

The Push/Pull model of transactions

Eric Koskinen¹ and Matthew Parkinson²

¹New York University

²Microsoft Research

Abstract

We present a general theory of serializability, unifying a wide range of transactional algorithms, including some that are yet to come. To this end, we provide a compact semantics in which concurrent transactions *push* their effects into the shared view (or *unpush* to recall effects) and *pull* the effects of potentially uncommitted concurrent transactions into their local view (or *unpull* to detangle). Each operation comes with simple side-conditions given in terms of commutativity (Lipton’s left-movers and right-movers [24]).

The benefit of this model is that most of the elaborate reasoning (coinduction, simulation, subtle invariants, etc.) necessary for proving the serializability of a transactional algorithm is already proved within the semantic model. Thus, proving serializability (or opacity) amounts simply to mapping the algorithm on to our rules, and showing that it satisfies the rules’ side-conditions.

1 Introduction

Recent years have seen an explosion of research on methods of providing **atomic** sections in modern programming languages, typically implemented via transactional memory (TM). The **atomic** keyword provides programmers with a powerful concurrent programming building block: the ability to specify when a thread’s operations on shared memory should appear to take place instantly when viewed by another thread.

To support such a construct, we must be able to reason about atomicity. Implementations typically achieve this by dynamically detecting conflicts between concurrent threads. This can be done tracking memory operations in hardware [15, 17, 16] or software [14, 6, 8, 25, 4]. Meanwhile, an alternate approach exploits abstract-level notions of conflict over linearizable data-structure operations such as commutativity [11, 28, 21, 20]. Both levels of abstraction also chose between optimistic execution, pessimistic execution, or mixtures of the two. Finally, there are multiple notions of correctness, and circumstances under which one may be preferable to another.

Unfortunately, we lack a unified way of formally describing this myriad of models, implementations and correctness criteria. This leads to confusion when trying to understand comparative advantages/disadvantages and how/when models can be combined or are interoperable. For example, with hardware support for transactions now available (*e.g.* Intel Haswell [17]), we need to understand how one can combine *memory-level* hardware transactions for unstructured memory operations with *abstract-level* data-structure operations (*e.g.* transactional boosting [11]). Today, at best, we have two custom semantics for reasoning about the models individually, but no unified view.

We present a simple calculus that illuminates the core of transactional memory systems. In our model concurrent transactions *push* their effects into the shared log (or *unpush* to roll-back) and *pull* in the effects of potentially uncommitted concurrent transactions (or *unpull* to detangle). Moreover, transactions can push or pull operations in non-chronological orders, provided certain commutativity (Lipton left/right-movers [24]) conditions hold. The benefit of this semantic model is that most of the elaborate reasoning (coinduction, simulation relations, subtle invariants, etc.) necessary for proving the correctness of a transactional algorithm is contained within the semantic model, and need only be proved once.

Our work formulates an expressive class of transactions and we have applied it to a wide range of TM systems including: optimistic read/write software TMs [6, 8], hardware transactional memories (Intel [17], IBM [16]), pessimistic TMs [25, 4, 11], hybrid optimistic/pessimistic TMs such as irrevocability [34], open nested transactions [28], and abstract-level techniques such as boosting [11].

Our choice of expressiveness includes transactions that are not opaque [10]: transactions *may* share their uncommitted effects. This choice carves out a design space for implementations to take advantage of the full spectrum of possibilities (*e.g.* dependent transactions [30]) and is relatively unrestrictive in terms of TM correctness criteria. However, despite expressive power, the model also gives the appropriate criteria to ensure serializability [29]. Meanwhile, we can also identify restrictions on the model for which opacity is recovered.

In our experience we have found that our model provides a mathematically rigorous foundation for intuitive concepts (*e.g.* *push* and *pull*) used in colloquial conversations contrasting TM systems.

Contributions. Our work includes the following:

- A general model of concurrent transactions in Section 4, capable of expressing a wide range of implementations, with only a few intuitive rules. Our model is parameterized by a coinductive sequential specification and the operational semantics of the programming language.
- We have proved that this model is serializable, discussed in Section 5. To cope with the non-monotonic nature of the model (arising from *unpush*, *unpull*, etc.), we devised a novel preservation invariant that is closed under rewinding both the local and global logs. The serializability proof shows simulation with an uninterleaved machine.
- We have shown conditions under which we can restrict the model to obtain a sub-model that satisfies opacity.
- In Section 6 we describe how our model accounts for the serializability of many transactional memory systems that range from software to hardware, pessimistic to optimistic, read/write-conflict to abstract-conflict, nesting and non-opaque features such as dependent transactions.

Limitations. The work presented in this paper models safety properties of transactions (*i.e.* serializability, opacity). A direction for future work is to consider liveness/progress issues.

Related work. In previous work [20] we provided a formal semantics for transactions that perform abstract-level data-structure operations. Our work involves two *separate* semantics: one for pessimistic transactions and one for optimistic transactions. There are several distinctions of our work: (i) The model presented in this paper is more expressive because it permits mixtures of these two flavors. This is particularly useful when combining hardware transactions [17, 16] with and abstract-level reasoning [11] for data-structures. (ii) In our mode, transactions may observe the effects of uncommitted, non-commutative transactions as seen in dependent transactions [30] and open nesting [28]. (iii) We have a simulation result which involves nontrivial formal groundwork such as a coinductive definition of state equality and left-mover.

Lesani *et al.* [22] describe a method of specifying and verifying TM algorithms. They specify some transactional algorithms in terms of I/O automata and this choice of language enables them to fully verify those specifications in PVS. In our work, we have aimed at a more abstract goal: to uncover the fundamental nature of transactions in the form of a general-purpose model. We leave the goal of full algorithm verification to future work.

There are other works in the literature that are focused on a variety of orthogonal semantic issues, including the privatization problem [32, 26, 1], correctness criteria such as dynamic/static/hybrid atomicity [33], and message passing within transactions [23]. These works are concerned with models that are restricted to read/write STMs and limited in expressive power (*e.g.* restricted to opacity [10]). Semantics also exist for other programming models that are similar to transactions [2, 3] but are not serializable. Finally, Cohen *et al.* [5] described some small hand proofs for particular transactional memory algorithms.

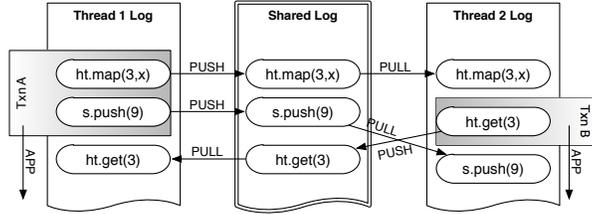


Figure 1: A diagram of PUSH/PULL

2 Overview

In this paper we distill the essence of reasoning about transactional implementations into a semantic model we call PUSH/PULL transactions. The model consists of a few simple rules—named PUSH, PULL, etc.—that correspond to natural stages in a transactional memory algorithm. For example, after a transaction applies an effect locally it then may PUSH this effect out into the shared view, where other transactions may PULL the effect into their local view.

The PUSH/PULL model has no concrete state, only a shared log of the object operations that have been applied, as well as per-thread local logs. An illustration is given in Figure 1. The full formal detail of the model is given in Section 4. We will now discuss this model informally.

Once a transaction (logically) applies an operation in its local log via the APP rule, it may PUSH the operation to the shared log. Note that, at this stage, the transaction may not have committed. Meanwhile, other threads may PULL the operation into their local log. The PULL case enables transactions to update their local view with operations that are permanent (that is, that correspond to committed transactions) or even to view the effects of another uncommitted transaction (*e.g.* for early conflict detection [14] or to establish a dependency [30]). PUSH/PULL also includes an UNPULL rule which discards a transaction’s knowledge of an effect due to another thread, and an UNPUSH rule which removes a thread’s operation from the shared view, perhaps implemented as an inverse. The UNAPPLY rule is useful for rewinding a transaction’s local state. Finally, there is a simple commit rule CMT that, roughly, stipulates that all operations must have been PUSHed and all PULled operations must have been committed.

Different algorithms will use different combinations of these rules (*cf.* Section 6). PUSH/PULL is expressive enough to describe a wide range of transactional implementations, all with only a few simple, tangible rules. Pessimistic algorithms [11, 25, 4] PUSH immediately after a local APP, optimistic algorithms [6, 8] PUSH their operations on commit, and hybrid [34] algorithms do a mixture of the two. Opaque [10] transactions do not PULL uncommitted effects. Non-opaque algorithms, such as dependent transactions [30], permit a transaction to PULL in uncommitted effects. From different patterns of PUSH/PULL rule usage one can derive correctness proofs for many transactional memory algorithms.

Example. Consider the transactional boosting [11, 12] hashtable implementation given in Figure 2. Recall that a boosted transaction uses a linearizable base object (in this case a `ConcurrentSkipListMap`), along with abstract locking to ensure that only commutative operations occur concurrently. In this example a thread executing the `atomic` block in `put` or `get` acquires a lock corresponding to the `key` of interest. In this way, no two transactions will conflict because if they try to access the same `key` one will block. Within the `put` method there are two scenarios depending on whether `key` is already defined in the `map` and, consequently, there are two cases for how to handle an abort. Finally, `put` ends by updating `map` and unlocking the `abstractLock`.

We can describe this algorithm intuitively, in terms of rules in the PUSH/PULL model. We have decomposed the code accordingly in Figure 2. After the transaction begins, it implements a PULL implicitly because, in transactional boosting, modifications are made directly to the shared state so the local view is the same as the shared view. Skipping the abort cases for the moment, the transaction then performs an

<pre> import java.util.concurrent.ConcurrentSkipListMap class BoostedSkipListMap[Key,Value] { val abstractLock = new AbstractLock() val map = new ConcurrentSkipListMap[Key, Value]() def put(key: Key, value: Value, t = Tx.current) { atomic { abstractLock lock key if (map contains key) { var oldValue = map(key) Tx.onAbort(() => map.put(key, oldValue) abstractLock unlock key) } else { Tx.onAbort(() => map.remove(key) abstractLock unlock key) } map.put(key,value) abstractLock unlock key } } def get(key: Key) { atomic { abstractLock lock key Tx.onExit(() => abstractLock.unlock(key)) key k = map.get(key) return k } } } </pre>	<pre> import java.util.concurrent.ConcurrentSkipListMap class BoostedSkipListMap[Key,Value] { val abstractLock = new AbstractLock() val map = new ConcurrentSkipListMap[Key, Value]() def put(key: Key, value: Value, t = Tx.current) { atomic { BEGIN abstractLock lock key; PULL(map) (implicit because shared access to map) if (map contains key) { var oldValue = map(key) Transaction.onAbort(() => UNPUSH(map.put(key,value)) and UNAPP(map.put(key,value)) map.put(key, oldValue) abstractLock unlock key) } else { Transaction.onAbort(() => UNPUSH(map.put(key,value)) and UNAPP(map.put(key,value)) map.remove(key) abstractLock unlock key) } APPLY(map.put(key,value)) and PUSH(map.put(key,value)) map.put(key, value) CMT abstractLock unlock key } } def get(key:Key) { . . . } } </pre>
---	---

Figure 2: On the left, an implementation of *transactional boosting* [11] which uses abstract locking and commutativity to safely access a shared Set, implemented as a `ConcurrentSkipList`. On the right, decomposition of the boosting implementation into PUSH/PULL rules such as APP, PUSH, PULL, CMT, as well as UNAPP, UNPUSH used for aborting a transaction. Each of these rules comes with a correctness criteria (see Figure 5) which, if proved to hold, implies that the implementation is serializable.

APPLY (effecting the local view) and a PUSH (sending the effect to the shared view) by updating the map. If the transaction aborts, it performs the opposite of these two PUSH/PULL rules in reverse: UNPUSH and then UNAPP by performing the appropriate inverse operation. The abstract locking of boosting ensures that the only operations in the shared log belong to committed transactions or else commute with all other uncommitted operations in the shared log.

Proofs of serializability. Each rule in PUSH/PULL comes with a few *correctness criteria*. In Section 5 we prove that if an implementation satisfies these criteria, then it is serializable. In this sense we have done the hard work of reasoning about transactional memory algorithms. The PUSH/PULL model encapsulates the

difficult components involved in a correctness argument (*e.g.* simulation proofs, coinduction, etc.) while, on the outside, offering rules that are simple and intuitive. For a user to prove the correctness of their algorithm they must simply:

1. Demarcate the algorithm into fragments: PUSH, PULL, etc.
2. Prove the implementation satisfies the respective correctness criteria.

Proofs of correctness criteria typically do not involve elaborate simulation relations or coinductive reasoning, but rather algebraic (*i.e.* commutative) properties of sequential code.

Synchronization in Figure 2 is accomplished using abstract locks to ensure that only commutative operations proceed in parallel. One correctness criterion of `PUSH(map.put(key, value))` is that the `put` operation must be able to commute with (more precisely: move to the right of) concurrent, uncommitted operations. We can show that this correctness criterion holds by showing that the two sequential sequences:

$$\begin{array}{c} \text{map.put}(\text{key1}, \text{value1}); \text{map.put}(\text{key2}, \text{value2}) \\ \text{vs.} \\ \text{map.put}(\text{key2}, \text{value2}); \text{map.put}(\text{key1}, \text{value1}) \end{array}$$

lead to the same final state, provided that `key1 ≠ key2`. Moreover, such proofs involving commutativity can be aided by recent works in the literature [7, 18]. Note that, without PUSH/PULL, the full formal argument would say that the data structure is atomic because there is a simulation relation between any configuration of concurrent transactions and a sequential history. The benefit of the PUSH/PULL semantic model is that the simulation relation has been identified and the difficult aspects of the correctness argument have already been proved. We believe that our work will cleanup transactional correctness proofs because it suffices to show that the implementation satisfies the correctness criteria given in our proof rules.

