Transactions in the Jungle

Rachid Guerraoui   Thomas A. Henzinger   Michal Kapalka   Vasu Singh

École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

{rachid.guerraoui,tah,michal.kapalka,vasu.singh}@epfl.ch

Abstract

Transactional memory (TM) has shown potential to simplify the task of writing concurrent programs. However, the semantics of interactions between transactions managed by a TM and non-transactional operations, while widely studied, lacks a clear formal specification. Those interactions can vary, sometimes in subtle ways, between TM implementations and underlying memory models.

We propose a new correctness condition for TMs, *parametrized opacity*, which captures the two following intuitive requirements: first, every transaction appears as if it is executed instantaneously with respect to other transactions and non-transactional operations, and second, non-transactional operations conform to a given memory model. Parametrized opacity corresponds to the well-studied strong atomicity property, which lacks a formal definition and is, in fact, ambiguous.

We use our formalization to theoretically investigate the inherent cost of implementing parametrized opacity. We first prove that parametrized opacity requires either instrumenting non-transactional operations (for most memory models) or writing to memory by transactions using potentially expensive read-modify-write instructions (such as compare-and-swap). Then, we show that for a class of relaxed memory models, parametrized opacity can indeed be implemented with constant-time instrumentation of non-transactional writes and no instrumentation of non-transactional reads. We show that, in practice, parametrizing the notion of correctness allows to develop more efficient TM implementations.

1. Introduction

Transactional memory (TM) [14, 26] offers a promising paradigm for concurrent programming in the multi-core era. A TM allows a programmer to think in terms of coarse-grained code blocks—transactions—that appear to be executed atomically, and at the same time, a TM yields a high level of parallelism. Ideally, a program could execute all operations on shared data within transactions, and have non-transactional operations on thread-local data. In practice, however, not all operations on shared data can be wrapped in transactions. For example, a programmer may wish to make shared data local to a thread, operate non-transactionally upon it for a while, and make it shared again [20, 30]. It is thus not surprising to see a large body of research dedicated to exploring the various models of interactions between transactions and non-transactional code [1, 6, 8, 30, 19, 21], and building TM implementations that implement those models [27, 20].

The mutual interaction between transactions is formalized by the notion of *serializability* [22] or *opacity* [12]. Serializability, a

Thread 1	Thread 2
atomic {	
$\quad x := 1$	
$\quad y := 1$	$r_1 := x$
}	$r_2 := y$

Figure 1. Can $r_1 = 1$ and $r_2 = 0$? It depends on the memory model (initially $x = y = 0$).

common correctness notion in database transactions, requires that committed transactions look as if they executed sequentially. But, it was observed [7, 13] that serializability is not sufficient for memory transactions, where it is important that even aborted transactions do not observe an inconsistent state of the memory. This gave rise to a new correctness criterion called opacity. The idea of opacity is to give programmers an illusion that there is no concurrency between transactions, whether committed or aborted. That is, opacity requires that all transactions "look like" they executed in some sequential order, consistent with their real-time order. Most of the TM implementations [7, 13, 23] satisfy opacity. The formal definition of opacity has allowed for a clear understanding of transactions [12] and also development of model checking techniques for TM algorithms [9, 10].

The interaction between transactions and non-transactional operations has been defined using a notion of *strong atomicity* [19, 17] in the literature. The intuition behind strong atomicity is that transactions execute atomically with respect to other transactions and non-transactional operations. Unfortunately, strong atomicity has not been formally defined, which has thus led to multiple interpretations [29]. Consider, for example, the execution depicted in Figure 1 (adapted from [8]). The transaction executed by thread 1 updates variables x and y. Thread 2 reads the variables x and y non-transactionally. Is it possible that thread 2 reads x as 0 and y as 1? According to the definition by Martin et al. [19], strong atomicity allows this result. But, according to the definition of strong atomicity by Larus et al. [17], this result is not allowed. The ambiguity in this definition can be attributed to an implicit assumption on the interaction between non-transactional operations, which in turn, depends upon the underlying memory model [2]. A *memory model* specifies the set of allowed behaviors of memory accesses in a shared memory program. A relaxed memory model allows the underlying system to reorder memory instructions, while a stringent memory model like sequential consistency enforces instructions to execute in program order. While Martin et al. [19] assume a relaxed memory model which allows to reorder independent reads (for example, RMO [31]), Larus et al. [17] assume a sequentially consistent memory model.

We provide a general formal framework for describing the interactions between transactions and non-transactional operations. We consider opacity as a correctness condition for transactions, and parametrize it by a memory model. We claim that while a TM can be implemented in a way to ensure opacity for transactions, there is little one can do (on a given platform or run-time environment) to change the underlying memory model. Hence, it is desirable to define opacity *parametrized* by a memory model.

Thread 1	Thread 2
$atomic$ {	
$\quad x := 1$	
$\quad x := 2$	$atomic$ {
}	$\quad z := x - y$
$atomic$ {	}
$\quad y := 2$	
}	

(a) Can $z < 0$?

Thread 1	Thread 2
$x := 1$	$r_1 := y$
$y := 1$	$r_2 := x$

(b) Can $r_1 = 1$ and $r_2 = 0$?

Thread 1	Thread 2
$atomic$ {	$z := x$
$\quad x := 1$	
$\quad x := 2$	
}	
$atomic$ {	
$\quad r_1 := z$	
$\quad r_2 := z$	
}	

(c) Can $z = 1$ or $r_1 \neq r_2$?

Figure 2. Motivating examples for the definition of parametrized opacity (initially $x = y = z = 0$ in every case)

Moreover, we want the definition of parametrized opacity to be implementation-agnostic (like opacity), so that it allows for transactional objects with semantics richer than that of simple read-write variables (which could help describing, for example, TMs that implement transactional boosting [15] or similar techniques).

We present a definition of a memory model which is general enough to capture a variety of memory models. Intuitively, we formalize a memory model as a function which, depending upon the sequence of operations, gives the set of possible order of operations. Our formalism can capture common memory models like TSO, PSO, RMO, and Alpha. Moreover, we allow different processes to observe different order of operations, which allows us to capture memory models with non-atomic stores, like IA-32. We classify memory models on the basis of the possible reorderings of operations they allow.

Our formalization of parametrized opacity is guided by the following intuition:

1. *Opacity for transactions.* Whatever the memory model is, executions that are purely transactional must ensure opacity. Indeed, the semantics of transactions should be intuitive and strong—in the end, we want TMs to be as easy to use as coarse-grained locking. For example, consider Figure 2(a). Thread 2 should observe x as 0 or 2, because the intermediate state of a transaction ($x = 1$) is not visible to other transactions. Also, y can be observed as 0 or 2. Moreover, y can be observed as 2 only if x is observed as 2, because the effect of transactions is visible in real-time order. Thus, the possible values of z are 0 and 2. Note that even if thread 2 aborts, opacity requires that z is 0 or 2.

2. *Efficiency of non-transactional operations.* Executions that are purely non-transactional have to adhere to the given memory model. In particular, parametrized opacity should not strengthen the semantics of non-transactional operations. The motivation here is to avoid a framework that would inherently require non-transactional operations to be instrumented with additional memory fences or software barriers, even for very weak memory models. For example, in Figure 2(b), a memory model may relax the order of write operations in Thread 1 or the read operations in Thread 2, resulting in $r_1 = 1$ and $r_2 = 0$.

3. *Isolation of transactions from non-transactional operations.* Transactions should appear, both to other transactions and non-transactional operations, as if they were executed instantaneously. In particular, isolation of transactions should be respected, regardless of the memory model. That is, first, the intermediate computations of transactions, or updates by aborted transactions, should never be visible to non-transactional operations, and, second, the non-transactional operations concurrent to a transaction should appear as if they happened before or after this transaction. For example, in Figure 2(c), Thread 2 cannot observe an intermediate state of a transaction, and thus

$z \neq 1$. Moreover, the effect of a non-transactional operation cannot show up in the middle of a transaction. Thus, $r_1 = r_2$.

Note that opacity parametrized by sequential consistency gives the notion of strong atomicity as proposed by Larus et al. [17]. On the other hand, parametrized opacity for a relaxed memory model like RMO matches the notion of strong atomicity given by Martin et al. [19].

In practice, TM implementations that guarantee strong atomicity require that non-transactional operations, instead of accessing memory directly, adhere to an access protocol as defined by the TM implementation. This modification of the semantics of non-transactional operations is known as *instrumentation*. For example, Tabatabai et al. [27] propose a TM implementation, where the non-transactional read and write operations follow the locking discipline as done by the transactions. The formal definition of opacity parametrized by a memory model allows us to theoretically analyze the cost of creating TM implementations that guarantee parametrized opacity. While parametrized opacity is the intuitive correctness property for transactional programs with non-transactional operations, we show that, without instrumentation, it is impossible to achieve on most memory models. Even for the small class of idealized memory models, where parametrized opacity can be achieved without instrumentation, we show that a TM implementation *must* use expensive read-modify-write operations for each object modified by a transaction.

