

# Safety of Live Transactions in Transactional Memory: TMS is Necessary and Sufficient

Hagit Attiya\*      Alexey Gotsman†      Sandeep Hans\*      Noam Rinetzky‡

## Abstract

The main challenge of stating the correctness of transactional memory (TM) systems is the need to provide guarantees on the system state observed by *live* transactions, i.e., those that have not yet committed or aborted. A TM correctness condition should not be too restrictive, to allow flexibility in implementation, yet strong enough to disallow undesirable TM behavior, which can lead to run-time errors in live transactions. The latter feature is formalized by *observational refinement* between TM implementations, stating that properties of a program using a concrete TM implementation can be established by analyzing its behavior with an abstract TM, serving as a specification of the concrete one.

We show that a variant of *transactional memory specification (TMS)* is equivalent to observational refinement for the common programming model in which local variables are rolled back upon a transaction abort and, hence, is the weakest consistency condition for this case. This is challenging due to the nontrivial formulation of TMS, which allows different aborted and live transactions to have different views on the system state. Our proof reveals some natural, but subtle, assumptions on the TM required for the equivalence result.

**This is a submission to the regular track.**

**It can be considered for the best student paper award: Sandeep Hans is a full-time student.**

---

\*Technion - Israel Institute of Technology, Israel, {hagit,sandeep}@cs.technion.ac.il

†IMDEA Software Institute, Spain, alexey.gotsman@imdea.org

‡Tel Aviv University, Israel, maon@cs.tau.ac.il

## 1 Introduction

*Transactional memory (TM)* eases the task of writing concurrent applications by letting the programmer designate certain code blocks as *atomic*. TM allows developing a program and reasoning about its correctness as if each atomic block executes as a *transaction*—in one step and without interleaving with others—even though in reality the blocks can be executed concurrently. Figure 1 shows how atomic blocks yield simple code for computations involving several shared *transactional objects*  $X$ ,  $Y$  and  $Z$ , access to which is mediated by the TM.

The common approach to stating TM correctness is through a *consistency condition* that restricts the possible TM executions. The main subtlety of formulating such a condition is the need to provide guarantees on the state of transactional objects observed by *live* transactions, i.e., those that have not yet committed or aborted. Because live transactions can always be aborted, one might think it unnecessary to provide any guarantees for them, as in fact done by common database consistency conditions [12]. However, in the setting of transactional memory, this is often unsatisfactory. For example, in Figure 1 the programmer may rely on the fact that  $X \neq Y$ , and, correspondingly, make sure that every committing transaction preserves this invariant. If we allow the transaction to read values of  $X$  and  $Y$  violating the invariant (counting on it to abort later, due to inconsistency), this will lead to the program *faulting* due to a division by zero.

The question of which TM consistency condition to use is far from settled, with several candidates having been proposed [2, 3, 6, 10]. An ideal condition should be as weak as possible, to allow flexibility in TM implementations, yet strong enough to satisfy the intuitive expectations of the programmer and, in particular, to disallow undesirable behaviors such as the one described above. *Observational refinement* [7, 8] allows formalizing the programmer’s expectations and thereby evaluating consistency conditions systematically. Consider two TM implementations—a *concrete* one, such as an efficient TM, and an *abstract* one, such as a TM executing every atomic block atomically. Informally, the concrete TM observationally refines the abstract one for a given programming language if every behavior a user can observe of a program  $P$  in this language linked with the concrete TM can also be observed when  $P$  is linked with the abstract TM instead. This allows the programmer to reason about the behavior of  $P$  (e.g., the preservation of the invariant  $X \neq Y$ ) using the expected intuitive semantics formalized by the abstract TM; the observational refinement relation implies that the conclusions (e.g., the safety of the division in Figure 1) will carry over to the case when  $P$  uses the concrete TM.

In prior work [1] we showed that a variant of the *opacity* condition [6] coincides with observational refinement for a particular programming language and, hence, is the weakest consistency condition for this language. Roughly speaking, a concrete TM implementation is in the opacity relation with an abstract one if for any sequence of interactions with the concrete TM, dubbed a *history*, there exists a history of the abstract TM where: (i) the actions of every separate thread are the same as in the original history; and (ii) the order of non-overlapping transactions present in the original history is preserved. However, our result considered a programming language in which local variables modified by a transaction are not rolled back upon an abort. Although this assumption holds in some situations (e.g., for Scala STM [13]), it is non-standard and most TM systems do not satisfy it. In this paper, we consider a variant of *transactional memory specification (TMS)* [3], a condition weaker than opacity,<sup>1</sup> and show that, under some natural assumptions on the TM, it coincides with observational refinement for a programming language in which local variables do get rolled back upon an abort.

This result is not just a straightforward adjustment of the one about opacity to a more realistic setting: TMS weakens opacity in a nontrivial way, which makes reasoning about its relationship with observational refinement much more intricate. In more detail, the key feature of opacity is that the behavior of *all* transactions in a history of the concrete TM, including aborted and live ones, has to be explained by a single history of the abstract TM. TMS relaxes this requirement by requiring only committed transactions in the concrete history to be explained by a single abstract one obeying (i)–(ii) above; every response obtained from the TM in an aborted or live transaction may be explained by a separate abstract history. The constraints on the choice of the abstract history are subtle: on one

```
result := abort;
while (result == abort) {
  result := atomic {
    x = X.read();
    y = Y.read();
    z = 42 / (x - y);
    Z.write(z);
  }
}
```

Figure 1: TM usage

<sup>1</sup>The condition we present here is actually called TMS1 in [3, 11]. These papers also propose another condition, called TMS2, but it is stronger than opacity [11] and therefore not considered here.

hand, somewhat counter-intuitively, TMS allows it to include transactions that aborted in the concrete history, with their status changed to committed, and exclude some that committed; on the other hand, this is subject to certain carefully chosen constraints. The flexibility in the choice of the abstract history is meant to allow the concrete TM implementation to perform as many optimizations as possible. However, it is not straightforward to establish that this flexibility does not invalidate observational refinement (and hence, the informal guarantees that programmers expect from a TM) or that the TMS definition cannot be weakened further.

Our results ensure that this is indeed the case. Informally, if local variables are not rolled back when transactions abort, threads can communicate to each other the observations they make inside aborted transactions about the state of transactional objects. This requires the TM to provide a consistent view of this state across all transactions, as formalized by the use of a single abstract history in opacity. However, if local variables are rolled back upon an abort, no information can leak out of an uncommitted transaction, possibly apart from the fact that the code in the transaction has faulted, stopping the computation. To get observational refinement in this case we only need to make sure that a fault in the transaction occurring with the concrete TM could be reproduced with the abstract one. For this it is sufficient to require that the state of transactional objects seen by every live transaction can be explained by some abstract history; different transactions can be explained by different histories.

Technically, we prove that TMS is sufficient for observational refinement by establishing a nontrivial property of the set of computations of a program, showing that a live transaction cannot notice the changes in the committed/aborted status of transactions concurrent with it that are allowed by TMS (Lemma 9, Section 6.1). Proving that TMS is necessary for observational refinement is challenging as well, as this requires us to devise multiple programs that can observe whether the subtle constraints governing the change of transaction status in TMS are fulfilled by the TM. We have identified several closure properties on the set of histories produced by the abstract TM required for these results to hold. Although intuitive, these properties are not necessarily provided by an arbitrary TM, and our results demonstrate their importance.

To concentrate on the core goal of this paper, the programming language we consider does not allow explicit transaction aborts or transaction nesting and assumes a static separation of transactional and non-transactional shared memory. Extending our development to lift these restrictions is an interesting avenue for future work. Also, due to space constraints, we defer some of the proofs to Appendix D.

## 2 Programming Language Syntax

We consider a language where a program  $P = C_1 \parallel \dots \parallel C_m$  is a parallel composition of *threads*  $C_t$ ,  $t \in \text{ThreadID} = \{1, \dots, m\}$ . Every thread  $t \in \text{ThreadID}$  has a set of *local variables*  $\text{LVar}_t = \{x, y, \dots\}$  and threads share a set of *global variables*  $\text{GVar} = \{g, \dots\}$ , all of type integer. We let  $\text{Var} = \text{GVar} \uplus \biguplus_{t=1}^m \text{LVar}_t$  be the set of all program variables. Threads can also access a transactional memory, which manages a fixed collection of *transactional objects*  $\text{Obj} = \{o, \dots\}$ , each with a set of *methods*  $\text{Method} = \{f, \dots\}$  that threads can call. For simplicity, we assume that each method takes one integer parameter and returns an integer value, and that all objects have the same set of methods. The syntax of commands  $C$  is standard:

$$C ::= c \mid C;C \mid \text{while}(b) \text{ do } C \mid \text{if}(b) \text{ then } C \text{ else } C \mid x := \text{atomic}\{C\} \mid x := o.f(e)$$

where  $b$  and  $e$  denote Boolean and integer expressions over local variables, left unspecified. The syntax includes *primitive commands*  $c$  from a set  $\text{PComm}$ , sequential composition, conditionals, loops, `atomic` blocks and object method invocations. Primitive commands execute atomically, and they include assignments to local and global variables and a special `fault` command, which stops the execution of the program in an error state. Thus, `fault` encodes illegal computations, such as division by zero or dereference of `null`-valued pointers.

An *atomic block*  $x := \text{atomic}\{C\}$  executes  $C$  as a *transaction*, which the TM can *commit* or *abort*. The system's decision is returned in the local variable  $x$ , which gets assigned distinguished values `committed` or `aborted`. We do not allow programs in our language to abort a transaction explicitly and forbid nested atomic blocks and, hence, nested transactions. We also assume that a program can invoke methods on transactional objects only inside atomic blocks and access global variables only outside them. Local variables can be accessed in both cases; however, threads cannot access local variables of other threads. Due to space constraints, we defer the formalisation of

the rules on variable accesses to Appendix A. When we later define the semantics of our programming language, we mandate that, if a transaction is aborted, local variables are rolled back to the values they had at its start, and hence, the values written to them by the transaction cannot be observed by the following non-transactional code.

### 3 Model of Computations

To define the notion of observational refinement for our programming language and the TMS consistency condition, we need a formal model for program computations. To this end, we introduce *traces*, which are certain finite sequences of *actions*, each describing a single computation step (we do not consider infinite computations).

DEFINITION 1. Let  $\text{ActionId}$  be a set of action identifiers. A **TM interface action**  $\psi$  has one of the following forms:

| Request actions               | Matching response actions                                |
|-------------------------------|--|
| $(a, t, \text{txbegin})$      | $(a, t, \text{OK}) \mid (a, t, \text{aborted})$          |
| $(a, t, \text{txcommit})$     | $(a, t, \text{committed}) \mid (a, t, \text{aborted})$   |
| $(a, t, \text{call } o.f(n))$ | $(a, t, \text{ret}(n') o.f) \mid (a, t, \text{aborted})$ |

where  $a \in \text{ActionId}$ ,  $t \in \text{ThreadID}$ ,  $o \in \text{Obj}$ ,  $f \in \text{Method}$  and  $n, n' \in \mathbb{Z}$ . A **primitive action**  $\phi$  has the form  $(a, t, c)$ , where  $c \in \text{PComm}$  is a primitive command. We use  $\varphi$  to range over actions of either type.

TM interface actions denote the control flow of a thread  $t$  crossing the boundary between the program and the TM: **request** actions correspond to the control being transferred from the former to the latter, and **response** actions, the other way around. A `txbegin` action is generated upon entering an `atomic` block, and a `txcommit` action when a transaction tries to commit upon exiting an `atomic` block. Actions `call` and `ret` denote a call to and a return from an invocation of a method on a transactional object and are annotated with the method parameter or return value. The TM may abort a transaction at any point when it is in control; this is recorded by an `aborted` response action.

A **trace**  $\tau$  is a finite sequence of actions satisfying certain natural well-formedness conditions (stated informally due to space constraints; see Appendix B): every action in  $\tau$  has a unique identifier; no action follows a `fault`; request and response actions are properly matched; actions denoting the beginning and end of transactions are properly matched; `call` and `ret` actions occur only inside transactions; and commands in  $\tau$  do not access local variables of other threads and do not access global variables when inside a transaction. We denote the set of traces by  $\text{Trace}$ . A **history** is a trace containing only TM interface actions; we use  $H, S$  to range over histories. We specify the behavior of a TM implementation by the set of possible interactions it can have with programs: a **transactional memory**  $\mathcal{T}$  is a set of histories that is prefix-closed and closed under renaming action identifiers.

We denote irrelevant expressions by  $\_$  and use the following notation:  $\tau(i)$  is the  $i$ -th element of  $\tau$ ;  $\tau|_t$  is the projection of  $\tau$  onto actions of the form  $(\_, t, \_)$ ;  $|\tau|$  is the length of  $\tau$ ;  $\tau_1\tau_2$  is the concatenation of  $\tau_1$  and  $\tau_2$ . We say that an action  $\varphi$  is in  $\tau$ , denoted by  $\varphi \in \tau$ , if  $\tau = \_ \varphi \_$ . The empty sequence of actions is denoted  $\varepsilon$ .

A **transaction**  $T$  is a non-empty trace such that it contains actions by the same thread, begins with a `txbegin` action and only its last action can be a committed or an aborted action. We use the following terminology for transactions. A transaction  $T$  is: **committed** if it ends with a committed action, **aborted** if it ends with aborted, **commit-pending** if it ends with `txcommit`, and **live**, in all other cases. We refer to this as  $T$ 's **status**. A transaction  $T$  is **completed** if it is either committed or aborted, and **visible** if it contains a `txcommit` action. A transaction  $T$  is **in a trace**  $\tau$ , written  $T \in \tau$ , if  $\tau|_t = \tau_1 T \tau_2$  for some  $t, \tau_1$  and  $\tau_2$ , where either  $T$  is completed or  $\tau_2$  is empty. We denote the set of all transactions in  $\tau$  by  $\text{tx}(\tau)$  and use self-explanatory notation for various subsets of transactions:  $\text{committed}(\tau)$ ,  $\text{aborted}(\tau)$ ,  $\text{pending}(\tau)$ ,  $\text{live}(\tau)$ ,  $\text{visible}(\tau)$ . For  $\varphi \in \tau$ , the **transaction of  $\varphi$  in  $\tau$** , denoted  $\text{txof}(\varphi, \tau)$ , is the subsequence of  $\tau$  comprised of all actions that are in the same transaction as  $\varphi$  (undefined if  $\varphi$  does not belong to a transaction).

### 4 Transactional Memory Specification (TMS)

In this section we formalise the TMS [3] correctness condition in our setting. TMS was originally formulated in an operational manner, using I/O automata; here we present a more abstract definition appropriate for our goals (we

provide further comparison in Section 7). As is common in consistency conditions for shared-memory concurrency, such as opacity [6] or linearizability [9], a crucial building block in the TMS definition is the following notion of the *real-time order*, which captures the order between non-overlapping transactions in a history.

**DEFINITION 2.** Let  $\psi = (\_, t, \_)$  and  $\psi' = (\_, t', \_)$  be two actions in a history  $H$ ;  $\psi$  is **before**  $\psi'$  **in the real-time order** in  $H$ , denoted by  $\psi \prec_H \psi'$ , if  $H = H\psi H_2 H_2' \psi' H_3$  and either (i)  $t = t'$  or (ii)  $(\_, t', \text{txbegin}) \in H_2' \psi'$  and either  $(\_, t, \text{committed}) \in \psi H_2$  or  $(\_, t, \text{aborted}) \in \psi H_2$ . A transaction  $T$  is **before** an action  $\psi'$  **in the real-time order** in  $H$ , denoted by  $T \prec_H \psi'$ , if  $\psi \prec_H \psi'$  for any  $\psi \in T$ . A transaction  $T$  is **before** a transaction  $T'$  **in the real-time order** in  $H$ , denoted by  $T \prec_H T'$ , if  $T \prec_H T'(1)$ .

Given a history  $H$  of program interactions with a concrete TM, TMS requires us to explain the behavior of all committed transactions in  $H$  by a single history  $S$  of the abstract TM, and to explain every response action  $\psi$  in  $H$  by an abstract history  $S_\psi$ . As we show in this paper, the existence of such explanations ensures that TMS imply observational refinement between the two TMs: the behavior of a program during some transaction in the history  $H$  of the program's interactions with the concrete TM can be reproduced when the program interacts with the abstract TM according to the history  $S$  or  $S_\psi$ . Below we use this insight when explaining the rationale for key TMS features.