3 Language and Atomic Semantics

In this section we describe a generic language of transactions and define an idealized semantics for concurrent transactions called the atomic semantics in which there are no interleaved effects on the shared state. We later introduce the PUSH/PULL semantics and show that it simulates the atomic semantics.

Language. We assume a set M of method calls (*e.g.* `ht.put('a', 5)`). Threads execute code from a programming language that includes transactions `tx c`, method names such as m , and a `skip` statement. Our first trick is to abstract away the threads' programming language c with two functions:

step(c): Pair $(m, c') \in \text{step}(c)$ if m is a next reachable method in the reduction of c , with remaining code c' .

fin(c): This predicate is true if there is a reduction of c to `skip` that does not encounter a method call.

These two functions allow us to obtain a simple semantics, despite an expressive input language, by introducing functions to resolve nondeterminism between method operation names and at the end of a transaction. We assume that code is well-formed in that a single operation name m is always contained within a transaction (orthogonal, is this issue of isolation [26]).

Example 1. *One could use the generic language:*

$$c ::= c_1 + c_2 \mid c_1 ; c_2 \mid (c)^* \mid \text{skip} \mid \text{tx } c \mid m$$

This grammar additionally consists of nondeterministic choice, sequential composition, and nondeterministic

looping. The corresponding functions for this example are:

$$\begin{aligned}
\text{step}(\text{skip}) &\equiv \emptyset \\
\text{step}(c_1 ; c_2) &\equiv (\text{step}(c_1) ; c_2) \cup (\text{fin}(c_1) ; \text{step}(c_2)) \\
\text{step}(c_1 + c_2) &\equiv \text{step}(c_1) \cup \text{step}(c_2) \\
\text{step}((c)^*) &\equiv \text{step}(c) ; (c)^* \\
\text{step}(\text{tx } c) &\equiv \text{step}(c) \\
\text{step}(m) &\equiv \{(m, \text{skip})\} \\
\text{fin}(\text{skip}) &\equiv \text{true} \\
\text{fin}(c_1 ; c_2) &\equiv \text{fin}(c_1) \wedge \text{fin}(c_2) \\
\text{fin}(c_1 + c_2) &\equiv \text{fin}(c_1) \vee \text{fin}(c_2) \\
\text{fin}((c)^*) &\equiv \text{true} \\
\text{fin}(\text{tx } c) &\equiv \text{fin}(c) \\
\text{fin}(m) &\equiv \text{false}
\end{aligned}$$

$$\begin{aligned}
S ; c_1 &\equiv \{(m, c_1; c_2) \mid (m, c_1) \in S\} \\
B ; S &\equiv \{(m, c_1) \mid B \wedge (m, c_1) \in S\}
\end{aligned}$$

Thus, if $c = \text{tx}(\text{skip} ; (c_1 + (m + n)) ; c_2)$, then one path through c reaches method n with a continuation of c_2 . Hence, $(n, c_2) \in \text{step}(c)$.

To make things more concrete, we instantiate our semantics with the above language through the remainder of this paper. We ignore nested transactions¹, however our model permits threads to roll backwards to any execution point [19] (thus modeling the partial abort nature of nested transactions).

Operations and logs. State is represented in terms of logs of operation records. An operation record (or, simply, an “operation”) $op = \langle m, \sigma_1, \sigma_2, id \rangle$ is a tuple consisting of the operation name m , a thread-local post-stack σ_1 (method arguments), a thread-local post-stack σ_2 (method return values), and a unique identifier id . We assume a predicate $\text{fresh}(id)$ that holds provided that id is globally unique (details omitted for lack of space). In the atomic semantics defined below, the shared state $\ell : \text{list } op$ is an ordered list of operations (more information is needed in the PUSH/PULL semantics, discussed later).

Parameter 3.1 (Sequential specification: allowed). *The sequential specification is a predicate on operation lists: allowed ℓ . We require that it be prefix closed.*

For convenience we will also write ℓ allows $\langle m, \sigma_1, \sigma_2, id \rangle$ which simply means allowed $\ell \cdot \langle m, \sigma_1, \sigma_2, id \rangle$. For example, if we have a simple TM based on memory read/write operations we might specify allowed $\ell \cdot \langle a := x, [x \mapsto 5], [x \mapsto 5, a \mapsto 5], id \rangle$, but \neg allowed $\ell \cdot \langle a := x, [x \mapsto 5], [x \mapsto 5, a \mapsto 3], id \rangle$ or more elaborate specifications that involve multiple tasks.

Ultimately, we expect the allowed predicate to be induced by the implementation’s operations on the state, $\llbracket op \rrbracket : \mathcal{P}(\text{State} \times \text{State})$, and the initial states, I . If we give a denotation to logs as $\llbracket \ell \cdot op \rrbracket \equiv \llbracket \ell \rrbracket ; \llbracket op \rrbracket$, and $\llbracket \epsilon \rrbracket \equiv I$, where $S; R \equiv \{s' \mid \exists s \in S. (s, s') \in R\}$. Then we can define allowed ℓ simply by checking if the denotation is non-empty, ($\llbracket \ell \rrbracket \neq \emptyset$).

We define a precongruence over operation logs $\ell_1 \preceq \ell_2$ coinductively, by requiring that all allowed extensions of the log ℓ_1 , are also allowed extension to the log ℓ_2 . This definition will ultimately be used in the simulation between PUSH/PULL and an atomic machine. We use a coinductive definition so that the precongruence can be defined up to all infinite suffixes.

Definition 3.1 (Shared log precongruence \preceq). *For all ℓ_1, ℓ_2 ,*

$$\frac{\text{allowed } \ell_1 \Rightarrow \text{allowed } \ell_2 \quad \forall op. (\ell_1 \cdot op) \preceq (\ell_2 \cdot op)}{\ell_1 \preceq \ell_2}$$

We use a double-line here to indicate greatest fixpoint.

¹For a discussion, see [19].

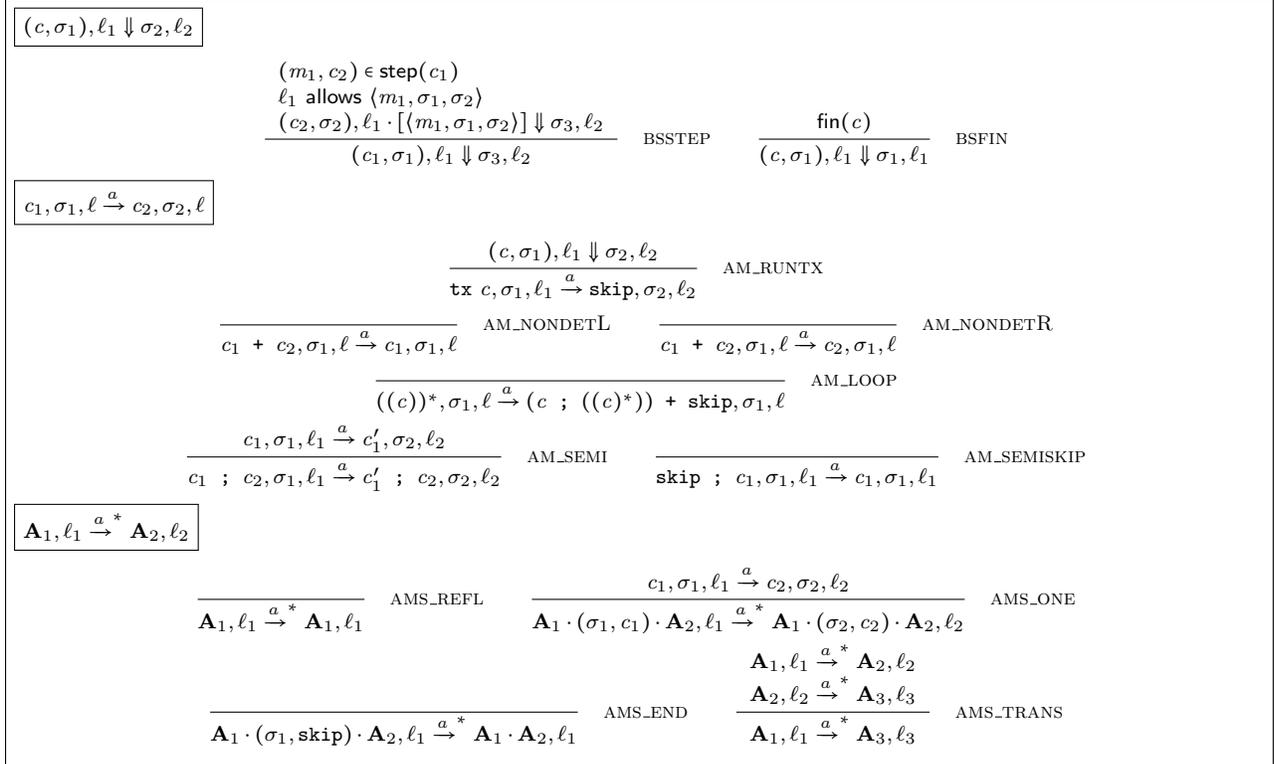


Figure 3: Atomic semantics of concurrent threads.

Informally, the above definition says that there is no sequence of observations we can make of ℓ_2 , that we can't also make of ℓ_1 . This is more general than just considering the set of states reached from executing the first log is included in the second: unobservable state differences are also permitted.

Atomic semantics. We define a simple atomic semantics, given in Figure 3 in which transactions are executed instantly, without interruption from concurrent threads. The semantics is a relation $\xrightarrow{a^*}$ over pairs consisting of a list of concurrent threads \mathbf{A} and a shared state ℓ . A single thread $(\sigma, c) \in \mathbf{A}$ is a local stack and code c . The relation $\xrightarrow{a^*}$ is reflexive, transitive, and permits a thread to complete (rules AMS_REFL, AMS_TRANS, AMS_END, respectively).

According to the rule AMS_ONE, a single thread can be reduced using the \xrightarrow{a} relation which is defined inductively over the structure of c . The rules for $+$, $;$, c^* and **skip** AM_NONDET L, AM_NONDET R, AM_LOOP, AM_SEMI, and AM_SEMISKIP are standard. The rule AM_RUNTX atomically executes the entire transaction **tx** c via the \Downarrow reduction (*i.e.* big step semantics). The big step semantics \Downarrow uses **step**() and **fin**() (rules BSSTEP and BSFIN, respectively) to scan through the nondeterminism in **tx** c to find a next operation name m or a path to **skip** denoting the end of the transaction. BSSTEP can be taken provided that the operation op is permitted by the sequential specification and that c_2 can be entirely reduced.

4 The PUSH/PULL Model

In this section we describe PUSH/PULL, an expressive model of serializable transactions. Concurrent threads execute the language described in the previous section but now transaction interleavings are possible. Moreover, we describe reductions APP, UNAPP, PUSH, UNPUSH, PULL, UNPULL, CMT, which can be made by a

$\frac{\mathbf{T}_1, G_1 \rightarrow^* \mathbf{T}_2, G_2}{\mathbf{T}_1, G_1 \rightarrow^* \mathbf{T}_3, G_3}$	MS_TRANS	$\frac{}{\mathbf{T}_1, G_1 \rightarrow^* \mathbf{T}_1, G_1}$	MS_REFL	$\frac{\mathbf{T}_1, G_1 \xrightarrow{\text{rt}} \mathbf{T}_2, G_2}{\mathbf{T}_1, G_1 \rightarrow^* \mathbf{T}_2, G_2}$	MS_ONE
$\frac{T_1, G_1 \xrightarrow{\text{rt}} T_2, G_2}{\mathbf{T}_1 \cdot T_1 \cdot \mathbf{T}_2, G_1 \xrightarrow{\text{rt}} \mathbf{T}_1 \cdot T_2 \cdot \mathbf{T}_2, G_2}$	MS_SELECT	$\frac{}{\mathbf{T}_1 \cdot \{\text{skip}, \sigma_1, L_1\} \cdot \mathbf{T}_2, G_1 \rightarrow^* \mathbf{T}_1 \cdot \mathbf{T}_2, G_1}$			MS_END

Figure 4: The machine reductions of PUSH/PULL. See Figure 5 for individual steps.

given transaction to control how its effects are shared with the environment or view the effects made by the environment.

As in the atomic semantics, the PUSH/PULL semantics has a reflexive, transitive reduction $\mathbf{T}, G \rightarrow^* \mathbf{T}', G'$ that reduces a list of threads $\mathbf{T} : \text{list}(c \times \sigma \times L)$ and a global log G to \mathbf{T}', G' . L and G are local and global operation logs, respectively, described in more detail below.

The reductions of the form $\mathbf{T}, G \rightarrow^* \mathbf{T}', G'$ are given in Figure 4. As in the atomic machine, \rightarrow^* is transitive (MS_TRANS). and reflexive (MS_REFL). The MS_END rule removes a completed thread (*i.e.* a thread that has reached `skip`) from the list of threads. The MS_SELECT rule reduces a single thread via the \rightarrow relation. Finally, the MS_ONE rule incorporates this reduction into the \rightarrow^* relation.

