The Semantics of Transactions and Weak Memory in x86, Power, ARM, and C++

Nathan Chong
ARM Ltd., UK

Tyler Sorensen
Imperial College London, UK

John Wickerson
Imperial College London, UK

Abstract

Weak memory models provide a complex, system-centric semantics for concurrent programs, while transactional memory (TM) provides a simpler, programmer-centric semantics. Both have been studied in detail, but their *combined* semantics is not well understood. This is problematic because such widely-used architectures and languages as x86, Power, and C++ all support TM, and all have weak memory models.

Our work aims to clarify the interplay between weak memory and TM by extending existing axiomatic weak memory models (x86, Power, ARMv8, and C++) with new rules for TM. Our formal models are backed by automated tooling that enables (1) the synthesis of tests for validating our models against existing implementations and (2) the model-checking of TM-related transformations, such as lock elision and compiling C++ transactions to hardware. A key finding is that a proposed TM extension to ARMv8 currently being considered within ARM Research is incompatible with lock elision without sacrificing portability or performance.

CCS Concepts • Theory of computation → Parallel computing models; Program semantics;

Keywords Shared Memory Concurrency, Weak Memory, Transactional Memory, Program Synthesis

ACM Reference Format:

Nathan Chong, Tyler Sorensen, and John Wickerson. 2018. The Semantics of Transactions and Weak Memory in x86, Power, ARM, and C++. In *Proceedings of 39th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'18)*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3192366.3192373>

1 Introduction

Transactional memory [28] (TM) is a concurrent programming abstraction that promises scalable performance without programmer pain. The programmer gathers instructions into *transactions*, and the system guarantees that each appears to be performed entirely and instantaneously, or not

at all. To achieve this, a typical TM system tracks each transaction's memory accesses, and if it detects a conflict (i.e., another thread concurrently accessing the same location, at least one access being a write), resolves it by aborting the transaction and rolling back its changes.

1.1 Motivating Example: Lock Elision in ARMv8

One important use-case of TM is *lock elision* [22, 46], in which the lock/unlock methods of a mutex are skipped and the critical region (CR) is instead executed speculatively inside a transaction. If two CRs do not conflict, this method allows them to be executed simultaneously, rather than serially. If a conflict is detected, the transaction is rolled back and the system resorts to acquiring the mutex as usual.

Lock elision may not apply to all CRs, so an implementation must ensure mutual exclusion between transactional and non-transactional CRs. This is typically done by starting each transactional CR with a read of the lock variable (and self-aborting if it is taken) [31, §16.2.1]. If the mutex is subsequently acquired by a non-transactional CR then the TM system will detect a conflict on the lock variable and abort the transactional CR.

Thus, reasoning about lock elision requires a concurrency model that accounts for both modes, transactional and non-transactional. In particular, systems with memory models weaker than *sequential consistency* (SC) [39] must ensure that the non-transactional lock/unlock methods synchronise sufficiently with transactions to provide mutual exclusion.

In their seminal paper introducing lock elision, *Rajwar and Goodman* argued that “correctness is guaranteed without any dependence on memory ordering” [46, §9]. In fact, by drawing on a decade of weak memory formalisations [5, 24, 45] and by extending state-of-the-art tools [4, 40, 55], we show it is straightforward to contradict this claim *automatically*.

Example 1.1 (Lock elision is unsound under ARMv8). Consider the program below, in which two threads use CRs to update a shared location x .

Initially: $[X0] = x = 0$	
lock() LDR W5, [X0] ADD W5, W5, #2 STR W5, [X0] unlock()	lock() MOV W7, #1 STR W7, [X0] unlock()
Test: $x = 2$

This work is licensed under a Creative Commons Attribution-ShareAlike International 4.0 License.

PLDI'18, June 18–22, 2018, Philadelphia, PA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5698-5/18/06.

<https://doi.org/10.1145/3192366.3192373>

It must not terminate with $x = 2$, for this would violate mutual exclusion. Now, let us instantiate the lock/unlock calls with two possible implementations of those methods.