The history  $S_\psi$  used to explain a response action  $\psi$  includes the transaction of  $\psi$  and a subset of transactions from  $H$  whose actions justify the response  $\psi$ . The following notion of a *possible past* of a history  $H = H_1\psi$  defines all sets of transactions from  $H$  that can form  $S_\psi$ . Note that, if a transaction selected by this definition is aborted or commit-pending in  $H$ , its status is changed to committed when constructing  $S_\psi$ , as formalized later in Definition 4. Informally, the response  $\psi$  is given as if all the transactions in its possible past have taken effect and all the others have not. We first give the formal definition of a possible past, and then explain it on an example.

**DEFINITION 3.** A history  $H_\psi = H_1'\psi$  is a **possible past** of a history  $H = H_1\psi$ , where  $\psi$  is a response action, if:

- (i)  $H_1'$  is a subsequence of  $H_1$ ;
- (ii)  $H_\psi$  is comprised of the transaction of  $\psi$  and some of the visible transactions in  $H$ :

$$\text{tx}(H_\psi) \subseteq \{\text{txof}(\psi, H)\} \cup \text{visible}(H).$$

- (iii) for every transaction  $T \in H_\psi$ , out of all transactions preceding  $T$  in the real-time order in  $H$ , the history  $H_\psi$  includes exactly the committed ones:

$$\forall T \in \text{tx}(H_\psi). \forall T'. T' \prec_{H_\psi} T \iff T' \prec_H T \wedge T' \in \text{committed}(H).$$

We denote the set of possible pasts of  $H$  by  $\text{TMSpast}(H)$ .

We explain the definition using the history  $H$  of the trace shown in Figure 2; one of its possible pasts  $H_\psi$  consists of the transactions  $T_1$ ,  $T_4$  and  $T_5$ . According to (ii), the transaction of  $\psi$  ( $T_5$  in Figure 2) is always included into any possible past, and live transactions are excluded: since they have not made an attempt to commit, they should not have an effect on  $\psi$ . Out of the visible transactions in  $H$ , we are allowed to select which ones to include (and, hence, treat as committed), subject to (iii): if we include a transaction  $T$  then, out of all transactions preceding  $T$  in the real-time order in  $H$ , we have to include exactly the committed ones. For example, since  $T_4$  and  $T_5$  are included in  $H_\psi$ ,  $T_1$  must also be included and  $T_3$  must not. This condition is necessary for TMS to imply observational refinement. Informally,  $T_3$  cannot be included into  $H_\psi$  because, in a program producing  $H$ , in between  $T_3$  aborting and  $T_5$  starting, thread  $t_2$  could have communicated to thread  $t_3$  the fact that  $T_3$  has aborted, e.g., using a global variable  $g$ , as illustrated in Figure 2. When executing  $\psi$ , the code in  $T_5$  may thus expect that  $T_3$  did not take effect; hence, the result of  $\psi$  has to reflect this, so that the code behavior is preserved when replacing the concrete TM by an abstract one in observational refinement. This is a key idea used in our proof that TMS is necessary for observational refinement (Section 6.2). In contrast to  $T_3$ , we can include  $T_4$  into  $H_\psi$  even if it is aborted or commit-pending. Since our language does not allow accessing global variables inside transactions, there is no way for the code in  $T_5$  to find out about the status of  $T_4$  from thread  $t_2$ , and hence, this code will not notice if the status of  $T_4$  is changed to committed when replacing the concrete TM by an abstract one in observational




Figure 2: Transactions  $T_1$ ,  $T_4$  and  $T_5$  form one possible past of the history  $H$  of the trace shown. Allowed status of transactions in  $H$  is denoted as follows: committed – C, aborted – A, commit-pending – CP, live – L. The transaction  $T_5$  executes only primitive actions after  $\psi$  in the trace.

refinement. For similar reasons, we can exclude  $T_2$  from  $H_\psi$  even if it is committed. This idea is used in our proof that TMS is sufficient for observational refinement (Section 6.1).

Before giving the definition of TMS, we introduce operations used to change the status of transactions in a possible past of a history to committed. *Suffix commit completion* below converts commit-pending transactions into committed; then *completed possible past* defines a possible past with all transactions committed.

DEFINITION 4. A history  $H^c$  is a **suffix completion** of a history  $H\psi$  if  $H^c = H\psi H'$ , any action in  $H'$  is either committed or aborted, and every transaction in  $H^c$  except possibly that of  $\psi$ , is completed. It is a **suffix commit completion** of  $H$  if  $H'$  consists of committed actions only. The sets of suffix completions and suffix commit completions of  $H$  are denoted  $\text{comp}(H)$  and  $\text{comcomp}(H)$ , respectively.

A history  $H_\psi^c$  is a **completed possible past** of a history  $H = H_1\psi$ , if  $H_\psi^c$  is a suffix commit completion of a history obtained from a possible past  $H_1'\psi$  of  $H$  by replacing all the aborted actions in  $H_1'$  by committed actions. The set of completed possible pasts of  $H$  is denoted  $\text{cTMSpast}(H)$ :

$$\text{cTMSpast}(H_1\psi) = \{H_\psi^c \mid \exists H_1'. H_1'\psi \in \text{TMSpast}(H_1\psi) \wedge H_\psi^c \in \text{comcomp}(\text{com}(H_1')\psi)\},$$

where  $|\text{com}(H_1')| = |H_1'|$  and  $\text{com}(H_1')(i) = (\text{if } (H_1')(i) = (a, t, \text{aborted})) \text{ then } (a, t, \text{committed}) \text{ else } H_1'(i)$ .

For example, one completed possible past of the history in Figure 2 consists of the transactions  $T_1$ ,  $T_4$  and  $T_5$ , with the status of the latter changed to committed if it was previously aborted or commit-pending. Note that a history  $H$  has a suffix completion only if  $H$  is of the form  $H = H_1\psi$  where all the transactions in  $H_1\psi$ , except possibly that of  $\psi$ , are commit-pending or completed. Also,  $\text{cTMSpast}(H_1\psi) \neq \emptyset$  only if, in addition,  $\psi$  is a response action.

The following definition of the *TMS relation* between TMs matches a history  $H$  arising from a concrete TM with a similar history  $S$  of an abstract TM. As part of this matching, we require that  $S$  be a permutation of  $H$  preserving the real-time order, which is formalised by the *opacity relation* [1, 6]. As in Definition 3(iii), this requirement is necessary to ensure observational refinement between the TMs: preserving the real-time order is necessary to preserve communication between threads when replacing the concrete TM with the abstract one.

DEFINITION 5. A history  $H$  is in the **opacity relation** with a history  $S$ , denoted by  $H \sqsubseteq_{\text{op}} S$ , if

$$\forall \psi, \psi'. (\psi \in S \iff \psi \in H) \wedge (\psi \prec_H \psi' \implies \psi \prec_S \psi').$$

DEFINITION 6. A history  $H$  is in the **TMS relation** with TM  $\mathcal{T}$ , denoted  $H \sqsubseteq_{\text{tms}} \mathcal{T}$ , if:

- (i)  $\exists H^c \in \text{comp}(H|_{\neg \text{live}})$ ,  $S \in \mathcal{T}$ .  $H^c|_{\text{com}} \sqsubseteq_{\text{op}} S$ , where  $\cdot|_{\neg \text{live}}$  and  $\cdot|_{\text{com}}$  are the projections to actions by transactions that are not live and by committed transactions, respectively; and
- (ii) for every response action  $\psi$  such that  $H = H_1\psi H_2$ :  $\exists H_\psi^c \in \text{cTMSpast}(H_1\psi)$ .  $\exists S_\psi \in \mathcal{T}$ .  $H_\psi^c \sqsubseteq_{\text{op}} S_\psi$ .

A TM  $\mathcal{T}_C$  is in the **TMS relation** with a TM  $\mathcal{T}_A$ , denoted by  $\mathcal{T}_C \sqsubseteq_{\text{tms}} \mathcal{T}_A$ , if  $\forall H \in \mathcal{T}_C$ .  $H \sqsubseteq_{\text{tms}} \mathcal{T}_A$ .

## 5 Observational Refinement

Our main result relates TMS to *observational refinement*, which we introduce in this section. This requires defining the semantics of the programming language, i.e., the set of traces that computations of programs produce. A **state** of a program records the values of all its variables:  $s \in \text{State} = \text{Var} \rightarrow \mathbb{Z}$ . The semantics of a program  $P = C_1 \parallel \dots \parallel C_m$  is given by the set of traces  $\llbracket P, \mathcal{T} \rrbracket(s) \subseteq \text{Trace}$  it produces when run with a TM  $\mathcal{T}$  from an

initial state  $s$ . To define this set, we first define the set of traces  $\llbracket P \rrbracket(s) \subseteq \text{Trace}$  that a program can produce when run from  $s$  with the behavior of the TM unrestricted, i.e., considering all possible values the TM can return to object method invocations and allowing transactions to commit or abort arbitrarily. We then restrict to the set of traces produced by  $P$  when run with  $\mathcal{T}$  by selecting those traces that interact with the TM in a way consistent with  $\mathcal{T}$ :  $\llbracket P, \mathcal{T} \rrbracket(s) = \{\tau \mid \tau \in \llbracket P \rrbracket(s) \wedge \text{history}(\tau) \in \mathcal{T}\}$ , where  $\text{history}(\cdot)$  projects to TM interface actions. Due to space constraints, we defer the formal definition of  $\llbracket P \rrbracket(s)$  to Appendix C. The definition follows the intuitive semantics of our programming language. In particular, it mandates that local variables be rolled back upon a transaction abort and includes traces corresponding to incomplete program computations into  $\llbracket P \rrbracket(s)$ .

We can now define the notions of *observations* and *observational refinement*. Informally, given a trace  $\tau$  of a client program, we consider observable: (i) the sequence of actions performed outside transactions in  $\tau$ ; (ii) the per-thread sequence of actions in  $\tau$  that are not part of uncommitted transactions; and (iii) whether a  $\tau$  ends with `fault` or not. Then observational refinement between a concrete TM  $\mathcal{T}_C$  and an abstract one  $\mathcal{T}_A$  states that every observable behavior of a program  $P$  using  $\mathcal{T}_C$  can be reproduced when  $P$  uses  $\mathcal{T}_A$ . Hence, any conclusion about its observable behavior that a programmer makes assuming  $\mathcal{T}_A$  will carry over to  $\mathcal{T}_C$ . Since our notion of observations excludes actions performed inside aborted or live transactions other than faulting, the programmer cannot make any conclusions about them. But, crucially, the programmer can be sure that, if a program is non-faulting under  $\mathcal{T}_A$ , it will stay so under  $\mathcal{T}_C$ . Let us call an action  $\varphi \in \tau$  **transactional** if  $\varphi \in T$  for some  $T \in \tau$ , and **non-transactional** otherwise. We denote by  $\tau|_{\text{trans}}$  and  $\tau|_{\neg\text{trans}}$  the projections of  $\tau$  to transactional and non-transactional actions.

**DEFINITION 7.** *The **thread-local observable behavior** of thread  $t$  in a trace  $\tau$ , denoted by  $\text{observable}_t(\tau)$ , is  $\downarrow$  if  $\tau|_t$  ends with a `fault` action, and  $(\tau|_t)|_{\text{obs}}$  otherwise, where  $\cdot|_{\text{obs}}$  denotes the projection to non-transactional actions and actions by committed transactions. A TM  $\mathcal{T}_C$  **observationally refines** a TM  $\mathcal{T}_A$ , denoted by  $\mathcal{T}_C \preceq \mathcal{T}_A$ , if for every program  $P$ , state  $s$  and trace  $\tau \in \llbracket P, \mathcal{T}_C \rrbracket(s)$  we have: (i)  $\exists \tau' \in \llbracket P, \mathcal{T}_A \rrbracket(s). \tau'|_{\neg\text{trans}} = \tau|_{\neg\text{trans}}$ ; and (ii)  $\forall t. \exists \tau'_t \in \llbracket P, \mathcal{T}_A \rrbracket(s). \text{observable}_t(\tau'_t) = \text{observable}_t(\tau)$ .*

## 6 Main Result

The main result of this paper is that the TMS relation can be characterized in terms of observational refinement for a class of TMs that enjoy certain natural closure properties.

- A TM  $\mathcal{T}$  is **closed under immediate aborts** if it allows adding arbitrary **immediately aborted** transactions to any history  $H \in \mathcal{T}$ : if  $H_1 H_2 H_3 \in \mathcal{T}$ ,  $H_2|_t = \varepsilon$  and  $H_1(\_, t, \text{txbegin}) H_2(\_, t, \text{aborted}) H_3$  is a well-formed history, then the latter belongs to  $\mathcal{T}$ .
- A TM  $\mathcal{T}$  is **closed under removing transaction responses** if whenever  $H_1(\_, t, \text{aborted}) H_2 \in \mathcal{T}$  or  $H_1(\_, t, \text{committed}) H_2 \in \mathcal{T}$  for  $H_2$  not containing actions by  $t$ , we also have  $H_1 H_2 \in \mathcal{T}$ .
- A TM  $\mathcal{T}$  is **closed under removing live and aborted transactions** if whenever  $H \in \mathcal{T}$ , we also have  $H' \in \mathcal{T}$  for any history  $H'$  which is a subsequence of  $H$  such that  $\text{committed}(H') = \text{committed}(H)$ ,  $\text{pending}(H') = \text{pending}(H)$ ,  $\text{live}(H') \subseteq \text{live}(H)$  and  $\text{aborted}(H') \subseteq \text{aborted}(H)$ .
- A TM  $\mathcal{T}$  is **closed under commit-pending transactions completion** if whenever  $H \in \mathcal{T}$  and  $\text{comp}(H) \neq \emptyset$ , we have  $\text{comp}(H) \cap \mathcal{T} \neq \emptyset$ .
- A TM  $\mathcal{T}$  is **closed under aborting live transactions** if whenever  $H_1 \psi \in \mathcal{T}$  and  $\psi \in \{(\_, t, \text{OK}), (\_, t, \text{ret}_\_) \}$ , we have  $H_1 \psi \psi_c(\_, t, \text{aborted}) \in \mathcal{T}$ , where  $\psi_c$  is an arbitrary request action by  $t$ .

These properties are satisfied by the expected TM specification that executes every transaction atomically [1].

**THEOREM 8.** *Let  $\mathcal{T}_C$  and  $\mathcal{T}_A$  be transactional memories.*

- If  $\mathcal{T}_A$  is closed under immediate aborts and removing transaction responses, then  $\mathcal{T}_C \sqsubseteq_{\text{tms}} \mathcal{T}_A \implies \mathcal{T}_C \preceq \mathcal{T}_A$ .*
- If  $\mathcal{T}_A$  is closed under removing live and aborted transactions, aborting live transactions and commit-pending transactions completion, then  $\mathcal{T}_C \preceq \mathcal{T}_A \implies \mathcal{T}_C \sqsubseteq_{\text{tms}} \mathcal{T}_A$ .*

### 6.1 Proof of Theorem 8(i) (Sufficiency)

Let us fix a program  $P = C_1 \parallel \dots \parallel C_m$ , state  $s$  and TMs  $\mathcal{T}_C$  and  $\mathcal{T}_A$  such that  $\mathcal{T}_C \sqsubseteq_{\text{tms}} \mathcal{T}_A$  and  $\mathcal{T}_A$  satisfies the closure conditions stated in the theorem. As we have noted before, the main subtlety of TMS lies in explaining the

behavior of a live transaction under  $\mathcal{T}_C$  by a history of  $\mathcal{T}_A$  where the committed/aborted status of some transactions is changed, as formalized by the use of  $\text{cTMSpast}$  in Definition 6(ii). Correspondingly, the most challenging part of the proof is to show that a trace from  $\llbracket P, \mathcal{T}_C \rrbracket(s)$  with a `fault` inside a live transaction can be transformed into a trace with the `fault` from  $\llbracket P, \mathcal{T}_A \rrbracket(s)$ . The following lemma describes the first and foremost step of this transformation: it converts a trace  $\tau \in \llbracket P \rrbracket(s)$  with a live transaction into another trace from  $\llbracket P \rrbracket(s)$  that contains the same live transaction, but whose history of non-aborted transactions is a given element of  $\text{cTMSpast}(\text{history}(\tau))$ . In other words, this establishes that the live transaction cannot notice the changes of the status of other transactions done by  $\text{cTMSpast}$ . Let  $\tau|_{\neg\text{abortedtx}}$  be the projection of  $\tau$  excluding aborted transactions.