Next, we focus on TM implementations that instrument non-transactional write operations, without any instrumentation for non-transactional read operations. We start with a basic result which shows that for a class of memory models which allows to reorder independent reads, it is possible to achieve parametrized opacity without instrumenting the read operations, and treating every non-transactional write as a transaction in itself. Note that this might not provide a practical solution, as we do not want a non-transactional operation to carry the overhead associated with a transaction. Moreover, we want non-transactional operations to finish in bounded time, while a transaction, in general, may take arbitrarily long to finish. The next question we ask is whether we can obtain parametrized opacity for a class of memory models with constant-time instrumentation on the writes? We show that for a class of memory models which allows to reorder a read of a variable following a read or write of another variable (like Alpha [28]), it is indeed possible to achieve parametrized opacity with constant-time instrumentation for non-transactional write operations. We also discuss how to adapt the constant-time write instrumentation solution for memory models which do not allow to reorder data-dependent reads (like RMO [31] and Java [18]).

Using our theoretical framework, we examine existing TM implementations that guarantee parametrized opacity. TM implementations in the literature that satisfy strong atomicity [27], in fact, satisfy parametrized opacity with respect to sequential consistency. We observe that the TM implementation can be designed to be more

efficient, if it is to satisfy opacity parametrized by a weaker memory model. We discuss examples of modifications that can be applied to the TM implementation by Tabatabai et al. [27] which remove overhead, while keeping it parametrized opaque with respect to a wide class of relaxed memory models. We also show that our theoretical framework can be used to specify weaker notions of correctness. As an example, we formalize single global lock atomicity (SGLA) [8, 17, 20]. Moreover, we show that the impossibility results we obtain for parametrized opacity do not hold for SGLA.

2. Preliminaries

We first describe a framework of a shared memory system consisting of shared objects and operations on those objects.

Operations. Let Obj denote a set of shared objects. We consider a shared-memory system consisting of a set P of processes, which communicate by executing commands on (shared) objects. Let C be a set of commands on shared objects, where arguments and return values are treated as part of a command. For example, in a system that supports only reading and writing of shared (natural number) variables, we have $C = \{\mathsf{rd}, \mathsf{wr}\} \times \mathbb{N}$. We define *operations* $O \subseteq C \times Obj$ as the set of all allowed command-object pairs.

Besides operations that issue commands on shared objects, every process $p \in P$ can execute the following special operations: start to start a new transaction, operation commit to commit a transaction, and abort to abort a transaction. Let $\hat{O} = O \cup \{\mathsf{start}, \mathsf{commit}, \mathsf{abort}\}$.

Histories. We define an *operation instance* as (o, p, k), where $o \in \hat{O}$ is an operation, $p \in P$ is a process which issues the operation, and $k \in \mathbb{N}$ is a natural number representing the identifier of the operation instance.

A *history* $h \in (\hat{O} \times P \times \mathbb{N})^*$ is a sequence of operation instances, such that for every pair $(o, p, i), (o', p', j)$ of operation instances in h, we have $i \neq j$. Intuitively, we want each operation instance in a history to have a unique identifier. For a natural number k, when we say "operation k", we mean "operation instance with identifier k", i.e., the element of h of the form (o, p, k), where $o \in \hat{O}$ and $p \in P$.

A *transaction* of a process p is a subsequence $(o_1, p, i_1) \ldots (o_n, p, i_n)$ of a history h such that (i) o_1 is a start operation, (ii) either operation i_n is the last operation instance of p in h, or we have $o_n \in \{\mathsf{commit}, \mathsf{abort}\}$, and (iii) all operations $o_2, \ldots, o_{n-1}$ belong to set O.

A transaction T is *committed* (resp. *aborted*) in a history h if the last operation instance of T has a commit operation (resp. an abort operation). A transaction T is *completed* if T is committed or aborted.

Given a history h and a natural number k, we say that operation k is *transactional* in h if operation k is part of a transaction in h. Otherwise, operation k is said to be *non-transactional*. We assume that every history h is *well-formed*: that is, every non-transactional operation in h belongs to set O. Intuitively, well-formedness of a history requires that every commit and abort of a transaction matches with a corresponding start, and there are no nested transactions. We denote by H the set of all histories.

Given a history h, we define the *real-time* partial order relation $\prec_h \subset \mathbb{N} \times \mathbb{N}$ of the operation identifiers in h, such that for two natural numbers i and j, we have $i \prec_h j$ if:

1. operations i and j belong to transactions T and T', respectively, where T is completed in h and the last operation instance of T precedes the first operation instance of T' in h, or

2. operation i precedes operation j in h, both operations are executed by the same process, and at least one of those operation instances is transactional.

For example, consider the history h illustrated in Figure 3(a). The transaction of process p_1 finishes before the transaction of process p_3 starts. The precedence relation $\prec_h$ consists of elements $(1, 2), (5, 7)$, and $(1, 9)$. On the other hand, $(1, 6)$ and $(6, 9)$ are *not* in $\prec_h$.

Object semantics. We use the concept of a *sequential specification* [16, 32] to describe the semantics of objects. Given an object $x \in Obj$, we define the semantics $[\![x]\!] \subseteq C^*$ as the set of all sequences of commands on x that could be generated by a single process accessing x.

For example, let x be a shared variable that supports only the commands to read and write its value (with initial value 0). Then, $[\![x]\!]$ is a subset of $(\{\mathsf{rd}, \mathsf{wr}\} \times \mathbb{N})^*$, such that, for every sequence $c_1 \ldots c_n$ in $[\![x]\!]$, and for all $i, v \in \mathbb{N}$, if $c_i = (\mathsf{rd}, v)$ then either (a) the latest write operation preceding c_i in $c_1 \ldots c_n$ is (wr, v), or (b) $v = 0$ and no write operation precedes c_i in $c_1 \ldots c_n$.

Sequential histories. We say that a history h is *sequential* if, for every transaction T in h, every operation instance between the start operation instance of T and the last operation instance of T in h is a part of T. That is, intuitively, no transaction T overlaps with another transaction or with any non-transactional operation in h.

We say that a sequential history s *respects* a partial or a total order $\prec \subseteq \mathbb{N} \times \mathbb{N}$ if, for every pair (i, j), if $i \prec j$ then operation i precedes operation j in s.

Let s be a sequential history. We denote by $s|_x$ the longest subsequence of all commands invoked on object x in s. We say that s is *legal* if for every object x, we have $s|_x \in [\![x]\!]$.

We denote by $visible(s)$ the longest subsequence of s that does not contain any operation instance of a non-committed transaction T, except if the T is not followed in s by any other transaction or non-transactional operation instance. We say that an operation k in s is *legal* in s if history $visible(s')$ is legal, where s' is the prefix of s that ends with operation k.

Two examples of sequential histories are given in Figure 3(b) and Figure 3(c). Note that s_1 and s_2 respect $\prec_h$ (where h is the history shown in Figure 3(a)). History s_1 is legal if $v = v' = 1$. Similarly, s_2 is legal if $v = 0$ and $v' = 1$.

3. Parametrized Opacity

We first formalize a memory model in our framework. Then, we define opacity parametrized by a memory model.

3.1 Memory models

A memory model describes the semantics of memory accesses in a shared memory system. We formalize a memory model as a transformation followed by per-process reordering of the history. The reordering function is defined in such a way that it allows different processes to have different views of the history, similar to the formalism by Sarkar et al. [25]. The transformation function allows to have complex operations at the level of the history, which may need a sequence of operations at the level of the implementation. Moreover, the transformation function allows to capture memory models which do not provide out-of-thin-air guarantees (e.g., in languages like C/C++ for unsynchronized code).

We define a *process view* $v : P \rightarrow 2^{\mathbb{N} \times \mathbb{N}}$ as a function such that $v(p)$ is a partial order on the set of natural numbers for every process $p \in P$. Let V be the set of all process views. A *memory model* is a pair $M = (\tau, R)$, where

- $\tau : O \rightarrow O^*$ is a *transformation* function that maps an operation to a sequence of operations, and

p_1	p_2	p_3
$((\mathsf{wr}, x, 1), 1)$		
$((\mathsf{start}), 2)$		
		$((\mathsf{rd}, y, 1), 3)$
$((\mathsf{wr}, y, 1), 4)$		
$((\mathsf{commit}), 5)$		
		$((\mathsf{rd}, x, v), 6)$
	$((\mathsf{start}), 7)$	
	$((\mathsf{commit}), 8)$	
	$((\mathsf{rd}, x, v'), 9)$	

(a) History h

p_1	p_2	p_3
$((\mathsf{wr}, x, 1), 1)$		
$((\mathsf{start}), 2)$		
$((\mathsf{wr}, y, 1), 4)$		
$((\mathsf{commit}), 5)$		
	$((\mathsf{rd}, y, 1), 3)$	
	$((\mathsf{rd}, x, v), 6)$	
	$((\mathsf{start}), 7)$	
	$((\mathsf{commit}), 8)$	
	$((\mathsf{rd}, x, v'), 9)$	

(b) Sequential history s_1

p_1	p_2	p_3
	$((\mathsf{rd}, x, v), 6)$	
$((\mathsf{wr}, x, 1), 1)$		
$((\mathsf{start}), 2)$		
$((\mathsf{wr}, y, 1), 4)$		
$((\mathsf{commit}), 5)$		
	$((\mathsf{rd}, y, 1), 3)$	
		$((\mathsf{start}), 7)$
		$((\mathsf{commit}), 8)$
		$((\mathsf{rd}, x, v'), 9)$

(c) Sequential history s_2

Figure 3. Examples of histories and sequential histories. Every history is read top to bottom. Notation: $(\mathsf{wr}, x, 1), 1)$ under a column marked with process p stands for the operation instance $(((\mathsf{wr}, 1), x), p, 1)$.