The single-thread reduction relation $\xrightarrow{\text{rt}}$ is defined inductively over c and has three types:

$$\xrightarrow{\text{rt}} ::= \xrightarrow{\text{struct}} \mid \xrightarrow{\text{fwd}} \mid \xrightarrow{\text{back}}$$

The structural $\xrightarrow{\text{struct}}$ reductions depend on the language. For the example language mentioned earlier, there are $\xrightarrow{\text{struct}}$ rules for nondeterministic choice, nondeterministic looping and sequential composition as in Figure 6 (NONDET_L, NONDET_R, LOOP, SEMI, SEMISKIP). The four $\xrightarrow{\text{fwd}}$ rules APP, CMT, PUSH, and PULL pertain to transactions making forward progress and the $\xrightarrow{\text{back}}$ rules UNAPP, UNPUSH, UNPULL pertain to transactions rewinding.

Figure 5 lists the seven proof rules that form the core of PUSH/PULL. These rules pertain to a thread performing a transaction `tx` c and manipulate the local stack, local log, and shared log in various ways. The local log $L : \text{list}(op \times l)$ is a list of operations, along with an additional flag l per operation, as to the status of the operation:

$$l ::= \begin{array}{ll} \text{npshd } c & \text{(local operation)} \\ | \text{ pshd } c & \text{(local operation shared to global view)} \\ | \text{ pld} & \text{(some other txn's operation)} \end{array}$$

The `npshd` and `pshd` flags save the code c that was active when the log entry was created. There is also a global log $G : \text{list}(op \times g)$ with flag g that distinguishes between operations that have or have not been committed: $g ::= \text{gUCmt} \mid \text{gCmt}$. Each proof rule comes with criteria, labeled as APP criterion (*i*), APP criterion (*ii*), etc. Note that we lift $\epsilon, \setminus, \subseteq$ to logs as follows, using id for equality:

$$\begin{aligned} \langle m_1, \sigma_1, \sigma'_1, id_1 \rangle \in L &\equiv \exists i. \text{ let } L[i] = [\langle m_2, \sigma_2, \sigma'_2, id_2 \rangle, l] \text{ in } id_1 = id_2 \\ G \setminus L &\equiv \text{filter } (\lambda (\langle m_1, \sigma_1, \sigma'_1, id_1 \rangle, g). \langle m_1, \sigma_1, \sigma'_1, id_1 \rangle \notin L) G \\ L \subseteq G &\equiv \forall i. \text{ let } L[i] = [\langle m_1, \sigma_1, \sigma'_1, id_1 \rangle, l] \text{ in } \langle m_1, \sigma_1, \sigma'_1, id_1 \rangle \in G \end{aligned}$$

Here the notation $L[i]$ refers to the i th list element of L .

The APP rule. APP is similar to the BSSTEP rule in the atomic semantics: it applies if there is a nondeterministic path in code c_1 that reaches a method m_1 (with continuation code c_2). APP criterion (*ii*) specifies that method m_1 must be *allowed* by the sequential specification with post-stack σ_2 . If so, the new operation

PUSH/PULL Rules

$\frac{\begin{array}{l} (i)\text{- } (m_1, c_2) \in \text{step}(c_1) \\ (ii)\text{- } L_1 \text{ allows } \langle m_1, \sigma_1, \sigma_2 \rangle \\ (iii)\text{- } \text{fresh}(id) \end{array}}{\{\text{tx } c_1, \sigma_1, L_1\}, G_1 \xrightarrow{\text{fwd}} \{\text{tx } c_2, \sigma_2, L_1 \cdot [(m_1, \sigma_1, \sigma_2, id), \text{npshd } c_1]\}, G_1} \text{ APP}$	
$\frac{}{\{\text{tx } c_1, \sigma_1, L_1 \cdot [(m_1, \sigma_2, \sigma_3, id), \text{npshd } c_2]\}, G_1 \xrightarrow{\text{back}} \{\text{tx } c_2, \sigma_2, L_1\}, G_1} \text{ UNAPP}$	
$\frac{\begin{array}{l} (i)\text{- } op \blacktriangleleft [L_1]_{\text{npshd}} \\ (ii)\text{- } [G_1]_{\text{gUCmt}} \setminus [L_1 \cdot L_2]_{\text{pshd}} \blacktriangleleft op \\ (iii)\text{- } G_1 \text{ allows } op \end{array}}{\{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{npshd } c_2] \cdot L_2\}, G_1 \xrightarrow{\text{fwd}} \{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{pshd } c_2] \cdot L_2\}, G_1 \cdot [op, \text{gUCmt}]}$	PUSH
$\frac{\begin{array}{l} (i)\text{- } \text{allowed } G_1 \cdot G_2 \\ (ii)\text{- } [L_2]_{\text{pshd}} \blacktriangleleft op \end{array}}{\{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{pshd } c_2] \cdot L_2\}, G_1 \cdot [op, g] \cdot G_2 \xrightarrow{\text{back}} \{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{npshd } c_2] \cdot L_2\}, G_1 \cdot G_2}$	UNPUSH
$\frac{\begin{array}{l} (i)\text{- } op \notin L \\ (ii)\text{- } L \text{ allows } op \\ (iii)\text{- } op \blacktriangleleft [L]_{\text{pshd}} \cup [L]_{\text{npshd}} \end{array}}{\{\text{tx } c_1, \sigma_1, L\}, G_1 \cdot [op, g] \cdot G_2 \xrightarrow{\text{fwd}} \{\text{tx } c_1, \sigma_1, L \cdot [op, \text{pld}]\}, G_1 \cdot [op, g] \cdot G_2}$	PULL
$\frac{(i)\text{- } \text{allowed } L_1 \cdot L_2}{\{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{pld}] \cdot L_2\}, G \xrightarrow{\text{back}} \{\text{tx } c_1, \sigma_1, L_1 \cdot L_2\}, G} \text{ UNPULL}$	
$\frac{\begin{array}{l} (i)\text{- } \text{fin}(c_1) \\ (ii)\text{- } L_1 \subseteq G_1 \\ (iii)\text{- } [L_1]_{\text{pld}} \subseteq [G_1]_{\text{gCmt}} \\ (iv)\text{- } \text{cmt}(G_1, L_1, G_2) \end{array}}{\{\text{tx } c_1, \sigma_1, L_1\}, G_1 \xrightarrow{\text{fwd}} \{\text{skip}, \sigma_1, []\}, G_2} \text{ CMT}$	<p>The CMT rule involves predicate <code>cmt</code>, defined as follows:</p> $\text{cmt}(G_1, L_1, G_2) \equiv G_2 = \text{map} \left(\lambda (op, g). \begin{cases} (op, \text{gCmt}) & \text{if } op \in [L_1]_{\text{pshd}} \\ (op, g) & \text{otherwise} \end{cases} \right)_{G_1}$

Figure 5: The PUSH/PULL rules. Notations $\setminus, \notin, \cdot, \subseteq$ are all lifted to lists where equality is given by *ids*, as discussed below. We will refer to the premise criteria of each rule as, for example, “PUSH criterion (ii).” Standard rules for reducing nondeterminism in the input language are displayed in Figure 6. Criteria that are written in gray font are not strictly necessary. See inline discussion.

$\frac{}{\{\text{c}_1 + \text{c}_2, \sigma_1, L\}, G \xrightarrow{\text{struct}} \{\text{c}_1, \sigma_1, L\}, G} \text{ NONDET\text{L}}$	$\frac{}{\{\text{c}_1 + \text{c}_2, \sigma_1, L\}, G \xrightarrow{\text{struct}} \{\text{c}_2, \sigma_1, L\}, G} \text{ NONDETR}$
$\frac{}{\{((c))^*, \sigma_1, L\}, G \xrightarrow{\text{struct}} \{(c ; ((c))^*) + \text{skip}, \sigma_1, L\}, G} \text{ LOOP}$	
$\frac{\{c_1, \sigma_1, L_1\}, G_1 \xrightarrow{\text{rt}} \{c'_1, \sigma_2, L_2\}, G_2}{\{c_1 ; c_2, \sigma_1, L_1\}, G_1 \xrightarrow{\text{rt}} \{c'_1 ; c_2, \sigma_2, L_2\}, G_2} \text{ SEMI}$	$\frac{}{\{\text{skip} ; c_1, \sigma_1, L_1\}, G_1 \xrightarrow{\text{struct}} \{c_1, \sigma_1, L_1\}, G_1} \text{ SEMISKIP}$

Figure 6: PUSH/PULL transaction reductions for standard input language features. Notice that the type of the SEMI reduction is inductive.

is appended to the local log L_1 with fresh operation id_1 (formalization of **fresh** in APP criterion (iii) is omitted). Intuitively, APP applies some next method m_1 locally but does not yet share it by sending it to the global log; it is marked as such with flag **npshd**. The APP rule also records the pre-code c_1 in the local log so that the transaction can later be reversed (*i.e.* aborted or undone). Indeed, the rule UNAPP moves backwards by taking the last item in the local log and, provided that it is still **npshd**, recalls the previous local stack and code.

The PUSH rule. A transaction may choose to share its effects with the global view via the PUSH rule. This reduction changes an operation’s flag from **npshd** to **pshd** in the local log and appends the operation to the global log, provided three conditions hold. These conditions use the notion of *left-mover* which is an algebraic property of operations and due to Lipton [24]. We provide a novel coinductive definition of left-mover that builds upon log precongruence (*i.e.* a form of observational equivalence):

Definition 4.1 (Left-mover [24], over logs). *For all op_1, op_2*

$$op_1 \blacktriangleleft op_2 \equiv \forall \ell. \ell \cdot \{op_1, op_2\} \preceq \ell \cdot \{op_2, op_1\}.$$

Intuitively, operation op_1 can move to the left of operation op_2 provided that whenever we are allowed to do $op_1 \cdot op_2$, we are also allowed to do $op_2 \cdot op_1$ and the resulting log is the same (precongruent). The proof of serializability involves several fairly straight-forward lemmas pertaining to **allowed** and left/right moverness, omitted for lack of space.

PUSH criterion (i) specifies that the pushed operation op is able to move to the left of all unpushed operations in the local log. This, intuitively, means that we can publish op as if it was the next thing to happen after all the operations published thus far by the current transaction. Formally, we lift \blacktriangleleft to lists and define projections such as $[L_1]_{\text{npshd}}$ using:

$$[L_1]_l \equiv \text{map fst (filter } (\lambda(op, l'). l = l') L_1)$$

and similar for $[G]_{\text{gUCmt}}$.

Application example: To our knowledge, all existing implementations satisfy this trivially because operations are PUSHed in the same order that they are APPLIED.

PUSH criterion (ii) is that *all uncommitted operations in the shared log $[G]_{\text{gUCmt}}$ —except those due to the current transaction—can move to the right of the current operation op .* This condition ensures that if the transaction commits at any point, it can serialize before all concurrent uncommitted transactions. (Recall that we have lifted \succ to lists where equality is given by the operation IDs and the order is determined by the first operand (in this case, G_1 .)

Application examples: A boosted transaction immediately performs a PUSH at the linearization point because it modifies the shared state in place. Optimistic STMs don’t perform PUSH until commit-time (unless there is some early conflict detection [13] which involves a form of PUSH). In boosted [11] and open nested [28] transactions, a commutativity requirement is sufficient to ensure this condition.

PUSH criterion (iii) is that op is **allowed** by the sequential specification of the global log. (Here we have lifted **allowed** to global logs.)

The UNPUSH rule. An operation op that has been PUSHed to the shared log can be UNPUSHed. This amounts to swapping the local flag from **pshd** to **npshd** and removing the corresponding global log entry for op . PUSH criterion (i) ensures that G_2 does not depend on op and PUSH criterion (ii) is that everything pushed chronologically after op could still have been pushed if op hadn’t been pushed. Note that PUSH criterion (i) is not strictly necessary because we can prove that it must hold whenever an UNPUSH occurs.

Application: When a boosted transaction aborts (*e.g.* due to deadlock) it must undo its effects on the shared state. This is modeled via the UNPUSH rule and typically implemented via inverse operations (such as **remove** on an element that had been **added**).

The PULL rule. Transactions can learn about the published effects of other transactions by PULLing operations from the global log into their local logs. An operation op can be pulled from the global log provided that it wasn't pulled before (PULL criterion (i)) and that the local log allows it (PULL criterion (ii)) according to the sequential specification. A transaction can only learn about the shared state through PULLing. In most applications, a transaction will PULL operations in chronological order. However, there are many examples for which this is not true. In a transaction that operates over two shared data-structures a and b , it may PULL in the effects on a even if they occurred after the effects on b because the transaction is only interested in modifying a . When the PULL rule occurs, the operation is appended to the local log L and marked as `pld`.

Finally, PULL criterion (iii) is that *everything that the current transaction has currently done locally must be able to move to the right of op* . This ensures that the transaction can behave as if the pulled effect preceded the transaction. We have marked this criterion in gray, indicating that it is not strictly necessary. One could imagine allowing transactions to PULL uncommitted, conflicting effects. However, we don't believe such behaviors to be particularly interesting or realistic.

Application: Many traditional STMs are opaque [10] (transactions cannot view the effects of other uncommitted transactions). Such systems never execute PULL operations marked as `gUCmt` and can only view operations that have been marked `gCmt`.