**LEMMA 9 (Live transaction insensitivity).** *Let  $\tau = \tau_1\psi\tau_2 \in \llbracket P \rrbracket(s)$  be such that  $\psi$  is a response action by thread  $t_0$  that is not a committed or aborted action and  $\tau_2$  contains only primitive actions by thread  $t_0$ . Consider  $H_\psi^c \in \text{cTMSpast}(\text{history}(\tau))$ . There exists  $\tau_\psi \in \llbracket P \rrbracket(s)$  such that  $\text{history}(\tau_\psi)|_{\neg\text{abortedtx}} = H_\psi^c$  and  $\tau_\psi|_{t_0} = \tau|_{t_0}$ .*

**PROOF.** We first show how to construct  $\tau_\psi$  and then prove that it satisfies the required properties. We illustrate the idea of its construction using the trace  $\tau$  in Figure 2. Let  $\text{history}(\tau) = H_1\psi$ . Since  $H_\psi^c \in \text{cTMSpast}(H)$ , by Definition 4 there exist histories  $H'_1$ ,  $H''_1$ , and  $H^{cc}$  such that

$$H'_1\psi \in \text{TMSpast}(H_1\psi) \quad \wedge \quad H''_1 = \text{com}(H'_1) \quad \wedge \quad H_\psi^c = H''_1\psi H^{cc} \in \text{comcomp}(H''_1\psi).$$

Recall that, for the  $\tau$  in Figure 2,  $H'_1\psi$  consists of the transactions  $T_1$ ,  $T_4$  and  $T_5$ . Then  $H''_1$  is obtained from  $H'_1$  by changing the last action of  $T_4$  to committed if it was aborted;  $H_\psi^c$  is obtained by completing  $T_4$  with a committed action if it was commit-pending. The trickiness of the proof comes from the fact that just mirroring these transformations on  $\tau$  may not yield a trace of the program  $P$ : for example, if  $T_4$  aborted, the code in thread  $t_2$  following  $T_4$  may rely on this fact, communicated to it by the TM via a local variable. Fortunately, we show that it is possible to construct the required trace by erasing certain suffixes of every thread and therefore getting rid of the actions that could be sensitive to the changes of transaction status, such as those following  $T_4$ . This erasure has to be performed carefully, since threads can communicate via global variables: for example, the value written by the assignment to  $g'$  in the code following  $T_4$  may later be read by  $t_1$ , and, hence, when erasing the the former, the latter action has to be erased as well. We now explain how to truncate  $\tau$  consistently.

Let  $\psi^b$  be the last `txbegin` action in  $H'_1\psi$ ; then for some traces  $\tau_1^b$  and  $\tau_2^b$  we have  $\tau = \tau_1^b\psi^b\tau_2^b\psi\tau_2$ . For the  $\tau$  in Figure 2,  $\psi^b$  is the `txbegin` action of  $T_4$ . Our idea is, for every thread other than  $t_0$ , to erase all its actions that follow the last of its transactions included into  $H'_1\psi$  or its last non-transactional action preceding  $\psi^b$ , whichever is later. Formally, for every thread  $t$ , let  $\tau_t^I$  denote the prefix of  $\tau|_t$  that ends with the last TM interface action of  $t$  in  $H'_1\psi$ , or  $\varepsilon$  if no such action exists. For example, in Figure 2,  $\tau_{t_1}^I$  and  $\tau_{t_2}^I$  end with the last TM interface actions of  $T_1$  and  $T_4$ , respectively. Similarly, let  $\tau_t^N$  denote the prefix of  $\tau|_t$  that ends in the last non-transactional action of  $t$  in  $\tau_1^b$ , or  $\varepsilon$  if no such action exists. For example, in Figure 2,  $\tau_{t_1}^N$  and  $\tau_{t_2}^N$  end with  $\phi_1$  and  $\phi_2$ , respectively. Let  $\tau_{t_0} = \tau|_{t_0}$  and for any  $t \neq t_0$  let  $\tau_t$  be  $\tau_t^I$ , if  $|\tau_t^N| < |\tau_t^I|$ , and  $\tau_t^N$ , otherwise. We then let the truncated trace  $\tau'$  be the subsequence of  $\tau$  such that  $\tau'|_t = \tau_t$  for any  $t$ . Thus, for the  $\tau$  in Figure 2, in the corresponding trace  $\tau'$  the actions of  $t_1$  end with  $\phi_1$  and those of  $t_2$  with the last action of  $T_4$ ; note that this erases both operations on  $g'$ . To construct  $\tau_\psi$  from  $\tau'$ , we mirror the transformations of  $H'_1$  into  $H''_1$  and  $H_\psi^c$ . Let  $\tau''$  be defined by  $|\tau''| = |\tau'|$  and

$$\tau''(i) = (\text{if } (\tau'(i) = (a, t, \text{aborted}) \wedge \tau'(i) \in H'_1)) \text{ then } (a, t, \text{committed}) \text{ else } \tau'(i)).$$

Then we let  $\tau_\psi = \tau''H^{cc}$ .

We first prove that  $\tau_\psi|_{t_0} = \tau|_{t_0}$ . Let  $T = \text{txof}(\psi, H_1\psi)$ ; then by Definition 3(ii),  $T \in H'_1\psi$ . Hence, by Definition 3(iii) we have

$$\forall T'. T' \prec_{H'_1\psi} T \iff T' \prec_{H_1\psi} T \wedge T' \in \text{committed}(H_1\psi), \quad (1)$$

so that  $(H'_1\psi)|_{t_0}$  does not contain aborted transactions and  $\tau''|_{t_0} = \tau'|_{t_0} = \tau|_{t_0}$ . Besides,  $H^{cc}|_{t_0} = \varepsilon$  and, hence,  $\tau_\psi|_{t_0} = \tau''|_{t_0} = \tau|_{t_0}$ .

We now sketch the proof that  $\tau_\psi \in \llbracket P \rrbracket(s)$ , appealing to the intuitive understanding of the programming language semantics. To this end, we show that  $\tau'$  and then  $\tau''$  belong to  $\llbracket P \rrbracket(s)$ . We start by analyzing how the






Figure 3: Cases in the proof of Lemma 9. \* all actions by  $t$  are transactional; \*\* all actions by  $t$  come from a single transaction, started before or by  $\psi^b$

trace  $\tau|_t$  is truncated to  $\tau_t$  for every thread  $t \neq t_0$ . Let us make a case split on the relative positions of  $\tau_t^N$ ,  $\tau_t^I$  and  $\psi^b$  in  $\tau$ . There are three cases, shown in Figure 3. Either  $\tau_t = \tau_t^N$  (a, thread  $t_1$  in Figure 2) or  $\tau_t = \tau_t^I$  (b, c). In the former case,  $\psi^b$  has to come after the end of  $\tau_t^N$ . In the latter case, either  $\psi^b$  comes after the end of  $\tau_t^I$  (b) or is its last action or precedes the latter (c, thread  $t_2$  in Figure 2).

By the choice of  $\tau_t^N$ , in (a) and (b) the fragment of  $\tau$  in between the end of  $\tau_t^N$  and  $\psi^b$  can contain only those actions by  $t$  that are transactional ( $T_2$  in Figure 2). By the choice of  $\tau_t^I$  and  $\psi^b$ , in (c) the fragment of  $\tau$  in between  $\psi^b$  and the end of  $\tau_t^I$  cannot contain a `txbegin` action by  $t$  and, hence, can contain only those actions by  $t$  that are transactional. Furthermore, these have to come from a single transaction, started either by  $\psi^b$  or before it ( $T_4$  in Figure 2). Finally, by the choice of  $\psi^b$  the actions of  $t_0$  following  $\psi^b$  are transactional and come from the transaction of  $\psi$ , also started either by  $\psi^b$  or before it ( $T_5$  in Figure 2). Given this analysis, the transformation from  $\tau$  to  $\tau'$  can be viewed as a sequence of two: (i) erase all actions following  $\psi^b$ , except those in some of transactions that were already ongoing at this time; (ii) erase some suffixes of threads containing only transactional actions. Since transactional actions do not access global variables, they are not affected by the actions of other threads. Furthermore, as we noted in Section 5,  $\llbracket P \rrbracket(s)$  includes incomplete program computations. This allows us to conclude that  $\tau' \in \llbracket P \rrbracket(s)$ .

We now show that  $\tau''$  is valid, again referring to cases (a-c). Let  $T = \text{txof}(\psi^b, H_1\psi)$ ; then  $T \in H_1'\psi$  by the choice of  $\psi^b$  and by Definition 3(iii) we get (1). Hence, for threads  $t$  falling into cases (a) or (b),  $\tau'|_t$  may not contain aborted transactions that are also in  $H_1'\psi$ . For threads  $t$  falling into case (c), the only aborted transaction included into  $H_1'\psi$  can be the last one in  $\tau'|_t$ . Finally, above we established that  $(H_1'\psi)|_{t_0}$  does not contain aborted transactions. Hence, transactions in  $\tau'$  whose status is changed from aborted to committed when switching to  $\tau''$  do not have any actions following them in  $\tau'$ . Furthermore,  $\llbracket P \rrbracket(s)$  allows committing or aborting transactions arbitrarily. This allows us to conclude that  $\tau'' \in \llbracket P \rrbracket(s)$ . For the same reason, we get  $\tau_\psi \in \llbracket P \rrbracket(s)$ .

Finally, we show that  $\text{history}(\tau_\psi)|_{\text{-abortedtx}} = H_\psi^c$ . It is sufficient to show that  $\text{history}(\tau'')|_{\text{-abortedtx}} = H_1''\psi$ ; since  $\tau_\psi = \tau''H^{cc}$  and  $H^{cc}$  contains only committed actions, this would imply

$$\text{history}(\tau_\psi)|_{\text{-abortedtx}} = \text{history}(\tau''H^{cc})|_{\text{-abortedtx}} = \text{history}(\tau'')|_{\text{-abortedtx}}H^{cc} = H_1''\psi H^{cc} = H_\psi^c.$$

By the choice of  $\tau_t^I$  for  $t \neq t_0$ , every transaction in  $(H_1'\psi)|_t$  is also in  $\tau_t^I$ . Hence,  $H_1'\psi$  is a subsequence of  $\text{history}(\tau')$ . By the definition of  $\tau''$  and  $H_1''$ ,  $H_1''\psi$  is a subsequence of  $\text{history}(\tau'')$ . Then since  $H_1''\psi$  does not contain aborted transactions,  $H_1''\psi$  is a subsequence of  $\text{history}(\tau'')|_{\text{-abortedtx}}$ .

Thus, to prove  $\text{history}(\tau'')|_{\text{-abortedtx}} = H_1''\psi$  it remains to show that every non-aborted transaction in  $\text{history}(\tau'')$  is in  $H_1''\psi$ . Since the construction of  $\tau''$  from  $\tau'$  changes the status of only those transactions that belong to  $H_1'\psi$ , it is sufficient to show that every non-aborted transaction in  $\text{history}(\tau')$  is in  $H_1'\psi$ . Here we only consider the case when such a transaction is by a thread  $t \neq t_0$  and  $\tau'|_t = \tau_t^N \neq \varepsilon$ ; we cover the other cases in Appendix D. Let  $\phi_t^N$  be the last action in  $\tau_t^N$  and  $T = \text{txof}(\psi^b, H_1\psi) \in H_1'\psi$ . Then by Definition 3(iii) we get (1). Since  $\phi_t^N$  comes before  $\psi^b$  in  $H_1\psi$ , any transaction  $T'$  in  $\tau'|_t$  is such that  $T' \prec_{H_1\psi} T$ , which together with (1) implies the required. This concludes the proof that  $\text{history}(\tau'')|_{\text{-abortedtx}} = H_1''\psi$ . ■

We now give the other lemmas necessary for the proof. Definition 6 matches a history of  $\mathcal{T}_C$  with one of  $\mathcal{T}_A$  using the opacity relation, possibly after transforming the former with `cTMSpast`. The following lemma is used to transform a trace of  $P$  accordingly. The lemma shows that, if we consider only traces where aborted

transactions abort immediately (i.e., are of the form  $(\_, \_, \text{txbegin}) (\_, \_, \text{aborted})$ ), then the opacity relation implies observational refinement with respect to observing non-transactional actions and thread-local trace projections. This result is a simple adjustment of the one about the sufficiency of opacity for observational refinement to our setting [1, Theorem 16] (it was proved in [1] for a language where local variables are *not* rolled back upon a transaction abort; this difference, however, does not matter if aborted transactions abort immediately).

LEMMA 10. *Consider  $\tau \in \llbracket P \rrbracket(s)$  such that all the aborted transactions in  $\tau$  abort immediately. Let  $S$  be such that  $\text{history}(\tau) \sqsubseteq_{\text{op}} S$ . Then there exists  $\tau' \in \llbracket P \rrbracket(s)$  such that  $\text{history}(\tau') = S$ ,  $\tau|_{\text{-trans}} = \tau'|_{\text{-trans}}$  and  $\forall t. \tau'|_t = \tau|_t$ .*

Let  $\tau|_{\text{-abortact}}$  be the trace obtained from  $\tau$  by removing all actions inside aborted transactions, so that every such transaction aborts immediately. We can benefit from Lemma 10 because local variables are rolled back if a transaction aborts, and, hence, applying  $\cdot|_{\text{-abortact}}$  to a trace preserves its validity.

PROPOSITION 11.  $\forall \tau. \tau \in \llbracket P \rrbracket(s) \implies \tau|_{\text{-abortact}} \in \llbracket P \rrbracket(s)$ .

Finally, Definition 6 matches only histories of committed transactions, but the histories of the traces in Lemma 10 also contain aborted transactions. Fortunately, the following lemma allows us to add empty aborted transactions into the abstract history while preserving the opacity relation.

LEMMA 12. *Let  $H$  be a history where all aborted transactions abort immediately and  $S$  be such that  $H|_{\text{-abortedtx}} \sqsubseteq_{\text{op}} S$ . There exists a history  $S'$  such that  $S'|_{\text{-abortedtx}} = S$  and  $H \sqsubseteq_{\text{op}} S'$ .*

Using Definition 6(i), Proposition 11 and Lemmas 10 and 12, we can show that the TMS relation preserves non-transactional actions and thread-local observable behavior of threads whose last action is not a `fault`.

LEMMA 13.  $\forall \tau \in \llbracket P, \mathcal{T}_C \rrbracket(s). \exists \tau' \in \llbracket P, \mathcal{T}_A \rrbracket(s). (\tau'|_{\text{-trans}} = \tau|_{\text{-trans}}) \wedge (\forall t. (\tau'|_t)|_{\text{obs}} = (\tau|_t)|_{\text{obs}})$ .

PROOF OF THEOREM 8(I). Given Lemma 13, we only need to establish the preservation of faults inside transactions. Consider  $\tau_0 \in \llbracket P, \mathcal{T}_C \rrbracket(s)$  such that  $\tau_0 = \tau_1 \psi \tau_2 \phi$ , where  $\phi = (\_, t_0, \text{fault})$  is transactional and  $\psi$  is the last TM interface action by thread  $t_0$ . Then  $\tau_2|_{t_0}$  consists of transactional actions and thus does not contain accesses to global variables. Hence,  $\tau = \tau_1 \psi (\tau_2|_{t_0}) \phi \in \llbracket P, \mathcal{T}_C \rrbracket(s)$ . By our assumption,  $\mathcal{T}_C \sqsubseteq_{\text{tms}} \mathcal{T}_A$ . Then there exists  $H_\psi^c \in \text{cTMSpast}(\text{history}(\tau))$  and  $S \in \mathcal{T}_A$  such that  $H_\psi^c \sqsubseteq_{\text{op}} S$ . By Lemma 9, for some trace  $\tau_\psi$  we have  $\tau_\psi \in \llbracket P \rrbracket(s)$ ,  $\text{history}(\tau_\psi)|_{\text{-abortedtx}} = H_\psi^c$  and  $\tau_\psi|_{t_0} = \tau|_{t_0}$ . By Proposition 11,  $\tau_\psi|_{\text{-abortact}} \in \llbracket P \rrbracket(s)$ . Using Lemma 12, we get a history  $S'$  such that  $\text{history}(\tau_\psi|_{\text{-abortact}}) \sqsubseteq_{\text{op}} S'$  and  $S'|_{\text{-abortedtx}} = S$ . Since  $S \in \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under immediate aborts, we get  $S' \in \mathcal{T}_A$ . Hence, by Lemma 10, for some  $\tau' \in \llbracket P, \mathcal{T}_A \rrbracket(s)$  we have  $\tau'|_{t_0} = \tau_\psi|_{t_0} = \tau|_{t_0} = \_ \phi$ , as required. ■

## 6.2 Proof Sketch for Theorem 8(ii) (Necessity)

Consider  $\mathcal{T}_C$  and  $\mathcal{T}_A$  such that  $\mathcal{T}_C \preceq \mathcal{T}_A$  and  $\mathcal{T}_A$  satisfies the closure conditions stated in the theorem. To show that for any  $H_0 \in \mathcal{T}_C$  we have  $H_0 \sqsubseteq_{\text{tms}} \mathcal{T}_A$ , we have to establish conditions (i) and (ii) from Definition 6. We sketch the more interesting case of (ii) (cf. Lemma 17), in which  $H_0 = H_1 \psi H_2 = H H_2 \in \mathcal{T}_C$ , where  $\psi$  is a response action by a thread  $t_0$ . We assume that  $\psi$  is not a committed or aborted action: the full proof considers the case of aborted actions separately (Lemma 18). We need to find  $H^c \in \text{cTMSpast}(H)$  and  $S \in \mathcal{T}_A$  such that  $H^c \sqsubseteq_{\text{op}} S$ .

