

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

Hagit Attiya<sup>1</sup>, Alexey Gotsman<sup>2</sup>, Sandeep Hans<sup>1</sup>, and Noam Rinetzky<sup>3</sup>

<sup>1</sup> Technion - Israel Institute of Technology, Israel

<sup>2</sup> IMDEA Software Institute, Spain

<sup>3</sup> Tel Aviv University, Israel

**Abstract.** One of the main challenges in 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 be weak enough 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)*, a TM correctness condition, 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 acceptable 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 of the system state. Our proof reveals some natural, but subtle, assumptions on the TM required for the equivalence result.

## 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 are used to manipulate 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 done by common database consistency conditions [1]. 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

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

**Fig. 1.** TM usage

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 [2–5]. An ideal condition should be weak enough 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* [6, 7] 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 any program  $P$  in this language using the concrete TM can also be observed when  $P$  uses 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 [8] we showed that a variant of the *opacity* condition [2] is equivalent to observational refinement for a particular programming language and, hence, is the weakest acceptable 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., Scala STM [9]), it is non-standard and most TM systems do not satisfy it. In this paper, we consider a variant of *transactional memory specification (TMS)* [5], a condition weaker than opacity,<sup>1</sup> and show that, under some natural assumptions on the TM, it is equivalent to 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 justified by a single history of the abstract TM. TMS relaxes this requirement by requiring only committed transactions in the concrete history to be justified by a single abstract one obeying (i)–(ii) above; every response obtained from the TM in an aborted or live transaction may be justified by a separate abstract history. The constraints on the choice of the abstract history are subtle: on one 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 al-

<sup>1</sup> The condition we present here is actually called TMS1 in [5, 10]. These papers also propose another condition, TMS2, but it is stronger than opacity [10] and therefore not considered here.

low 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 justified by some abstract history; different transactions can be justified 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 1, 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 [11, 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  $\text{Var} = \text{GVar} \uplus \bigsqcup_{t=1}^m \text{LVar}_t$  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* 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  $\text{Method} = \{f, \dots\}$ . The syntax of commands  $C$  is standard:  $C$  can be of the forms

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

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.

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 [11, 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 *ActionId* be a set of action identifiers. A **TM interface action**  $\psi$  has one of the following forms:

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

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**  $\chi$  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 [11, Appendix B]): every action in  $\tau$  has a unique identifier; no action follows a `fault`; request and response actions are properly matched; for every thread  $t$ ,  $\tau|_t$  cannot contain a request action

immediately followed by a primitive action; 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 nonempty 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. 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 in  $\tau$  as  $\varphi$  (undefined if  $\varphi$  does not belong to a transaction).

## 4 Transactional Memory Specification (TMS)

In this section we define the TMS [5] correctness condition in our setting. TMS was originally formulated using I/O automata; here we define it in a different style appropriate for our goals (we provide further comparison in Section 7). Since threads may communicate through global variables outside of transactions, they may observe the *real-time order* between non-overlapping transactions in a history. Therefore, this order is a crucial building block in the TMS definition, as is common in consistency conditions for shared-memory concurrency, such as opacity [2] or linearizability [12].

**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 every  $\psi \in T$ . A transaction  $T$  is **before a transaction  $T'$  in the real-time order** in  $H$ , denoted by  $T \prec_H T'$ , if  $T \prec_H T'(1)$ .

The following *opacity relation* [2, 8]  $H \sqsubseteq_{\text{op}} S$  ensures that  $S$  is a permutation of  $H$  preserving the real-time order.

**Definition 3.** A history  $H$  is in the **opacity relation** with a history  $S$ , denoted by  $H \sqsubseteq_{\text{op}} S$ , if  $\forall \psi, \psi'. (\psi \in S \iff \psi \in H) \wedge (\psi \prec_H \psi' \implies \psi \prec_S \psi')$ .

Given a history  $H$  of program interactions with a concrete TM, TMS requires us to justify the behavior of all committed transactions in  $H$  by a single history  $S$  of

the abstract TM, and to justify each response action  $\psi$  inside a transaction in  $H$  by an abstract history  $S_\psi$ . As we show in this paper, the existence of such justifications ensures that TMS implies 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 justify 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 5. 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 using an example.


