epsilon', \emptyset)$$

6. Traditional (Fine-grained) STM

Most transactional memory systems [7, 28] use the optimistic semantics at the fine-grained level of memory read/write operations. Transactions concurrently execute operations and when they try to commit they either succeed, in which case their effects atomically become visible, or abort, in which case their effects are discarded and the transactions restarts.

Formally, our optimistic semantics uses per-transaction copies of the shared state. For performance reasons, however, many transactional memory implementations do in-place modifications of the shared state. They rely on *undo*-logs (rather than our *intention*-logs) to reverse the effects of aborted transactions. Since the operations (read and write memory locations) are simple, it is easy to keep undo-logs which consist of previous memory values.

We now formalize traditional transactional memory using our semantics. Let $\mathcal{L} : \mathbb{N}$ be a set of locations and $\mathcal{V} : \mathbb{N}$ be a set of values. The shared store σ_{sh} consists of a single object η representing a memory, as well as per-transaction return values: $\sigma_{\text{sh}} = \langle \eta, r_{\tau_1}, r_{\tau_2}, \dots \rangle$. We define the two method calls as follows:

$$\begin{aligned} \llbracket \tau_i : \eta.\text{read } x \rrbracket \sigma_{\text{sh}} &= \sigma_{\text{sh}}[r_{\tau_i} \mapsto \eta[x]] \\ \llbracket \tau_i : \eta.\text{write } x \ v \rrbracket \sigma_{\text{sh}} &= \sigma_{\text{sh}}[\eta[x \mapsto v]] \end{aligned}$$

When transaction τ_i reads x from the shared store, the r_{τ_i} component of the shared store is updated to contain the current value of x in η . The write operation updates η such that the value of x is v . Moverness is intuitive (let x, y and v, w each be distinct and means neither \triangleleft nor \triangleright):

$$\eta.\text{read } x \ \triangleleft \ \eta.\text{read } y \quad (11)$$

$$\eta.\text{write } x \ v \ \triangleleft \ \eta.\text{write } y \ v \quad (12)$$

$$\eta.\text{read } x \ \triangleleft \ \eta.\text{write } y \ v \quad (13)$$

$$\eta.\text{read } x \ \triangleleft \ \eta.\text{read } x \quad (14)$$

$$\eta.\text{write } x \ v \ \triangleleft \ \eta.\text{write } x \ v \quad (15)$$

$$\eta.\text{write } x \ v \ \cdot \ \eta.\text{write } x \ v \quad (16)$$

$$\eta.\text{read } x \ \cdot \ \eta.\text{write } x \ v \quad (17)$$

Existing transactional memory implementations mostly abide by the above rules. However they typically do not consider Eqn. 15 to be a both-mover.

Invisible reads Invisible reads [12, 33] and other forms of read-only access in transactional memory implementations exploit Eqns. 11 and 14. Since read operations are always both movers with respect to each other, it is always sound to allow read-only workloads to interleave in any order. Implementations have various ways of handling the advent of write operations in such workloads.

Eager Conflict Detection in TL2 The TL2 implementation involves an added optimization: eager conflict detection. Rather than keeping a thread-local view of memory and checking conflicts at commit, with each read/write operation, TL2 checks the location timestamp and compares it with the transaction-local begin timestamp. Transactions are aborted if a conflict is detected – conflicts correspond to any instance of read/write operations on the same location (Eqns. 14, 15, 16, 17). Eager conflict detection can be over-approximated in our semantics as simply a non-deterministic choice to take the OABT rule.

Checkpoints and Nested Transactions Our semantics is amenable to the “checkpoints” alternative [17] to nested transactions. We intentionally designed our semantics to be non-deterministic: the UNDO rule can be invoked at any point during a transaction, allowing the implementation to decide how many times to apply the UNDO rule. Checkpoints allow us to model (i) nested transactions [17] and (ii) mechanisms which escape deadlock [18] as described next.

Deadlock The pessimistic semantics may deadlock. Consider a configuration in which T includes $\langle \tau, o.m; C, \{p.n\}, \sigma \rangle$ and $\langle \tau', p.n'; C', \{o.m'\}, \sigma' \rangle$. If neither $o.m \triangleleft o.m'$ nor $p.n \triangleleft p.n'$, then neither transaction will be able to take the PAPP step in the semantics.