To this end, we construct a program  $P_H$  (as we explain further below) where every thread  $t$  performs the sequence of transactions specified in  $H|_t$ . The program monitors certain properties of the TM behavior, e.g., checking that the return values obtained from methods of transactional objects in committed transactions correspond to those in  $H$  and that the real-time order between actions includes that in  $H$ . If these properties hold, thread  $t_0$  ends by executing the `fault` command. Let  $s$  be a state with all variables set to distinguished values. We next construct a trace  $\tau \in \llbracket P_H, \mathcal{T}_C \rrbracket(s)$  such that  $\text{history}(\tau) = H$  and  $t_0$  faults in  $\tau$ . By Definition 7, there exists  $\tau' \in \llbracket P_H, \mathcal{T}_A \rrbracket(s)$  such that  $t_0$  faults in  $\tau'$ . However, the program  $P_H$  is constructed so that  $t_0$  can fault in  $\tau'$  only if the properties of the TM behaviour the program monitors hold, and thus  $H$  is related to  $\text{history}(\tau')$  in a certain way. This relationship allows us to construct  $H^c \in \text{cTMSpast}(H)$  from  $H$  and  $S \in \mathcal{T}_A$  from  $\text{history}(\tau')$  such that  $H^c \sqsubseteq_{\text{op}} S$ .

In more detail, thread  $t_0$  in  $P_H$  monitors the return status of every transaction and the return values obtained inside the atomic blocks corresponding to transactions committed in  $H|_{t_0}$  and the (live) transaction of  $\psi$ . If there

is a mismatch with  $H|_{t_0}$ , this is recorded in a special local variable. At the end of the transaction of  $\psi$ ,  $t_0$  checks the variable and faults if the TM behavior matched  $H|_{t_0}$ . This construction is motivated by the fact that faulting is the only observation Definition 7 allows us to make about the behavior of the live transaction of  $\psi$ . Since the definition does not correlate actions by threads  $t$  other than  $t_0$  between  $\tau$  and  $\tau'$ , such threads monitor TM behavior differently: if there is a mismatch with  $H|_t$ , a thread  $t$  faults immediately. Since a trace can have at most one fault and  $t_0$  faults in  $\tau'$ , this ensures that any committed transaction in  $\tau'$  behaves as in  $H$ .

To check whether an execution of  $P_H$  complies with the real-time order in  $H$ , for each transaction in  $H$ , we introduce a global variable  $g$ , which is initially 0 and is set to 1 by the thread executing the transaction right after the transaction completes, by a command following the corresponding atomic block. Before starting a transaction, each thread checks whether all transactions preceding this one in the real-time order in  $H$  have finished by reading the corresponding  $g$  variables. Thread  $t_0$  records the outcome in the special local variable checked at the end; all other threads fault upon detecting a mismatch.

This construction of  $P_H$  allows us to infer that: (i) the projection of  $\text{history}(\tau')|_{t_0}$  to committed transactions and  $\text{txof}(\varphi, \tau')$  is equal to the corresponding projection of  $H|_{t_0}$ ; (ii) for all other threads  $t$  a similar relationship holds for the prefix of  $\text{history}(\tau')|_t$  ending with the last transaction preceding  $\text{txof}(\varphi, \tau')$  in the real-time order; (iii) the real-time order in  $\text{history}(\tau')$  includes that in  $H$ . Transactions concurrent with  $\text{txof}(\varphi, \tau')$  in  $\tau'$  may behave differently from  $H$ . However, checks done by  $P_H$  inside these transactions ensure that, if such a transaction  $T$  is visible in  $\tau'$ , then the return values inside  $T$  match those in  $H$ . The checks on the global variables  $g$  done right before  $T$  also ensure that all transactions preceding  $T$  in the real-time order in  $H$  commit or abort in  $\tau'$  as prescribed by  $H$ . This relationship between  $H$  and  $\text{history}(\tau')$  allows us to establish the requirements of Definition 6(ii). ■

## 7 Related Work

When presenting TMS [3], Doherty et al. discuss why it allows programmers to think only of serial executions of their programs, in which the actions of a transaction appear consecutively. This discussion—corresponding to our sufficiency result—is informal, since the paper lacks a formal model for programs and their semantics. Most of it explains how Definition 6(i) ensures the correctness of committed transactions. The discussion of the most challenging case of live transactions—corresponding to Definition 6(ii) and our Lemma 9—is one paragraph long. It only roughly sketches the construction of a trace with an abstract history allowed by TMS and does not give any reasoning for why this trace is a valid one, but only claims that constraints in Definition 6(ii) ensure this. This reasoning is very delicate, as indicated by our proof of Lemma 9, which carefully selects which actions to erase when transforming the trace. Moreover, Doherty et al. do not try to argue that TMS is the weakest condition possible, as we established by our necessity result.

Another TM consistency condition, weaker than opacity but incomparable to TMS, is *virtual world consistency* (VWC) [10]. Like TMS, VWC allows every operation in a live or aborted transaction to be explained by a separate abstract history. However, it places different constraints on the choice of abstract histories, which do not take into account the real-time order between actions. Because of this, VWC does not imply observational refinement for our programming language: taking into account the real-time order is necessary when threads can communicate via global variables outside transactions.

Our earlier paper [1] has laid the groundwork for relating TM consistency and observational refinement, and it includes a detailed comparison with related work on opacity and observational refinement. The present paper considers a much more challenging case of a language where local variables are rolled back upon an abort. To handle this case, we have developed new techniques, such as establishing the live transaction insensitivity property (Lemma 9) to prove sufficiency and proposing monitor programs for the nontrivial constraints used in the TMS definition to prove necessity. Similarly to [1] and other papers using observational refinement to study consistency conditions [4, 5], we reformulate TMS so that it is not restricted to an abstract TM  $\mathcal{T}_A$  that executes transactions atomically. This generality, not allowed by the original TMS definition, has two benefits. First, our reformulation can be used to compare two TM implementations, e.g., an optimized and an unoptimized one. Second, dealing with the general definition forces us to explicitly state the closure properties required from the abstract TM, rather than having them follow implicitly from its atomic behavior.

## References

- [1] H. Attiya, A. Gotsman, S. Hans, and N. Rinetzky. A programming language perspective on transactional memory consistency. In *PODC*, pages 309–318, 2013.
- [2] H. Attiya, S. Hans, P. Kuznetsov, and S. Ravi. Safety of deferred update in transactional memory. In *ICDCS*, 2013.
- [3] S. Doherty, L. Groves, V. Luchangco, and M. Moir. Towards formally specifying and verifying transactional memory. *Formal Aspects of Computing*, 25(5):769–799, 2013.
- [4] I. Filipovic, P. W. O’Hearn, N. Rinetzky, and H. Yang. Abstraction for concurrent objects. In *ESOP*, pages 252–266, 2009.
- [5] A. Gotsman and H. Yang. Liveness-preserving atomicity abstraction. In *ICALP (2)*, pages 453–465, 2011.
- [6] R. Guerraoui and M. Kapalka. On the correctness of transactional memory. In *PPOPP*, pages 175–184, 2008.
- [7] J. He, C. Hoare, and J. Sanders. Data refinement refined. In *ESOP*, pages 187–196, 1986.
- [8] J. He, C. Hoare, and J. Sanders. Prespecification in data refinement. *Information Processing Letters*, 25(2):71–76, 1987.
- [9] M. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463–492, 1990.
- [10] D. Imbs and M. Raynal. Virtual world consistency: A condition for STM systems (with a versatile protocol with invisible read operations). *Theor. Comput. Sci.*, 444:113–127, 2012.
- [11] M. Lesani, V. Luchangco, and M. Moir. Putting opacity in its place. In *WTTM*, 2012.
- [12] C. H. Papadimitriou. The serializability of concurrent database updates. *J. ACM*, 26:631–653, October 1979.
- [13] Scala STM Expert Group. Scala STM quick start guide, 2012. [http://nbronson.github.io/scala-stm/quick\\_start.html](http://nbronson.github.io/scala-stm/quick_start.html).

## A Restrictions on Variable Accesses

To formalise restrictions on accesses to variables by primitive commands, we partition the set  $\text{PComm} - \{\text{fault}\}$  into  $2m$  classes:  $\text{PComm} - \{\text{fault}\} = \bigsqcup_{t=1}^m (\text{LPcomm}_t \uplus \text{GPcomm}_t)$ . The intention is that commands from  $\text{LPcomm}_t$  can access only the local variables of thread  $t$  ( $\text{LVar}_t$ ); commands from  $\text{GPcomm}_t$  can additionally access global variables ( $\text{LVar}_t \uplus \text{GVar}$ ). We formalize these restrictions in Appendix C. To ensure that a thread  $t$  does not access local variables of other threads, we require that the thread cannot mention such variables in the conditions of `if` and `while` commands and can only use primitive commands from  $\text{LPcomm}_t \uplus \text{GPcomm}_t$ .

## B Formal Definition of Traces

DEFINITION 14 (Traces). A **trace**  $\tau$  is a finite sequence of actions, satisfying the following conditions:

- (i) every action in  $\tau$  has a unique identifier: if  $\tau = \tau_1(a_1, -, -)\tau_2(a_2, -, -)\tau_3$  then  $a_1 \neq a_2$ .
- (ii) no action follows a `fault`: if  $\tau = \tau'\varphi$  then  $\tau'$  does not contain a `fault` action.
- (iii) request and response actions are properly matched: for every thread  $t$ ,  $\text{history}(\tau)|_t$  consists of alternating request and corresponding response actions, starting from a request action;
- (iv) actions denoting the beginning and end of transactions are properly matched: for every thread  $t$ , in the projection of  $\text{history}(\tau)|_t$  to `txbegin`, `committed` and `aborted` actions, `txbegin` alternates with `committed` or `aborted`, starting from `txbegin`;
- (v) call and ret actions occur only inside transactions: for every thread  $t$ , if  $\tau|_t = \tau_1\psi\tau_2$  for a call or ret action  $\psi$ , then  $\tau_1 = \tau'_1\psi'\tau''_1$  for some `txbegin` action  $\psi'$ , and  $\tau'_1$  and  $\tau''_1$  such that  $\tau''_1$  does not contain `committed` or `aborted` actions;
- (vi) commands in  $\tau$  do not access local variables of other threads: if  $(-, t, c) \in \tau$  then  $c \in \text{LPcomm}_t \uplus \text{GPcomm}_t \uplus \{\text{fault}\}$ ;
- (vii) commands in  $\tau$  do not access global variables inside a transaction: if  $\tau = \tau_1(-, t, c)\tau_2$  for  $c \in \text{GPcomm}_t$ , then it is not the case that  $\tau_1 = \tau'_1(-, t, \text{txbegin})\tau''_1$ , where  $\tau''_1$  does not contain `committed` or `aborted` actions.

## C Formal Definition of the Semantics of the Programming Language

This section formally defines the set  $\llbracket P \rrbracket(s)$ . It is computed in two stages. First, we compute a set  $A(P)$  of traces that resolves all issues regarding sequential control flow and interleaving. Intuitively, if one thinks of each thread  $C_t$  in  $P$  as a control-flow graph, then  $A(P)$  contains all possible interleavings of paths in the graphs of  $C_t$ ,  $t \in \text{ThreadID}$  starting from their initial nodes. The set  $A(P)$  is a superset of all the traces that can actually be executed: e.g., if a thread executes the command “ $x := 1$ ; `if` ( $x = 1$ )  $y := 1$  `else`  $y := 2$ ” where  $x, y$  are local variables, then  $A(P)$  will contain a trace where  $y := 2$  is executed instead of  $y := 1$ . To filter out such nonsensical traces, we *evaluate* every trace to determine whether it is **valid**, i.e., whether its control flow is consistent with the effect of its actions on program variables. This is formalized by a function  $\text{eval} : \text{State} \times \text{Trace} \rightarrow \mathcal{P}(\text{State}) \cup \{\zeta\}$  that, given an initial state and a trace, produces the set of states resulting from executing the actions in the trace, an empty set if the trace is invalid, or a special state  $\zeta$  if the trace contains a `fault` action. Thus,  $\llbracket P \rrbracket(s) = \{\tau \in A(P) \mid \text{eval}(s, \tau) \neq \emptyset\}$ .

When defining the semantics, we encode the evaluation of conditions in `if` and `while` statements with `assume` commands. More specifically, we expect that the sets  $\text{LPcomm}_t$  contain special primitive commands `assume`( $b$ ), where  $b$  is a Boolean expression over local variables of thread  $t$ , defining the condition. We state their semantics formally below; informally, `assume`( $b$ ) does nothing if  $b$  holds in the current program state, and stops the computation otherwise. Thus, it allows the computation to proceed only if  $b$  holds. The `assume` commands are only used in defining the semantics of the programming language; hence, we forbid threads from using them directly.

**The trace set  $A(P)$ .** The function  $A'(\cdot)$  in Figure 4 maps commands and programs to sequences of actions they may produce. Technically,  $A'(\cdot)$  might contain sequences that are not traces, e.g., because they do not have unique identifiers or continue beyond a `fault` command. This is resolved by intersecting the set  $A'(P)$  with the set of all traces to define  $A(P)$ .  $A'(C)t$  gives the set of action sequences produced by a command  $C$  when it is executed

$$\begin{aligned}
A'(c)t &= \{(\_, t, c)\} \\
A'(C_1; C_2)t &= \{\tau_1 \tau_2 \mid \tau_1 \in A'(C_1)t \wedge \tau_2 \in A'(C_2)t\} \\
A'(\text{if } (b) \text{ then } C)t &= \{(\_, t, \text{assume}(b)) (A'(C)t)\} \\
A'(\text{if } (b) \text{ then } C_1 \text{ else } C_2)t &= \{(\_, t, \text{assume}(b)) \tau_1 \mid \tau_1 \in A'(C_1)t\} \cup \{(\_, t, \text{assume}(\neg b)) \tau_2 \mid \tau_2 \in A'(C_2)t\} \\
A'(\text{while } (b) \text{ do } C)t &= \{((\_, t, \text{assume}(b)) (A'(C)t))^* (\_, t, \text{assume}(\neg b))\} \\
A'(x := o.f(e))t &= \{(\_, t, \text{assume}(e = n)) (\_, t, \text{call } o.f(n)) (\_, t, \text{ret}(n') o.f) (\_, t, x := n') \mid n, n' \in \mathbb{Z}\} \cup \\
&\quad \{(\_, t, \text{assume}(e = n)) (\_, t, \text{call } o.f(n)) (\_, t, \text{aborted}) \mid n \in \mathbb{Z}\} \\
A'(x := \text{atomic } \{C\})t &= \{(\_, t, \text{txbegin}) (\_, t, \text{aborted}) (\_, t, x := \text{aborted})\} \cup \\
&\quad \{(\_, t, \text{txbegin}) (\_, t, \text{OK}) \tau (\_, t, \text{aborted}) (\_, t, x := \text{aborted}) \mid \tau \in A'(C)t \wedge (\_, t, \text{aborted}) \notin \tau\} \cup \\
&\quad \{(\_, t, \text{txbegin}) (\_, t, \text{OK}) \tau (\_, t, \text{txcommit}) (\_, t, r) (\_, t, x := r) \mid \tau \in A'(C)t \wedge \\
&\quad \quad (\_, t, \text{aborted}) \notin \tau \wedge (r = \text{committed} \vee r = \text{aborted})\} \\
A'(C_1 \parallel \dots \parallel C_m) &= \text{prefix}(\bigcup \{\text{interleave}(\tau_1, \dots, \tau_m) \mid \forall t. 1 \leq t \leq m \implies \tau_t \in A'(C_t)t\}) \\
A(P) &= A'(P) \cap \text{Trace}
\end{aligned}$$

Figure 4: The definition of  $A(P)$ .

by thread  $t$ . To define  $A'(P)$ , we first compute the set of all the interleavings of action sequences produced by the threads constituting  $P$ . Formally,  $\tau \in \text{interleave}(\tau_1, \dots, \tau_m)$  if and only if every action in  $\tau$  is performed by some thread  $t \in \{1, \dots, m\}$ , and  $\tau|_t = \tau_t$  for every thread  $t \in \{1, \dots, m\}$ . We then let  $A'(P)$  be the set of all prefixes of the resulting sequences, as denoted by the prefix operator. We take prefix closure here to account for incomplete program computations as well as those in which the scheduler preempts a thread forever.