**Definition 4.** A history  $H_\psi = H'_1\psi$  is a **possible past** of a history  $H = H_1\psi$ , where  $\psi$  is a response action that it is not a committed or aborted 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:

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

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



**Fig. 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.

$T_4$  is changed to committed when replacing the concrete TM by an abstract one in observational 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 5.** A history  $H^c$  is a *suffix completion* of a history  $H\psi$  if  $H^c = H\psi H'$ , every 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  $\text{ccomp}(H)$ , 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  $\text{cTMSpast}(H)$ :

$$\text{cTMSpast}(H_1\psi) = \{H_\psi^c \mid \exists H'_1. H'_1\psi \in \text{TMSpast}(H_1\psi) \wedge H_\psi^c \in \text{ccomp}(\text{com}(H'_1)\psi)\},$$

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  $\psi$  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  $S$  preserves the real-time order of  $H$ . As in Definition 4(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 6.** A history  $H$  is in the **TMS relation** with TM  $\mathcal{T}$ , denoted  $H \sqsubseteq_{\text{tms}} \mathcal{T}$ , if:

- (i)  $\exists H^c \in \text{comp}(H|_{\text{-live}})$ ,  $S \in \mathcal{T}$ .  $H^c|_{\text{com}} \sqsubseteq_{\text{op}} S$ , where  $\cdot|_{\text{-live}}$  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 it is not a committed or aborted action and  $H = H_1\psi H_2$ , we have  $\exists H_\psi^c \in \text{cTMSpast}(H_1\psi)$ .  $\exists S_\psi \in \mathcal{T}$ .  $H_\psi^c \sqsubseteq_{\text{op}} S_\psi$ .

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. Due to space constraints, we defer its formal definition to [11, Appendix C] and describe only its high-level structure. 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 executed 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 executed 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 executed 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. The definition of  $\llbracket P \rrbracket(s)$  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 *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  $\tau$  excluding 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  $\mathcal{T}_C$ . An action  $\varphi \in \tau$  is **transactional** if  $\varphi \in T$  for some  $T \in \tau$ , and **non-transactional** otherwise. We denote by  $\tau|_{\text{trans}}$  and  $\tau|_{\text{-trans}}$  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  $\frac{1}{2}$  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)  $\exists \tau' \in \llbracket P, \mathcal{T}_A \rrbracket(s)$ .  $\tau'|_{\text{-trans}} = \tau|_{\text{-trans}}$ ; 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 is equivalent to observational refinement for abstract TMs that enjoy certain natural closure properties. Their formulation relies on the following notions.

A history  $H_a$  is an **immediate abort extension of a history**  $H$  if  $H$  is a subsequence of  $H_a$ , and whenever  $\psi \in H_a$  and  $\psi \notin H$  we have: (i)  $\psi = (\_, \_, \text{txbegin})$  or  $\psi = (\_, \_, \text{aborted})$ , (ii) if  $\psi = (\_, t, \text{txbegin})$  then  $H_a = H'_a \psi (\_, t, \text{aborted}) \_$ , where  $H'_a \in \{\varepsilon, \_ (\_, \_, \text{committed}), \_ (\_, \_, \text{aborted})\}$ , and (iii) if  $\psi = (\_, \_, \text{aborted})$  then there exists  $\psi' \notin H$  such that  $H_a = \_ \psi' \psi \_$ . We denote by  $\text{addab}(H)$  the set of all immediate abort extensions of  $H$ . Informally, a history  $H_a \in \text{addab}(H)$  is an extension of  $H$  with transactions that abort immediately after their invocation. Note that the added transactions are placed either right before other transactions begin or right after they complete.

A history  $H_c$  is a **non-interleaved completion** of a history  $H$  if  $H$  is a subsequence of  $H_c$ ,  $\text{pending}(H_c) = \emptyset$  and whenever  $\psi \in H_c$  and  $\psi \notin H$  we have  $H_c = \_ (\_, t, \text{txcommit}) \psi \_$  and either  $\psi = (\_, t, \text{committed})$  or  $\psi = (\_, t, \text{aborted})$ . We denote the set of non-interleaved completions of  $H$  by  $\text{nicomp}(H)$ . Informally,  $H' \in \text{nicomp}(H)$  completes each commit-pending transaction in  $H$  by adding a committed or aborted action at its end.