There are implementation techniques for detecting such deadlock [18]. These techniques can recover from deadlock by performing either a full abort (reverting all of a transaction’s operations through multiple instances of UNDO) or partial abort (reverting a subset of operations). Ideally, recovery would only revert the operations necessary to allow another transaction to make progress.

6.1 Opacity

In addition to serializability, recent work has proposed a new correctness condition called *opacity* [9] which holds for most modern implementations of transactional memory [7, 28]. Opacity involves serializability as well as two new requirements. We now show that both of our execution semantics are opaque.

Definition 6.1. (*Opacity*) An execution ϵ is opaque, denoted *opaque* ϵ if and only if

1. ϵ is serial (*serial* ϵ).
2. Operations performed by an aborted transaction are never visible to other transactions.
3. Every transaction observes a consistent version of the state.

Proposition 6.1. (*Opacity of Pessimistic Execution*) For any execution ϵ given by \rightarrow_P , *opaque* ϵ .

Proof: Intuitively, the proposition holds because transactions only perform PAPP if all concurrent operations (including those of transactions which have or may abort) are left-movers, the transaction logically only observes the state of the system at commit-time. For details see our Technical Report [19].

Proposition 6.2. (*Opacity of Optimistic Execution*) For any execution ϵ given by \rightarrow_O , *opaque* ϵ .

Proof: Follows from the fact that the shared state is updated atomically in OCMT and that transactions duplicate the shared state in OBEG. For details see our Technical Report [19].

7. Coarse-grained Implementations

7.1 Pessimistic Coarse-Grained (Transactional Boosting)

In recent work we have given a methodology for enabling transactional access to concurrent data structures, say a ConcurrentSkipList, which already have built-in thread synchronization [11]. The methodology is based on pessimism, and allows us to use such off-the-self concurrent objects in a transactional setting without the pitfalls that would come with fine-grained synchronization.

The methodology also makes one particular choice about how coarse-grained synchronization is accomplished: abstract locks. The user defines an abstract locking discipline such that, for any two transactions executing concurrent method calls, the locks conflict if the method calls do not commute. Here is an example:

```
global SkipList Set;
global AbsLock keys[10];

T1: atomic {
  with(lock(keys[5])) {
    Set.add(5);
    with(lock(keys[3])) {
      Set.remove(3);
    }
  }
}

T2: atomic {
  with(lock(keys[6])) {
    if (Set.contains(6)) {
      with(lock(keys[7])) {
        Set.remove(7);
      }
    }
  }
}
```

The above example involves a SkipList implementation of a Set of integers. Set has the standard operations $\text{add}(x)$, $\text{remove}(y)$, and $\text{contains}(z)$. $\text{add}(x)$ commutes with $\text{add}(y)$ if x and y are distinct. An example locking discipline (as above) is to define one abstract lock per key (i.e. x, y, z). Thus, the abstract locks for $\text{add}(x)$ and $\text{add}(y)$ do not conflict when x and y are distinct, so these calls can proceed in parallel.

With our semantics in place, we can formalize the examples we described earlier [11]. Moreover, we can be more precise and consider \triangleleft and \triangleright in addition to \bowtie (commutativity):

- *Multi-set* – The moverness of a multi-set is similar to the commutativity of the Set given in Figure 1 of our previous work [11]. However, a multi-set has the following additional movers:

$$\begin{aligned} \text{contains}(x)/\text{true} &\triangleright \text{insert}(x)/\text{true} \\ \text{remove}(x)/\text{true} &\triangleright \text{insert}(x)/\text{true} \end{aligned}$$

- *Priority Queue (Heap)* – We gave the commutativity of a priority queue in Figure 4 of our previous work [11]. The method $\text{insert}(x)$ inserts an element into the heap and the method $\text{removeMin}()/y$ returns the minimal element. Figure 4 of [11] can be refined:

$$\text{insert}(x) \triangleleft \text{removeMin}()/x$$

In fact we noticed a mistake. Figure 4 of [11] claims that

$$\text{insert}(x) \bowtie \text{removeMin}()/y \quad \forall y \leq x$$

but when $x = y$, only \triangleleft holds.