Application: Some (non-opaque) transaction A may become *dependent* [30] on another transaction B if the effects of B are released to A before B commits. This is captured by B performing a PUSH of some effects that are then PULLED by A even though B has not committed.

The UNPULL rule. A PULLED operation may be removed from the local log. UNPULL criterion (i) is that the local log is *allowed* without operation op . Informally, this means that the transaction must not have done anything that depended on op . Without this criterion the local log might become invalid with respect to the sequential specification.

Applications: Breaking dependencies [30].

The CMT rule. If there is a path through `tx` c that reaches `skip` (CMT criterion (i)), then the transaction can commit. There are three additional conditions: CMT criterion (ii) is that the local log L_1 must be contained within the global log G_1 , indicating that *all of the transaction's operations have been pushed*. CMT criterion (iii) says that *all pulled operations correspond to transactions that have been committed*. Finally, CMT criterion (iv) is that the global log is updated to G_2 in which all of the transaction's operations are marked as committed. This is achieved with the `cmt(G_1, L_1, G_2)` predicate, defined at the bottom of Figure 5. The CMT rule serves as the instantaneous moment when all of a transaction's effects become permanent. Note that a transaction does not have to PULL all committed operations. Instead, transactions check whether they conflict with other transactions' operations each time they PUSH an operation.

5 Serializability

In this section we present the proof of serializability of PUSH/PULL. The proof is achieved by a simulation between a PUSH/PULL machine and the atomic semantics.

5.1 Mnemonics

There are a few helpful mnemonics to keep in mind for this proof. The first is a relationship between \blacktriangleleft and \preceq . When there is a relationship of the form $op_1 \blacktriangleleft op_2$, and we have $\ell \cdot op_1 \cdot op_2$ then, after swapping the ops, we have the relationship $\ell \cdot op_1 \cdot op_2 \preceq \ell \cdot op_2 \cdot op_1$. That is, the order of the operations in the expression $op_1 \blacktriangleleft op_2$ is the order they will be in the log on the left-hand side of \preceq . Generally speaking, we will refer to logs on the right-hand side of \preceq as the hypothetical log (the log of the atomic machine we simulate).

5.2 Lemmas

We develop lemmas that later help us establish the simulation relation between PUSH/PULL and the atomic semantics.

Lemma 5.1. *For all ℓ_1, ℓ_2, op , we have that $\ell_2 \blacktriangleleft op \wedge \text{allowed } \ell_1 \cdot \ell_2 \cdot op \Rightarrow \text{allowed } \ell_1 \cdot op$.*

Proof. First we consider the case where ℓ_2 is a singleton operation op' . This holds easily from the definition of $op' \blacktriangleleft op$ and prefix closure of **allowed**. We then proceed by induction on ℓ_2 . \square

Lemma 5.2 (Transitivity of precongruence). *For all ℓ_a, ℓ_b, ℓ_c , we have that $\ell_a \leq \ell_b \wedge \ell_b \leq \ell_c \Rightarrow \ell_a \leq \ell_c$*

Proof. By coinduction, using the fact that **allowed** $\ell_a \Rightarrow$ **allowed** ℓ_c holds by transitivity of implication. \square

Lemma 5.3 (Precongruence over append). *For all ℓ_a, ℓ_b, ℓ_c , we have that $\ell_a \leq \ell_b \Rightarrow \ell_a \cdot \ell_c \leq \ell_b \cdot \ell_c$.*

Proof. By induction on ℓ_c , using the definition of \leq . \square

Lemma 5.4 (Precongruence and BSSTEP). *For all $c, \sigma, \sigma', \ell_1, \ell'_1$ and ℓ_2 ,*

$$(c, \sigma), \ell_1 \Downarrow \sigma', \ell'_1 \wedge \ell_2 \leq \ell_1 \Rightarrow \exists \ell'_2. (c, \sigma), \ell_2 \Downarrow \sigma', \ell'_2 \wedge \ell'_2 \leq \ell'_1$$

Proof. We proceed by rule induction on the derivation of $(c, \sigma), \ell_1 \Downarrow \sigma', \ell'_1$. In the base case, $\ell'_1 = \ell_1$. \square

5.3 Invariants

In order to prove simulation, numerous invariants were necessary. We say that a predicate $P(T, G)$ is *invariant w.r.t.* I_1, \dots, I_n provided that

$$\begin{aligned} \forall \mathbf{T} \ G \ \mathbf{T}' \ G'. \quad & \mathbf{T}, G \xrightarrow{\text{tt}} \mathbf{T}', G' \Rightarrow \\ & (\forall T \in \mathbf{T}. P(T, G) \wedge I_1(T, G) \wedge \dots \wedge I_n(T, G)) \Rightarrow \\ & (\forall T' \in \mathbf{T}'. P(T', G')) \end{aligned}$$

Lemma 5.5. *For all \mathbf{T}, G , if $P(T, G)$ is invariant and $\mathbf{T}, G \rightarrow^* \mathbf{T}', G'$, then $P(T', G')$ for \mathbf{T}' .*

Proof. Induction on \rightarrow^* . \square

We use the following lemma to prove properties are invariant.

Lemma 5.6. *If for all reductions $T, G \rightarrow T', G'$,*

1. *if $P(T, G)$ and $\bigwedge_{i \in [1, n]} I_i(T, G)$ implies $P(T', G')$; and*
2. *for all \hat{T} , if $P(T, G)$ and $\bigwedge_{i \in [1, n]} I_i(T, G)$ and $P(\hat{T}, G)$ and $\bigwedge_{i \in [1, n]} I_i(\hat{T}, G)$ implies $P(\hat{T}, G')$*

then P is invariant wrt I_1, \dots, I_n .

Part of proving that serializability holds, involves some log manipulations that arise from the commutativity conditions imposed by the PUSH/PULL model. This includes a few invariants. First, we have an invariant that ensures that the **pshd**/**npshd** flags in the local log are consistent with whether or not the op is in the global log:

Lemma 5.7. *The following is invariant:*

$$\begin{aligned} I_{LG}(\{c, \sigma, L\}, G) \equiv & \forall [(m, \sigma, \sigma', id), l] \in L \\ & \left\{ \begin{array}{l} l = \text{pshd } c \Rightarrow \langle m, \sigma, \sigma', id \rangle \in G \\ l = \text{npshd } c \Rightarrow \langle m, \sigma, \sigma', id \rangle \notin G \end{array} \right\} \end{aligned}$$

Proof. By Lemma 5.6 and case analysis on reduction relation. \square

Next, we have two invariants that ensure your uncommitted operations can move to the right of other transactions' uncommitted operations:

Lemma 5.8. *The following is invariant wrt I_{LG} :*

$$\begin{aligned} I_{slideR}(\{c, \sigma, L\}, G) &\equiv \forall [\langle m_1, \sigma_1, \sigma'_1, id_1 \rangle, pshd] \in L \text{ and } [\langle m_2, \sigma_2, \sigma'_2, id_2 \rangle, pshd \mid npshd] \notin L. \\ G &= G_1 \cdot [(\langle m_1, \sigma_1, \sigma'_1, id_1 \rangle, gUCmt)] \cdot G_2 \cdot [(\langle m_2, \sigma_2, \sigma'_2, id_2 \rangle, g)] \cdot G_3 \\ \Rightarrow &\langle m_1, \sigma_1, \sigma'_1, id_1 \rangle \blacktriangleleft \langle m_2, \sigma_2, \sigma'_2, id_2 \rangle \end{aligned}$$

Here the notion $pshd \mid npshd$ means that the label can be either $pshd$ or $npshd$, but not pld .

Proof. To prove that I_{slideR} is invariant there are two cases to consider:

1. Let $T \equiv \{c, \sigma, L\} \in \mathbf{T}$ be the thread that took a step to $T' \equiv \{c', \sigma', L'\} \in \mathbf{T}'$. Assume $I_{slideR}(\{c, \sigma, L\}, G)$.
Claim: $I_{slideR}(\{c', \sigma', L'\}, G')$.

Pf. All cases (APP, UNAPP, PUSH, UNPUSH, PULL, UNPULL, CMT) are trivial.

2. Let $\{c, \sigma, L\} \in \mathbf{T}$ be the thread that took a step and let $\{\hat{c}, \hat{\sigma}, \hat{L}\} \in \mathbf{T}$ be another thread. Assume $I_{slideR}(\{c, \sigma, L\}, G)$ and $I_{slideR}(\{\hat{c}, \hat{\sigma}, \hat{L}\}, G)$.

Claim: $I_{slideR}(\{\hat{c}, \hat{\sigma}, \hat{L}\}, G')$.

Pf. Proof by case analysis on the step taken by $\{c, \sigma, L\}$, cases APP, UNAPP, UNPUSH, PULL, UNPULL, CMT are trivial. For case PUSH we rely on PUSH criterion (ii). □

Lemma 5.9. *The following is invariant wrt I_{LG} and I_{slideR} :*

$$I_{slidePushed}(\{c, \sigma, L\}, G) \equiv G \preceq (G \setminus [L]_{pshd}) \cdot (G \cap [L]_{pshd})$$

Note that \setminus and \cap preserves the order of their first arguments.

Proof. By induction on L , using Lemma 5.8. □

The next invariants intuitively mean that if a transaction PUSHes operations out of order, the resulting log bares some precongruence to a log in which the operations were pushed in the correct order.

Lemma 5.10. *For all c, σ, L, G , the following is invariant:*

$$\begin{aligned} I_{reorderPUSH}(\{c, \sigma, L_1 \cdot L_2\}, G) &\equiv \forall [\langle m_1, \sigma_1, \sigma'_1, id_1 \rangle, pshd \mid npshd] \in L_1 \text{ and } [\langle m_2, \sigma_2, \sigma'_2, id_2 \rangle, pshd \mid npshd] \in L_2 \\ G &= G_1 \cdot [(\langle m_2, \sigma_2, \sigma'_2, id_2 \rangle, gUCmt)] \cdot G_2 \cdot [(\langle m_1, \sigma_1, \sigma'_1, id_1 \rangle, gUCmt)] \cdot G_e \\ \Rightarrow &\langle m_2, \sigma_2, \sigma'_2, id_2 \rangle \blacktriangleleft \langle m_1, \sigma_1, \sigma'_1, id_1 \rangle \end{aligned}$$

Proof. To prove that $I_{reorderPUSH}$ is invariant there are two cases to consider:

1. Let $t \equiv \{c, \sigma, L\} \in \mathbf{T}$ be the thread that took a step to $T' \equiv \{c', \sigma', L'\} \in \mathbf{T}'$. Assume $I_{reorderPUSH}(\{c, \sigma, L\}, G)$.
Claim: $I_{reorderPUSH}(\{c', \sigma', L'\}, G')$.

Pf. Cases APP, UNAPP, PULL, UNPULL, CMT are trivial. Case UNPUSH uses UNPUSH criterion (ii) and case PUSH uses PUSH criterion (i).

2. Let $\{c, \sigma, L\} \in \mathbf{T}$ be the thread that took a step and let $\{\hat{c}, \hat{\sigma}, \hat{L}\} \in \mathbf{T}$ be some other thread. Assume $I_{reorderPUSH}(\{c, \sigma, L\}, G)$ and $I_{reorderPUSH}(\{\hat{c}, \hat{\sigma}, \hat{L}\}, G)$.

Claim: $I_{reorderPUSH}(\{\hat{c}, \hat{\sigma}, \hat{L}\}, G')$.

Pf. Trivial because $I_{reorderPUSH}$ only pertains to a transaction's own operations. □

The following Lemma indicates a log equivalence between one in which operations have been pushed in non-chronological order, and a log in which they have been pushed chronologically.

Lemma 5.11. *The following is invariant wrt $I_{reorderPUSH}$*

$$I_{chronPush}(\{c, \sigma, L\}, G) \equiv (G \setminus \lfloor L \rfloor_{pshd}) \cdot (G \cap \lfloor L \rfloor_{pshd}) \preceq (G \setminus \lfloor L \rfloor_{pshd}) \cdot \lfloor L \rfloor_{pshd}$$

Proof. We prove a slightly stronger property

$$\begin{aligned} & (G \setminus \lfloor L_1 \cdot L_2 \rfloor_{pshd}) \cdot (G \cap \lfloor L_1 \cdot L_2 \rfloor_{pshd}) \preceq \\ & (G \setminus \lfloor L_1 \cdot L_2 \rfloor_{pshd}) \cdot (G \cap \lfloor L_1 \rfloor_{pshd}) \cdot \lfloor L_2 \rfloor_{pshd} \end{aligned}$$

By induction on the size of L_2 . □

The next invariants intuitively mean that any operation $[\langle m_1, \sigma_1, \sigma'_1, id_1 \rangle, npshd \ c_1]$ can move to the left of some $[\langle m_2, \sigma_2, \sigma'_2, id_2 \rangle, pshd \ c_2]$ provided that the first operation is earlier than the second operation in the local log.