$A'(c)t$  returns a singleton set with the action corresponding to the primitive command  $c$  (primitive commands execute atomically).  $A'(C_1; C_2)t$  concatenates all possible action sequences corresponding to  $C_1$  with those corresponding to  $C_2$ . The set of action sequences of a conditional considers cases where either branch is taken. We record the decision using an assume action; at the evaluation stage, this allows us to ensure that this decision is consistent with the program state. The set of action sequences for a loop is defined using the Kleene closure operator  $*$  to produce all possible unfoldings of the loop body. Again, we record branching decisions using assume actions.

The set of action sequences of a method invocation  $x := f(e)$  includes both sequences where the method executes successfully and where the current transaction is aborted. The former set is constructed by nondeterministically choosing two integers  $n$  and  $n'$  to describe the parameter  $n$  and the return value  $n'$  for the method call. To ensure that  $e$  indeed evaluates to  $n$ , we insert  $\text{assume}(e = n)$  before the call action, and to ensure that  $x$  gets the return value  $n'$ , we add the assignment  $x := n'$  after the  $\text{ret}$  action. Note that some of the choices here might not be feasible: the chosen  $n$  might not be the value of the parameter expression  $e$  when the method is invoked, or the method might never return  $n'$  when called with  $n$ . Such infeasible choices are filtered out at the following stages of the semantics definition: the former in the definition of  $\llbracket P \rrbracket(s)$  by the use of evaluation and the semantics of assume, and the latter in the definition of  $\llbracket P, \mathcal{T} \rrbracket(s)$  by selecting the sequences from  $\llbracket P \rrbracket(s)$  that interact with the transactional memory correctly. The set of action sequences of  $x := \text{atomic } \{C\}$  contains those in which  $C$  is aborted in the middle of its execution (at an object operation or right after it begins) and those in which  $C$  executes until completion and then the transaction commits or aborts.

**Semantics of primitive commands.** To define evaluation, we assume a semantics of every command  $c \in \text{PComm} - \{\text{fault}\}$ , given by a function  $\llbracket c \rrbracket$  that defines how the program state is transformed by executing  $c$ . As we noted before, different classes of primitive commands are supposed to access only certain subsets of variables. To ensure that this is indeed the case, we define  $\llbracket c \rrbracket$  as a function of only those variables that  $c$  is allowed to access. Namely, the semantics of  $c \in \text{LPComm}_t$  is given by

$$\llbracket c \rrbracket : (\text{LVar}_t \rightarrow \mathbb{Z}) \rightarrow \mathcal{P}(\text{LVar}_t \rightarrow \mathbb{Z}).$$

The semantics of  $c \in \text{GPcomm}_t$  is given by

$$\llbracket c \rrbracket : ((\text{LVar}_t \uplus \text{GVar}) \rightarrow \mathbb{Z}) \rightarrow \mathcal{P}((\text{LVar}_t \uplus \text{GVar}) \rightarrow \mathbb{Z}).$$

Note that we allow  $c$  to be non-deterministic.

For a valuation  $q$  of variables that  $c$  is allowed to access,  $\llbracket c \rrbracket(q)$  yields the set of their valuations that can be obtained by executing  $c$  from a state with variable values  $q$ . For example, an assignment command  $x := g$  has the following semantics:

$$\llbracket x := g \rrbracket(q) = \{q[g \mapsto q(g)]\}.$$

We define the semantics of assume commands following the informal explanation given at the beginning of this section: for example,

$$\llbracket \text{assume}(x = n) \rrbracket(q) = \begin{cases} \{q\}, & \text{if } q(x) = n; \\ \emptyset, & \text{otherwise.} \end{cases} \quad (2)$$

Thus, when the condition in assume does not hold of  $q$ , the command stops the computation by not producing any output.

We lift functions  $\llbracket c \rrbracket$  to full states by keeping the variables that  $c$  is not allowed to access unmodified and producing  $\not\downarrow$  if  $c$  faults. For example, if  $c \in \text{LPcomm}_t$ , then

$$\llbracket c \rrbracket(s) = \{s|_{\text{LVar} \setminus \text{LVar}_t} \uplus q \mid q \in \llbracket c \rrbracket(s|_{\text{LVar}_t})\},$$

where  $s|_V$  is the restriction of  $s$  to variables in  $V$ . Finally, we let

$$\llbracket \text{fault} \rrbracket(s) = \not\downarrow,$$

so that the only way a program can fault is by executing the `fault` command.

**Trace evaluation.** Using the semantics of primitive commands, we first define the evaluation of a single action on a given state:

$$\begin{aligned} \text{eval} &: \text{State} \times \text{Action} \rightarrow \mathcal{P}(\text{State}) \cup \{\not\downarrow\} \\ \text{eval}(s, (\_, t, c)) &= \llbracket c \rrbracket(s); \\ \text{eval}(s, \psi) &= \{s\}. \end{aligned}$$

Note that this does not change the state  $s$  as a result of TM interface actions, since their return values are assigned to local variables by separate actions introduced when generating  $A(P)$ . We then lift `eval` to traces as follows:

$$\begin{aligned} \text{eval} &: \text{State} \times \text{Trace} \rightarrow \mathcal{P}(\text{State}) \cup \{\not\downarrow\} \\ \text{eval}(s, \tau) &= \begin{cases} \emptyset, & \text{if } \tau = \tau' \varphi \wedge \text{eval}(s, \tau') = \emptyset; \\ \text{evalna}(s, \tau|_{\neg \text{abortact}}), & \text{otherwise,} \end{cases} \end{aligned}$$

where

$$\text{evalna}(s, \tau) = \begin{cases} \{s\}, & \text{if } \tau = \varepsilon; \\ \{s'' \in \text{eval}(s', \varphi) \mid s' \in \text{evalna}(s, \tau')\}, & \text{if } \tau = \tau' \varphi. \end{cases}$$

The set of states resulting from evaluating trace  $\tau$  from state  $s$  is effectively computed by the helper function  $\text{evalna}(s, \tau)$ , which ignores actions inside aborted transactions to model local variable roll-back. However, ignoring the contents of aborted transactions completely poses a risk that we might consider traces including sequences of actions inside aborted transactions that yield an empty set of states. To mitigate this,  $\text{eval}(s, \tau)$  recursively evaluates every prefix of  $\tau$ , thus ensuring that sequences of actions inside aborted transaction are valid.

As we explained in Section 5, we define  $\llbracket P \rrbracket(s)$  as the set of those traces from  $A(P)$  that can be evaluated from  $s$  without getting stuck, as formalized by `eval`. Note that this definition enables the semantics of assume defined by (2) to filter out traces that make branching decisions inconsistent with the program state. For example, consider again the program “ $x := 1$ ; if  $(x = 1)$   $y := 1$  else  $y := 2$ ”. The set  $A(P)$  includes traces where both branches are explored. However, due to the semantics of the assume actions added to the traces according to Figure 4, only the trace executing  $y := 1$  will result in a non-empty set of final states after the evaluation and, therefore, only this trace will be included into  $\llbracket P \rrbracket(s)$ .

## D Additional Proofs

### D.1 Remaining Cases from the Proof of Lemma 9

- $t \neq t_0$  is such that  $\tau'|_t = \tau_t^I \neq \varepsilon$ . Let  $\psi_t^I$  be the last action in  $\tau_t^I$ . Let  $T = \text{txof}(\psi_t^I, H_1\psi)$ . By the choice of  $\tau_t^I$  we have  $T \in H_1\psi$ ; then by Definition 3(iii) we get (1). Since any transaction  $T'$  in  $\text{history}(\tau'|_t)$  is either  $T$  or is such that  $T' \prec_{(H_1\psi)|_t} T$ , this implies the required.
- $t = t_0$ . Let  $T = \text{txof}(\psi, H_1\psi) \in H_1\psi$ . Then by Definition 3(iii) we get (1). Since any transaction  $T'$  in  $\text{history}(\tau'|_{t_0})$  is either  $T$  or is such that  $T' \prec_{(H_1\psi)|_{t_0}} T$ , this implies the required.

### D.2 Proof of Lemma 12

Let  $n$  be the number of aborted transactions in  $H$ . To construct the desired  $S'$ , we inductively construct a sequence of histories  $S_i, i = 0..n$  such that

$$\begin{aligned} |\text{aborted}(S_i)| = i; \quad S_i|_{\text{-abortedtx}} = S; \quad \{\psi \mid \psi \in S_i\} \subseteq \{\psi \mid \psi \in H\}; \\ \forall \psi_1, \psi_2 \in S_i. \psi_1 \prec_H \psi_2 \implies \psi_1 \prec_{S_i} \psi_2. \end{aligned} \quad (3)$$

We then let  $S' = S_n$ , so that  $H \sqsubseteq_{\text{op}} S'$ .

For  $i = 0$ , we take  $S_0 = S$ , and all the requirements in (3) hold vacuously. Assume a history  $S_i$  satisfying (3) was constructed; we get  $S_{i+1}$  from  $S_i$  by the following construction. Let  $H = H_1\psi_b H_2\psi_a H_3$ , where  $\psi_b = (\_, t, \text{txbegin}), \psi_a = (\_, t, \text{aborted}), H_2|_t = \varepsilon$  and

$$\neg \exists \psi'. \psi' = (\_, \_, \text{txbegin}) \in H_1 \wedge \text{txof}(\psi', H) \in \text{aborted}(H) \wedge \psi' \notin S_i.$$

That is, out of all aborted transactions in  $H$  that are not in  $S_i$ ,  $\psi_b\psi_a$  is the one with the earliest  $\text{txbegin}$ . We now consider two cases.

**Case I:**  $H_1$  does not contain a committed or an aborted action.

In this case, let  $S_{i+1} = \psi_b\psi_a S_i$ . We only need to show that for any  $\psi' \in S_i$  we have  $\psi' \prec_H \psi_b \implies \psi' \prec_{S_{i+1}} \psi_b$  and  $\psi_a \prec_H \psi' \implies \psi_a \prec_{S_{i+1}} \psi'$ . The latter holds by the construction of  $S_{i+1}$ . To show the former, observe that, since  $H_1$  does not contain a committed or aborted action, it cannot contain actions by thread  $t$ . Hence, we cannot have  $\psi' \prec_H \psi_b$  for any  $\psi'$ .

**Case II:**  $H_1$  contains a committed or an aborted action.

Let  $\psi$  be the last committed or aborted action in  $S_i$  that is also in  $H_1$  and let  $S_i = S'\psi S''$ . We then let  $S_{i+1} = S'\psi\psi_b\psi_a S''$ . We again need to show that for any  $\psi' \in S_i$  we have  $\psi' \prec_H \psi_b \implies \psi' \prec_{S_{i+1}} \psi_b$  and  $\psi_a \prec_H \psi' \implies \psi_a \prec_{S_{i+1}} \psi'$ .

Assume  $\psi' \prec_H \psi_b$  for some  $\psi' \in S_i$ ; then  $\psi' \in H_1$ . By the choice of  $\psi_b$  and  $\psi_a$ , all the committed and aborted actions in  $H_1$  are in  $S_i$ , and by the choice of  $\psi$ , all such actions are in  $S'\psi$ . Hence, if  $\psi'$  is a committed or an aborted action, then  $\psi' \in S'\psi$  and, hence,  $\psi' \prec_{S_{i+1}} \psi_b$ . If  $\psi'$  is by thread  $t$ , then it is either a committed or an aborted action (and, hence,  $\psi' \prec_{S_{i+1}} \psi_b$ ) or it precedes such an action  $\psi'' \in S_i$  by  $t$  in  $H_1$ :  $\psi' \prec_{H_1} \psi''$ . Then  $\psi' \prec_{S_{i+1}} \psi''$  and  $\psi'' \prec_{S_{i+1}} \psi_b$ , which implies  $\psi' \prec_{S_{i+1}} \psi_b$ .

Now assume  $\psi_a \prec_H \psi'$  for some  $\psi' \in S_i$ ; then  $\psi' \in H_3$ . If  $\psi'$  is a  $\text{txbegin}$  action, then  $\psi \prec_H \psi'$ . Hence,  $\psi \prec_{S_i} \psi'$ , i.e.,  $\psi' \in S''$ , which implies  $\psi_a \prec_{S_{i+1}} \psi'$ . If  $\psi'$  is by thread  $t$ , then it is either a  $\text{txbegin}$  action (and, hence,  $\psi_a \prec_{S_{i+1}} \psi'$ ) or it follows such an action  $\psi'' \in S_i$  by thread  $t$  in  $H_3$ :  $\psi'' \prec_{H_3} \psi'$ . Then  $\psi'' \prec_{S_{i+1}} \psi'$  and  $\psi_a \prec_{S_{i+1}} \psi''$ , which implies  $\psi_a \prec_{S_{i+1}} \psi'$ . ■

### D.3 Proof of Lemma 13

Let  $H = \text{history}(\tau)$ . By assumption,  $\mathcal{T}_C \sqsubseteq_{\text{tms}} \mathcal{T}_A$ . Hence, there exist histories  $H^c \in \text{comp}(H|_{\text{-live}})$  and  $S \in \mathcal{T}_A$  such that  $H^c|_{\text{com}} \sqsubseteq_{\text{op}} S$ . Then  $H^c = (H|_{\text{-live}})H'$  for some  $H'$ . Let  $\tau^c$  be the trace obtained from  $\tau$  in the same way as  $H^c$  is obtained from  $H$ :  $\tau^c = \tau_0 H'$ , where  $\tau_0$  is obtained from  $\tau$  by discarding all live transactions; thus,  $\text{history}(\tau^c) = H^c$ . It is easy to see that  $\tau^c \in \llbracket P \rrbracket(s)$ . Besides,  $\tau^c|_{\text{-trans}} = \tau|_{\text{-trans}}$  and



$(\tau|_t)|_{\text{obs}}$  is a prefix of  $(\tau^c|_t)|_{\text{obs}}$  for any  $t$ . Let  $\tau^{na} = \tau^c|_{\text{-abortact}}$ . By Proposition 11 we get  $\tau^{na} \in \llbracket P \rrbracket(s)$ . Since  $(H^c|_{\text{-abortact}})|_{\text{-abortedtx}} = H^c|_{\text{com}} \sqsubseteq_{\text{op}} S$ , by Lemma 12, for some history  $S'$  we have  $\text{history}(\tau^{na}) = H^c|_{\text{-abortact}} \sqsubseteq_{\text{op}} S'$  and  $S'|_{\text{-abortedtx}} = S$ . Since  $\mathcal{T}_A$  is closed under immediate aborts and  $S \in \mathcal{T}_A$ , we have  $S' \in \mathcal{T}_A$ . We have  $\tau^{na} \in \llbracket P \rrbracket(s)$ ; hence, by Lemma 10, there exists a trace  $\tau'' \in \llbracket P \rrbracket(s)$  such that  $\text{history}(\tau'') = S' \in \mathcal{T}_A$ ,

$$\tau''|_{\text{-trans}} = \tau^{na}|_{\text{-trans}} = \tau^c|_{\text{-trans}} = \tau|_{\text{-trans}},$$

and  $\tau''|_t = \tau^{na}|_t$  for any  $t$ . Let  $\tau'$  be the history obtained from  $\tau''$  by discarding the actions in  $H^c$ , which are last by the corresponding threads. Then

$$\tau'|_{\text{-trans}} = \tau''|_{\text{-trans}} = \tau|_{\text{-trans}},$$

$\tau' \in \llbracket P \rrbracket(s)$  and, since  $\mathcal{T}_A$  is closed under removing transaction responses,  $\text{history}(\tau') \in \mathcal{T}_A$ . Given  $\tau''|_t = \tau^{na}|_t$ , it is also easy to check that  $(\tau'|_t)|_{\text{obs}} = (\tau|_t)|_{\text{obs}}$ , as required.  $\blacksquare$

#### D.4 Proof of Theorem 8(ii) (Necessity)

Let  $\tau|_i$  denote the prefix of a trace  $\tau$  containing  $i$  actions and let  $\tau|_{\varphi}$  denote the prefix of  $\tau$  until (and including) the action  $\varphi$ .