- $R : H \to 2^V$ is a *reordering* function that maps a history to a set of process views.

Intuitively, τ maps every operation to its internal representation (e.g., in hardware or a run-time environment). For instance, a write to a 64-bit memory word might be, on some systems, executed as two writes to its 32-bit parts. Moreover, if a memory model does not give out-of-thin-air guarantees, τ can map every write to a special *havoc* operation followed by the write. A read which follows the havoc operation and precedes the write operation can return any value. A process view allows different processes to observe different orders of operations in a given history. This can capture memory models, like IA-32 [5] which allow non-atomic stores.

If h is a well-formed history, then $\tau(h)$ denotes a well-formed history obtained from h by replacing every operation instance (o, p_j, k) of h with a sequence $(o_1, p_j, k_1), \ldots, (o_m, p_j, k_m)$, where $\tau(o) = o_1, \ldots, o_m$ and $k_1, \ldots, k_m \in \mathbb{N}$. (Note that identifiers of operation instances of $\tau(h)$ can be arbitrary, as long as they are unique.)

Well-formed memory models. A transformation function τ is *well-formed* if for every well-formed, sequential and legal history s, sequence $\tau(s)$ is also a well-formed, sequential, and legal history. Intuitively, a transformation is well-formed if it does not change the semantics of transactional operations.

A reordering function R is *well-formed* if for every history $h \in H$, for every view $v \in R(h)$, for every process $p \in P$: (i) if $(i, j) \in v(p)$, then operations i and j are non-transactional in h, (ii) if $(i, j) \in v(p)$, then i precedes j in h, and (iii) if $(i, j) \in v(p)$, then $(j, i) \notin v(p')$ for all processes $p' \in P$. Intuitively, condition (i) requires that a reordering function can impose an order only on non-transactional operations, condition (ii) requires that a view does not force a process to observe operations of some process in out of program order, and condition (iii) requires that a view should not force two processes to observe operations to occur in different orders. A memory model $M = (\tau, R)$ is well-formed if τ and R are well-formed. All memory models we know of are indeed well-formed.

Capturing dependence of operations. Often, a memory model [18, 31] allows to reorder operations unless the latter operation is control-dependent or data-dependent on the former operation. We need to distinguish dependent reads from independent reads in our framework to capture these memory models. We can capture these dependencies in our framework using additional commands: $\{\mathsf{cdrd}, \mathsf{ddrd}, \mathsf{cdwr}, \mathsf{ddwr}\} \times \mathbb{N} \times 2^\mathbb{N}$. For example, an operation instance (o, p, k) in h with $o = (\mathsf{cdrd}, v, \{k_1, \ldots, k_n\}), x)$ denotes a read operation which is control-dependent on operations $k_1 \ldots k_n$ in h. Well-formedness of a history with control and data

dependent operations requires that if an operation k is dependent on $k_1 \ldots k_n$ in h, then all operations $k_1 \ldots k_n$ precede k in h.

3.2 Examples of memory models

We now give examples of memory models for histories on read and write operations on shared variables. We first define an *identity transformation* function τ_I such that $\tau_I(o) = o$ for every operation $o \in O$. We say that a view v is identical across processes, if $v(p) = v(p')$ for all processes $p, p' \in P$.

- *Sequential consistency* requires that the order of operations of a process in a history is preserved in every view, and all processes view an identical order of operations of different processes. Formally, $M_{SC} = (\tau_I, R_{SC})$ such that for all histories h, we have $v \in R_{SC}(h)$ if (a) v is identical across processes, and (b) for every process $p \in P$, for every pair $(o_1, p, i), (o_2, p, j)$ of operations such that operation i precedes operation j in h, we have $(i, j) \in v(p)$.

- *Total store order* allows a write operation to forward the value of a variable to a following read operation, and allows to reorder a write operation followed a read operation to a different variable. Formally, the memory model $M_{tso} = (\tau_I, R_{tso})$ such that for all histories h, we have $v \in R_{tso}(h)$ if (a) v is identical across processes, and (b) for every pair $(o_1, p, i), (o_2, p, j)$ of operations such that operation i precedes operation j in h, we have $(i, j) \in v(p)$ if one of following conditions holds:

  - o_2 is a write operation[1],

  - o_1 and o_2 are to the same object x, or

  - o_1 is a read operation of the form (rd, x, v) such that (wr, x, v) is the last preceding write operation to x by process p in h.

  The intuition for the last case is to allow two read operations to different variables to reorder if the first read obtains the value from a store buffer.

- *Relaxed memory order* allows to reorder read and write operations to different variables, unless the first operation is a read, and the second operation is either a write control/data-dependent on the first operation, or a read data-dependent on the first operation. RMO is specified by the memory model $M_{rmo} = (\tau_I, R_{rmo})$ such that for all histories h, we have $v \in R_{rmo}(h)$ if (a) v is identical across processes, and (b) for every process p, for every pair $(o_1, p, i), (o_2, p', j)$ of operations such that operation i precedes operation j in h, we have $(i, j) \in v(p)$ if one of the following conditions holds:

---

[1] We use the term "write operation" as a general term for a simple write, or a control/data dependent write. Similarly, for read operation.

- o_1 and o_2 are to the same object x,
- o_1 is a read of a variable x and $o_2 = ((\mathsf{cdwr}, v, K), y)$ or $o_2 = ((\mathsf{ddwr}, v, K), x)$ for some $v \in V$ and $K \subseteq \mathbb{N}$ such that $i \in K$, or
- o_1 is a read of a variable x and $o_2 = ((\mathsf{ddrd}, v, K), y)$ for some $v \in V$ and $K \subseteq \mathbb{N}$ such that $i \in K$.

- *Junk-SC* is a memory model we describe here to show how memory models that allow junk (out-of-thin-air) values can be expressed in our formalism. Junk-SC requires sequential consistency, but if there exist a read and a write to a variable, such that they are not real-time ordered with respect to each other, then the read can observe any junk value. We denote Junk-SC as $M_{junk} = (\tau, R_{SC})$, where τ is given as: $\tau(\mathsf{wr}, x, v) = havoc(x) \cdot (\mathsf{wr}, x, v)$ and $\tau(o) = o$ otherwise.

Classes of memory models. We now present a classification of memory models on the basis of reorderings they allow. We build the following four classes depending upon the restrictions posed by the memory model.

1. We first describe the class M_{rr} which represents the *read-read restrictive memory models*. We let $M_{rr} = M_{rr}^i \cup M_{rr}^c \cup M_{rr}^d$, where:

   M_{rr}^i is the set of memory models $M = (\tau_I, R)$ such that for all histories h, for all $i, j \in \mathbb{N}$, if operation i is a read operation to x and operation j is a read operation to y such that $x \neq y$, and both operations i, j are by process p, then for all view orders $v \in R(h)$, for all $p' \in P$, we have $(i, j) \in v(p')$.

   M_{rr}^c be the set of memory models $M = (\tau_I, R)$ such that for all histories h, for all $i, j \in \mathbb{N}$, if operation i is a read operation to x and operation j is of the form $(((\mathsf{cdrd}, v, K), y), p, j)$ for some v and $K \subseteq \mathbb{N}$ such that $x \neq y$ and $i \in K$, and both operations i, j are by process p, then for all view orders $v \in R(h)$, for all $p' \in P$, we have $(i, j) \in v(p')$.

   Similarly, we define M_{rr}^d as the set of memory models which do not allow to reorder two read operations to different variables, if the second operation is data-dependent on the first operation. Generally, if a memory model M is in M_{rr}^i, then $M \in M_{rr}^c$ and $M \in M_{rr}^d$.

2. We define the class of *read-write restrictive memory models* as $M_{rw} = M_{rw}^i \cup M_{rw}^c \cup M_{rw}^d$, where:

   M_{rw}^i is the set of memory models $M = (\tau_I, R)$ such that for all histories h, for all $i, j \in \mathbb{N}$, if operation i is a read operation to x and operation j is a write operation to y such that $x \neq y$, and both operations i, j are by process p, then for all view orders $v \in R(h)$, for all $p' \in P$, we have $(i, j) \in v(p')$.

   M_{rw}^c is the set of memory models $M = (\tau_I, R)$ such that for all histories h, for all $i, j \in \mathbb{N}$, if operation i is a read operation to x and operation j is of the form $(((\mathsf{cdwr}, v, K), y), p, j)$ for some v and $K \subseteq \mathbb{N}$ such that $x \neq y$ and $i \in K$, and both operations i, j are by process p, then for all view orders $v \in R(h)$, for all $p' \in P$, we have $(i, j) \in v(p')$.

   Similarly, we define M_{rw}^d as the set of memory models which do not allow to reorder a read followed by a write operation to a different variable, if the second operation is data-dependent on the first operation.

3. We define the class of *write-read restrictive memory models* as M_{wr}, where $M = (\tau_I, R)$ belongs to M_{wr} if for all histories h, for all $i, j \in \mathbb{N}$, if operation i is a write operation to x and operation j is a read operation to y such that $x \neq y$, and both operations i, j are by process p, then for all view orders $v \in R(h)$, for all $p' \in P$, we have $(i, j) \in v(p')$.

4. We define the class of *write-write restrictive memory models* as M_{ww}, where $M = (\tau_I, R)$ belongs to M_{ww} if for all histories h, for all $i, j \in \mathbb{N}$, if operations i and j are write operations to x and y such that $x \neq y$, and both operations i, j are by process p, then for all view orders $v \in R(h)$, for all $p' \in P$, we have $(i, j) \in v(p')$.