Lemma 5.12. *The following is invariant*

$$I_{localOrder}(\{c, \sigma, L\}, G) \equiv \left\{ \begin{array}{l} L = L_1 \cdot [\langle m_2, \sigma_2, \sigma'_2, id_2 \rangle, npshd \ c_2] \cdot L_2 \cdot [\langle m_1, \sigma_1, \sigma'_1, id_1 \rangle, pshd \ c_1] \cdot L_3 \\ \Rightarrow \langle m_1, \sigma_1, \sigma'_1, id_1 \rangle \blacktriangleleft \langle m_2, \sigma_2, \sigma'_2, id_2 \rangle \end{array} \right\}$$

Proof. To prove that $I_{localOrder}$ is invariant there are two cases to consider:

1. Let $t \equiv \{c, \sigma, L\} \in \mathbf{T}$ be the thread that took a step to $T' \equiv \{c', \sigma', L'\} \in \mathbf{T}'$. Assume $I_{localOrder}(\{c, \sigma, L\}, G)$.
Claim: $I_{localOrder}(\{c', \sigma', L'\}, G')$.

Pf. Cases APP, UNAPP, PULL, UNPULL, CMT are trivial. Case PUSH uses PUSH criterion (i). Case UNPUSH uses UNPUSH criterion (i).

2. Let $\{c, \sigma, L\} \in \mathbf{T}$ be the thread that took a step and let $\{\hat{c}, \hat{\sigma}, \hat{L}\} \in \mathbf{T}$ be some other thread. Assume $I_{localOrder}(\{c, \sigma, L\}, G)$ and $I_{localOrder}(\{\hat{c}, \hat{\sigma}, \hat{L}\}, G)$.
Claim: $I_{localOrder}(\{\hat{c}, \hat{\sigma}, \hat{L}\}, G')$.

Pf. Trivial because $I_{localOrder}$ only pertains to a transaction's own operations. □

The following lemma indicates that we can reorder a given thread's operations to match the order they were applied in the local log.

Lemma 5.13. *The following is invariant wrt $I_{localOrder}$:*

$$I_{localReorder}(\{c, \sigma, L\}, G) \equiv (G \setminus \lfloor L \rfloor_{pshd}) \cdot \lfloor L \rfloor_{pshd} \cdot \lfloor L \rfloor_{npshd} \preceq (G \setminus \lfloor L \rfloor_{pshd}) \cdot \lfloor L \rfloor_{npshd}^{pshd}$$

Proof. We prove the stronger property

$$\begin{aligned} & (G \setminus \lfloor L_1 \cdot L_2 \rfloor_{pshd}) \cdot \lfloor L_1 \cdot L_2 \rfloor_{pshd} \cdot \lfloor L_1 \cdot L_2 \rfloor_{npshd} \preceq \\ & (G \setminus \lfloor L \rfloor_{pshd}) \cdot \lfloor L_1 \rfloor_{pshd}^{npshd} \cdot \lfloor L_2 \rfloor_{pshd} \cdot \lfloor L_2 \rfloor_{npshd} \end{aligned}$$

by induction on L_1 . □

5.4 Commit preservation invariant

The heart of the simulation requires that we prove an invariant of the system that the shared log is equivalent to what it would be if concurrently executing transactions removed their effects and applied them atomically. More precisely, imagine that at a given moment there is a shared log G , and a given thread $T = \{c, \sigma, L\}$ atomically marks all of its pushed operations as committed, reaching a shared log of G_{post} . Note that T may still have unpushed operations $[L]_{\text{npshd}}$. The invariant states that, there is a precongruence between the shared log reached by completing T from $G_{\text{post}} \cdot [L]_{\text{npshd}}$ and the shared log that would have been reached if T rewound itself and atomically ran the entire transaction from G (that is, $G \setminus L$, *i.e.* the previous shared log, with all operations belonging to T filtered out).

As described so far, the commit preservation invariant would look like the following:

$$\begin{aligned} & \forall G_{\text{post}}. \text{cmt}(G, L, G_{\text{post}}) \Rightarrow \\ & \forall \sigma', \ell_a. (c, \sigma), G_{\text{post}} \cdot [L]_{\text{npshd}} \Downarrow \sigma', \ell_a \Rightarrow \\ & \exists \ell_b. \text{otx}(\{c, \sigma, L\}), G \setminus L \Downarrow \sigma', \ell_b \quad \wedge \quad \ell_a \leq \ell_b \end{aligned}$$

where otx rewinds the transaction to its original state/code, recorded in L as follows:

$$\begin{aligned} \text{otx}(\{c, \sigma, []\}) & \equiv \{c, \sigma, []\} \\ \text{otx}(\{c, \sigma, [(m_1, \sigma_1, \sigma'_1, id_1), \text{npshd } 'c] \cdot L\}) & \equiv \{c, \sigma_1, []\} \\ \text{otx}(\{c, \sigma, [(m_1, \sigma_1, \sigma'_1, id_1), \text{pshd } 'c] \cdot L\}) & \equiv \{c, \sigma_1, []\} \\ \text{otx}(\{c, \sigma, [(m_1, \sigma_1, \sigma'_1, id_1), \text{pld}] \cdot L\}) & \equiv \text{otx}(\{c, \sigma, L\}) \end{aligned}$$

Partial rewind. This is not enough to give us the simulation result as the property is not an invariant. As the system makes steps, which undo operations from the logs, the property must be closed with respect to these *backwards* steps. Thus we need the above to hold after any partial rewinding of the local log and/or partial removal of other transactions' uncommitted operations in the shared log.

Definition 5.1 (Self-rewind). *A transaction's self-rewind denoted $\{c, \sigma, L\}, G \circ_{\text{self}} \{c, \sigma, L\}, G$ is defined as:*

$$\begin{aligned} & \overline{\{c, \sigma, L \cdot \langle m, \sigma_1, \sigma_2, id, \text{npshd } 'c \rangle\}, G \circ_{\text{self}} \{c, \sigma_1, L\}, G} \text{ PRU} \\ & \overline{\{c, \sigma, L \cdot \langle m_1, \sigma_1, \sigma'_1, id_1, \text{pshd } 'c \rangle\}, G_1 \cdot [(m_1, \sigma_1, \sigma'_1, id_1), \text{gUCmt}] \cdot G_2 \circ_{\text{self}} \{c, \sigma_1, L\}, G_1 \cdot G_2} \text{ PRP} \\ & \overline{\{c, \sigma, L\}, G \circ_{\text{self}} \{c, \sigma, L'\}, G} \text{ PRM} \quad \overline{\{c, \sigma, L\}, G \circ_{\text{self}} \{c, \sigma, L\}, G} \text{ PRR} \end{aligned}$$

The self-rewind allows us to cope with the fact that a transaction may have PULLED operations from another uncommitted transaction. In particular, we preserve the fact that the transaction *may* be able to detangle from the uncommitted transaction, and atomically commit. In the definition above of \circ_{self} , the first two cases describe a transaction rewinding its log to an operation that has been APPLIED but may or may not have been PUSHED. The transactional may pass beyond a PULLED operation via the third rule. Finally, the relation is reflexive.

The *shared log* partial rewind permits uncommitted operations of other transactions to be dropped from the shared log, and is defined as follows:

$$\overline{G_1 \circ_L G_1 \quad \langle m, \sigma, \sigma', id \rangle \notin L \quad G_2 \circ_L G_2} \quad \overline{G_1 \cdot [(m, \sigma, \sigma', id), \text{gUCmt}] \cdot G_2 \quad \circ_L G_1 \cdot G_2} \quad \overline{G \quad \circ_L G}$$

We can now state the commit preservation invariant as follows:

Definition 5.2 (Commit preservation invariant). For all G ,

$$\begin{aligned}
\text{cmtpres}(c, \sigma, L, G) &\equiv \left\{ \begin{array}{l} \forall \text{``}G. G \circlearrowleft \text{``}G. \Rightarrow \quad (0) \\ \forall \{c, \sigma, L\}. \{c, \sigma, L\}, \text{``}G \circlearrowleft_{\text{self}} \{c, \sigma, L\}, \text{``}G \Rightarrow \quad (1) \\ \forall G_{\text{post}}. \text{cmt}(\text{``}G, \text{``}L, G_{\text{post}}) \Rightarrow \quad (2) \\ \forall \sigma', \ell_a. (c, \sigma), G_{\text{post}} \cdot [L]_{\text{npshd}} \Downarrow \sigma', \ell_a \Rightarrow \quad (3) \\ \exists \ell_b. \text{otx}(\{c, \sigma, L\}), G \setminus L \Downarrow \sigma', \ell_b \wedge \ell_a \leq \ell_b \quad (4) \end{array} \right.
\end{aligned}$$

Intuitively, this invariant means that under any dropping of others' uncommitted operations (Line 0) and after partially rewinding $\circlearrowleft_{\text{self}}$ your local transaction to some local log L (Line 1), if you are now able to atomically “commit” by swapping your commit flags (Line 2) and running the rest of your transaction (Line 3), then the shared log reached ℓ_a , is contained within a shared log ℓ_b that would have been reached if the thread appended its entire transaction to G atomically.

Lemma 5.14. For all $\ell, \ell', \ell \leq \ell' \Rightarrow (\text{cmtpres}(\{c, \sigma, L\}, G) \Rightarrow \text{cmtpres}(\{c, \sigma, L'\}, G))$

Proof. Trivial. □

The following lemma indicates that partial rewind behaviors are included in the PUSH/PULL transition system:

Lemma 5.15. *Invariance of*

$$I_{\subseteq}(\{c, \sigma, L\}, G) \equiv \left\{ \begin{array}{l} (\{c, \sigma, L\}, G \circlearrowleft_{\text{self}} \{c, \sigma, L\}, \text{``}G) \Rightarrow \\ (\{c, \sigma, L\}, G \xrightarrow{\text{rt}} \{c, \sigma, L\}, \text{``}G) \end{array} \right\}$$

Proof. Case analysis. □

Lemma 5.16. *cmtpres is invariant w.r.t. $I_{\text{slidePushed}}$, $I_{\text{chronPush}}$, $I_{\text{localReorder}}$, and I_{\subseteq} .*

Proof. We assume that cmtpres holds for some \mathbf{T}, G , that $\mathbf{T}, G \xrightarrow{\text{rt}} \mathbf{T}', G'$ and then show that cmtpres holds for \mathbf{T}', G' . We split the proof into two cases, proving each one by induction on $\xrightarrow{\text{rt}}$.

1. Let $t \equiv \{c, \sigma, L\} \in \mathbf{T}$ be the thread that took a step to $T' \equiv \{c', \sigma', L'\} \in \mathbf{T}'$. (Generally speaking, we reserve the primed notation for components of \mathbf{T}', G' , for example when we apply the invariant to T' obtaining variables such as ℓ'_a .) Assume $\text{cmtpres}(\{c, \sigma, L\}, G)$.

Claim: $\text{cmtpres}(\{c', \sigma', L'\}, G')$.

Pf. By induction, with the inductive hypothesis

$$\begin{aligned}
\forall \text{``}G. G \circlearrowleft \text{``}G. &\Rightarrow \quad (0) \\
\forall \{c, \sigma, L\}. \{c, \sigma, L\}, \text{``}G \circlearrowleft_{\text{self}} \{c, \sigma, L\}, \text{``}G &\Rightarrow \quad (1) \\
\forall G_{\text{post}}. \text{cmt}(\text{``}G, \text{``}L, G_{\text{post}}) &\Rightarrow \quad (2) \\
\forall \sigma', \ell_a. (c, \sigma), G_{\text{post}} \cdot [L]_{\text{npshd}} \Downarrow \sigma', \ell_a &\Rightarrow \quad (3) \\
\exists \ell_b. \text{otx}(\{c, \sigma, L\}), G \setminus L \Downarrow \sigma', \ell_b \wedge \ell_a \leq \ell_b &\quad (4)
\end{aligned}$$

Note in Line 4 that \setminus does not remove operations from G that have been pld into L .

Case APP: $L' = L \cdot [(m, \sigma, \sigma', id), \text{npshd } c]$. From $\xrightarrow{\text{fwd}}$, L allows $\langle m, \sigma, \sigma', id \rangle$ and $(m, c') \in \text{step}(c)$. Also, note that $\text{otx}(\{c, \sigma, L\}) = \text{otx}(\{c', \sigma', L'\})$. We must show that the invariant holds for \mathbf{T}', G' where L' has one more unpushed operation. That is, we must show that

$$\begin{aligned}
\forall \text{``}G. G \circlearrowleft \text{``}G &\Rightarrow \quad (0) \\
\forall \{c, \sigma, L\}. \{c', \sigma', L'\}, \text{``}G \circlearrowleft_{\text{self}} \{c, \sigma, L\}, \text{``}G &\Rightarrow \quad (1) \\
\forall G_{\text{post}}. \text{cmt}(\text{``}G, \text{``}L', G_{\text{post}}) &\Rightarrow \quad (2) \\
\forall \sigma', \ell_a. (c', \sigma'), G_{\text{post}} \cdot [L']_{\text{npshd}} \Downarrow \sigma', \ell_a &\Rightarrow \quad (3) \\
\exists \ell_b. \text{otx}(\{c', \sigma', L'\}), G \setminus L' \Downarrow \sigma', \ell_b \wedge \ell_a \leq \ell_b &\quad (4)
\end{aligned}$$

Note that $G' = G$. So we can chose the same “ G from the inductive hypothesis. Consequently, Lines 2,3,4 hold when we rewind to every *strictly* previous $\{c, \sigma, L\}$. However, we must show that those lines hold without any rewind:

$$\forall G_{\text{post}}. \text{cmt}(\langle G, L', G_{\text{post}} \rangle) \Rightarrow \quad (2)$$

$$\forall \sigma', \ell_a. (c', \sigma'), G_{\text{post}} \cdot [L']_{\text{nps hd}} \Downarrow \sigma', \ell_a \Rightarrow \quad (3)$$

$$\exists \ell_b. \text{otx}(\{c', \sigma', L'\}), G \setminus L' \Downarrow \sigma', \ell_b \wedge \ell_a \leq \ell_b \quad (4)$$

G_{post} is the same as in the inductive hypothesis because the new operation is yet unpushed. What remains is to show that

$$\forall \sigma', \ell_a. (c, \sigma), G_{\text{post}} \cdot [L]_{\text{nps hd}} \Downarrow \sigma', \ell_a \Rightarrow \quad (i)$$

$$\exists \ell_b. \text{otx}(\{c, \sigma, L\}), G \setminus L \Downarrow \sigma', \ell_b \wedge \ell_a \leq \ell_b \quad (ii)$$

implies

$$\forall \sigma', \ell_a. (c', \sigma'), G_{\text{post}} \cdot [L]_{\text{nps hd}} \cdot [\langle m, \sigma, \sigma', id \rangle] \Downarrow \sigma', \ell_a \Rightarrow \quad (iii)$$

$$\exists \ell_b. \text{otx}(\{c, \sigma, L\}), G \setminus (L \cdot [\langle m, \sigma, \sigma', id \rangle]) \Downarrow \sigma', \ell_b \wedge \ell_a \leq \ell_b \quad (iv)$$

where $\text{cmt}(\langle G, L', G_{\text{post}} \rangle)$. Above Line (ii) is equivalent to (iv) because id is not in G . We instantiate \Downarrow in Line (i) for the case where $(m, c') \in \text{step}(c)$, which gives us precisely Line (iii).