DEFINITION 15. Two traces  $\tau$  and  $\tau'$  are **equivalent up to identifiers**, denoted  $\tau \equiv \tau'$ , if  $|\tau| = |\tau'|$  and for every  $i = 1 \dots |\tau|$ , actions  $\tau(i)$  and  $\tau'(i)$  may differ only in their action identifiers.

LEMMA 16. Let  $\mathcal{T}_C$  and  $\mathcal{T}_A$  be transactional memories such that  $\mathcal{T}_C \preceq \mathcal{T}_A$ , and  $\mathcal{T}_A$  is closed under removing live and aborted transactions and commit-pending transactions completion. Let  $H$  be a history in  $\mathcal{T}_C$ . There exists a history  $H^c \in \text{comp}(H|_{\text{-live}})$  and a history  $S \in \mathcal{T}_A$ , such that  $H^c|_{\text{com}} \sqsubseteq_{\text{op}} S$ .

**Proof:** Let us choose an integer value  $u \neq 1$  which does not appear in  $H$ . We use the following shorthands:

- We let  $m$  be the largest thread identifier occurring in  $H$ .
- We denote by  $k^t$  the number of transactions started by thread  $t$  in  $H$ , i.e., the number of  $(\_, t, \text{txbegin})$  actions in  $H$ .
- We partition  $H|_t$  into  $k^t$  subsequences:  $H|_t = H_1^t \dots H_{k^t}^t$ , where  $H_i^t$  is comprised of the actions in the  $i$ -th transaction of  $t$ . Specifically,  $H_i^t(1) = (\_, t, \text{txbegin})$ .
- We let  $c_i^t$  be the outcome of the  $i$ -th transaction of thread  $t$ , i.e.,  $c_i^t = \text{committed}$  or  $c_i^t = \text{aborted}$ . If the transaction is not completed,  $c_i^t$  is undefined.
- We denote by  $q_i^t$  the number of call actions of thread  $t$  in its  $i$ -th transaction, i.e., in  $H_i^t$ .
- We let  $(\_, t, \text{call } o_{i,j}^t \cdot f_{i,j}^t(n_{i,j}^t))$  be the  $j$ -th call action of thread  $t$  in its  $i$ -th transaction.
- We let  $(\_, t, \text{ret}(r_{i,j}^t) o_{i,j}^t \cdot f_{i,j}^t(n_{i,j}^t))$  be the  $j$ -th ret action of thread  $t$  in its  $i$ -th transaction. If the response to  $(\_, t, \text{call } o_{i,j}^t \cdot f_{i,j}^t(n_{i,j}^t))$  is an aborted action, we let  $r_{i,j}^t = \text{aborted}$ . If there is no response to  $(\_, t, \text{call } o_{i,j}^t \cdot f_{i,j}^t(n_{i,j}^t))$ , i.e., the transaction is live and  $(\_, t, \text{call } o_{i,j}^t \cdot f_{i,j}^t(n_{i,j}^t))$  is its last action, then we let  $r_{i,j}^t = u$ .
- We denote by  $\text{lasttx}(t, i, t')$  the number of transactions of thread  $t'$  in  $H$  that either committed or aborted before the  $i$ -th transaction of thread  $t$  started, i.e., the number of  $(\_, t', \text{committed})$  and  $(\_, t', \text{aborted})$  actions preceding the  $i$ -th  $(\_, t, \text{txbegin})$  action in  $H$ .

For every thread  $t = 1..m$  we construct a straight-line command

$$C_H^t = GP_1^t; CP_1^t; GP_2^t; CP_2^t; \dots; GP_{k^t}^t; CP_{k^t}^t,$$

where  $GP_i^t$  and  $CP_i^t$  are defined in Figure 5. Here the  $g_i^t$  variables are global and all others are local. The  $g_i^t$  variables are used to monitor the real-time order:  $g_i^t$  is written only by thread  $t$  and is used to signal that the  $i$ -th transaction of thread  $t$  ended. The variable  $z_{i,t'}^t$  is used to record whether the  $\text{lasttx}(t, i, t')$ -th transaction of thread  $t'$  signaled that it had ended before the  $i$ -th transaction of thread  $t$  started. As  $\text{lasttx}(t, i, t')$  might be 0, we add a dummy variable  $g_0^t$  for every thread  $t$ . Later in the proof we execute the program from a state in which  $g_0^t$  is initialised to 1. The variable  $w_i^t$  records whether the  $i$ -th transaction of thread  $t$  committed or aborted. The variables  $y_{i,j}^t$  record the return value of the  $j$ -th object method invocation in the  $i$ -th transaction of thread  $t$ .

$$\begin{aligned}
GP_i^t &= z_{i,1}^t := g_{\text{lasttx}(t,i,1)}^1; \text{if}(z_{i,1}^t \neq 1) \text{ then fault}; \\
&\dots \\
&z_{i,m}^t := g_{\text{lasttx}(t,i,m)}^m; \text{if}(z_{i,m}^t \neq 1) \text{ then fault};
\end{aligned}$$

- If  $H_i^t$  is a committed or aborted transaction, then

$$\begin{aligned}
CP_i^t &= w_i^t := \text{atomic} \{ y_{i,1}^t := o_{i,1}^t \cdot f_{i,1}^t(n_{i,1}^t); \text{if}(y_{i,1}^t \neq r_{i,1}^t) \text{ then fault}; \\
&\dots; \\
&y_{i,q_i}^t := o_{i,q_i}^t \cdot f_{i,q_i}^t(n_{i,q_i}^t); \text{if}(y_{i,q_i}^t \neq r_{i,q_i}^t) \text{ then fault}; \} \\
&\text{if}(w_i^t \neq c_i^t) \text{ then fault else } g_i^t := 1;
\end{aligned}$$

- If  $H_i^t$  is a live transaction, then

$$\begin{aligned}
CP_i^t &= w_i^t := \text{atomic} \{ y_{i,1}^t := o_{i,1}^t \cdot f_{i,1}^t(n_{i,1}^t); \\
&\dots; \\
&y_{i,q_i}^t := o_{i,q_i}^t \cdot f_{i,q_i}^t(n_{i,q_i}^t); \\
&\text{fault}; \}
\end{aligned}$$

- If  $H_i^t$  is a commit-pending transaction, then

$$\begin{aligned}
CP_i^t &= w_i^t := \text{atomic} \{ y_{i,1}^t := o_{i,1}^t \cdot f_{i,1}^t(n_{i,1}^t); \text{if}(y_{i,1}^t \neq r_{i,1}^t) \text{ then fault}; \\
&\dots; \\
&y_{i,q_i}^t := o_{i,q_i}^t \cdot f_{i,q_i}^t(n_{i,q_i}^t); \text{if}(y_{i,q_i}^t \neq r_{i,q_i}^t) \text{ then fault}; \}
\end{aligned}$$

Figure 5: The construction of  $CP_i^t$  and  $GP_i^t$  for Lemma 16 and Lemma 17 (case of  $t \neq t_0$ ).

Thus, the command  $GP_i^t; CP_i^t$  begins by reading the signals of the last transaction of every thread that, according to  $H$ , should end before the  $i$ -th transaction of thread  $t$  starts. It then performs an `atomic` block in which it invokes the sequence of object method invocations induced by  $H_i^t$ . After the `atomic` block ends, the command signals the end of the transaction. The desired program  $P_H$  is:

$$P_H = C_H^1 \parallel \dots \parallel C_H^m.$$

We now construct a particular trace  $\tau$  of  $P_H$ . We build  $\tau$  by first constructing a trace  $\tau^t$  for every sequential command  $C_H^t$  and then interleaving the traces  $\tau^1, \dots, \tau^m$  in a particular way. Consider the set of traces  $A(C_H^t)t$  of the sequential command  $C_H^t$ , and let  $\tau^t \in A(C_H^t)t$  be the maximal trace without any `fault` actions such that  $H|_t = \text{history}(\tau^t)$ . The trace  $\tau^t$  exists, since by construction of  $C_H^t$  and the definition of the trace set of a sequential command (Figure 4), there is a trace in  $A(C_H^t)t$  for every possible parameter and return value of object method invocations and `atomic` blocks in  $C_H^t$ ; in particular,  $A(C_H^t)t$  contains a trace where the parameters and return values of object method invocations and the return values of transactions are as in  $H|_t$ . We now partition every  $\tau^t$  into  $|H|_t|$  subsequences that we later interleave to create  $\tau$ :

$$\tau^t = \tau_1^t \dots \tau_{|H|_t}^t.$$

Formally, for every  $i = 1..|H|_t|$  there is exactly one TM interface action  $\psi_i^t$  in  $\tau_i^t$  and the conditions in Figure 6 hold. This defines  $\tau_i^t$  uniquely and ensures that, if  $\tau_i^t$  ends with  $\psi_i^t = (\_, t, \text{txbegin})$ , then it contains all the actions that are used to read the global signaling variables that precede  $\psi_i^t$  in  $\tau^t$ . The desired trace is constructed by interleaving the subsequences of the traces  $\tau^1, \dots, \tau^m$  according to the order induced by  $H$ . Formally,

$$\tau = \tau_{j_1}^{t_1} \dots \tau_{j_{|H|}}^{t_{|H|}},$$

$$\tau_i^t = \begin{cases} \neg\psi_i^t, & \psi_i^t \in \{(\_, t, \text{txbegin}), (\_, t, \text{call } o.f(n)), (\_, t, \text{txcommit})\}; \\ \psi_i^t, & \psi_i^t = (\_, t, \text{OK}); \\ \psi_i^t(\_, t, y_{i,-}^t := n), & \psi_i^t = (\_, t, \text{ret}(n) \text{ o.f}) \wedge \text{txof}(\psi_i^t, \tau^t) \in \text{live}(\tau^t); \\ \psi_i^t(\_, t, y_{i,-}^t := n) (\_, t, \text{assume}(y_{i,-}^t = r_{i,-}^t)), & \psi_i^t = (\_, t, \text{ret}(n) \text{ o.f}) \wedge \text{txof}(\psi_i^t, \tau^t) \notin \text{live}(\tau^t); \\ \psi_i^t(\_, t, w_i := c_i^t) (\_, t, \text{assume}(w_i = c_i^t)) (\_, t, g_i^t := 1), & \psi_i^t \in \{(\_, t, \text{committed}), (\_, t, \text{aborted})\}. \end{cases}$$

Figure 6: The construction of  $\tau_i^t$  for Lemma 16 and Lemma 17.

where  $H(i) = (\_, t_i, \_)$  and  $j_i = |(H|_i)|_{t_i}|$ . Note that by construction  $\text{history}(\tau) = H$ .

Since  $\tau^t \in A(C_H^t)t$ , we have  $\tau \in A(P_H)$ . Let  $s$  be the state where all the local variables are set to  $u$  and for all  $t$ ,  $g_i^t = 0$  for  $i \neq 0$  and  $g_0^t = 1$ . By the construction of  $\tau$ , we have  $\text{eval}(s, \tau) \neq \emptyset$ . Then, since  $\text{history}(\tau) = H \in \mathcal{T}_C$ , we have  $\tau \in \llbracket P_H \rrbracket(s, \mathcal{T}_C)$ . Since  $\tau \in \llbracket P_H \rrbracket(s, \mathcal{T}_C)$  and  $\mathcal{T}_C \preceq \mathcal{T}_A$ , by Definition 7 there exists a trace  $\tau' \in \llbracket P_H \rrbracket(s, \mathcal{T}_A)$  such that  $\tau' \in A(P_H)$ ,  $\tau'|_{\neg\text{trans}} = \tau|_{\neg\text{trans}}$  and  $S' = \text{history}(\tau') \in \mathcal{T}_A$ .

We now analyze the relationship between a transaction in  $\tau$  and the one in  $\tau'$  arising from the execution of the same commands. The construction in Figure 5 ensures that:

- If a transaction in  $\tau$  is completed, then so is the corresponding transaction in  $\tau'$ , since every completed transaction in  $\tau$  is followed by a non-transactional action assigning to  $g_i^t$  and  $\tau'|_{\neg\text{trans}} = \tau|_{\neg\text{trans}}$ . Furthermore, a transaction in  $\tau$  is committed if and only if so is the corresponding transaction in  $\tau'$ , and in this case the return values for transactional actions inside the transactions match.
- If a transaction in  $\tau$  is live, then the corresponding transaction in  $\tau'$  is live, aborted or does not appear in  $\tau'$  at all.
- If a transaction in  $\tau$  is commit-pending, then the corresponding transaction in  $\tau'$  can have any status or may not appear in  $\tau'$  at all. However, if the transaction in  $\tau'$  is visible (in particular, commit-pending or committed), then the return values for transactional actions inside the transactions match.

Let  $H_1 = (H|_{\neg\text{live}})|_{\neg\text{abortedtx}}$  and  $S^r = (S'|_{\neg\text{live}})|_{\neg\text{abortedtx}}$ . Since  $S' \in \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under removing live and aborted transactions,  $S^r \in \mathcal{T}_A$ . Let  $p'_t$ , respectively,  $p''_t$  be the index of last txcommit or committed action in  $H_1|_t$ , respectively,  $S^r|_t$ ; if there is no such action, the corresponding index is 0. Let  $p_t = \min(p'_t, p''_t)$ . From the above analysis it follows that  $(H_1|_t)|_{p_t} \equiv (S^r|_t)|_{p_t}$ .

Let  $S_1$  be a history obtained from  $S^r$  by renaming the action identifiers such that  $S_1 \equiv S^r$ ,  $(H_1|_t)|_{p_t} = (S_1|_t)|_{p_t}$  and all actions in  $S_1$  that do not belong to  $(S^r|_t)|_{p_t}$  for some  $t$  have identifiers that do not appear in  $H$ . Since  $S^r \in \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under renaming action identifiers, we get  $S_1 \in \mathcal{T}_A$ .

The history  $S_1$  does not contain live transactions so that  $\text{comp}(S_1) \neq \emptyset$ . Then, since  $S_1 \in \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under commit-pending transactions completion, we get that there exists a history  $S_2 \in \text{comp}(S_1) \cap \mathcal{T}_A$ . Then  $S_2 = S_1 S'_1$  for some history  $S'_1$  that contains only committed and aborted actions; without loss of generality we can assume that identifiers of actions in  $S'_1$  do not appear in  $H$ . Let  $S''_1$  be the subsequence of committed actions in  $S_1$  that are not in  $H_1$ . The history  $H_1 S''_1 S'_1$  contains only visible transactions. We construct the desired history  $H^c$  by aborting all the commit-pending transactions in  $H_1 S''_1 S'_1$ . Let  $H^c = H_1 S''_1 S'_1 H_a$ , where  $H_a$  consists of actions  $(\_, t, \text{aborted})$  for every thread  $t$  ending with a commit-pending transaction in  $H_1 S''_1 S'_1$ ; these actions have unique identifiers that do not appear in  $H$ . We have that  $H^c \in \text{comp}(H|_{\neg\text{live}})$ , for the following reasons:

- $S''_1$  completes the commit-pending transactions in  $H$  that get committed in  $S'$ .
- $S'_1$  completes the commit-pending transactions in  $H$  that stay commit-pending in  $S'$ .
- $H_a$  aborts the commit-pending transactions in  $H$  that become live or aborted in  $S'$  or do not appear there at all.

Let  $S = S_2|_{\text{com}}$ . Since  $S_2 \in \mathcal{T}_A$  has only completed transactions and  $\mathcal{T}_A$  is closed under removing live and aborted transactions, we get that  $S = S_2|_{\text{com}} \in \mathcal{T}_A$ . We now show  $H^c|_{\text{com}} \sqsubseteq_{\text{op}} S$ , thus completing the proof. We first show that  $\forall t \in \text{ThreadID}. (H^c|_t)|_{\text{com}} = S|_t$  by considering several cases.