Initially: $[X0] = x = 0$, $[X1] = m = 0$		
1 Loop:		3 TXBEGIN } begin txn
LDAXR W2, [X1]	atomically	LDR W6, [X1] } load m
CBNZ W2, Loop	update m	CBZ W6, L1 } and abort
4 MOV W3, #1	from 0	TXABORT } if non-
STXR W4, W3, [X1]	to 1	L1: } zero
CBNZ W4, Loop		MOV W7, #1 } $x \leftarrow 1$
2 LDR W5, [X0]		STR W7, [X0]
5 ADD W5, W5, #2	$x \leftarrow x + 2$	TXEND } end txn
STR W5, [X0]		
STLR WZR, [X1]	$m \leftarrow 0$	
Test: $x = 2$		

The left thread executes its CR non-transactionally, using the recommended ARMv8 spinlock [7, K9.3], while the right thread uses lock elision (with unofficial but representative TM instructions). This program *can* terminate with $x = 2$, thus witnessing the unsoundness of lock elision, as follows:

- 1 The left thread reads the lock variable m as 0 (free). LDAXR indicates an *acquire* load, which means that the read cannot be reordered with any later event in program-order.
- 2 The left thread reads x as 0. This load can execute speculatively because ARMv8 does not require that the earlier store-exclusive (STXR) completes first [45].
- 3 The right thread starts a transaction, sees the lock is still free, updates x to 1, and commits its transaction.
- 4 The left thread updates m to 1 (taken). This is a store-exclusive (STXR) [36], so it only succeeds if m has not been updated since the last load-exclusive (LDAXR). It does succeed, because the right thread only *reads* m .
- 5 Finally, the left thread updates x to 2, and m to 0. STLR is a *release* store, which means that the update to m cannot be reordered with any earlier event in program-order.

The crux of our counterexample is that a (non-transactional) CR can start executing after the lock has been observed to be free, but before it has actually been taken. Importantly, this relaxation is safe if all CRs are mutex-protected (i.e., the spinlock *in isolation* is correct), since every lock acquisition involves writing to the lock variable and at most one store-exclusive can succeed. Rather, the counterexample only arises when this relaxation is *combined* with any reasonable TM extension to ARMv8. This includes a proposed extension currently being considered within ARM Research.

Furthermore, there appears to be no easy fix. Re-implementing the spinlock by appending a DMB fence to the lock() implementation would prevent the problematic reordering, but would also inhibit compatibility with code that uses the ARM-recommended spinlock, and may decrease performance when lock elision is not in use. Otherwise, if software

portability is essential, transactional CRs could be made to *write* to the lock variable (rather than just read it), but this would induce serialisation, and thus nullify the potential speedup from lock elision.

1.2 Our Work

In this paper, we use formalisation to tame the interaction between TM and weak memory. Specifically, we propose axiomatic models for how transactions behave in x86 [31], Power [29], ARMv8 [7], and C++ [34]. As well as the lock elision issue already explained, our formalisations revealed:

- an ambiguity in the specification of Power TM (§5.2),
- a bug in a register-transfer level (RTL) prototype implementation of ARMv8 TM (§6.2),
- a simplification to the C++ TM proposal (§7.2), and
- that coalescing transactions is unsound in Power (§8.1).

Although TM is conceptually simple, it is notoriously challenging to implement correctly, as exemplified by Intel repeatedly having to disable TM in its processors due to bugs [26, 30], IBM describing adding TM to Power as “arguably the single-most invasive change ever made to IBM’s RISC architecture” [1], and the C++ TM Study Group listing “conflict with the C++ memory model and atomics” as one of their hardest challenges [56]. To cope with the combined complexity of transactions and weak memory that exist in real systems, we build on several recent advances in automated tooling to help develop and validate our models. In the x86 and Power cases, we use the SAT-based Memalloy tool [55], extended with an exhaustive enumeration mode à la Lustig et al. [40], to automatically synthesise exactly the ‘minimally forbidden’ tests (up to a bounded size) that distinguish our TM models from their respective non-TM baselines. We then use the Litmus tool [4] to check that these tests are never observed on existing hardware (i.e., that our models are sound). We also generate a set of ‘maximally allowed’ tests, which we use to assess the completeness of our models (i.e., how many of the behaviours our models allow are empirically observable).