We classify some well-known memory models in these classes.

- SC memory model M_{SC} is in $M_{rr}^i \cap M_{rw}^i \cap M_{wr} \cap M_{ww}$.
- TSO memory model M_{tso} is in $M_{rr}^i \cap M_{rw}^i \cap M_{ww}$ and $M_{tso} \notin M_{wr}$.
- Partial store order (PSO) memory model M_{pso} is in $M_{rr}^i \cap M_{rw}^i$ and $M_{pso} \notin M_{ww} \cup M_{wr}$.
- Relaxed memory order (RMO) M_{rmo} is in $M_{rr}^d \cap M_{rw}$ and $M_{rmo} \notin M_{ww} \cup M_{wr}$. Moreover, note that $M_{rmo} \notin M_{rr}^i$ and $M_{rmo} \notin M_{rw}^i$.
- Alpha memory model M_{alpha} is in M_{rw} and $M_{alpha} \notin M_{rr} \cup M_{wr} \cup M_{ww}$.

Note that these classes do not impose any restrictions on views of different processes, and thus memory models which allow non-atomic stores (like IA-32 [5]) can also be classified under these classes. For example, the IA-32 memory model has a similar classification as TSO.

3.3 Parametrized opacity

We now define the notion of parametrized opacity, i.e., opacity parametrized by a given memory model M. Recall that, intuitively, parametrized opacity requires that (1) every transaction appears as if it took place instantaneously between its first and last operation, and (2) non-transactional operations ensure the requirements specified by the given memory model.

We say that a history h ensures *opacity parametrized by a memory model* $M = (\tau, R)$, if there exists a total order $\ll$ on the set of transactional operations in h and a process view $v \in R(\tau(h))$, such that, for every process $p \in P$, there exists a sequential history s that satisfies the following conditions:

1. s is a permutation of $\tau(h)$,
2. s respects relation $\ll \cup \prec_h \cup v(p)$, and
3. every operation is legal in s.

For example, the history h shown in Figure 3(a) is parametrized opaque with respect to memory model M_{SC} if $v = 1$. This is because: (a) operation 3 reads the value of y as 1 which is written by the transaction of process p_1, (b) operation 1 writes x as 1 before the transaction, and (c) SC requires that p_2 reads x after y. Moreover, h is parametrized opaque with respect to memory model M_{rmo} if $v = 0$ or $v = 1$. This is because RMO allows to reorder reads of different variables. h is parametrized opaque with respect to M_{junk} if $v = 1$ for the same reasons as for M_{SC}. But note that if operation 3 read y as 0, then opacity parametrized by M_{junk} allows operation 6 to read any value $v \in \mathbb{N}$.

4. TM Implementations

In this section, we define a TM implementation $\mathcal{I}$, and when a given TM implementation ensures opacity parametrized by a given memory model M. Intuitively, $\mathcal{I}$ ensures opacity parametrized by M if every history *generated* by $\mathcal{I}$ ensures opacity parametrized by M. We thus need to define what a TM implementation is, and precisely which histories are generated by a given TM implementation.

Instructions. We start by defining hardware primitives that a TM implementation is allowed to use. Let $Addr$ be a set of memory

p_1	p_2
$((\triangleright, \mathsf{start}), 1)$	
$(\langle \mathsf{cas}\ g, 0, 1 \rangle, 1)$	
	$((\triangleright, (\mathsf{rd}, x, 1)), 2)$
$((\triangleleft, \mathsf{start}), 1)$	
	$((\langle \mathsf{load}\ x, 1 \rangle), 2)$
$((\triangleright, (\mathsf{wr}, x, 1)), 3)$	
	$((\triangleleft, (\mathsf{rd}, x, 1)), 2)$
$((\langle \mathsf{store}\ a_x, 1 \rangle), 3)$	
$((\triangleleft, (\mathsf{wr}, x, 1)), 3)$	
$((\triangleright, \mathsf{commit}), 4)$	
$(\langle \mathsf{store}\ g, 0 \rangle, 4)$	
$((\triangleleft, \mathsf{commit}), 4)$	

(a) Trace r

p_1	p_2
$(\mathsf{start}, 1)$	
	$((\mathsf{rd}, x, 1), 2)$
$((\mathsf{wr}, x, 1), 3)$	
$(\mathsf{commit}, 4)$	

(b) History h_1

p_1	p_2
	$((\mathsf{rd}, x, 1), 2)$
$(\mathsf{start}, 1)$	
$((\mathsf{wr}, x, 1), 3)$	
$(\mathsf{commit}, 4)$	

(c) History h_2

Figure 4. A trace and corresponding histories. Notation: (in, k) under column p represents the instruction instance (in, p, k)

addresses. We define the set In of *instructions* as follows, where $a \in Addr$ and $v, v' \in \mathbb{N}$:

$$In ::= \langle \mathsf{load}\ a, v \rangle \mid \langle \mathsf{store}\ a, v \rangle \mid \langle \mathsf{cas}\ a, v, v' \rangle$$

We call the store and CAS instructions as update instructions. An operation of a history corresponds to a sequence of instructions. To know the begin and end points of an operation at the level of hardware, we use two special instructions $\triangleright, \triangleleft$ for each operation. For every operation $o \in \hat{O}$, we denote the *invocation* (instruction) of o as $(\triangleright, o)$, and the *response* (instruction) of o as $(\triangleleft, o)$. Let $\hat{In} = In \cup (\{\triangleright, \triangleleft\} \times \hat{O})$.

Traces. We define an *instruction instance* as (in, p, k), where $in \in \hat{In}$ is an instruction, $p \in P$ is the process that issues the instance, and $k \in \mathbb{N}$ is an operation identifier. Every instruction instance (in, p, k) is said to *correspond* to operation k.

A *trace* $r \in (\hat{In} \times P \times \mathbb{N})^*$ is a sequence of instruction instances. Let $k \in \mathbb{N}$ be an operation identifier. A *complete operation trace* of operation k is a sequence of the form $((\triangleright, o), p, k)$ $(in_1, p, k) \ldots (in_m, p, k)$ $((\triangleleft, o), p, k)$, where $p \in P$ is a process, $o \in \hat{O}$ is an operation, and $in_1 \ldots in_m \in In$ are instructions. An *incomplete operation trace* of operation k is a sequence of the form $((\triangleright, o), p, k)$ $(in_1, p, k) \ldots (in_m, p, k)$, where $p \in P$, $o \in \hat{O}$, and $in_1 \ldots in_m \in In$.

Let r be a trace. Given a process $p \in P$, we denote by $r|_p$ the longest subsequence of instruction instances in r issued by process p. We assume that every trace r satisfies the following property: for every process $p \in P$, the sequence $r|_p$ is a sequence of complete operation traces, possibly ending with an incomplete operation trace. We say that a history h *corresponds* to a trace r if:

1. Given any natural number k, an operation k is in h if and only if there is an instruction instance that corresponds to operation k in r, and

2. Operation k occurs before operation j if some instruction instance that corresponds to operation k occurs in r before some instruction instance that corresponds to operation j.

Intuitively, a history h that corresponds to a trace r represents the logical order of operations in r. If we assigned a point in time to every instruction in r, and every operation in h, then every operation (o, k) in h must be somewhere in between the corresponding invocation instruction $((\triangleright, o), p, k))$ and response instruction $((\triangleleft, o), p, k)$ in r.

For example, consider the trace r shown in Figure 4(a). In r, the start operation issues a CAS instruction to address g to change the value from 0 to 1, and the commit instruction stores value 0 to g.

Histories h_1 and h_2 are two examples of histories that correspond to r.

A trace r is *well-formed* if every history h corresponding to r is well-formed. We assume that every trace is well-formed.

Let r be a trace, and p be a process. A *transaction* T of p in r is a sequence in $r|_p$ of the form $((\triangleright, \mathsf{start}), p, k)\ (in_1, p, k_1) \ldots (in_m, p, k_m)$, where the following conditions are satisfied:

- $in_m \in \{(\triangleleft, \mathsf{commit}), (\triangleleft, \mathsf{abort})\}$, or (in_m, p, k_m) is the last instruction instance of r_p, and

- for all j s.t. $1 \le j < m$, we have $in_j \notin \{(\triangleright, \mathsf{start}), (\triangleleft, \mathsf{commit}), (\triangleleft, \mathsf{abort})\}$

Moreover, T is committed (resp. aborted) if the last instruction instance in T is a response of a commit (resp. abort) operation.

An instance $((\triangleright, o), p, k)$ of an invocation is said to be *transactional* in a trace r if it belongs to some transaction in r. Otherwise, the instance is said to be *non-transactional*. For example, consider the trace r shown in Figure 4(a). The (single) invocation instance of process p_2 is non-transactional, while all invocation instances of process p_1 are transactional in r.

TM implementations. A *TM implementation* $\mathcal{I} = \langle \mathcal{I}_T, \mathcal{I}_N \rangle$ is a pair, where $\mathcal{I}_T : \hat{O} \to 2^{In^*}$ is the implementation for transactional operations, and $\mathcal{I}_N : O \to 2^{In^*}$ is the implementation for non-transactional operations.