Case PUSH: In this case, we let $L = L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{nps hd } c_\alpha] \cdot L_2$. The post-log L' is identical, except the nps hd flag is replaced with ps hd . To prove that the invariant still holds, we must again consider all rewindings of the logs. We chose the same $G \in \mathcal{C}_L$ “ G as in the inductive hypothesis. If the local log is rewound to L_1 (or some prefix thereof), then we have passed the operation in question, and Lines 2–4 of the invariant hold directly from the inductive hypothesis. What remains is to prove that the invariant holds after rewinding to some $\hat{L} \equiv L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{nps hd } c_\alpha] \cdot \hat{L}_2$ where \hat{L}_2 is a prefix of L_2 . We must now show that

$$\forall G'_{\text{post}}. \text{cmt}(\langle G, L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{ps hd } c_\alpha] \cdot \hat{L}_2, G'_{\text{post}} \rangle) \Rightarrow \quad (2)$$

$$\forall \sigma_{\text{post}}, \ell_a. (c, \sigma), G'_{\text{post}} \cdot [\hat{L}]_{\text{nps hd}} \Downarrow \sigma_{\text{post}}, \ell_a \Rightarrow \quad (3)$$

$$\exists \ell_b. \text{otx}(\{c, \sigma, L\}), G \setminus \hat{L} \Downarrow \sigma_{\text{post}}, \ell_b \wedge \ell_a \leq \ell_b \quad (4)$$

Let G_{post} be from the inductive hypothesis. Note that $G'_{\text{post}} = G_{\text{post}} \cdot [\langle m, \sigma, \sigma', id \rangle, \text{gCmt}]$. Distributing the $[\]_{\text{nps hd}}$ over append , we can drop the op in Line 3, as well as replacing G'_{post} . What remains is to show:

$$\forall \sigma_{\text{post}}, \ell_a. (c, \sigma), G_{\text{post}} \cdot [\langle m, \sigma, \sigma', id \rangle] \cdot [L_1 \cdot \hat{L}_2]_{\text{nps hd}} \Downarrow \sigma_{\text{post}}, \ell_a \Rightarrow \quad (3)$$

$$\exists \ell_b. \text{otx}(\{c, \sigma, L\}), G \setminus \hat{L} \Downarrow \sigma_{\text{post}}, \ell_b \wedge \ell_a \leq \ell_b \quad (4)$$

We now perform the following rewrites of the log in Line 3, with each successive log, having precongruence with the previous:

$$\begin{aligned} & G_{\text{post}} \cdot [\langle m, \sigma, \sigma', id \rangle] \cdot [L_1 \cdot \hat{L}_2]_{\text{nps hd}} \\ \leq & (G_{\text{post}} \setminus [L_1 \cdot \hat{L}_2]_{\text{ps hd}}) \cdot (G_{\text{post}} \cap [L_1 \cdot \hat{L}_2]_{\text{ps hd}}) \cdot [\langle m, \sigma, \sigma', id \rangle] \cdot [L_1 \cdot \hat{L}_2]_{\text{nps hd}} && \text{Lm 5.9} \\ \leq & (G_{\text{post}} \setminus [L_1 \cdot \hat{L}_2]_{\text{ps hd}}) \cdot (G_{\text{post}} \cap [L_1 \cdot \hat{L}_2]_{\text{ps hd}}) \cdot [L_1]_{\text{nps hd}} \cdot [\langle m, \sigma, \sigma', id \rangle] \cdot [L_2]_{\text{nps hd}} && \text{PUSH (i)} \\ \leq & (G_{\text{post}} \setminus [L_1 \cdot \hat{L}_2]_{\text{ps hd}}) \cdot [L_1 \cdot L_2]_{\text{ps hd}} \cdot [L_1]_{\text{nps hd}} \cdot [\langle m, \sigma, \sigma', id \rangle] \cdot [L_2]_{\text{nps hd}} && \text{Lm 5.11} \\ \leq & (G_{\text{post}} \setminus [L_1 \cdot \hat{L}_2]_{\text{ps hd}}) \cdot [L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, --] \cdot \hat{L}_2]_{\text{ps hd}}^{\text{nps hd}} && \text{Lm 5.13} \end{aligned}$$

Finally, we use Lemma 5.4 to substitute this rewritten log into place on the left side of \Downarrow on Line 3, and we have proved that the invariant still holds.

Cases UNPUSH: In this case, we let

$$\begin{aligned} L &= L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{ps hd } c_\alpha] \cdot L_2 \\ G &= G_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{gUCmt}] \cdot G_2 \end{aligned}$$

The post-log L' is identical to L , except the `pshd` flag is replaced with `npshd`. $G' = G_1 \cdot G_2$. To prove that the invariant still holds, we must again consider all rewindings of the logs. From the inductive hypothesis we have

$$G_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{gUCmt}] \cdot G_2 \quad \mathcal{C}_L \quad \text{“}G_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{gUCmt}] \cdot \text{“}G_2.$$

If the local log is rewound to L_1 (or some prefix thereof), then we have passed the operation in question, and Lines 2–4 of the invariant hold directly from the inductive hypothesis. What remains is to prove that the invariant holds after local rewinding to some

$$\hat{L} \equiv L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{pshd } c_\alpha] \cdot \hat{L}_2 \text{ and } \text{“}G_1 \cdot \text{“}G_2$$

where \hat{L}_2 is a prefix of L_2 . We must now show that

$$\forall G'_{\text{post}}. \text{cmt}(\text{“}G_1 \cdot \text{“}G_2, L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{npshd } c_\alpha] \cdot \hat{L}_2, G'_{\text{post}}) \Rightarrow \quad (2)$$

$$\forall \sigma_{\text{post}}, \ell_a. (\text{“}c, \text{“}\sigma), G'_{\text{post}} \cdot [L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{npshd } c_\alpha] \cdot \hat{L}_2]_{\text{npshd}} \Downarrow \sigma_{\text{post}}, \ell_a \Rightarrow \quad (3)$$

$$\exists \ell_b. \text{otx}(\{\text{“}c, \text{“}\sigma, \text{“}L\}, G \setminus \hat{L} \Downarrow \sigma_{\text{post}}, \ell_b) \wedge \ell_a \leq \ell_b \quad (4)$$

Here, G'_{post} is identical to G_{post} , except that $\langle m, \sigma, \sigma', id \rangle$ has been dropped. We now perform the following rewrites of the log in Line 3, with each successive log having precongruence with its previous:

$$\begin{aligned} & G'_{\text{post}} \cdot [L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{npshd } c_\alpha] \cdot \hat{L}_2]_{\text{npshd}} \\ \leq & (G'_{\text{post}} \setminus [L_1 \cdot \hat{L}_2]_{\text{pshd}}) \cdot (G'_{\text{post}} \cap [L_1 \cdot \hat{L}_2]_{\text{pshd}}) \cdot [L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{npshd } c_\alpha] \cdot \hat{L}_2]_{\text{npshd}} \quad \text{Lm 5.9} \\ \leq & (G'_{\text{post}} \setminus [L_1 \cdot \hat{L}_2]_{\text{pshd}}) \cdot [L_1 \cdot L_2]_{\text{pshd}} \cdot [L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{npshd } c_\alpha] \cdot \hat{L}_2]_{\text{npshd}} \quad \text{Lm 5.11} \\ \leq & (G'_{\text{post}} \setminus [L_1 \cdot \hat{L}_2]_{\text{pshd}}) \cdot [L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{npshd } c_\alpha] \cdot \hat{L}_2]_{\text{pshd}}^{\text{npshd}} \quad \text{Lm 5.13} \end{aligned}$$

Finally, we use Lemma 5.4 to substitute this rewritten log in to place on the left side of \Downarrow on Line 3, and prove that the invariant still holds.

Case CMT: We again pick the same “ G from the inductive hypothesis. Now, since the transaction has performed a CMT, $L' = []$ and “ $G = \text{“}G$. So we need only show that

$$\forall G'_{\text{post}}. \text{cmt}(\text{“}G, [], G'_{\text{post}}) \Rightarrow \quad (2)$$

$$\forall \sigma_{\text{post}}, \ell_a. (\text{“}c', \text{“}\sigma'), G'_{\text{post}} \Downarrow \sigma_{\text{post}}, \ell_a \Rightarrow \quad (3)$$

$$\exists \ell_b. \text{otx}(\{c', \sigma', []\}, G \Downarrow \sigma_{\text{post}}, \ell_b) \wedge \ell_a \leq \ell_b \quad (4)$$

Now, $G = G_{\text{post}}$ and $\text{otx}(\{c', \sigma', []\}) = (c', \sigma')$. Thus, Lines 3 and 4 are identical, and we can use ℓ_a as a witness for ℓ_b .

Case PULL: After PULL, $L' = L \cdot [\langle m, \sigma, \sigma', id \rangle, \text{pld}]$. For all rewinds of L' that produce prefixes of L , the invariant holds due to the inductive hypothesis. So we need only show that it holds for the unrewind L' . However, in that case this pulled operation is filtered out by the $[_{\text{npshd}}$ in Line 3 of the invariant, so it holds from the inductive hypothesis.

Case UNPULL: Let $L = L_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{pld}] \cdot L_2$. Due to the UNPULL rule, $L = L_1 \cdot L_2 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{pld}]$ is allowed. This case then follows from the fact that our inductive hypothesis is closed under rewind, Lemma 5.14, and that Lines 2,3,4 of the inductive hypothesis don't depend on PULLED operations.

Case UNAPP: This case follows from the fact that our inductive hypothesis is closed under rewind.

Cases NONDET \bar{L} , NONDET \bar{R} , LOOP, SEMI, SEMISKIP: Details omitted because they are uninteresting.

2. Let $\{c, \sigma, L\} \in \mathbf{T}$ be the thread that took a step and let $\{\hat{c}, \hat{\sigma}, \hat{L}\} \in \mathbf{T}$ be some other thread. Assume $\text{cmtpres}(\{c, \sigma, L\}, G)$ and $\text{cmtpres}(\{\hat{c}, \hat{\sigma}, \hat{L}\}, G)$.

Claim: $\text{cmtpres}(\{\hat{c}, \hat{\sigma}, \hat{L}\}, G')$.

Pf. We use the notations T, T' , and \hat{T} . Proof by induction on the step taken by $\{c, \sigma, L\}$, with inductive hypothesis

$$\forall \text{“}G. G \circ_{\hat{L}} \text{“}G \Rightarrow \quad (0)$$

$$\forall \{c, \sigma, L\} \text{“}G. \{c, \sigma, L\}, \text{“}G \circ_{\text{self}} \{c, \sigma, L\}, \text{“}G \Rightarrow \quad (1)$$

$$\forall G_{\text{post}}. \text{cmt}(G, L, G_{\text{post}}) \Rightarrow \quad (2)$$

$$\forall \sigma_{\text{post}}, \ell_a. (c, \sigma), G_{\text{post}} \cdot [L]_{\text{npshd}} \Downarrow \sigma_{\text{post}}, \ell_a \Rightarrow \quad (3)$$

$$\exists \ell_b. \text{otx}(\{c, \sigma, L\}), G \setminus L \Downarrow \sigma_{\text{post}}, \ell_b \wedge \ell_a \leq \ell_b \quad (4)$$

and

$$\forall \text{“}G. G \circ_{\hat{L}} \text{“}G \Rightarrow \quad (5)$$

$$\forall \{c, \sigma, L\} \text{“}G. \{\hat{c}, \hat{\sigma}, \hat{L}\}, \text{“}G \circ_{\text{self}} \{c, \sigma, L\}, \text{“}G \Rightarrow \quad (6)$$

$$\forall G_{\text{post}}. \text{cmt}(G, L, G_{\text{post}}) \Rightarrow \quad (7)$$

$$\forall \sigma_{\text{post}}, \ell_a. (c, \sigma), G_{\text{post}} \cdot [L]_{\text{npshd}} \Downarrow \sigma_{\text{post}}, \ell_a \Rightarrow \quad (8)$$

$$\exists \ell_b. \text{otx}(\{c, \sigma, L\}), G \setminus L \Downarrow \sigma_{\text{post}}, \ell_b \wedge \ell_a \leq \ell_b \quad (9)$$

Cases APP, UNAPP, PULL, UNPULL: In all of these cases, $G' = G$, so $\text{cmtpres}(\{\hat{c}, \hat{\sigma}, \hat{L}\}, G')$ is trivial.