Moreover, we investigate several properties of our models. For instance, C++ offers ‘atomic’ transactions and ‘relaxed’ transactions; we prove that atomic transactions are strongly isolated, and that race-free programs with no non-SC atomics and no relaxed transactions enjoy ‘transactional SC’. Other properties of our models we verify up to a bound using Memalloy; these are that introducing, enlarging, or coalescing transactions introduces no new behaviours, and that C++ transactions compile soundly to x86, Power, and ARMv8.

Finally, we show how Memalloy can be used to check a library implementation against its specification by encoding it as a program transformation. We apply our technique to check that x86 and Power lock elision libraries correctly

implement mutual exclusion – but that this is not so, as we have seen, in ARMv8.

Summary Our contributions are as follows:

- a fully-automated toolflow for generating tests from an axiomatic memory model and using them to validate the model's soundness, its completeness, and its metatheoretical properties (§4);
- formalisations of TM in the SC (§3), x86 (§5), Power (§5), ARMv8 (§6), and C++ (§7) memory models;
- proofs that the transactional C++ memory model guarantees strong isolation for atomic transactions, and transactional SC for race-free programs with no non-SC atomics or non-atomic transactions (§7);
- the automatic, bounded verification of transactional monotonicity and compilation from C++ transactions to hardware (§8); and
- a technique for validating lock elision against hardware TM models, which is shown to be effective through the discovery of the serious flaw of Example 1.1 (§8).

Companion Material We provide all the models we propose (in the .cat format [5]), the automatically-generated litmus tests used to validate our models, litmus tests corresponding to all the executions discussed in our paper, and Isabelle proofs of all statements marked with the $\clubsuit$ symbol.

2 Background: Axiomatic Memory Models

Here we give the necessary background on the formal framework we use for reasoning about programs, which is standard across several recent works [5, 40, 55].

A *memory model* defines how threads interact with shared memory. An *axiomatic* memory model consists of constraints (i.e., axioms) on *candidate executions*. An execution is a graph representing a runtime behaviour, whose structure is defined below. The candidate executions of a program are obtained by assuming a non-deterministic memory system: each load can observe a store from anywhere in the program. After filtering away the candidates that fail the constraints, we are left with the *consistent* executions; i.e., those that are allowed in the presence of the actual memory system.

2.1 Executions

Let $\mathbb{X}$ be the set of all executions. Each execution is a graph whose vertices, E , represent runtime memory-related events and whose labelled edges represent various relations between them. The events are partitioned into R , W , and F , the sets of read, write, and fence events.¹ Events in an execution are connected by the following relations:

- po , program order (a.k.a. sequenced-before);

¹ We encode fences as *events* (rather than edges) because this simplifies execution minimisation (§4.2). We then derive architecture-specific fence relations that connect events separated by fence events, which we use in our models and execution graphs.

Figure 1. An execution and its litmus test

- $addr/ctrl/data$, an address/control/data dependency;
- rmw , to indicate read-modify-write operations;
- $sloc$, between events that access the same location;
- rf , the ‘reads-from’ relation; and
- co , the ‘coherence’ order in which writes hit memory.

We restrict our attention to executions that are *well-formed* as follows: po forms, for each thread, a strict total order over that thread's events; $addr$, $ctrl$, and $data$ are within po and always originate at a read; rmw links the read of an RMW operation to its corresponding write; rf connects writes to reads accessing the same location, with no read having more than one incoming rf edge; and co connects writes to the same location and forms, for each location, a strict total order over the writes to that location.