- *Queue* – Similar to the multi-set, a queue has an operation which is a left-mover but not a right-mover: $\text{enqueue}(x) \triangleleft \text{dequeue}()$
- *Semaphore* – The standard semaphore is represented as an integer whose value is initially 0 with methods incr (increments the value) and decr (decrements the value). Applying decr when the value is 0 is either undefined or blocks. Therefore, $\text{incr} \triangleleft \text{decr}$.

7.2 Optimistic Coarse-Grained (Baldassin and Burckhardt)

Recently, Baldassin and Burckhardt proposed a new programming model which resembles coarse-grained transactions [3]. In their experience adapting a game application to transactions, they describe a technique based on user-provided specifications (i.e. moverness) to resolve conflicts.

We now discuss their implementation, using parenthetical comments to refer back to the corresponding rule in our optimistic execution semantics in Section 3.2. Baldassin and Burckhardt describe transactions which begin (OBEG) with a replica (σ_τ) of the shared state (σ_{sh}). Transactions then execute performing operations (OAPP) on their local copies, and when they commit their local updates are “propagated to all other replicas” (logically equivalent to $\text{merge}(\sigma_{\text{sh}}, \ell_\tau)$ in OCMT) and when transactional operations are read-only, a single shared state is used (direct equivalence to $\text{merge}(\sigma_{\text{sh}}, \ell_\tau)$).

We now give a formal treatment for the objects described [3]. In the following, the first two are base objects where any write is treated as a conflict (i.e. non-moverness) while concurrent reads are acceptable (i.e. both-moverness). The remaining three use higher-order moverness.

- *Values* (base)

$$\text{get} \bowtie \text{get} \quad \text{set} \cdot \text{get} \quad \text{set} \cdot \text{set}$$
- *Collections* (base)

$$\text{GetRO} \bowtie \text{GetRO} \quad \text{GetRW} \cdot \text{GetRO} \quad \text{GetRW} \cdot \text{GetRW}$$
- *SharedGrowingList* (wrapper) – This object appears to be a multiset where the order of element insertion does not matter: $\text{insert}(x) \bowtie \text{insert}(y)$ for all x, y .
- *SharedObjectWithPriority* (wrapper) – This object allows transactions to associate a priority with the operations performed on an object and at commit, conflicts are resolved according to priority. Note that this violates serializability, and may rely on application-specific invariants or behaviors for soundness.
- *SharedObjectWithMerge* (wrapper) – This general wrapper allows a user to define a commit-time merge procedure which computes new values in the shared store (σ_{sh}) given the values desired by committing transactions ($\sigma_\tau, \sigma_{\tau'}$, etc.). This can be thought of as a specific form of applying the operational log ℓ_τ to the shared state in $\text{merge}(\sigma_{\text{sh}}, \ell_\tau)$. Since this merge procedure is provided, the assumption is that all transactional operations on these objects are both-movers or, in case of conflicts, operations can be inverted without aborting the transaction.

In the game application concurrent transactions are aligned on a per-frame basis, all beginning and committing (roughly) simultaneously. They must assume that any begin order was possible and thus all operations must be right movers with respect to each other – i.e. every operation must be a both-mover. By adding begin timestamps this requirement could be relaxed, and operations need only move to the right of committed operations. In a more general setting (not aligned to frames) right moverness may be useful.

8. Related Work

Our work is the first formal treatment of *coarse-grained transactions*, which captures the trend in the recent literature away from low-level memory synchronization towards more efficient, higher-level synchronization. Our semantics, encompassing both pessimistic and optimistic execution models, unifies a wide range of ad-hoc techniques: traditional transactional memory [7, 28], transactional boosting [11], and recent optimistic techniques [3]. Moreover, our work provides a formal basis for an emerging new wave of research [2, 27] which uses automated techniques to discover moverness.