The required closure properties are formulated as follows:

CLP1 A TM  $\mathcal{T}$  is **closed under immediate aborts** if whenever  $H \in \mathcal{T}$  and  $\text{aborted}(H) = \emptyset$ , we also have  $H' \in \mathcal{T}$  for any history  $H' \in \text{addab}(H)$ .

CLP2 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  $H_1H_2 \in \mathcal{T}$ .

CLP3 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  $\text{aborted}(H') \subseteq \text{aborted}(H)$ .

CLP4 A TM  $\mathcal{T}$  is **closed under completing commit-pending transactions** if whenever  $H \in \mathcal{T}$ , we have  $\text{nicomp}(H) \cap \mathcal{T} \neq \emptyset$ .

These properties are satisfied by the expected TM specification that executes every transaction atomically [8].

**Theorem 1.** *Let  $\mathcal{T}_C$  and  $\mathcal{T}_A$  be transactional memories.*

(i) *If  $\mathcal{T}_A$  satisfies CLP1 and CLP2, then  $\mathcal{T}_C \sqsubseteq_{\text{tms}} \mathcal{T}_A \implies \mathcal{T}_C \preceq \mathcal{T}_A$ .*

(ii) *If  $\mathcal{T}_A$  satisfies CLP3 and CLP4, then  $\mathcal{T}_C \preceq \mathcal{T}_A \implies \mathcal{T}_C \sqsubseteq_{\text{tms}} \mathcal{T}_A$ .*

### 6.1 Proof of Theorem 1(i) (Sufficiency)

Let us fix a program  $P = C_1 \parallel \dots \parallel C_m$  and a state  $s$ . As we have noted before, the main subtlety of TMS lies in justifying 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: given a

trace  $\tau \in \llbracket P \rrbracket(s)$  with a live transaction and a history  $H_\psi^c \in \text{cTMSpast}(\text{history}(\tau))$ , the lemma converts  $\tau$  into another trace from  $\llbracket P \rrbracket(s)$  that contains the same live transaction, but whose history of non-aborted transactions is  $H_\psi^c$ . In other words, this establishes that the live transaction cannot notice changes in the committed/aborted status of other transactions done by  $\text{cTMSpast}$ . Let  $\tau|_{\text{-abortedtx}}$  be the projection of  $\tau$  excluding aborted transactions.

**Lemma 1 (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  $\tau_2$  is a sequence of primitive actions by thread  $t_0$ . For any  $H_\psi^c \in \text{cTMSpast}(\text{history}(\tau))$  there exists  $\tau_\psi \in \llbracket P \rrbracket(s)$  such that  $\text{history}(\tau_\psi)|_{\text{-abortedtx}} = H_\psi^c$  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 5 there exist histories  $H'_1$ ,  $H''_1$ , and  $H^{cc}$  such that

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

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 committing. 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 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  $H'_1\psi$ , or  $\varepsilon$  if no such action exists. For example, in Figure 2,  $\tau_{t_1}^N$  and  $\tau_{t_2}^N$  end with  $\chi_1$  and  $\chi_2$ , respectively. Let  $\tau_{t_0} = \tau|_{t_0}$  and for each  $t \neq t_0$  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 each  $t$ . Thus, for the  $\tau$  in Figure 2, in the corresponding trace  $\tau'$  the actions of  $t_1$  end with  $\chi_1$  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

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



**Fig. 3.** Cases in the proof of Lemma 1. \* all actions by  $t$  are transactional; \*\* all actions by  $t$  come from a single transaction, started before or by  $\psi^b$ .

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 4(ii),  $T \in H_1^1\psi$ . Hence, by Definition 4(iii) we have

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

so that  $(H_1^1\psi)|_{t_0}$  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 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  $t$ ; hence, by the choice of  $\tau_t^N$  it 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  $T \in H_1^1\psi$  by the choice of  $\psi^b$  and by Definition 4(iii) we get (1). Hence, for threads  $t$  falling into cases (a) or (b),  $\tau'|_t$  does not contain aborted transactions that are also in  $H_1^1\psi$ . For threads  $t$  falling into case (c), an aborted transaction by  $t$  included into  $H_1^1\psi$  can only be the last one in  $\tau'|_t$ . Finally, above we established that  $(H_1^1\psi)|_{t_0}$  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  $\text{history}(\tau_\psi)|_{\neg\text{abortedtx}} = H_\psi^c$ . It is sufficient to show that  $\text{history}(\tau'')|_{\neg\text{abortedtx}} = H_1''\psi$ ; since  $\tau_\psi = \tau''H^{cc}$  and  $H^{cc}$  contains only committed actions, this would imply

$$\begin{aligned} \text{history}(\tau_\psi)|_{\neg\text{abortedtx}} &= \text{history}(\tau''H^{cc})|_{\neg\text{abortedtx}} = \\ &= \text{history}(\tau'')|_{\neg\text{abortedtx}}H^{cc} = H_1''\psi H^{cc} = H_\psi^c. \end{aligned}$$