Notation Given a relation r , r^{-1} is its inverse, $r^?$ is its reflexive closure, r^+ is its transitive closure, and r^* is its reflexive transitive closure. We use $\neg$ for the complement of a set or relation, implicitly with respect to the set of all events or event pairs in the execution. We write ‘;’ for relational composition: $r_1; r_2 = \{(x, z) \mid \exists y. (x, y) \in r_1 \wedge (y, z) \in r_2\}$. We write $[-]$ to lift a set to a relation: $[s] = \{(x, x) \mid x \in s\}$. To restrict a relation r to being inter-thread or intra-thread, we use $r_e = r \setminus (po \cup po^{-1})^*$ or $r_i = r \cap (po \cup po^{-1})^*$, respectively. Similarly, $r_{loc} = r \cap sloc$.

Derived Relations The *from-read* (fr) relation relates each read event to all the write events on the same location that are *co*-later than the write the read observed [40]. The *com* relation captures three ways events can ‘communicate’ with each other.

$$\begin{aligned} fr &= ([R]; sloc; [W]) \setminus (rf^{-1}; (co^{-1})^*) \\ com &= rf \cup co \cup fr \end{aligned}$$

Visualising Executions We represent executions using diagrams like the one in Fig. 1 (left). Here, the po -edges separate the execution's events into two threads, each drawn in one column. Each event is labelled with the sets it belongs to, such as R and W . We use location names such as x to identify the $sloc$ -classes.

2.2 From Executions to Litmus Tests

In order to test whether an execution of interest is observable in practice, it is necessary to convert it into a *litmus test* (i.e., a program with a postcondition) [18]. This litmus test is constructed so that the postcondition only passes when the particular execution of interest has been taken [3, 55].

Figure 2. A transactional execution and its litmus test

As an example, the execution on the left of Fig. 1 corresponds to the pseudocode litmus test on the right. Read events become loads, writes become stores, and the *po*-edges induce the order of instructions and their partitioning into threads. To ensure that the litmus test passes only when the intended *rf*-edges are present, we arrange that each store writes a unique non-zero value, and then check that each local register holds the value written by the store it was intended to observe – this corresponds to the $r0 = 2$ in the postcondition. To ensure that the intended *co*-edges are present, we check the final value of each memory location – this corresponds to the $x = 2$ in the postcondition.²

3 Axiomatising Transactions

Transactional memory (TM) can be provided either at the architecture level (x86, Power, ARMv8) or in software (C++). Since we are concerned only with the *specification* of TM, and not its implementation, we can formalise both forms within a unified framework. In this section, we describe how program executions can be extended to express transactions (§3.1) and how we can derive litmus tests to test for these executions (§3.2). We then propose axioms for capturing the *isolation* of transactions (§3.3), and for strengthening the SC memory model to obtain *transactional* SC (§3.4).

3.1 Transactional Executions

To enable transactions in an axiomatic memory modelling framework, we extend executions with an *stxn* relation that relates events in the same successful (i.e., committed) transaction. For an execution to be well-formed, *stxn* must be a partial equivalence relation (i.e., symmetric and transitive), and each *stxn*-class must coincide with a contiguous subset of *po*. When generating the candidate executions of a program with transactions, each transaction is assumed to succeed or fail non-deterministically. That is, each either gives rise to a *stxn*-class of events, or vanishes as a no-op.

Diagrammatically, we represent *stxn* using boxes. For instance, events *a* and *b* in Fig. 2 form a successful transaction.

² When there are more than two writes to a location, extra constraints on executions are needed to fix all the *co*-edges [55].

Figure 3. Four SC executions that are allowed by weak isolation but forbidden by strong isolation

Remark 3.1. To study the behaviour of unsuccessful transactions in more detail, one could add an explicit representation of them in executions, perhaps using dashed boxes. However, the behaviour of unsuccessful transactions is tricky to ascertain on hardware because of the rollback mechanism. Moreover, it is unclear how they should interact with *co*, since *co* is the order in which writes hit the memory, which writes in unsuccessful transactions never do.