Let $\mathcal{I} = \langle \mathcal{I}_T, \mathcal{I}_N \rangle$ be a TM implementation. We say that a complete operation trace $((\triangleright, o), p, k)\ (in_1, p, k) \ldots (in_m, p, k)$ $((\triangleleft, o), p, k)$ is *transactionally* (resp. *non-transactionally*) *generated* by $\mathcal{I}$, if sequence $in_1 \ldots in_m$ is in $\mathcal{I}_T(o)$ (resp. in $\mathcal{I}_N(o)$). We say that an incomplete operation trace $((\triangleright, o), p, k)\ (in_1, p, k) \ldots (in_m, p, k)$ is *transactionally* (resp. *non-transactionally*) *generated* by $\mathcal{I}$, if sequence $in_1 \ldots in_m$ is a prefix of some element in $\mathcal{I}_T(o)$ (resp. in $\mathcal{I}_N(o)$).

Given a trace r and a TM implementation $\mathcal{I}$, we say that r is *generated* by $\mathcal{I}$ if for every transactional (resp. non-transactional) operation k in r, the complete or incomplete operation trace of operation k in r is transactionally (resp. non-transactionally) generated by $\mathcal{I}$.

Instrumentation. We say that a TM implementation $\mathcal{I} = \langle \mathcal{I}_T, \mathcal{I}_N \rangle$ is *uninstrumented* if for every variable x, we have $\mathcal{I}_N(\mathsf{rd}, x, v) = \{\langle \mathsf{load}\ a_x, v \rangle\}$ and $\mathcal{I}_N(\mathsf{wr}, x, v) = \{\langle \mathsf{store}\ a_x, v \rangle\}$, where a_x is the address of variable x. Otherwise, the TM implementation is *instrumented*. Note that these terms refer only to the implementation of non-transactional operations.

Languages. We define the *language* $L(\mathcal{I})$ of a TM implementation as the set of all traces generated by a TM implementation. We define that a TM implementation $\mathcal{I}$ *guarantees* opacity parametrized by a memory model M if, for every trace $r \in L(\mathcal{I})$, there is a history h that corresponds to r such that h ensures opacity parametrized by M.

Note that our definition of a trace does not require that after an instruction of the form $\langle \mathsf{store}\ a_x, v \rangle$, the subsequent load of a_x is of the form $\langle \mathsf{load}\ a_x, v \rangle$. Indeed, the underlying hardware may execute a relaxed memory model (which, in principle, may be different from the programmer's memory model at the level of operations). For example, a programmer may wish to guarantee opacity parametrized by sequential consistency on a hardware with memory model RMO. But for the sake of simplicity, in this following sections, we assume that the underlying hardware guarantees a strong memory model equivalent to linearizability, that is, every instruction is executed to completion when it is issued. Note that our impossibility results hold even when the underlying hardware executes a weaker memory model.

5. Achieving Parametrized Opacity

We use our framework to investigate the inherent cost of achieving opacity parametrized by a memory model.

5.1 Uninstrumented TM implementations

We first study uninstrumented TM implementations. We show that for most of the practical memory models, uninstrumented TM implementations cannot achieve parametrized opacity. Moreover, we show that even to achieve opacity parametrized by very relaxed memory models, it is required that transactional write operations are implemented as expensive compare-and-swap instructions.

We start with a lemma which states that if a committed transaction consists of a write operation, then the transaction must consist of a store or a compare-and-swap instruction.

Lemma 1. For every memory model M, an uninstrumented TM implementation $\mathcal{I}$ guarantees parametrized opacity only if for all traces $r \in L(\mathcal{I})$, for every committing transaction T in r, if there is an operation (wr, x, v) in T, then the transaction T in r consists of an update instruction to a_x with value v.

Proof. Let the value of a_x be initially 0. Consider a trace r with a single process p (shown in Figure 5(a)). Let T be a committed transaction of process p in r, such that T consists of an operation $o_1 = (\mathsf{wr}, x, v)$. Let r consist of a non-transactional operation $o_2 = (\mathsf{rd}, x, v')$, such that the invocation of o_2 occurs after the last instruction of T, that is, the response of the commit operation. As the TM implementation $\mathcal{I}$ is uninstrumented, we have $\mathcal{I}_N(\mathsf{rd}, x, v') = \{\langle \mathsf{load}\ a_x, v' \rangle\}$. Note that as r has a single process, there is only one history h corresponding to r. By definition of $\prec_h$, we know that $o_1 \prec_h o_2$. Thus, to guarantee parametrized opacity, we must have $v = v'$. Thus, T must issue an update instruction to a_x with value v. $\square$

Using the above lemma, we now prove that for memory models which restrict the order of some pair of read or write operations to different variables, parametrized opacity cannot be achieved.

Theorem 1. Given a memory model M such that $M \in (M_{rr} \cup M_{rw} \cup M_{wr} \cup M_{ww})$, there does not exist an uninstrumented TM implementation that guarantees opacity parametrized with respect to M.

Proof. We prove the theorem in four parts, where each part corresponds to one of the cases of ordering restrictions between read and write operations. In every case, we construct a trace r with two processes p_1 and p_2, where p_1 executes a transaction T, and p_2 issues non-transactional (and possibly transactional) operations in such a way that for every history corresponding to r, there is no sequential history s which respects the restriction posed by the memory model, and at the same time, every operation in s is legal. In every case, a_x and a_y are initialized to 0.

Case 1. Let $M \in M_{rr}$. That is, M does not allow to reorder two read operations to different variables. The trace r is shown in Figure 5(b). Let T consist of operations (wr, x, v_1) and (wr, y, v_2). From Lemma 1, we know that T updates addresses a_x and a_y with values v_1 and v_2 respectively. Without loss of generality, we assume that a_x is updated before a_y.[2] Let the trace r consist of two non-transactional operations (rd, x, v_3) and (rd, y, v_4) issued by process p_2 (with identifiers j and k respectively). As $\mathcal{I}$ is uninstrumented, the non-transactional read operations are implemented as load instructions. Let the two load instructions $\langle \mathsf{load}\ a, v_3 \rangle$ and $\langle \mathsf{load}\ a, v_4 \rangle$ of process p_2 execute between the updates of a_x and

---

[2] Figure 5(b) shows the updates as part of the commit operation. In general, the updates can happen anywhere during the transaction, but always as two separate instructions.

Figure 5. Different traces r constructed in Lemma 1, Theorem 1, and Theorem 2. An $\langle \mathsf{update}\ a, v \rangle$ instruction denotes a store or a successful cas instruction to address a with value v. We omit the instruction idenfiers. We use $\ldots$ as a shorthand for a sequence of instructions.

a_y. Thus, we have $v_3 = v_1$ and $v_4 = 0$. Note that if T is an aborted transaction in r, then there is no history h corresponding to r such that h satisfies opacity parametrized by M. This is because operation j observes the update to a_x. Thus, let T be a committed transaction in r. Consider an arbitrary history h corresponding to r. By definition of M_{rr}, every view function $v \in R(h)$ requires that $(j, k) \in v(p)$ for all $p \in P$. On the other hand, legality requires operation j to appear after T, and operation k to appear before T. Thus, there does not exist a sequential history which satisfies conditions 2 and 3 at the same time.

Case 2. Let $M \in M_{wr}$. That is, M does not allow to reorder a write followed by a read operation to a different variable. The trace r is shown in Figure 5(c). Let T consist of operations $o_1 = (\text{rd}, x, v_1)$ and $o_2 = (\text{wr}, y, v_2)$. From Lemma 1, we know that T updates address a_y with value v_2. Let r consist of two non-transactional operations of p_2 in the order (wr, x, v_3) (with identifier j) followed by $(\text{rd}, y, 0)$ (with identifier k), such that $v_3 \neq v_1$. As $\mathcal{I}$ is uninstrumented, p_2 executes $\langle \text{store } a_x, v_3 \rangle$ followed by $\langle \text{load } a_y, 0 \rangle$. Let these two instructions execute after the response of o_1 and immediately before the update of a_y by process p_1.[3] Note that as p_2 does not change the value of a_y, the use of CAS instruction by p_1 to update a_y will be successful. Moreover, the key point to note here is that once p_1 updates a_y with value v_2, the transaction T cannot abort. This is because r can be extended to trace r', where a non-transactional read by process p_2 executes $\langle \text{load } a_y, v_2 \rangle$. If T is an aborted transaction, then no history corresponding to r' is parametrized opaque with respect to M. Thus, T is a committed transaction. Consider an arbitrary history h corresponding to r. Legality requires that operation k occurs before T and operation j occurs after T. By definition of M_{wr}, every view $v \in R(h)$ requires that $(j, k) \in v(p)$ for all $p \in P$. Thus, h does not ensure opacity parametrized by M.