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  $\text{history}(\tau'')|_{\neg\text{abortedtx}}$ .

Thus, to prove  $\text{history}(\tau'')|_{\neg\text{abortedtx}} = H_1''\psi$  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 [11, Appendix D]. Let  $\chi_t^N$  be the last action in  $\tau_t^N$  and  $T = \text{txof}(\psi^b, H_1'\psi) \in H_1'\psi$ . Then by Definition 4(iii) we get (1). Since  $\chi_t^N$  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  $\text{history}(\tau'')|_{\neg\text{abortedtx}} = H_1''\psi$ .  $\square$

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  $\text{cTMSpast}$ . 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 [8, Theorem 16] (it was proved in [8] 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 2.** *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  $\text{history}(\tau') = S$ ,  $\tau|_{\neg\text{trans}} = \tau'|_{\neg\text{trans}}$  and  $\forall t. \tau'|_t = \tau|_t$ .*

Let  $\tau|_{\neg\text{abortact}}$  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 2 because local variables are rolled back if a transaction aborts, and, hence, applying  $\cdot|_{\neg\text{abortact}}$  to a trace preserves its validity.

**Proposition 1.**  $\forall \tau. \tau \in \llbracket P \rrbracket(s) \implies \tau|_{\neg\text{abortact}} \in \llbracket P \rrbracket(s)$ .

Finally, Definition 6 matches only histories of committed transactions, but the histories of the traces in Lemma 2 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 3.** *Let  $H$  be a history where all aborted transactions abort immediately and  $S$  be such that  $H|_{\neg\text{abortedtx}} \sqsubseteq_{\text{op}} S$ . There exists a history  $S' \in \text{addab}(S)$  such that  $H \sqsubseteq_{\text{op}} S'$ .*

Definition 6(i), Proposition 1 and Lemmas 2 and 3 can be used to prove that the TMS relation preserves non-transactional actions and thread-local observable behavior of threads whose last action is not a `fault`.

**Lemma 4.** *If  $\mathcal{T}_C \sqsubseteq_{\text{tms}} \mathcal{T}_A$  and  $\mathcal{T}_A$  satisfies CLP1 and CLP2, then*

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

**Proof of Theorem 1(i).** Given Lemma 4, we only need to establish the preservation of faults inside transactions. Consider  $\tau_0 \in \llbracket P, \mathcal{T}_C \rrbracket(s)$  such that  $\tau_0 = \tau_1\psi\tau_2\chi$ , where  $\chi = (\_, t_0, \text{fault})$  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,  $\tau = \tau_1\psi(\tau_2|_{t_0})\chi \in \llbracket P, \mathcal{T}_C \rrbracket(s)$ . 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 1, for some trace  $\tau_\psi$  we have  $\tau_\psi \in \llbracket P \rrbracket(s)$ ,  $\text{history}(\tau_\psi)|_{\neg\text{abortedtx}} = H_\psi^c$  and  $\tau_\psi|_{t_0} = \tau|_{t_0}$ . By Proposition 1,  $\tau_\psi|_{\neg\text{abortact}} \in \llbracket P \rrbracket(s)$ . Using Lemma 3, we get a history  $S'$  such that  $\text{history}(\tau_\psi|_{\neg\text{abortact}}) \sqsubseteq_{\text{op}} S'$  and  $S' \in \text{addab}(S)$ . Since  $S \in \mathcal{T}_A$  and  $\mathcal{T}_A$  is closed under immediate aborts (CLP1), we get  $S' \in \mathcal{T}_A$ . Hence, by Lemma 2, for some  $\tau' \in \llbracket P, \mathcal{T}_A \rrbracket(s)$  we have  $\tau'|_{t_0} = \tau_\psi|_{t_0} = \tau|_{t_0} = \_ \chi$ , as required.  $\square$

## 6.2 Proof Sketch for Theorem 1(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), in which  $H_0 = H_1\psi H_2 = HH_2 \in \mathcal{T}_C$ , where  $\psi$  is a response action by a thread  $t_0$  that is not a committed or aborted action. 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.