Case PUSH: $G' = G \cdot [\langle m, \sigma, \sigma', id \rangle, \text{gUCmt}]$. We must show that Lines 5–9 still hold for G' . In this case, G'_{post} is similar to what it was in the inductive hypothesis, but with the additional uncommitted operation $\langle m, \sigma, \sigma', id \rangle$ that was appended by T . So we must show that:

$$\forall \text{“}G. G' \circ_{\hat{L}} \text{“}G \Rightarrow \quad (5)$$

$$\forall \{c, \sigma, L\} \text{“}G. \{\hat{c}, \hat{\sigma}, \hat{L}\}, \text{“}G \circ_{\text{self}} \{c, \sigma, L\}, \text{“}G \Rightarrow \quad (6)$$

$$\forall G'_{\text{post}}. \text{cmt}(G \cdot [\langle m, \sigma, \sigma', id \rangle, \text{gUCmt}], L, G'_{\text{post}}) \Rightarrow \quad (7)$$

$$\forall \sigma_{\text{post}}, \ell_a. (c, \sigma), G'_{\text{post}} \cdot [L]_{\text{npshd}} \Downarrow \sigma_{\text{post}}, \ell_a \Rightarrow \quad (8)$$

$$\exists \ell_b. \text{otx}(\{c, \sigma, L\}), (G \cdot [\langle m, \sigma, \sigma', id \rangle, \text{gUCmt}]) \setminus L \Downarrow \sigma_{\text{post}}, \ell_b \wedge \ell_a \leq \ell_b \quad (9)$$

In the case where $\circ_{\hat{L}}$ removes the PUSHed op, Lines 5–9 follow directly from the inductive hypothesis. Otherwise, there is now a new operation $\langle m, \sigma, \sigma', id \rangle$ that is not in the local log \hat{L} . We now perform the following rewrites of the log in Line 8, with each successive log, having precongurence with the previous:

$$\begin{aligned} & \leq (G'_{\text{post}} \setminus [L]_{\text{pshd}}) \cdot (G'_{\text{post}} \cap [L]_{\text{pshd}}) \cdot [\langle m, \sigma, \sigma', id \rangle] \cdot [L]_{\text{npshd}} && \text{Lm 5.9} \\ & \leq (G'_{\text{post}} \setminus [L]_{\text{pshd}}) \cdot [\langle m, \sigma, \sigma', id \rangle] \cdot (G'_{\text{post}} \cap [L]_{\text{pshd}}) \cdot [L]_{\text{npshd}} && \text{PUSH (ii)} \\ & \leq ((G'_{\text{post}} \cdot [\langle m, \sigma, \sigma', id \rangle]) \setminus [L]_{\text{pshd}}) \cdot (G'_{\text{post}} \cap [L]_{\text{pshd}}) \cdot [L]_{\text{npshd}} && \text{via disjointness} \\ & \leq ((G'_{\text{post}} \cdot [\langle m, \sigma, \sigma', id \rangle]) \setminus [L]_{\text{pshd}}) \cdot [L]_{\text{pshd}} \cdot [L]_{\text{npshd}} && \text{Lm 5.11} \\ & \leq ((G'_{\text{post}} \cdot [\langle m, \sigma, \sigma', id \rangle]) \setminus [L]_{\text{pshd}}) \cdot L && \text{Lm 5.13} \end{aligned}$$

Finally, we use Lemma 5.4 to substitute this rewritten log into place on the left side of \Downarrow on Line 8, and we have proved that the invariant still holds.

Case UNPUSH: Let $G = G_1 \cdot [\langle m, \sigma, \sigma', id \rangle, \text{gUCmt}] \cdot G_2$ and $G' = G_1 \cdot G_2$. We must show that:

$$\forall \text{“}G. G' \circ_{\hat{L}} \text{“}G \Rightarrow \quad (5)$$

$$\forall \{c, \sigma, L\} \text{“}G. \{\hat{c}, \hat{\sigma}, \hat{L}\}, \text{“}G \circ_{\text{self}} \{c, \sigma, L\}, \text{“}G \Rightarrow \quad (6)$$

$$\forall G_{\text{post}}. \text{cmt}(G_1 \cdot G_2, L, G_{\text{post}}) \Rightarrow \quad (7)$$

$$\forall \sigma_{\text{post}}, \ell_a. (c, \sigma), G_{\text{post}} \cdot [L]_{\text{npshd}} \Downarrow \sigma_{\text{post}}, \ell_a \Rightarrow \quad (8)$$

$$\exists \ell_b. \text{otx}(\{c, \sigma, L\}), G_1 \cdot G_2 \setminus L \Downarrow \sigma_{\text{post}}, \ell_b \wedge \ell_a \leq \ell_b \quad (9)$$

This holds based on the fact that for all $\{c, \sigma, L\}, op, G_1, G_2$ that

$$op \notin L \wedge \text{cmtpres}(\{c, \sigma, L\}, G_1 \cdot [op, \text{gUCmt}] \cdot G_2) \Rightarrow \text{cmtpres}(\{c, \sigma, L\}, G_1 \cdot G_2)$$

which we prove by induction.

Case CMT: In this case, $\text{cmt}(G, L, G')$, so all of the operations in G belonging to T are now marked as gCmt . We must show that:

$$\forall \langle G, G' \rangle \in \mathcal{O}_L \langle G \rangle \Rightarrow \quad (5)$$

$$\forall \{c, \sigma, L\} \in G. \{\hat{c}, \hat{\sigma}, \hat{L}\} \in \mathcal{O}_{\text{self}} \{c, \sigma, L\} \Rightarrow \quad (6)$$

$$\forall G'_{\text{post}}. \text{cmt}(G', L, G'_{\text{post}}) \Rightarrow \quad (7)$$

$$\forall \sigma_{\text{post}}, \ell_a. (c, \sigma), G'_{\text{post}} \cdot [L]_{\text{npshd}} \Downarrow \sigma_{\text{post}}, \ell_a \Rightarrow \quad (8)$$

$$\exists \ell_b. \text{otx}(\{c, \sigma, L\}, G' \setminus L \Downarrow \sigma_{\text{post}}, \ell_b \wedge \ell_a \preceq \ell_b \quad (9)$$

This follows easily from the fact that $[G]_{\text{gCmt}} \subseteq [G']_{\text{gCmt}}$

Cases NONDETL, NONDETR, LOOP, SEMI, SEMISKIP: Details omitted because they are uninteresting. □

5.5 Main Theorem

We now prove the main theorem via simulation between the PUSH/PULL machine and an atomic machine. The theorem depends strongly on the cmtpres invariant.

Theorem 5.17 (Serializability). *PUSH/PULL is serializable.*

Proof. Via a simulation relation between \rightarrow^* and \xrightarrow{a}^* . The simulation relation is defined as follows:

$$\mathbf{T}, G \sim \mathbf{A}, \ell \equiv (\text{map rewind } \mathbf{T}) = \mathbf{A} \wedge [G]_{\text{gCmt}} \preceq \ell$$

where $\text{rewind } T$ rolls a transaction back to its original code (details pertaining to semi-colon treatment in rewind omitted for simplicity). We define $\mathbf{T} \sim \mathbf{A}$ and $G \sim \ell$ with the appropriate conjunct from above.

To prove simulation, we show that for every $\mathbf{T}, G \rightarrow^* \mathbf{T}', G'$ such that $\mathbf{T}, G \sim \mathbf{A}, \ell$, there exists some \mathbf{A}', ℓ' such that $\mathbf{A}, \ell \xrightarrow{a}^* \mathbf{A}', \ell'$ and that $\mathbf{T}', G' \sim \mathbf{A}', \ell'$. The reflexive, transitive and thread-end rules are straight-forward. What remains is aligning AMACH_ONE and MACH_ONE . In this case, we prove a helper lemma that shows that the simulation relation holds after each single step \rightarrow . That is, for every $\mathbf{T}, G \rightarrow \mathbf{T}', G'$ such that $\mathbf{T}, G \sim \mathbf{A}, \ell$, there exists some \mathbf{A}', ℓ' such that $\mathbf{A}, \ell \xrightarrow{a}^* \mathbf{A}', \ell'$ and that $\mathbf{T}', G' \sim \mathbf{A}', \ell'$.

So we must consider each \rightarrow step from Figures 5 and 6 and show that an appropriate \mathbf{A}', ℓ' can be found. In each case, the inductive hypothesis gives us that the simulation relation **rewinds** all uncommitted transactions in \mathbf{T} to obtain \mathbf{A} and **drops** all uncommitted operations from G to obtain ℓ . Moreover, we rely on all of the invariants holding for \mathbf{T}, G as well as \mathbf{T}', G' (most significantly, the cmtpres invariant).

Case APP: This step is straight-forward. We let $\mathbf{A}' = \mathbf{A}$ and $\ell' = \ell$. The inductive hypothesis tells us that $\mathbf{T}, G \sim \mathbf{A}, \ell$ and, after an APP step, \mathbf{T}' is very similar to \mathbf{T} , except with one more operation $\langle m, \sigma_1, \sigma_2, id, \text{npshd } c \rangle$ in the local log of a single thread.

Case UNAPP: Similar to APP.

Case PUSH: In this case a new operation has been PUSHed into the shared log. Since the new operation has not been committed, it will be filtered out by rewind so, again, we can let $\mathbf{A}' = \mathbf{A}$ and $\ell' = \ell$.

Case UNPUSH: In this case, we first use a simple invariant that the operation that is UNPUSHed was not yet committed. Then, we like the PUSH case, since the new operation has not been committed, it will be filtered out by rewind so, again, we can let $\mathbf{A}' = \mathbf{A}$ and $\ell' = \ell$.

Cases PULL, UNPULL: Let $\mathbf{A}' = \mathbf{A}$ and $\ell' = \ell$. In both PULL and UNPULL, the shared log is unchanged, so **dropping** leads to the same ℓ . Moreover, \sim is independent of the operations PULLED into the local logs in \mathbf{T} .

Cases NONDET_L, NONDET_R, LOOP: Let $\mathbf{A}' = \mathbf{A}$ and $\ell' = \ell$. These cases are straight-forward, and simply involve rewinding (with an additional lemma that threads not executing transactions have empty local logs).

Cases SEMI, SEMISKIP: (Details omitted because they are uninteresting.)

Case CMT: This is the only case in which we map \mathbf{T}', G' to new \mathbf{A}' and ℓ' . We must show that such an \mathbf{A}', ℓ' exists such that $\mathbf{T}', G' \sim \mathbf{A}', \ell'$. Let \mathbf{A}' be constructed from each of the transactions in \mathbf{A} , except the committing one. The committing one is defined by rewinding the transaction T entirely, and then atomically running the transaction via `AM_RUNTX`. $(\text{map rewind } \mathbf{T}') = \mathbf{A}'$ holds by construction.

Consider the committing transaction $T \equiv \{\text{tx } c_1, \sigma_1, L_1\}$. The CMT rule gives us that $\text{cmt}(G, L_1, G')$, where G' is the same as G except that the flag `gUCmt` has been swapped to `gCmt` for all operations that have been pushed in L_1 . \mathbf{T}' is defined by the CMT rule, where the committing transaction has been replaced with $\{\text{skip}, \sigma_1, []\}$.