Case 3. Let $M \in M_{rw}$. That is, M does not allow to reorder a read followed by a write operation to a different variable. The trace r is shown in Figure 5(d). Let T consist of operations $o_1 = (\text{wr}, x, v_1)$ and $o_2 = (\text{wr}, y, v_2)$. From Lemma 1, we know that T updates addresses a_x and a_y with values v_3 and v_4. Without loss of generality, we assume that a_x is updated before a_y. Process p_2 issues the following operations non-transactionally in the order: (rd, x, v_3), (wr, y, v_4), $(\text{wr}, y, 0)$. As $\mathcal{I}$ is uninstrumented, these three operations are executed as: $\langle \text{load } a_x, v_3 \rangle$, $\langle \text{store } a_y, v_4 \rangle$, and $\langle \text{store } a_y, 0 \rangle$. Let r be such that these three instructions occur immediately before the update of a_y with value v_2. As p_2 changes and restores the value of a_y to 0, process p_1 does not observe the change, and thus, the update of a_y with value v_2 is successful. As in case 2, T cannot be an aborting transaction, once T updates a_y. We now extend r as follows: after the response of the commit operation of T, let p_2 execute a transaction T' followed by two non-transactional operations: (rd, x, v_5), (rd, y, v_6). Note that $v_5 = v_1$ and $v_6 = v_2$. Consider an arbitrary history h corresponding to r. Legality requires that the first non-transactional read by process p_2 occurs after T, and the two non-transactional writes of p_2 occur before T. This contradicts the expected view order for a memory model $M \in M_{rw}$. Thus, the history h does not ensure opacity parametrized by M. We also show that T cannot update a_y with value 0 due to the following argument: consider a trace r' which is identical to r, except that in r', process p_2 does not issue the two non-transactional writes. At the point where p_1 updates a_y, it cannot observe a difference from trace r (as process p_2 uses two store instructions for the non-transactional write operations). Consider an arbitrary history h corresponding to r'. Legality requires

that the last two non-transactional operations by p_2 see value v_1 and v_2. Thus, p_1 cannot update a_y to 0 in transaction T.

Case 4. Let $M \in M_{ww}$. That is, M does not allow to reorder two write operations to different variables. The trace r is similar to the one for case 3, shown in Figure 5(d), except that p_1 has two read operations, and the first operation of p_2 is a write instead of a read operation. Let T consist of four operations: (rd, x, v_1), (rd, y, v_2), (wr, x, v_3) and (wr, y, v_4). As in case 3, we know that T updates addresses a_x and a_y with values v_3 and v_4, and we assume that a_x is updated before a_y. Let the trace r consist of three non-transactional operations of process p_2 in the order (wr, x, v_5), (wr, y, v_6), $(\text{wr}, y, 0)$. As $\mathcal{I}$ is uninstrumented, the non-transactional write operations are implemented as store instructions. Let the three store instructions: $\langle \text{store } a_x, v_5 \rangle$, $\langle \text{store } a_y, v_6 \rangle$, and $\langle \text{store } a_y, 0 \rangle$ by process p_2 occur immediately before the update of a_y with value v_4 by process p_1 in trace r. As in case 2, T cannot be an aborting transaction, after it updates a_y. We can extend the trace r exactly as in case 3 now, and show a contradiction between the legality and the view order required by a memory model $M \in M_{ww}$. Thus, for every history h corresponding to r, we know that h does not ensure opacity parametrized by M.
$\square$

We saw in classification of memory models (Section 3.2) that most of the practical memory models do restrict some order of operations. Thus, Theorem 1 gives an intuition that without instrumentation, it is not possible to achieve opacity parametrized by practical memory models. We now show that for an idealized memory model, which allows reordering all operations to different variables, achieving parametrized opacity is still expensive: a TM implementation must use cas instruction within a transaction to update every variable which is read and written by the transaction.

Theorem 2. For a memory model M, an uninstrumented TM implementation $\mathcal{I}$ guarantees parametrized opacity with respect to M only if for all traces $r \in L(\mathcal{I})$, for all variables x, if a committed transaction T consists of operations (rd, x, v) and (wr, x, v'), then T consists of a $\langle \text{cas } x, v, v' \rangle$ instruction.

Proof. Consider a trace r with two processes p_1 and p_2 as shown in Figure 5(e). Let T be a transaction of process p_1 in r, such that T consists of operations (rd, x, v) and (wr, x, v'). By Lemma 1, we know that T updates a_x with value v' using either a store or a compare-and-swap instruction. Suppose T changes the value of a_x using a store instruction. Let r consist of two non-transactional operations (wr, x, v_1) followed by (rd, x, v_2) of process p_2. As $\mathcal{I}$ is uninstrumented, we know that a read operation is implemented as a load, and a write as a store. Consider the trace r where $\langle \text{store } a_x, v_1 \rangle$ instruction of process p_2 occurs immediately before $\langle \text{store } a_x, v' \rangle$ instruction of process p_1 and the instruction $\langle \text{load } x, v_2 \rangle$ occurs immediately after $\langle \text{store } x, v' \rangle$, such that we get $v_2 = v'$. For a corresponding history of r to be parametrized opaque with respect to M, the transaction T has to be a committed transaction. After the response of the commit of T, let there be an empty transaction T' of process p_2 in r followed by a non-transactional read of x, with a load instruction $\langle \text{load } x, v_3 \rangle$. Note that we have $v_3 = v'$. We argue that irrespective of the memory model M, there does not exist a history h corresponding to r such that h ensures opacity parametrized by M. This is because the operation (wr, x, v_1) of process p_1 can appear neither before T (as the read of v in T returns 0), nor after T (as the read after T' returns v'). Thus, an uninstrumented TM implementation has to use a cas instruction to update a variable in a transaction T, if T consists of both read and write operations to T.

$\square$

Now, we establish a complementary result to Theorem 1. We show that with an idealized memory model which relaxes the order

---

[3] For simplicity, we assume that the transaction loads the value of a_x before the response of o_1. The argument holds in general, as the value of a_x is loaded before the update of a_y.

```
On transaction start:
lg := g
while (lg ≠ p) do
    ⟨cas g, lg, p⟩
endwhile
return

On transaction write of x with value v':
issue a transactional read of x
if x ∈ writeset(p) then
    update (x, v) to (x, v') in writeset(p)
    return
endif
add (x, v) to writeset(p)


On transaction read of x:
if x ∉ readset(p) then
    ⟨load a_x, v⟩
    add (x, v) to readset(p)
    return v
endif
return v s.t. (x, v) ∈ readset(p)

On transaction abort:
empty readset(p)
empty writeset(p)


On transaction commit:
while writeset(p) is not empty
    pick and remove (x, v') from writeset(p)
    pick and remove (x, v) from readset(p)
    ⟨cas x, v, v'⟩
endwhile
⟨store g, 0⟩
return
```

Figure 6. A global lock based TM implementation

of all operations to different variables, it is possible to obtain parametrized opacity with an uninstrumented TM implementation.

Theorem 3. Given a memory model $M \notin (M_{rr} \cup M_{rw} \cup M_{wr} \cup M_{ww})$, there exists an uninstrumented TM implementation which guarantees parametrized opacity with respect to M.

Proof. Consider an uninstrumented global lock based TM implementation $\mathcal{I}$ described in pseudo code in Figure 6. $\mathcal{I}$ acquires a global lock g during the start of each transaction, and releases the lock (using $\langle\text{store } g, 0\rangle$) immediately before the response of the commit or abort of the transaction. Moreover, for every variable x, every transaction T loads the value of a_x only at the first read or write operation to x within T (if such an operation exists in T), and if T writes to x, then $\mathcal{I}$ uses a CAS instruction to update a_x in T. Consider an arbitrary trace $r \in L(\mathcal{I})$. Consider a history h corresponding to the trace r obtained by choosing the following logical points of execution of an operation within an operation trace: start operation at the successful cas instruction, commit and abort operations at $\langle\text{store } g, 0\rangle$ instruction, every non-transactional read at the load instruction, every non-transactional write at the store instruction, and every transactional read or write at the invocation. First, we note that the history h obtained is such that for every pair T, T' of transactions in h, either $T \prec_h T'$ or $T' \prec_h T$. Now, consider a variable x and a transaction T in h. Let $k_1 \ldots k_n$ be a sequence of non-transactional operation instances to x in h, which occur between the first and last operations of the transaction T. We create a sequential history s from h by repeating the following two steps for all variables:
Step 1. Consider an non-transactional write operation k_i for some $1 \le i \le n$. If there is no load or cas of x in T after the store instruction with identifier k_i in r, we place all operations $k_i \ldots k_n$ after the transaction T in s. For all other non-transactional write operations k_i, we we place all operations $k_1 \ldots k_i$ before the transaction T in s. We now have no non-transactional write operation to x between the first and last operations of T.
Step 2. For all remaining non-transactional read operations o, we place o before T in s if the load instruction corresponding to o precedes the update to x by T in r, and after T otherwise.

Note that as the memory model freely allows to reorder instructions to independent variables, we can move operations of different variables freely with respect to each other. Moreover, the two steps do not change the order of non-transactional operations of a process to a variable. $\square$

5.2 Instrumented TM implementations

In practice, TM implementations use instrumentation of non-transactional operations to achieve parametrized opacity. We now investigate instrumented TM implementations. Generally, a his-

tory contains more read operations than write operations. So, it is worthwhile to study whether we can achieve parametrized opacity by just instrumenting the non-transactional write operations and leaving the non-transactional read operations uninstrumented.

We first show that it is indeed possible to achieve parametrized opacity with uninstrumented reads for a class of memory models that allow to reorder read operations. The construction implements non-transactional write operations as single operation transactions.

Theorem 4. There exists a TM implementation with uninstrumented reads that guarantees parametrized opacity for memory models $M \notin M_{rr}$.

Proof. Consider a global lock based TM implementation $\mathcal{I}$ that treats transactional operations as in the proof of Theorem 3, and shown in Figure 6. Moreover, $\mathcal{I}$ implements a non-transactional write operation (wr, x, v) as: acquire the global lock g using cas (as in transaction start), followed by the instruction $\langle\text{store } x, v\rangle$, followed by $\langle\text{store } g, 0\rangle$. Intuitively, $\mathcal{I}$ treats every non-transactional write operation as a transaction in itself. $\mathcal{I}$ implements non-transactional read operations as load instructions (no instrumentation).