Transactional Boosting In prior work [11] we proved the correctness of Transactional Boosting using histories of events [14, 32]. Boosting is pessimistic and the proof in this paper of the pessimistic execution semantics roughly corresponds to the previous proof [11] when labels are aligned with events. Nonetheless, our treatment here has novelties:

1. We define a coarse-grained language of programs (rather than a library-style methodology) and provide a formal proof of serializability for this language.
2. We have added support for non-determinism, showing equivalence by containment of an execution in a set of serial executions, rather than by finding a single serial execution.
3. We have added a second, optimistic semantics (boosting is purely pessimistic) which models modern read/write STMs such as TL2 [7] and Intel’s STM compiler [28], as well as recent attempts at larger granularity [3] as described in Section 7.2.

Previous Semantics of STM In a similar vein, the work by Jagannathan *et al.* gives a semantics for both pessimistic and optimistic read/write transactional memory [15]. Our work can be seen as a generalization in which the operations are over abstract data-types rather than simple memory operations.

Other previous works on semantics of transactions are orthogonal and have focused instead on *isolation* of transactions (i.e. the privatization problem [30]). These semantics are purely fine-grained, and tackle the issue of ensuring shared memory is not accessed outside of a transaction. By contrast, in our semantics the shared/local distinction is assumed and we focus on the correctness criteria for coarser granularities.

For example, the ATOMS FAMILY presented by Moore and Grossman [23] involves a taxonomy of semantics which vary based on degrees of isolation, availability of abort, and allowed parallel nesting. Another example is the semantics of Automatic Mutual Exclusion (AME) given by Abadi *et al.* [1]. The programming model of AME is substantially different from ours; its execution model involves spawning asynchronous method calls which are scheduled such that only non-conflicting code runs concurrently. This programming model can be used to model fine-grained optimistic implementations, with “eager” updates. We next discuss the distinction between eager and lazy.

Another important distinction is the choice of *which form of optimism* to model: eager or lazy updates to memory. The eager version may perform better but is subject to transactions observing inconsistent states (i.e. not opaque [9] as discussed in Section 6.1), which appears to have significant drawbacks. AME discusses distinction and why they chose to model lazy updates. The ATOMS FAMILY explores both lazy (StrongBasic) and eager (Weak), using a type system to ensure serializability in the latter case.

In our work, lazy updates are fundamentally part of the model (in \rightarrow_O the shared store is only updated in the OCMT rule). Our model matches that of TL2 [7] and Intel’s transactional memory compiler [28].

Commutativity The notion of commutativity, more precisely moverness [22], is fundamental in databases and parallel programming. The literature includes Bernstein [5], Steele [31], Weihl [32], Schwartz and Spector [29], Beeri [4], and Korth [16]. More recently, the Jade programming language [21, 26] has been developed which includes support for programmer-provided specifications as to whether certain functions are commutative.

In the very recent literature there have been some techniques which attempt to detect commutativity. This work is inspired by Ri-nard and Diniz [27] whose *commutativity analysis* checked for concrete equivalence mostly via heap disjointness, and leveraged some limited arithmetic identities. Recently, Aleen and Clark presented a commutativity analysis [2]. Their work is a whole-program analysis for checking whether context-sensitive procedure invocations commute. To check equivalence of re-orderings they use random interpretation. Their work is orthogonal to ours, and could be used to detect or prove moverness of certain objects. Some effort may be needed since (1) their work is for parallelizing compilers, not transactions and (2) precision could be improved by considering left- and right-moverness rather than simply both-moverness.

9. Conclusion

In the recent literature there has been a new trend of broadening transactional operations from fine-grained memory operations to more coarse-grained operations, often drastically improving performance. We have presented a semantics which unifies a broad range of such ad-hoc techniques, encompassing both pessimistic (e.g. transactional boosting) and optimistic (e.g. traditional STM and recent coarse-grained work [3]) execution semantics. We have proven serializability and found that the choice of execution semantics imposes different requirements, expressed in terms of left- and right-movers, on the coarse-grained operations themselves. Our work serves also as a formal foundation for recent automated techniques for discovering moverness [2, 27]. We expect that this will be a promising research direction for the future.

Acknowledgments We thank Byron Cook, Tim Harris, Michael Hicks and Peter Sewell for their comments on an earlier draft. We acknowledge funding from a Gates scholarship (Koskinen), a Royal Academy of Engineering Research Fellowship (Parkinson), and NSF grant 0811289 (Herlihy).