- $H|_t$  does not contain commit-pending transactions. Then  $H_1|_t = S_1|_t$ ,  $H_1|_t$  contains only committed transactions and  $S'_1 = S''_1 = H_a = \varepsilon$ . Hence,

$$(H^c|_t)|_{\text{com}} = (H_1|_t)|_{\text{com}} = H_1|_t = S_1|_t = (S_1|_t)|_{\text{com}} = (S_2|_t)|_{\text{com}} = S|_t.$$

- $H|_t$  ends with a commit-pending transaction and  $(S''_1 S'_1 H_a)|_t = S''_1|_t = (\_, t, \text{committed})$ . Then  $H_1|_t = (H_1|_t)|_{p_t}$  and  $S_1|_t = ((S_1|_t)|_{p_t})(S''_1|_t)$ . Hence,

$$H^c|_t = (H_1|_t)(S''_1|_t) = ((H_1|_t)|_{p_t})(S''_1|_t) = ((S_1|_t)|_{p_t})(S''_1|_t) = S_1|_t = S_2|_t.$$

Then  $(H^c|_t)|_{\text{com}} = (S_2|_t)|_{\text{com}} = S|_t$ .

- $H|_t$  ends with a commit-pending transaction and  $(S'_1 S'_1 H_a)|_t = S'_1|_t$ . Then  $H_1|_t = S_1|_t$ . Hence,

$$H^c|_t = (H_1|_t)(S'_1|_t) = (S_1|_t)(S'_1|_t) = S_2|_t.$$

Then  $(H^c|_t)|_{\text{com}} = (S_2|_t)|_{\text{com}} = S|_t$ .

- $H|_t$  ends with a commit-pending transaction and  $(S''_1 S'_1 H_a)|_t = H_a|_t = (\_, t, \text{aborted})$ . Then  $((H_1|_t)(H_a|_t))|_{\text{com}} = ((H_1|_t)|_{p_t})|_{\text{com}}$  and  $S_1|_t = (S_1|_t)|_{p_t}$ . Hence,

$$(H^c|_t)|_{\text{com}} = ((H_1|_t)(H_a|_t))|_{\text{com}} = ((H_1|_t)|_{p_t})|_{\text{com}} = ((S_1|_t)|_{p_t})|_{\text{com}} = S_1|_t = S_2|_t.$$

This shows  $\forall t \in \text{ThreadID}. (H^c|_t)|_{\text{com}} = S|_t$ . Finally, for every completed transaction  $T$  in  $\tau$ , the value of  $g_t^t$  is set to 1 after the transaction completes and is read before every transaction  $T'$  that begins after the completion of  $T$ . Hence, the transaction in  $\tau'$  corresponding to  $T$  completes before the transaction corresponding to  $T'$  begins. Then the real-time order in  $H^c$  is preserved in  $S$ , which implies  $H^c|_{\text{com}} \sqsubseteq_{\text{op}} S$ . ■

**LEMMA 17.** *Let  $\mathcal{T}_C$  and  $\mathcal{T}_A$  be transactional memories such that  $\mathcal{T}_C \preceq \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under live transactions omission and pending transactions completion. Let  $H = H'\psi \in \mathcal{T}_C$ , where  $\psi$  is a response action and  $\psi \notin \{(\_, \_, \text{committed}), (\_, \_, \text{aborted})\}$ . There exists a history  $S \in \mathcal{T}_A$  such that  $H^c \sqsubseteq_{\text{op}} S$ , where  $H^c \in \text{cTMSpact}(H)$ .*

**Proof:** Let  $\psi$  be an action of thread  $t_0$ . As in Lemma 16, we build a program  $P_H$  for  $H$  that records the interaction of every thread  $t$  with the transactional memory and monitors whether the real-time order includes that in  $H$ . The sequential commands  $C_H^i$  constructed for all the threads except  $t_0$  are the same as in Lemma 16, but  $C_H^0$ , the sequential command built for  $t_0$ , is different.

As in Lemma 16, the command  $C_H^0$  records the arguments passed to method calls and the return values during a transaction the program using local variables that are never rewritten. However,  $C_H^0$  observes these values *after* the transaction completes using conditional command. This makes the sequence of actions inside a committed transaction observable to thread  $t_0$ . As a result, if  $P_H$  is evaluated with a history  $H'$  which has a different per-thread sequence of actions inside committed transactions than that of  $H$ , then the trace would reflect the difference by either a `fault` by some thread other than  $t_0$  or by the non-transactional actions of  $t_0$ . Histories  $H$  and  $H'$  might perform different actions inside aborted transactions, but this is fine for our purposes because the actions inside aborted transactions are not observable. Monitoring the real-time order is done using the same monitoring mechanism as in Lemma 16.

The challenge here is to construct a program that can observe the sequence of actions inside the transaction of  $\psi$ , which is live. The only one way to get this observation is to use faults. Hence, given an incomplete history  $H$ , we construct a program which faults *if* the local observations of  $t_0$  are the same as in  $H$ .

Now we formalize the intuitive proof sketch. To make the proof self-contained we repeat some parts of the constructions that were already discussed in Lemma 16. Let us choose an integer value  $u \neq 1$  which does not appear in  $H$ . We use the following shorthands:

- We denote by  $m$  the number of threads that have an action in  $H$ , and for simplicity, let  $m$  be the largest thread identifier occurring in  $H$ .

$$\begin{aligned}
GP_i^{t_0} &= z_{i,1}^{t_0} := g_{\text{lasttx}(t_0,i,1)}^1; \text{if}(z_{i,1}^{t_0} \neq 1) \text{ then mismatch} := 1; \\
&\dots \\
z_{i,m}^{t_0} &:= g_{\text{lasttx}(t_0,i,m)}^m; \text{if}(z_{i,m}^{t_0} \neq 1) \text{ then mismatch} := 1;
\end{aligned}$$

- For the live transaction of  $\psi$ ,  $CP_{k^{t_0}}^{t_0}$  is constructed as follows:

$$\begin{aligned}
CP_{k^{t_0}}^{t_0} &= w_{k^{t_0}}^{t_0} := \text{atomic} \{ y_{k^{t_0},1}^{t_0} := o_{k^{t_0},1}^{t_0} \cdot f_{k^{t_0},1}^{t_0}(n_{k^{t_0},1}^{t_0}); \\
&\quad \text{if}(y_{k^{t_0},1}^{t_0} \neq r_{k^{t_0},1}^{t_0}) \text{ then mismatch} := 1; \\
&\quad \dots; \\
&\quad y_{k^{t_0},q_{k^{t_0}}^{t_0}}^{t_0} := o_{k^{t_0},q_{k^{t_0}}^{t_0}}^{t_0} \cdot f_{k^{t_0},q_{k^{t_0}}^{t_0}}^{t_0}(n_{k^{t_0},q_{k^{t_0}}^{t_0}}^{t_0}); \\
&\quad \text{if}(y_{k^{t_0},q_{k^{t_0}}^{t_0}}^{t_0} \neq r_{k^{t_0},q_{k^{t_0}}^{t_0}}^{t_0}) \text{ then mismatch} := 1; \\
&\quad \text{if}(\text{mismatch} \neq 1) \text{ then fault}; \}
\end{aligned}$$

- For an aborted or a committed transaction by  $t_0$ ,  $CP_i^{t_0}$  is constructed as follows:

$$\begin{aligned}
CP_i^{t_0} &= w_i^{t_0} := \text{atomic} \{ y_{i,1}^{t_0} := o_{i,1}^{t_0} \cdot f_{i,1}^{t_0}(n_{i,1}^{t_0}); \text{if}(y_{i,1}^{t_0} \neq r_{i,1}^{t_0}) \text{ then mismatch} := 1; \\
&\quad \dots; \\
&\quad y_{i,q_i^{t_0}}^{t_0} := o_{i,q_i^{t_0}}^{t_0} \cdot f_{i,q_i^{t_0}}^{t_0}(n_{i,q_i^{t_0}}^{t_0}); \text{if}(y_{i,q_i^{t_0}}^{t_0} \neq r_{i,q_i^{t_0}}^{t_0}) \text{ then mismatch} := 1; \} \\
&\text{if}(w_i^{t_0} \neq c_i^{t_0}) \text{ then mismatch} := 1 \text{ else } g_i^{t_0} := 1;
\end{aligned}$$

Figure 7: The construction of  $CP_i^{t_0}$  and  $GP_i^{t_0}$  for Lemma 17.

- We denote by  $k^t$  the number of transactions started by thread  $t$  in  $H$ , i.e., the number of  $(\_, t, \text{txbegin})$  actions in  $H$ .
- We partition  $H|_t$  into  $k^t$  subsequences:  $H|_t = H_1^t \dots H_{k^t}^t$ , where  $H_i^t$  is comprised of the actions in the  $i$ -th transaction of  $t$ . Specifically,  $H_i^t(1) = (\_, t, \text{txbegin})$ .
- We let  $c_i^t$  be the outcome of the  $i$ -th transaction of thread  $t$ , i.e.,  $c_i^t = \text{committed}$  or  $c_i^t = \text{aborted}$ . If the transaction is not completed,  $c_i^t$  is undefined.
- We denote by  $q_i^t$  the number of call actions of thread  $t$  in its  $i$ -th transaction, i.e., in  $H_i^t$ .
- We let  $(\_, t, \text{call } o_{i,j}^t \cdot f_{i,j}^t(n_{i,j}^t))$  be the  $j$ -th call action of thread  $t$  in its  $i$ -th transaction.
- We let  $(\_, t, \text{ret}(r_{i,j}^t) \cdot o_{i,j}^t \cdot f_{i,j}^t)$  be the  $j$ -th ret action of thread  $t$  in its  $i$ -th transaction.
- If the response to  $(\_, t, \text{call } o_{i,j}^t \cdot f_{i,j}^t(n_{i,j}^t))$  is an aborted action, we let  $r_{i,j}^t$  be aborted. If there is no response to  $(\_, t, \text{call } o_{i,j}^t \cdot f_{i,j}^t(n_{i,j}^t))$ , i.e. the transaction is live and  $(\_, t, \text{call } o_{i,j}^t \cdot f_{i,j}^t(n_{i,j}^t))$  is the last action in the transaction, then  $r_{i,j}^t$  is set to  $u$ .
- We denote by  $\text{lasttx}(t, i, t')$  the number of transactions of thread  $t'$  in  $H$  that either committed or aborted before the  $i$ -th transaction of thread  $t$  started, i.e., the number of  $(\_, t', \text{committed})$  and  $(\_, t', \text{aborted})$  actions preceding the  $i$ -th  $(\_, t, \text{txbegin})$  action in  $H$ .

For every thread  $t = 1..m$  we construct a straight-line command

$$C_H^t = GP_1^t; CP_1^t; GP_2^t; CP_2^t; \dots; GP_{k^t}^t; CP_{k^t}^t.$$

$CP_i^t$  is a sequence of commands constructed according to the  $i$ -th transaction of thread  $t$  in  $H$ . The construction of  $CP_i^{t_0}$  and  $GP_i^{t_0}$  is shown in Figure 7, and the construction of  $CP_i^t$  and  $GP_i^t$ , when  $t \neq t_0$ , is shown in Figure 5.

Here the  $g$  variables are global and all others are local. The  $g$  and  $z$  variables are used to monitor the real-time order. The variable  $g_i^t$  is written only by thread  $t$  and is used to signal that the  $i$ -th transaction of thread  $t$  ended:  $g_i^t$  is set to 1 only after  $t$ 's  $i$ -th transaction either committed or aborted. The variable  $z_{i,t'}^t$  is local to thread  $t$ . It is used to record if the  $\text{lasttx}(t, i, t')$ -th transaction of thread  $t'$  signaled that it had ended before the  $i$ -th transaction of

thread  $t$  started. As  $\text{lasttx}(t, i, t')$  might be 0, we add a dummy variable  $g_0^t$  for every thread  $t$ . Later in the proof we execute the program from a state in which  $g_0^t$  is initialised to 1. The variable  $w_i^t$  records whether the  $i$ -th transaction of thread  $t$  committed or aborted. The variables  $y_{i,j}^t$  record the return value of the  $j$ -th object method invocation in the  $i$ -th transaction of thread  $t$ . Variable `mismatch` is a local variable of thread  $t_0$ . It is used to check that the return values for object methods by thread  $t_0$  are as expected: it is set to 1 if there is a mismatch. It will help in producing a different trace if there is a mismatch.

Thus, the command  $C_H^t$  begins by reading the signals of the last transaction of every thread that, according to  $H$ , should end before the  $i$ -th transaction of thread  $t$  starts. It then performs an atomic block in which it invokes the sequence of object method invocations induced by  $H_i^t$ . After the atomic block ends,  $CP_i^t$  signals the end of the transaction. The desired program  $P_H$  is the parallel composition of the straight-line commands of the different threads:

$$P_H = C_H^1 \parallel \dots \parallel C_H^m.$$

We now construct a particular trace  $\tau$  of  $P_H$ . We build  $\tau$  by first constructing a trace  $\tau^t$  for every sequential command  $C_H^t$  and then interleaving the traces  $\tau^1, \dots, \tau^m$  in a particular way.

Consider the set of traces  $A(C_H^t)t$  (cf. Figure 4) of the sequential command  $C_H^t$ , and let  $\tau^t \in A(C_H^t)t$  be the maximal trace such that  $H|_t = \text{history}(\tau^t)$ . The trace  $\tau^t$  exists, since by construction of  $C_H^t$  and the definition of the trace set of a sequential command, there is a trace in  $A(C_H^t)t$  for every possible parameter and return value of object method invocations and atomic blocks in  $C_H^t$ ; in particular,  $A(C_H^t)t$  contains a trace where the parameters and return values of object method invocations and the return values of transactions are as in  $H|_t$ .

We now partition every  $\tau^t$  into  $|H|_t|$  subsequences that we later interleave to create  $\tau$ :

$$\tau^t = \tau_1^t \dots \tau_{|H|_t}^t.$$

Formally, for every  $i = 1..|H|_t|$  there is exactly one TM interface action  $\psi_i^t$  in  $\tau_i^t$ . The construction of  $\tau_i^t$  is shown in Figure 6. Note that this defines  $\tau_1^t \dots \tau_{|H|_t}^t$  uniquely. The above definition also ensures that, if  $\tau_i^t$  ends with  $\psi_i^t = (\_, t, \text{txbegin})$ , then it contains all the actions that are used to read the global signaling variables that precede  $\psi_i^t$  in  $\tau^t$ .

The desired trace is constructed by interleaving the subsequences of the traces  $\tau^1, \dots, \tau^m$  according to the order induced by  $H$ . Formally,

$$\tau = \tau_{j_1}^{t_1} \dots \tau_{j_{|H|}}^{t_{|H|}},$$

where  $H(i) = (\_, t_i, \_)$  and  $j_i = |H|_i|_{t_i}|$ . Note that by construction  $\text{history}(\tau) = H$ .

This ends the construction, which is similar to Lemma 16. The rest of the proof diverges from the proof of Lemma 16.

Since  $\tau^t \in A(C_H^t)t$ , we have  $\tau \in A(P_H)$ . Since  $\text{history}(\tau) = H \in \mathcal{T}_C$ , we have  $\tau \in \llbracket P_H \rrbracket(s, \mathcal{T}_C)$ .

Since  $\tau \in \llbracket P_H \rrbracket(s, \mathcal{T}_C)$  has a `fault` after the action  $\psi$  and  $\mathcal{T}_C \preceq \mathcal{T}_A$ , by Definition 7 there exists a trace  $\tau' \in \llbracket P_H \rrbracket(s, \mathcal{T}_A)$  that ends with a fault by thread  $t_0$ . Hence,  $\tau' \in A(P_H)$  and  $S' = \text{history}(\tau')$  is in  $\mathcal{T}_A$ . Since  $\tau$  and  $\tau'$  have only one fault, the global variables in  $\tau$  and  $\tau'$  have same values, and thus the real-time order of actions in  $H$  is preserved in  $S'$ . Also, the actions inside committed transactions in  $S'$  are same as in  $H$  and all the committed transactions in  $\tau$  are committed in  $\tau'$  since we check the local values  $w_i^t$  outside the atomic block.

Now we create a history  $H^c$  using the history  $S'$ . Let  $H_1 = H|_{\vee \cup \psi}$  and  $S^r = S'|_{\vee \cup \psi}$ , respectively, where  $\cdot|_{\vee \cup \psi}$  is the projection to actions by all the visible transactions and the transaction of  $\psi$ .

Since the return values of object invocations are checked inside the completed transactions in  $H_1$  and that of  $\psi$ , all the return values inside completed transactions that completed before `txbegin` of `txof`( $\psi, S^r$ ) in  $S^r$  and the transaction of  $\psi$  are same as in  $H_1$ ; otherwise, the trace  $\tau$  would have a `fault` by a thread  $t \neq t_0$ , or would not have a `fault` after  $\psi$  by thread  $t_0$ .