Consider an arbitrary trace $r \in L(\mathcal{I})$. Note that there exists a history h corresponding to r (obtained using the logical points as in Theorem 3) such that no non-transactional write occurs between the start and commit/abort of a transaction, and for every pair T, T' of transactions, either $T \prec_h T'$ or $T' \prec_h T$. Consider a variable x. Given a transaction T in h, let $k_1 \ldots k_n$ be the identifiers of non-transactional operation instances in h, which occur between the first and last operation of the transaction T. We create a sequential history s from h as in Theorem 3: for all non-transactional read operations o to x, we place o before T in s if the load instruction corresponding to o precedes the cas instruction to x by T in r, and after T otherwise. Note that as the memory model freely allows to reorder read operations to different variables, we can move reads of different variables freely with respect to each other. $\square$

Note that this construction implements non-transactional write operations using a cas instruction to acquire locks. Thus, a non-transactional write operation is implemented in an expensive manner. Moreover, a non-transactional write may may take an arbitrarily long time to complete in this construction. Formally speaking, we have $(\langle\text{load } g, 0\rangle)^* \in \mathcal{I}(o)$ if o is a write operation. We would like to create a TM implementation with inexpensive instrumentation.

Given a TM implementation $\mathcal{I}$, we say that a non-transactional write operation has *constant-time instrumentation* in $\mathcal{I}$ if there exists a constant c such that for every operation $o \in (\text{wr} \times Obj \times \mathbb{N})$, every instruction sequence in $\mathcal{I}_N(o)$ has length at most c. We now prove an interesting result, where we present a TM implementation,

with no instrumentation on reads and constant-time instrumentation on writes, which guarantees opacity parametrized by a memory models that allow reordering read/write followed by a read of a different variable.

Theorem 5. There exists a TM implementation $\mathcal{I}$ with constant-time instrumentation of writes and no instrumentation of reads such that $\mathcal{I}$ guarantees opacity parametrized by a memory model M, where $M \notin M_{rr} \cup M_{wr}$.

Proof. We build a TM implementation $\mathcal{I}$ which uses a global lock to execute transactions (as in Theorem 3 and 4). Moreover, $\mathcal{I}$ uses a version number per process. When a process p issues a non-transactional write operation , p increments its version number, and writes the value, the process id, and the version number using a store instruction. $\mathcal{I}$ does not use instrumentation for non-transactional read operations. Now, we prove that $\mathcal{I}$ guarantees opacity parametrized by M, where $M \notin M_{rr} \cup M_{wr}$. Consider an arbitrary trace $r \in L(\mathcal{I})$. Consider a history h corresponding to the trace r obtained by choosing the logical points as in Theorem 3. We know thhat for every two transactions T, T' in h, we have $T \prec_h T'$ or $T' \prec_h T$. Consider an arbitrary transaction T in h and a variable x. Let $k_1 \ldots k_m$ be the identifiers of the non-transactional operations to x that occur between the first and last operation of T in h. Note that as $\mathcal{I}$ uses `cas` instruction for variables which are written in a transaction, no non-transactional write to x stores to a_x between a load and an update of a_x in T. We thus can obtain a sequential history from h by repeating the following for all variables x and for all transactions in h: for a non-transactional operation k_i to x such that operation k_i occurs between the start and end of transaction T, if the corresponding store (or load) to a_x occurs after the update of a_x in T, we place operation k_i after T in s, otherwise we place k_i before T in s. We place remaining non-transactional read operations before T. Note that it cannot happen that the corresponding store of a non-transactional write operation k_i to a_x occurs after a load and before an update to a_x.

If a process issues non-transactional reads of x and y, such that their corresponding loads occur between the updates by a transaction, we need to reorder the non-transactional reads for legality. Thus, we want $M \notin M_{rr}$. Similarly, note that if a process issues a non-transactional write of x followed by a read of y, such that the corresponding store to a_x and load of a_y occur between the updates to a_x and a_y, we need to reorder the non-transactional access for legality. Thus, we want $M \notin M_{wr}$. But interestingly, we built $\mathcal{I}$ in such a way that it first adds all variables to be updated in the writeset. Only then, $\mathcal{I}$ starts updating the variables using `cas` instruction. Hence, if a process issues a non-transactional read which loads a value updated by T, we know that every following non-transactional write can also occur after T. Thus, we do not require that $M \notin M_{rw}$. Similarly, if a process issues two non-transactional writes, then we know that either both the writes occur before T or after T. This implies that $\mathcal{I}$ guarantees parametrized opacity even for memory models which restrict the order of a read/write followed by a write to a different variable, but allow reordering read/write followed by a read to a different variable. □

Relaxed memory models like Alpha [28] are neither in M_{rr}, nor in M_{wr}. So, this construction provides an inexpensive way to guarantee opacity parametrized by memory models like Alpha. Moreover, memory models like RMO and Java are in M_{rr}^d, that is, they do not allow data dependent reads to reorder. But, these memory models do allow independent reads or control-dependent reads to reorder. This implies that if we use special synchronization

for data-dependent reads, we can use the result of Theorem 5 for a vast class of memory models.[4]

6. Discussion

We first discuss the practical implications of our work. Then, we discuss how our framework can be extended to formalize weaker notions of correctness, like single global lock atomicity.

6.1 Impact on practical TM implementations

The core contribution of our theoretical framework is to provide TM designers with a correctness property that exploits the inherent relaxations of the underlying memory model. For example, most memory models, like TSO, PSO, RMO, Alpha, and Java allow to reorder a write operation followed by a read/write operation. On the other hand, the only TM implementation [27] we know of which guarantees strong atomicity, indeed guarantees parametrized opacity with respect to sequential consistency. Given that a programmer expects behavior of non-transactional operations within the scope of behaviors under the given memory model, one can build a more efficient implementation which guarantees opacity with respect to the programmer's memory model.

A TM implementation for strong atomicity. We give a brief description of the TM implementation by Shpeisman [27]. Intuitively, the TM implementation maintains a transactional record for every variable. The record implements the locking discipline. Each record can be in four possible states: shared, exclusive, exclusive anonymous, and private. The shared state allows concurrent access to a number of reading transactions. The exclusive state allows read-write access to a single transaction. The exclusive anonymous state allows a process to have read-write access non-transactionally. Variables in private state are visible to only one process. In this TM implementation, a non-transactional read needs to check whether the variable is being written concurrently by a transaction. Moreover, non-transactional writes have to gain ownership (go to exclusive anonymous state) of a variable before performing the write.

We argue that depending upon the memory model M, this design can be optimized to guarantee opacity parametrized by M. Carrying our results from the previous section, if M allows to reorder a read/write with a following read of a different variables, then the non-transactional read operations can be uninstrumented.

6.2 Weaker notions of correctness

Apart from the intuitive strong correctness property, parametrized opacity, which requires that transactions be isolated from other transactions and non-transactional operations, weaker notions of correctness have also been considered in the literature. Some claim that transactions should behave as global locks, thus isolating transactions from other transactions, but not from non-transactional operations. This yields the correctness property called single global lock atomicity (SGLA). We now formalize SGLA and show that SGLA is strictly weaker than parametrized opacity for every memory model M. We also show that the impossibility result given in Theorem 1 for parametrized opacity does not hold for SGLA.

SGLA. We slightly modify the framework we developed in Section 2 and 3. SGLA requires a weaker notion of sequential history, where transactions execute sequentially, but non-transactional operations may be concurrent with transactions. A history h is *trans-*

---

[4] In practice, memory models like Java enforce strict ordering of operations marked with the *volatile* keyword. Indeed, these volatile accesses have to be treated differently than simple non-transactional accesses. For example, a volatile access may be considered as a single operation transaction. Volatile accesses are few and inherently expensive, even in non-transactional programs.

actionally sequential if for every transaction T in h, every operation instance between the start operation instance of T and the last operation instance of T in h is either an operation instance of T, or a non-transactional operation instance.

The behavior of locks in terms of reorderings with other operations of the process depends upon the memory model. Thus, we cannot enforce a strict ordering between non-transactional operations and following or preceding transactions. Instead of defining a precedence relation between operations in a history using $\prec_h$, we now define the partial order using the memory model. Recall that when we talk of a memory model for parametrized opacity, a view does not enforce an order on start, commit, and abort operations. Whereas, when we formalize SGLA, a start operation has the memory model semantics of a lock operation, whereas commit and abort operations have the memory model semantics of an unlock operation. Thus, a memory model for SGLA can be seen as an extension of a memory model for parametrized opacity.

We say that a history h ensures *SGLA parametrized by a memory model $M = (\tau, R)$*, if there exists an ordering function $v \in R(\tau(h))$, such that, for every process $p \in P$, there exists a transactionally sequential history s such that

1. s is a permutation of $\tau(h)$
2. s respects the total order $v(p)$, and
3. every operation in legal in s.