References

- [1] ABADI, M., BIRRELL, A., HARRIS, T., AND ISARD, M. Semantics of transactional memory and automatic mutual exclusion. In *The 35th ACM SIGPLAN SIGACT Symposium on Principles of Programming Languages (POPL’08)* (2008), G. C. Necula and P. Wadler, Eds., ACM, pp. 63–74.
- [2] ALEEN, F., AND CLARK, N. Commutativity analysis for software parallelization: letting program transformations see the big picture. In *Proceedings of the 14th international conference on Architectural support for programming languages and operating systems (ASPLOS-XII)* (2009), M. L. Soffa and M. J. Irwin, Eds., ACM, pp. 241–252.
- [3] BALDASSIN, A., AND BURCKHARDT, S. Lightweight software transactions for games. In *Proceedings of the First USENIX Workshop on Hot Topics in Parallelism (HotPar’09)* (2009).
- [4] BEERI, C., BERNSTEIN, P., GOODMAN, N., LAI, M.-Y., AND SHASHA, D. A concurrency control theory for nested transactions (preliminary report). In *Proceedings of the 2nd annual ACM symposium on Principles of distributed computing (PODC’83)* (New York, NY, USA, 1983), ACM Press, pp. 45–62.
- [5] BERNSTEIN, A. Analysis of programs for parallel processing. *IEEE Transactions on Electronic Computers* 15, 5 (1966), 757–763.
- [6] DAMRON, P., FEDOROVA, A., LEV, Y., LUCHANGCO, V., MOIR, M., AND NUSSBAUM, D. Hybrid transactional memory. In *Proceedings of the 12th international conference on Architectural support for*