What remains is to *show the existence an ℓ' such that $\lfloor G' \rfloor_{\text{gCmt}} \leq \ell'$* , as follows:

- We start by applying the `cmtpres` invariant to the configuration \mathbf{T}, G .
- We pick the “ G that has dropped all uncommitted operations due to other transactions (excluding T):

$${}^{\prime}G \equiv \text{filter } (\lambda[\{m, \sigma, \sigma', id\}, g].g = \text{gCmt}) G'$$

- Now, we do a 0-step local log rollback, keeping L_1 untouched and ${}^{\prime}G = {}^{\prime}G$.
- The remainder of the `cmtpres` invariant tells us that

$$\begin{aligned} \forall G_{\text{post}}. \text{cmt}({}^{\prime}G, L_1, G_{\text{post}}) &\Rightarrow \\ \forall \sigma', H. (c, \sigma), G_{\text{post}} \cdot [L_1]_{\text{npshd}} \Downarrow \sigma', H &\Rightarrow \\ \exists \hat{\ell}'. \text{otx}(\{c, \sigma, L_1\}), {}^{\prime}G \setminus L_1 \Downarrow \sigma', \hat{\ell}' \wedge H &\leq \hat{\ell}' \end{aligned}$$

- Note that $G_{\text{post}} = \lfloor G' \rfloor_{\text{gCmt}}$ due to a commutativity: G_{post} is obtained from G by dropping all operations due to other uncommitted transactions and then marking T operations as committed. Meanwhile, $\lfloor G' \rfloor_{\text{gCmt}}$ is obtained from G by committing T and then dropping uncommitted operations.
- Note also that $[L_1]_{\text{npshd}} = \emptyset$ because of CMT criterion (ii) and that $\text{fin}(c)$ due to CMT criterion (i). By the `BSFIN` rule, $H = G_{\text{post}}$. This leaves us with

$$\exists \hat{\ell}'. \text{otx}(\{c, \sigma, L_1\}), \lfloor G \rfloor_{\text{gCmt}} \Downarrow \sigma', \hat{\ell}' \wedge H \leq \hat{\ell}'$$

- By `BSSTEP`, we know that $\hat{\ell}' = \lfloor G \rfloor_{\text{gCmt}} \cdot \ell_{\text{new}}$. We use this as our successor atomic log in the simulation relation and the right conjunct gives us that $H \leq \lfloor G \rfloor_{\text{gCmt}} \cdot \ell_{\text{new}}$.
- Note that H contains only committed operations, so $H = \lfloor H \rfloor_{\text{gCmt}}$. Thus, we can conclude that $\lfloor H \rfloor_{\text{gCmt}} \leq \lfloor G \rfloor_{\text{gCmt}} \cdot \ell_{\text{new}}$ (i.e. $\lfloor H \rfloor_{\text{gCmt}} \leq \hat{\ell}'$),
- The inductive hypothesis gives us that $\lfloor G \rfloor_{\text{gCmt}} \leq \ell$, so $\lfloor H \rfloor_{\text{gCmt}} \leq \ell \cdot \ell_{\text{new}}$ and thus we have the witness to the step of the atomic machine so the simulation relation still holds. \square

6 Evaluation

We applied our model to reason about a wide variety of transactional memory implementations in the literature. In each case, we recast the implementation strategy in terms of the PUSH/PULL model and then show that the implementations satisfy the conditions of each rule in PUSH/PULL.

6.1 Opacity

For general PUSH/PULL transactions, opacity [10] does not necessarily hold: transactions may view the uncommitted effects of other concurrent transactions. However, there are several ways that we can characterize opacity as a fragment of PUSH/PULL transactions. For example, if transactions do not perform PULL operations during execution then they are *opaque*.

However, we can take things a step further. An active transaction T may PULL an operation m' that is due to an uncommitted transaction T' provided that T will never execute a method m that does not commute with m . This suggests an interesting way of ensuring opacity while PULLing uncommitted effects by examining (statically or dynamically) the set of all reachable operations that a transaction may perform.

6.2 Optimistic Models

STMs such as TL2 [6], TinySTM [8], Intel STM [31] are optimistic (or mostly-optimistic) and do not share their effects until they commit. Transactions begin by PULLing all operations (there are never uncommitted operations) by simply viewing the shared state. As they continue to execute, they APP locally and do not PUSH until an uninterleaved moment when they check the second PUSH condition on all of their effects (which is approximated via read/write sets) and, if it holds, PUSH everything and CMT. Effects are pushed in order so the first PUSH condition is trivial. If a transaction discovers a conflict, it can simply perform UNAPP repeatedly and needn't UNPUSH.

Transactions that use checkpoints [19] and (closed) nested transactions [27] do not share their effects until commit time. They are similar to the above optimistic models, except that placemarkers are set so that, if an abort is detected, UNAPP only needs to be performed for some operations.

6.3 Pessimistic Models

Matveev and Shavit [25] describe how pessimistic transactions can be implemented by delaying write operations until the commit phase. In this way, write transactions appear to occur instantaneously at the commit point: all write operations are PUSHed just before CMT, with no interleaved transactions. Consequently, read operations perform PULL only on committed effects.

Transactional boosting [11] is also a pessimistic model. It's PUSH/PULL representation is straightforward.

6.4 Mixed Models

For the irrevocable transactions of Welc *et al.* [34], there is at most one pessimistic (“irrevocable”) transaction and many optimistic transactions. The pessimistic transaction PUSHes its effects instantaneously after APP.

6.5 Reading Uncommitted Effects

As discussed in Section 4, the early release mechanism [14] and dependent transactions [30] can be modeled with PUSH/PULL. In early release, an executing transaction T communicates with T' to determine whether the transactions conflict. This is modeled as T' performing a PUSH(op) and T checking whether it is able to PULL(op). A *dependent transaction* T will PULL the effects of another transaction T' . This comes with the stipulation that T does not commit until T' has committed. If T' aborts, then T must abort. However, note that T must only move backwards (via $\xrightarrow{\text{back}}$) insofar as to detangle from T' .

7 Boosting/HTM interaction

The PUSH/PULL model is expressive, permitting transactions to announce their effects in orders different from the way they are done locally (see the PUSH rule). Moreover, transactions can undo their effects in

different orders from the order they were announced in (see the UNPUSH rule). This may seem a fairly obscure behavior which, to our knowledge, has not been realized in any transactional implementations. Nonetheless, in this section we demonstrate a simple example where these complex behaviors seem natural.

Consider the following example transaction that accesses a boosted [11] version of a `ConcurrentSkipList` and a boosted version of a `ConcurrentHashTable`, as well as integer variables `size`, `x`, and `y` that are controlled via a hardware transactional memory [17]:

```

1 BoostedConcurrentSkipList skiplist;
2 BoostedConcurrentHashTable hashT;
3 HTM int size;
4 HTM int x, y;
5
6 atomic {
7   skiplist.insert(foo);
8   size++;
9
10  hashT.map(foo => bar);
11  if (*)
12    x ++;
13  else
14    y ++;
15 }
```

Let us say that execution proceeds, modifying the `skiplist`, incrementing `size`, updating the `hashT`, and the following the `if` branch. At this underlined point when `x` is about to be incremented, let us say that the hardware transactional memory detects a conflict with a concurrent access to `x`.

The PUSH/PULL model shows that the implementation can rewind (UNPUSH) the effects of the HTM, but leave the effects of the boosted objects (which are expensive to replay) in the shared view. So the HTM can discard the effects to `x` and `size` with UNPUSHP, perform a partial rewind via UNAPP, then execute:

```

if (*)
  x ++;
else
  y ++;
}
```

In terms of the PUSH/PULL model, the transaction has performed the rules given in Figure 7. This figure decomposes the elaborate behavior into the simple PUSH/PULL rules. We can then construct a correctness argument for the example from the criteria of each rule, and the hard work of the simulation proof is done for us.

8 Conclusions and Future Work

We have described an expressive model of transactions and shown that it is capable of serving as proof of serializability for a wide variety of transactional memory algorithms. We work with pure logs and develop a model in which transactions pass around their effects by PUSHing to or PULLing from a shared log. The model gives rise to simple proof rules that allow us to more easily construct proofs for a wide range of transactional behaviors—optimism, pessimism, opacity, dependency, etc.—all within a unified treatment.

One potential avenue for future work is to consider weaker notions than serializability [9, 3].

<i>Transaction begins.</i>	PULL(all skiplist operations) APP(skiplist.insert(foo)), PUSH(skiplist.insert(foo)), APP(size++), PULL(all hashT operations) APP(hashT.map(foo=>bar)), PUSH(hashT.map(foo=>bar)), APP(x++),
<i>Push HTM ops:</i>	PUSH(size++), PUSH(x++),
<i>HTM signals abort.</i>	UNPUSH(x++), UNPUSH(size++),
<i>Rewind some code:</i>	UNAPP(x++),
<i>March forward again:</i>	APP(y++),
<i>Uninterleaved commit:</i>	PUSH(size++), PUSH(y++), CMT

Figure 7: Decomposing behavior in terms of PUSH/PULL rules.

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] 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).
- [3] BURCKHARDT, S., AND LEIJEN, D. Semantics of concurrent revisions. In *Programming Languages and Systems - 20th European Symposium on Programming, ESOP 2011, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2011, Saarbrücken, Germany, March 26-April 3, 2011. Proceedings* (2011), G. Barthe, Ed., vol. 6602 of *Lecture Notes in Computer Science*, Springer, pp. 116–135.
- [4] CHEREM, S., CHILIMBI, T. M., AND GULWANI, S. Inferring locks for atomic sections. In *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation, Tucson, AZ, USA, June 7-13, 2008* (2008), R. Gupta and S. P. Amarasinghe, Eds., ACM, pp. 304–315.
- [5] COHEN, A., PNUELI, A., AND ZUCK, L. D. Mechanical verification of transactional memories with non-transactional memory accesses. In *Computer Aided Verification, 20th International Conference, CAV 2008, Princeton, NJ, USA, July 7-14, 2008, Proceedings* (2008), A. Gupta and S. Malik, Eds., vol. 5123, Springer, pp. 121–134.
- [6] DICE, D., SHALEV, O., AND SHAVIT, N. Transactional locking II. In *Proceedings of the 20th International Symposium on Distributed Computing (DISC'06)* (September 2006).
- [7] DIMITROV, D., RAYCHEV, V., VECHEV, M., AND KOSKINEN, E. Commutativity race detection. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'14)*, Edinburgh, UK (2014).
- [8] FELBER, P., FETZER, C., AND RIEGEL, T. Dynamic performance tuning of word-based software transactional memory. In *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming* (New York, NY, USA, 2008), PPOPP'08, ACM, pp. 237–246.
- [9] FELBER, P., GRAMOLI, V., AND GUERRAOU, R. Elastic transactions. In *DISC'09* (Berlin, Heidelberg, 2009), pp. 93–107.
- [10] 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.
- [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., AND KOSKINEN, E. Composable transactional objects: A position paper. In *Proceedings of the 23rd European Symposium on Programming Languages and Systems (ESOP'14)* (2014), vol. 8410, pp. 1–7.
- [13] HERLIHY, M., LUCHANGCO, V., MOIR, M., AND III, W. N. S. Software transactional memory for dynamic-sized data structures. In *Proceedings of the Twenty-Second ACM Symposium on Principles of Distributed Computing, PODC 2003, Boston, Massachusetts, USA, July 13-16, 2003* (2003), E. Borowsky and S. Rajsbaum, Eds., ACM, pp. 92–101.
- [14] 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.
- [15] 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. San Diego, CA, May 1993* (1993), A. J. Smith, Ed., ACM, pp. 289–300.
- [16] IBM. Ibm transactional memory compiler https://www.ibm.com/developerworks/mydeveloperworks/blogs/5894415f-be62-4bc0-81c5-3956e82276f3/entry/ibm_s_alphaworks_software_transactional_memory_compiler?lang=en.
- [17] INTEL. Transactional synchronization in haswell. <http://software.intel.com/en-us/blogs/2012/02/07/transactional-synchronization-in-haswell>.
- [18] KIM, D., AND RINARD, M. C. Verification of semantic commutativity conditions and inverse operations on linked data structures. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2011, San Jose, CA, USA, June 4-8, 2011* (2011), ACM, pp. 528–541.
- [19] KOSKINEN, E., AND HERLIHY, M. Checkpoints and continuations instead of nested transactions. In *SPAA 2008: Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures* (2008), F. Meyer auf der Heide and N. Shavit, Eds., ACM, pp. 160–168.
- [20] KOSKINEN, E., PARKINSON, M. J., AND HERLIHY, M. Coarse-grained transactions. In *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (2010), M. V. Hermenegildo and J. Palsberg, Eds., ACM, pp. 19–30.
- [21] 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.
- [22] LESANI, M., LUCHANGCO, V., AND MOIR, M. A framework for formally verifying software transactional memory algorithms. In *CONCUR 2012 - Concurrency Theory - 23rd International Conference, CONCUR 2012, Newcastle upon Tyne, UK, September 4-7, 2012. Proceedings* (2012), M. Koutny and I. Ulidowski, Eds., vol. 7454, Springer, pp. 516–530.
- [23] LESANI, M., AND PALSBERG, J. Communicating memory transactions. In *Proceedings of the 16th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2011, San Antonio, TX, USA, February 12-16, 2011* (2011), C. Cascaval and P.-C. Yew, Eds., ACM, pp. 157–168.
- [24] LIPTON, R. J. Reduction: a method of proving properties of parallel programs. *Commun. ACM* 18, 12 (1975), 717–721.
- [25] MATVEEV, A., AND SHAVIT, N. Towards a fully pessimistic stm model. In *Proc. Workshop on transactional memory (TRANSACTION'12)* (2012).
- [26] 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.
- [27] 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.
- [28] NI, Y., MENON, V. S., ADL-TABATABAI, A.-R., HOSKING, A. L., HUDSON, R. L., MOSS, J. E. B., SAHA, B., AND SHPEISMAN, T. Open nesting in software transactional memory. In *Proceedings of the 12th ACM SIGPLAN symposium on Principles and practice of parallel programming (PPoPP'07)* (New York, NY, USA, 2007), ACM Press, pp. 68–78.
- [29] PAPADIMITRIOU, C. The serializability of concurrent database updates. *Journal of the ACM (JACM)* 26, 4 (1979), 631–653.
- [30] RAMADAN, H. E., ROY, I., HERLIHY, M., AND WITCHEL, E. Committing conflicting transactions in an stm. In *PPOPP* (2009), pp. 163–172.
- [31] 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.
- [32] 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.
- [33] 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.
- [34] WELC, A., SAHA, B., AND ADL-TABATABAI, A.-R. Irrevocable transactions and their applications. In *SPAA '08: Proceedings of the twentieth annual symposium on Parallelism in algorithms and architectures* (New York, NY, USA, 2008), ACM, pp. 285–296.