Note that the general definition of a memory model for SGLA may allow counterintuitive results, when a non-transactional access following a transaction is allowed to precede the transaction in a view, or two processes may view two transactions to execute in different orders. Based on this intuition, in the framework for SGLA, we define that $M' = (\tau, R')$ is a *well-formed extension* of $M = (\tau, R)$ if: (i) for every history h, for every view $v \in R'(h)$, for every pair $(o, p_1, i), (o', p_2, j)$ of operation instances such that o and o' belong to {start,commit,abort}, for all processes $p, p' \in P$, if $(i, j) \in v(p)$, then $(i, j) \in v(p')$, (ii) for every history h and for every process $p \in P$, if a non-transactional operation j of p precedes a transaction T of p in h, then there exists a $v \in R(h)$, such that for every process $p' \in P$, we have $(k, j) \in v$, where operation k is the start operation of transaction T. (iii) for every history h and for every process $p \in P$, if a non-transactional operation j of p precedes a transaction T of p in h, then there does not exist $v \in R(h)$, such that for some process $p' \in P$, we have $(k, j) \in v$, where operation k is the commit or abort operation of transaction T.

Theorem 6. *Given a history h and a memory model M, if h ensures parametrized opacity with respect to M, then h ensures SGLA with respect to M', where M' is a well-formed extension of M.*

Proof idea. Consider a history h. As h ensures parametrized opacity with respect to M, we know that there exists a view function $v \in R(h)$ such that for every process $p \in P$, there exists a sequential history s such that s satisfies the three requirements for parametrized opacity. Now, we note that a sequential history is, by definition, transactionally sequential. Moreover, as s respects the partial order $\ll \cup \prec_h$, we know that s respects the view function $v' \in R'(h)$, where the order of a transactional operation with respect to all other operations of a process is maintained in v'. We know that such a v' exists, as M' is a well-formed extension of M. This implies that h ensures SGLA with respect to M'. $\square$

Theorem 7. *Given a memory model M, there exists an uninstrumented TM implementation that guarantees SGLA parametrized by M.*

Proof idea. Consider an uninstrumented TM implementation $\mathcal{I}$ which executes every transaction using a global lock as shown in Figure 6 and explained in Theorem 3. Consider an arbitrary trace $r \in L(\mathcal{I})$. Note that there exists a history h corresponding to r (as explained in Theorem 3) such that for every pair T, T' of transactions, either $T \prec_h T'$ or $T' \prec_h T$. This means that h is transactionally sequential. Moreover, every operation in h is legal. Thus, h ensures SGLA parametrized by M. $\square$

7. Related Work

Various formalisms for memory models have been proposed in the literature [3, 4, 11, 24, 25]. Most of these formalisms are tailored to capture the intricacies of specific memory models. Our reason for building a formalism for memory models was to obtain a general framework, with the focus on classification of memory models on the basis of reordering of instructions they allow.

The interaction of transactions with non-transactional operations has been widely studied. The study was pioneered by Grossman et al. [8], where the authors raised several issues that need to be tackled in order to create TM implementations that handle non-transactional operations properly. The authors express the correctness property using sample executions, and thus it lacks a formal specification. Work by Scott et al. [29] focuses on providing a set of rules that should hold irrespective of the memory model. This work is restricted to memory models that do not allow out-of-thin-air values. Menon et al. [20] define correctness by mapping transactions to critical sections, thus providing an intuitive definition of single global lock atomicity. Type systems and operational semantics for transactional programs with non-transactional accesses have been proposed by Abadi et al. [1] and Grossman et al. [21]. We believe that while each of the above approaches sheds light on the intricacies involved in transactional programs with non-transactional operations, none provides a formal specification of the correctness property in the presence of relaxed memory models.

8. Conclusion

We formalized the correctness property of opacity parametrized by a memory model. Intuitively, parametrized opacity requires two things. Firstly, transactions are isolated from other transactions and non-transactional operations. Secondly, the behavior of non-transactional operations is governed by the underlying memory model. We used our formalism to prove several results on achieving parametrized opacity with instrumented or uninstrumented TM implementations under different memory models. In particular, we show that for most memory models, parametrized opacity cannot be achieved without instrumenting non-transactional operations. On the positive side, we show that with constant-time instrumentation for writes, and no instrumentation for reads, it is indeed possible to achieve opacity parametrized by a significant class of memory models. The practical relevance of our definition lies in the fact that while creating TM implementations, one may use the relaxations provided by the memory model to create more efficient TM implementations. In future, we plan to use our formal framework to verify different TM implementations in the literature.

References

[1] Martín Abadi, Andrew Birrell, Tim Harris, and Michael Isard. Semantics of transactional memory and automatic mutual exclusion. In *POPL*, pages 63–74, 2008.

[2] Sarita V. Adve and Kourosh Gharachorloo. Shared memory consistency models: A tutorial. *IEEE Computer*, 29(12):66–76, 1996.

[3] Gérard Boudol and Gustavo Petri. Relaxed memory models: An operational approach. In *POPL*, pages 392–403, 2009.

[4] Sebastian Burckhardt, Madanlal Musuvathi, and Vasu Singh. Verifying compiler transformations for concurrent programs. Technical Report MSR-TR-2008-171, Microsoft Research, 2008.

[5] Intel Corporation. *Intel 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*. 2006.

[6] Luke Dalessandro and Michael Scott. Strong isolation is a weak idea. In *TRANSACT*, 2009.

[7] D. Dice, O. Shalev, and N. Shavit. Transactional locking II. In *DISC*, pages 194–208. Springer, 2006.

[8] Dan Grossman, Jeremy Manson, and William Pugh. What do high-level memory models mean for transactions? In *Memory System Performance and Correctness*, 2006.

[9] Rachid Guerraoui, Thomas A. Henzinger, Barbara Jobstmann, and Vasu Singh. Model checking transactional memories. In *PLDI*, pages 372–382. ACM, 2008.

[10] Rachid Guerraoui, Thomas A. Henzinger, and Vasu Singh. Nondeterminism and completeness in transactional memories. In *CONCUR*. Springer, 2008.

[11] Rachid Guerraoui, Thomas A. Henzinger, and Vasu Singh. Software transactional memory on relaxed memory models. In *CAV*, pages 321–336, 2009.

[12] Rachid Guerraoui and Michał Kapałka. On the correctness of transactional memory. In *PPoPP*, 2008.

[13] M. Herlihy, V. Luchangco, M. Moir, and W. N. Scherer. Software transactional memory for dynamic-sized data structures. In *PODC*, pages 92–101, 2003.

[14] M. Herlihy and J. E. B. Moss. Transactional memory: Architectural support for lock-free data structures. In *ISCA*, pages 289–300. ACM Press, 1993.

[15] Maurice Herlihy and Eric Koskinen. Transactional boosting: A methodology for highly-concurrent transactional objects. In *PPoPP*, 2008.

[16] Maurice Herlihy and Jeannette M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems*, 12(3), 1990.

[17] J. R. Larus and R. Rajwar. *Transactional Memory*. Synthesis Lectures on Computer Architecture. Morgan & Claypool, 2007.

[18] Jeremy Manson, William Pugh, and Sarita V. Adve. The Java memory model. In *POPL*, pages 378–391. ACM, 2005.

[19] Milo M. K. Martin, Colin Blundell, and E. Lewis. Subtleties of transactional memory atomicity semantics. *Computer Architecture Letters*, 5(2), 2006.

[20] Vijay Menon, Steven Balensiefer, Tatiana Shpeisman, Ali-Reza Adl-Tabatabai, Richard L. Hudson, Bratin Saha, and Adam Welc. Practical weak-atomicity semantics for Java STM. In *SPAA*, 2008.

[21] Katherine F. Moore and Dan Grossman. High-level small-step operational semantics for transactions. In *POPL*, pages 51–62. ACM, 2008.

[22] C. H. Papadimitriou. The serializability of concurrent database updates. *Journal of the ACM*, pages 631–653, 1979.

[23] Bratin Saha, Ali-Reza Adl-Tabatabai, Richard L. Hudson, Chi Cao Minh, and Ben Hertzberg. McRT-STM: A high performance software transactional memory system for a multi-core runtime. In *PPoPP*, pages 187–197. ACM, 2006.

[24] Vijay A. Saraswat, Radha Jagadeesan, Maged Michael, and Christoph von Praun. A theory of memory models. In *PPoPP*, pages 161–172, New York, NY, USA, 2007. ACM.

[25] Susmit Sarkar, Peter Sewell, Francesco Zappa Nardelli, Scott Owens, Tom Ridge, Thomas Braibant, Magnus O. Myreen, and Jade Alglave. The semantics of x86-CC multiprocessor machine code. In *POPL*, pages 379–391, 2009.

[26] N. Shavit and D. Touitou. Software transactional memory. In *PODC*, pages 204–213, 1995.

[27] Tatiana Shpeisman, Vijay Menon, Ali-Reza Adl-Tabatabai, Steven Balensiefer, Dan Grossman, Richard L. Hudson, Katherine F. Moore, and Bratin Saha. Enforcing isolation and ordering in STM. In *PLDI*, pages 78–88. ACM, 2007.

[28] Richard L. Sites, editor. *Alpha Architecture Reference Manual*. Digital Press, 2002.

[29] Michael F. Spear, Luke Dalessandro, Virendra J. Marathe, and Michael L. Scott. Ordering-based semantics for software transactional memory. In *OPODIS*, 2008.

[30] Michael F. Spear, Virendra J. Marathe, Luke Dalessandro, and Michael L. Scott. Privatization techniques for software transactional memory. In *PODC*, pages 338–339, 2007.

[31] D. Weaver and T. Germond, editors. *The SPARC Architecture Manual (version 9)*. Prentice-Hall, Inc., 1994.

[32] William E. Weihl. Local atomicity properties: Modular concurrency control for abstract data types. *ACM Transactions on Programming Languages and Systems*, 11(2), 1989.