- programming languages and operating systems (ASPLOS-XII)* (New York, NY, USA, 2006), ACM Press, pp. 336–346.
- [7] DICE, D., SHALEV, O., AND SHAVIT, N. Transactional locking II. In *Proceedings of the 20th International Symposium on Distributed Computing (DISC'06)* (September 2006).
- [8] FLANAGAN, C., FREUND, S. N., AND YI, J. Velodrome: a sound and complete dynamic atomicity checker for multithreaded programs. In *PLDI'08: Proceedings of the 2008 ACM SIGPLAN conference on Programming language design and implementation* (New York, NY, USA, 2008), ACM, pp. 293–303.
- [9] GUERRAOU, R., AND KAPALKA, M. On the correctness of transactional memory. In *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming (PPoPP'08)* (New York, NY, USA, 2008), ACM, pp. 175–184.
- [10] HARRIS, T., MARLOW, S., JONES, S. L. P., AND HERLIHY, M. Composable memory transactions. *Commun. ACM* 51, 8 (2008), 91–100.
- [11] HERLIHY, M., AND KOSKINEN, E. Transactional boosting: A methodology for highly concurrent transactional objects. In *Proceedings of the 13th ACM SIGPLAN symposium on Principles and practice of parallel programming (PPoPP'08)* (2008).
- [12] HERLIHY, M., LUCHANGCO, V., MOIR, M., AND SCHERER, III, W. N. Software transactional memory for dynamic-sized data structures. In *Proceedings of the 22nd annual symposium on Principles of distributed computing (PODC'03)* (2003), ACM Press, pp. 92–101.
- [13] HERLIHY, M., AND MOSS, J. E. B. Transactional memory: architectural support for lock-free data structures. In *Proceedings of the 20th Annual International Symposium on Computer Architecture (ISCA'93)* (1993), ACM Press, pp. 289–300.
- [14] HERLIHY, M. P., AND WING, J. M. Linearizability: a correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 12, 3 (1990), 463–492.
- [15] JAGANNATHAN, S., VITEK, J., WELC, A., HOSKING, A. A Transactional object calculus. *Science of Computer Programming*. 57, 2 (August 2005), 164–186.
- [16] KORTH, H. F. Locking primitives in a database system. *J. ACM* 30, 1 (1983), 55–79.
- [17] KOSKINEN, E., AND HERLIHY, M. Checkpoints and continuations instead of nested transactions. In *Proceedings of the twentieth annual symposium on Parallelism in algorithms and architectures (SPAA'08)* (New York, NY, USA, 2008), ACM, pp. 160–168.
- [18] KOSKINEN, E., AND HERLIHY, M. Dreadlocks: efficient deadlock detection. In *Proceedings of the twentieth annual symposium on Parallelism in algorithms and architectures (SPAA'08)* (New York, NY, USA, 2008), ACM, pp. 297–303.
- [19] KOSKINEN, E., PARKINSON, M., AND HERLIHY, M. Coarse-grained transactions. Tech. Rep. 759, Computer Laboratory, University of Cambridge, 2009.
- [20] KULKARNI, M., PINGALI, K., WALTER, B., RAMANARAYANAN, G., BALA, K., AND CHEW, L. P. Optimistic parallelism requires abstractions. In *Proceedings of the ACM SIGPLAN 2007 Conference on Programming Language Design and Implementation (PLDI'07)* (2007), J. Ferrante and K. S. McKinley, Eds., ACM, pp. 211–222.
- [21] LAM, M., AND RINARD, M. Coarse-grain parallel programming in Jade. In *Proceedings of the third ACM SIGPLAN symposium on Principles and practice of parallel programming (PPoPP'91)* (1991), ACM New York, NY, USA, pp. 94–105.
- [22] LIPTON, R. J. Reduction: a method of proving properties of parallel programs. *Commun. ACM* 18, 12 (1975), 717–721.
- [23] MOORE, K. F., AND GROSSMAN, D. High-level small-step operational semantics for transactions. In *Proceedings of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages (POPL'08)* (New York, NY, USA, 2008), ACM, pp. 51–62.
- [24] MORAVAN, M. J., BOBBA, J., MOORE, K. E., YEN, L., HILL, M. D., LIBLIT, B., SWIFT, M. M., AND WOOD, D. A. Supporting nested transactional memory in logtm. In *Proceedings of the 12th international conference on Architectural support for programming languages and operating systems (ASPLOS-XII)* (New York, NY, USA, 2006), ACM Press, pp. 359–370.
- [25] MORAVAN, M. J., BOBBA, J., MOORE, K. E., YEN, L., HILL, M. D., LIBLIT, B., SWIFT, M. M., AND WOOD, D. A. Supporting nested transactional memory in logtm. *SIGOPS Oper. Syst. Rev.* 40, 5 (2006), 359–370.
- [26] RINARD, M., AND LAM, M. The design, implementation, and evaluation of Jade. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 20, 3 (1998), 483–545.
- [27] RINARD, M. C., AND DINIZ, P. C. Commutativity analysis: A new analysis technique for parallelizing compilers. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 19, 6 (November 1997), 942–991.
- [28] SAHA, B., ADL-TABATABAI, A.-R., HUDSON, R. L., MINH, C. C., AND HERTZBERG, B. McRT-STM: a high performance software transactional memory system for a multi-core runtime. In *Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming (PPoPP'06)* (New York, NY, USA, 2006), ACM, pp. 187–197.
- [29] SCHWARZ, P. M., AND SPECTOR, A. Z. Synchronizing shared abstract types. *ACM Transactions on Computer Systems* 2, 3 (1984), 223–250.
- [30] SPEAR, M. F., MARATHE, V. J., DALESSANDRO, L., AND SCOTT, M. L. Privatization techniques for software transactional memory. In *Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing (PODC'07)* (New York, NY, USA, 2007), ACM, pp. 338–339.
- [31] STEELE, JR, G. L. Making asynchronous parallelism safe for the world. In *Proceedings of the 17th ACM SIGPLAN-SIGACT symposium on Principles of programming languages (POPL'90)* (New York, NY, USA, 1990), ACM Press, pp. 218–231.
- [32] WEIHL, W. E. Data-dependent concurrency control and recovery (extended abstract). In *Proceedings of the second annual ACM symposium on Principles of distributed computing (PODC'83)* (New York, NY, USA, 1983), ACM Press, pp. 63–75.
- [33] WILLIAM N. SCHERER, I., AND SCOTT, M. L. Advanced contention management for dynamic software transactional memory. In *Proceedings of the 24th annual ACM symposium on Principles of distributed computing (PODC'05)* (New York, NY, USA, 2005), ACM, pp. 240–248.