Let  $H' = \text{history}(\tau')$ . This construction of  $P_H$  allows us to infer that: (i) the projection of  $H'|_{t_0}$  to committed transactions and  $\text{txof}(\psi, H')$  is equal to the corresponding projection of  $H|_{t_0}$ ; (ii) for all other threads  $t$  a similar relationship holds for the prefix of  $H'|_t$  ending with the last transaction preceding  $\text{txof}(\psi, H')$  in the real-time order; (iii) the real-time order in  $H'$  includes that in  $H$ . Transactions concurrent with  $\text{txof}(\psi, H')$  in  $H'$  may behave differently from  $H$ . However, checks done by  $P_H$  inside these transactions ensure that, if such a transaction  $T$  is visible in  $H'$ , 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  $H'$  as prescribed by  $H$ . This relationship between  $H$  and  $H'$  allows us to establish the requirements of Definition 6(ii).  $\square$

## 7 Related Work

When presenting TMS [5], 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 1—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 1, 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) [3]. Like TMS, VWC allows every operation in a live or aborted transaction to be justified 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 [8] 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 developed new techniques, such as establishing the live transaction insensitivity property (Lemma 1) to prove sufficiency and proposing monitor programs for the nontrivial constraints used in the TMS definition to prove necessity. Similarly to [8] and other papers using observational refinement to study consistency conditions [13, 14], we reformulate TMS so that it is not restricted to a particular abstract TM  $\mathcal{T}_A$ . 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.

**Acknowledgements.** We thank Mohsen Lesani, Victor Luchangco and the anonymous reviewers for comments that helped us improve the paper. This work was supported by EU FP7 project ADVENT (308830).

## References

1. Papadimitriou, C.H.: The serializability of concurrent database updates. *J. ACM* 26, 631–653 (1979)
2. Guerraoui, R., Kapalka, M.: On the correctness of transactional memory. In: PPOPP, pp. 175–184 (2008)
3. Imbs, D., Raynal, M.: Virtual world consistency: A condition for STM systems (with a versatile protocol with invisible read operations). *Theor. Comput. Sci.* 444, 113–127 (2012)
4. Attiya, H., Hans, S., Kuznetsov, P., Ravi, S.: Safety of deferred update in transactional memory. In: ICDCS (2013)
5. Doherty, S., Groves, L., Luchangco, V., Moir, M.: Towards formally specifying and verifying transactional memory. *Formal Aspects of Computing* 25, 769–799 (2013)
6. He, J., Hoare, C., Sanders, J.: Prespecification in data refinement. *Information Processing Letters* 25, 71–76 (1987)
7. He, J., Hoare, C., Sanders, J.: Data refinement refined. In: Robinet, B., Wilhelm, R. (eds.) *ESOP 1986. LNCS*, vol. 213, pp. 187–196. Springer, Heidelberg (1986)
8. Attiya, H., Gotsman, A., Hans, S., Rinetzky, N.: A programming language perspective on transactional memory consistency. In: PODC, pp. 309–318 (2013)
9. 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)
10. Lesani, M., Luchangco, V., Moir, M.: Putting opacity in its place. In: WTTM (2012)
11. Attiya, H., Gotsman, A., Hans, S., Rinetzky, N.: Safety of live transactions in transactional memory: Tms is necessary and sufficient. Technical Report CS-2014-02, Technion (2014)
12. Herlihy, M., Wing, J.M.: Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems* 12, 463–492 (1990)
13. Gotsman, A., Yang, H.: Liveness-preserving atomicity abstraction. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) *ICALP 2011, Part II. LNCS*, vol. 6756, pp. 453–465. Springer, Heidelberg (2011)
14. Filipović, I., O’Hearn, P., Rinetzky, N., Yang, H.: Abstraction for concurrent objects. In: Castagna, G. (ed.) *ESOP 2009. LNCS*, vol. 5502, pp. 252–266. Springer, Heidelberg (2009)