Thus, for thread  $t_0$ ,  $(S^r|_{\psi})|_{t_0} \equiv (H_1|_{\psi})|_{t_0}$  and for every thread  $t \neq t_0$ , if there is a `txcommit` action in  $S^r|_t$ , then  $(S^r|_{k_t})|_t \equiv (H_1|_{k_t})|_t$ , where  $k_t$  is the index of the  $k^{\text{th}}$  `txcommit` action,  $\psi_t$ , in  $H_1|_t$  and  $k$  is the number of `txcommit` actions in  $S^r|_t$ . Let  $S_1$  be a history with unique action identifiers obtained from  $S^r$  by renaming the

action identifiers in  $S^r$  such that  $S_1 \equiv S^r$  and  $\forall i = 1 \dots |(H_1 \downarrow_{\psi})|_{t_0}|, (S_1|_{t_0})(i) = (H_1|_{t_0})(i)$ , and for every thread  $t \neq t_0$ , if there is a txcommit action in  $S^r$ , then  $\forall i = 1 \dots k_t, (S_1|_t)(i) = (H_1|_t)(i)$  and all actions in  $S_1$  that are not in  $H_1$  do not have action identifier same as any action in  $H_1$ . Note that if there is not txcommit action in  $S^r|_t$ , then  $S_1|_t = S^r|_t = \varepsilon$ . Note that since  $S_1|_t$  contains only visible transactions, either  $S_1|_t = \varepsilon$  or  $S_1|_t = (S_1|_t) \downarrow_{\psi_t}$  or  $S_1|_t = (S_1|_t) \downarrow_{\psi_t} \psi'_t$ , where  $\psi'_t = (\_, t, \text{committed})$  or  $(\_, t, \text{aborted})$ .

Let  $H_p$  be the subsequence of  $H_1$  such that for thread  $t_0$ ,  $H_p|_{t_0} = H_1|_{t_0}$  and for all other threads, if  $S^r|_t$  does not contains a txcommit action then  $H_p|_t = \varepsilon$ , if  $S^r|_t$  contains a txcommit action then  $H_p|_t = H_1|_t$  if  $\psi_t$  is the last action in  $H_1|_t$ ; otherwise  $H_p|_t = (H_1|_t) \downarrow_{\psi_t} \psi''_t$ , where  $\psi''_t$  is the response of  $\psi_t$  in  $H_1$ . Let  $H'_1$  be a subsequence of  $H_1$  such that for thread  $t_0$ ,  $H'_1|_{t_0} = H_1|_{t_0}$  and for all other threads, if  $S^r|_t$  does not contains a txcommit action then  $H'_1|_t = \varepsilon$ , otherwise  $H'_1|_t = (H_1|_t) \downarrow_{\psi_t}$ . Note that for every thread  $t \neq t_0$ , either  $S_1|_t$  does not contain a txcommit action and  $H'_1|_t = H_p|_t = \varepsilon$ , or  $S_1|_t$  contains a txcommit action and  $H'_1|_t = (H_p|_t) \downarrow_{\psi_t} = (H_1|_t) \downarrow_{\psi_t}$ , where  $\psi_t$  is the last txcommit action in  $S_1|_t$ .

Note that a transaction in  $\tau$  that does not complete before txbegin of txof( $\psi, \tau$ ) may have different response in  $\tau'$ . Also note that a transaction that is in  $H_1$  but not in  $H'_1$  cannot complete before the txbegin of txof( $\psi, H_1$ ). The reason for this is that for every completed transaction, either committed or aborted, that completed in  $H_1$  before txbegin of txof( $\psi, H_1$ ), there is a completed transaction, with the same return value, in  $S_1$ . This is because for every committed and aborted transaction in  $\tau$  that completed before txbegin of txof( $\psi, \tau$ ),  $\tau'$  contains assignment to the variable  $g$  following the completion of the transaction, and the txof( $\psi, H_1$ ) has checks in  $\tau$  before it began on the number of transactions completed by other threads. However, for a transaction in  $\tau$  that does not complete before txbegin of txof( $\psi, \tau$ ) may have only a prefix or no action executed in  $\tau'$ , and if it is completed in both  $\tau$  and  $\tau'$ , it may have different response. Note that for any transaction in  $H'_1$ , all its preceding transactions are in  $S_1$  with the same response actions, for the same reason.

Since  $S' \in \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under removing live and aborted transactions,  $S^r \in \mathcal{T}_A$ . Since  $\mathcal{T}_A$  is closed under renaming action identifiers, we get  $S_1 \in \mathcal{T}_A$ . Since  $S_1 \in \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under commit-pending transactions completion, there exists a history  $S_2 \in \text{comp}(S_1) \cap \mathcal{T}_A$ . Let  $S_2 = S_1 S'_1 \in \mathcal{T}_A$  be a history such that every action in  $S'_1$  is either committed or aborted and has unique action identifiers with no action identifier that appear in  $S_1$  or  $H_1$ , and every transaction in  $S_2$ , except that of  $\psi$ , is complete. Let  $S = S_2|_{c \cup \psi}$ , where  $\cdot|_{c \cup \psi}$  is the projection to actions by all the committed transactions and the transaction of  $\psi$ . Let  $S''_1$  be the subsequence of committed and aborted actions in  $S_1$  which are not in  $H'_1$ . There are committed and aborted actions in  $S_1$  which are not in  $H'_1$  because some of the commit-pending transactions in  $\tau$  are executed to completion in  $\tau'$  after txbegin of txof( $\psi, \tau$ ).

Let  $H^c = (H'_1 S''_1 S'_1)|_{c \cup \psi}$ . Note that all the transactions in  $H^c$  and  $S$  are completed, except for the transaction of  $\psi$ . The reason for this is for every thread  $t \neq t_0$ , the last transaction in  $H'_1|_t$  is a completed transaction or a commit-pending transaction, which gets completed either by  $S''_1$  if it is completed in  $S_1$ , or it gets completed by  $S'_1$  if it is commit-pending in  $S_1$ . Since  $S_2 \in \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under removing live and aborted transactions, in particular removing aborted transactions, we get that  $S = S_2|_{c \cup \psi} \in \mathcal{T}_A$ . Now we show that  $H^c \in \text{cTMSpact}(H)$  (Claim 17.1) and then show Claim 17.2 and Claim 17.3 to prove that  $H^c \sqsubseteq_{\text{op}} S$ .

CLAIM 17.1.  $H^c \in \text{cTMSpact}(H)$ .

**Proof:** Let  $H''_1 \psi$  be the subsequence of  $H_p$  obtained by removing some transactions in the following way:

- a committed transaction  $T_c$  is removed if  $T_c$  is not in  $H^c$ .
- an aborted transaction  $T_a$  is removed if there is no txbegin action, by any thread in  $H''_1 \psi$ , after  $T_a$  completes.

Consider the following about the history  $H''_1 \psi$ :

- (i)  $H''_1 \psi$  is a subsequence of  $H_p$ , which is a subsequence of  $H$ .
- (ii)  $H''_1$  contains some of the visible transactions in  $H$  and also includes txof( $\psi, H$ ). This is because  $H_1$  contains all the visible transactions in  $H$  and txof( $\psi, H$ ).
- (iii) Since for any transaction in  $H_p$ , all its preceding transactions are in  $S_1$  with the same response actions of preceding transactions. (The response of the last transaction by a thread  $t \neq t_0$  in  $H_p$  might have different response in  $S_1$ ). Thus, we get that for a transaction in  $H_p$  (and in  $H$ ), if it is included in  $H_p$  then all of its preceding transactions in  $H_p$  (and in  $H$ ) are also included. Since all the committed transactions in  $H_p$ , which

have a transaction after them in real-time order in  $H_1''\psi$ , are included in  $H_1''\psi$ , it follows that

$$\forall T \in \text{tx}(H_1''\psi). \forall T'. T' \prec_{H_1''\psi} T \iff (T' \prec_H T \wedge T' \in \text{committed}(H)).$$

Therefore,  $H_1''\psi \in \text{TMSpast}(H)$ , by Definition 3.

Note that for every aborted transaction by a thread  $t \neq t_0$ , that is last in  $H_p$  and included in  $H_1''\psi$ , the response action is removed in  $H_1'$  and replaced by a committed or aborted action in  $S_1''$ . Therefore,  $(H_1'S_1''S_1')|_{\text{c}\cup\psi} \in \text{comcomp}(\text{com}(H_1''\psi))$ , and hence  $H^c \in \text{comcomp}(\text{com}(H_1''\psi))$ .

Since  $H_1''\psi \in \text{TMSpast}(H)$  and  $H^c \in \text{comcomp}(\text{com}(H_1''\psi))$ , Definition 4 implies that  $H^c \in \text{cTMSpast}(H)$ . ■

**CLAIM 17.2.** *For every thread  $t \neq t_0$ , if there is a txcommit action in  $S_1$ , then  $(S_1 \downarrow_{\psi_t})|_t = (H_1' \downarrow_{\psi_t})|_t$ , where  $\psi_t$  is the last txcommit action in  $S_1|_t$ .*

**Proof:** Since  $(S^r \downarrow_{k_t})|_t \equiv (H_1 \downarrow_{k_t})|_t$ , where  $H_1|_t(k_t) = \psi_t$ , and by the construction of  $S_1$ , for every thread  $t \neq t_0$  it holds that  $\forall i = 1 \dots k_t$ ,  $(S_1|_t)(i) = (H_1|_t)(i)$ , we get that  $(S_1 \downarrow_{\psi_t})|_t = (H_1 \downarrow_{\psi_t})|_t$ . By construction of  $H_1'$ ,  $H_1'|_t \downarrow_{\psi_t} = H_p|_t \downarrow_{\psi_t} = H_1|_t \downarrow_{\psi_t}$ . Therefore,  $(S_1 \downarrow_{\psi_t})|_t = (H_1' \downarrow_{\psi_t})|_t$ . ■

**CLAIM 17.3.**  *$H^c|_t = S|_t$ , for every thread  $t \neq t_0$ .*

**Proof:** If  $S_1|_t$  does not contain a txcommit, then  $S_1|_t = H_1'|_t = \varepsilon$ . Also,  $H^c|_t = \varepsilon$  and  $S|_t = \varepsilon$ , and the claim holds vacuously.

If there is a txcommit action in  $S_1|_t$ , Claim 17.2 implies that  $(S_1 \downarrow_{\psi_t})|_t = (H_1' \downarrow_{\psi_t})|_t$ , where  $\psi_t$  is the last txcommit action in  $S_1|_t$ . Let  $S_1|_t = (S_1 \downarrow_{\psi_t})|_t S_t$ , for some history  $S_t$ . Note that either  $S_t = \varepsilon$  or  $S_t = (\_, t, \text{committed})$  or  $(\_, t, \text{aborted})$ .

If  $S_t = (\_, t, \text{committed})$  or  $(\_, t, \text{aborted})$ , then  $S_1|_t$  is a completed history and hence,  $S_2|_t = S_1|_t$  and  $S_1'|_t = \varepsilon$ . Since  $S_1''$  contains all committed and aborted transactions in  $S_1$  which are not in  $H_1'$ , it follows that  $S_1''|_t = S_t$ . Hence,  $S_2|_t = S_1|_t = (H_1'S_1'')|_t$ , and

$$H^c|_t = ((H_1'S_1''S_1')|_{\text{c}\cup\psi})|_t = ((H_1'S_1'')|_{\text{c}\cup\psi})|_t = (S_1|_{\text{c}\cup\psi})|_t = (S_2|_{\text{c}\cup\psi})|_t = S|_t$$

If  $S_t = \varepsilon$ , then the last transaction in  $S_1|_t$  is either commit-pending or completed in both  $H_1'|_t$  and  $S_1|_t$ , and hence,  $S_2|_t = (S_1S_1')|_t$  and  $S_1''|_t = \varepsilon$ . Since  $(S_1 \downarrow_{\psi_t})|_t = (H_1' \downarrow_{\psi_t})|_t$ , it follows that  $(S_1S_1')|_t = (H_1'S_1')|_t$ , and

$$H^c|_t = ((H_1'S_1''S_1')|_{\text{c}\cup\psi})|_t = ((H_1'S_1')|_{\text{c}\cup\psi})|_t = ((S_1S_1')|_{\text{c}\cup\psi})|_t = (S_2|_{\text{c}\cup\psi})|_t = S|_t$$

Since thread  $t_0$  faults only if all the checks that ensures it succeeded, we get that  $H_1'|_{t_0} = S_1|_{t_0}$ . Therefore,

$$H^c|_{t_0} = (H_1'S_1''S_1')|_{t_0} = H_1'|_{t_0} = S_1|_{t_0} = S_2|_{t_0} = S|_{t_0}.$$

Claim 17.3 implies that  $H^c|_t = S|_t$ , for all other threads  $t \neq t_0$ .

Since for every complete transaction  $T$  in  $\tau$  the value of  $g_i^t$  is set to 1 after the transaction completes and is read by every transaction  $T'$  that begin after the completion of  $T$ , we get that  $T$  completes in  $\tau'$  before the transaction  $T'$  begins; otherwise, it would generate a `fault` by a thread  $t \neq t_0$ , and will not generate a `fault` by  $t_0$ . Thus, the real-time order in  $H^c$  is preserved in  $S$ , implying that  $H^c \sqsubseteq_{\text{op}} S$  and completing the proof. ■

**LEMMA 18.** *Let  $\mathcal{T}_C$  and  $\mathcal{T}_A$  be transactional memories such that  $\mathcal{T}_C \preceq \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under live transactions omission, pending transactions completion and aborting live transactions. Let  $H = H'\psi \in \mathcal{T}_C$ , where  $\psi = (t_0, \_, \text{aborted})$ . There exists a history  $S \in \mathcal{T}_A$  such that  $H^c \sqsubseteq_{\text{op}} S$ , where  $H^c \in \text{cTMSpast}(H)$ .*



**Proof:** Let  $H_0 = H|_{\psi_r}$ , where  $\psi_r$  is the last response action by  $\text{txof}(\psi, H)$ . Let  $H_1 = H_0\psi_c\psi$ , where  $\psi_c$  is the last request action in  $H$  by  $\text{txof}(\psi, H)$ . Note that  $\psi_c$  may not necessarily be a  $\text{txcommit}$  action. It can be any arbitrary request action that got aborted. Since every prefix of  $H$  is in  $\mathcal{T}_C$ , the history  $H_0$  is also in  $\mathcal{T}_C$ . Lemma 17 yields a history  $H_0^c \in \text{cTMSpast}(H_0)$  and a history  $S_0 \in \mathcal{T}_A$  such that  $H_0^c \sqsubseteq_{\text{op}} S_0$ , where  $H_0^c \in \text{cTMSpast}(H_\psi)$ .

Let  $H^c = H_0\psi_c\psi$  and  $S = S_0\psi_c\psi$ . Since  $S_0 \in \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under aborting live transactions, we get that  $S \in \mathcal{T}_A$ . Since  $H_0^c \in \text{cTMSpast}(H_0)$ , we get  $H^c \in \text{cTMSpast}(H)$ , implying that  $H^c \sqsubseteq_{\text{op}} S$ , since  $H_0^c \sqsubseteq_{\text{op}} S_0$ . ■

**PROOF OF THEOREM 8(II).** Let  $H \in \mathcal{T}_C$ . Since  $\mathcal{T}_C \preceq \mathcal{T}_A$ , Lemma 16 gives a history  $H^c \in \text{comp}(H|_{\text{-live}})$  and a history  $S \in \mathcal{T}_A$  such that  $H^c|_{\text{com}} \sqsubseteq_{\text{op}} S$ .

Since every prefix of  $H$  is in  $\mathcal{T}_C$ , for every response  $\psi$  action in  $H$ ,  $H_\psi = H|_\psi \in \mathcal{T}_C$ .

If  $\psi$  is a committed action, then Lemma 16 gives a history  $H^c \in \text{comp}(H_\psi|_{\text{-live}})$  and a history  $S \in \mathcal{T}_A$  such that  $H^c|_{\text{com}} \sqsubseteq_{\text{op}} S$ .

If  $\psi$  is a response action, but neither committed nor aborted, then Lemma 17 gives a history  $S_\psi \in \mathcal{T}_A$  such that  $H_\psi^c \sqsubseteq_{\text{op}} S_\psi$ , where  $H_\psi^c \in \text{cTMSpast}(H_\psi)$ .

If  $\psi$  is an aborted action, then Lemma 18 gives a history  $S_\psi \in \mathcal{T}_A$  such that  $H_\psi^c \sqsubseteq_{\text{op}} S_\psi$ , where  $H_\psi^c \in \text{cTMSpast}(H_\psi)$ .

In all cases,  $H \sqsubseteq_{\text{tms}} \mathcal{T}_A$ . ■