3.2 From Transactional Executions to Litmus Tests

A transactional execution can be converted into a litmus test by extending the construction of §2.2. As an example, the execution on the left of Fig. 2 corresponds to the litmus test on the right. The instructions in the transaction simply need enclosing in instructions that begin and end a transaction. We write these as *txbegin* and *txend* here; our tooling specialises these for each target architecture. The postcondition checks that the transaction succeeded using the ‘*ok*’ location, which is zeroed in the transaction’s fail-handler, L_{fail} , the label of which is provided with the *txbegin* instruction.

3.3 Weak and Strong Isolation

We now explain how the *isolation* property of transactions can be captured as a property of an execution graph. A TM system provides *weak* isolation if transactions are isolated from other transactions; that is, their intermediate state cannot affect or be affected by other transactions [12, 27]. It provides *strong* isolation if transactions are also isolated from non-transactional code.

The four 3-event SC executions in Fig. 3 illustrate the difference between strong and weak isolation. In each, the interfering event would need to be within a transaction to be forbidden by weak isolation; strong isolation does not make this distinction. Executions (a) and (d) correspond to what Blundell et al. call *non-interference* and *containment*, respectively, and (b) is similar to the standard axiom for RMW isolation (cf. *RMWISOL* in Fig. 5).

Failures of isolation can be characterised as communication cycles between transactions. To define these cycles, the following constructions are useful:

$$\begin{aligned} \text{weaklift}(r, t) &= t; (r \setminus t); t \\ \text{stronglift}(r, t) &= t^2; (r \setminus t); t^2. \end{aligned}$$

$\text{acyclic}(hb)$ where $hb = po \cup com$	(ORDER)
$\text{acyclic}(\text{stronglift}(hb, stxn))$	(TxnORDER)

Figure 4. SC axioms [49], with TSC extensions highlighted

If r relates events e_1 and e_2 in different transactions, then $\text{weaklift}(r, stxn)$ relates all the events in e_1 's transaction to all those in e_2 's transaction. The stronglift version also includes edges where the source and/or the target event are not in a transaction. Weak and strong isolation can then be axiomatised by treating all the events in a transaction as a single event whenever the transaction communicates with another transaction (WEAKISOL) or any other event (STRONGISOL).

$\text{acyclic}(\text{weaklift}(com, stxn))$	(WEAKISOL)
$\text{acyclic}(\text{stronglift}(com, stxn))$	(STRONGISOL)

3.4 Transactional Sequential Consistency

Although isolation is a critical property for transactions, it only provides a lower bound on the guarantees that real architectures provide. Meanwhile, an upper bound on the guarantees provided by a reasonable TM implementation is *transactional sequential consistency* (TSC) [19]. The models we propose in §5–7 all lie between these bounds.

TSC is a strengthening of the SC memory model in which consecutive events in a transaction must appear consecutively in the overall execution order. Where SC can be characterised axiomatically (Fig. 4) by forbidding cycles in program order and communication (ORDER) [49], we can obtain TSC by additionally forbidding such cycles between transactions and non-transactional events (TxnORDER). Note that TxnORDER subsumes the STRONGISOL axiom.

4 Methodology

We identify three components of a memory modelling methodology: (1) developing and refining axioms, (2) synthesising and running conformance tests, and (3) checking metatheoretical properties. In this section, we explain our approach to each of these components, and in particular, how we have extended the Memalloy tool [55] to support each task.

Background on Memalloy The original Memalloy tool, built on top of Alloy [35], was developed for comparing memory models. It takes two models (say, M and N), and searches for a single execution that distinguishes them (i.e., is inconsistent under M but consistent under N). Additionally, if Memalloy is supplied with a translation on executions (e.g., representing a compiler mapping or a compiler optimisation), then it searches for a witness that the translation is unsound. This translation is defined by a relation, typically named π , from ‘source’ events to ‘target’ events.

4.1 Developing and Refining Axioms

For each model, we make a first attempt at a set of axioms using information obtained from specifications, programming manuals, research papers, and discussions with designers. Then, for each proposed change to the model, we use Memalloy to generate tests that become disallowed or allowed as a result. We decide whether to accept the change based on discussing these tests with designers, and running them on existing hardware (where available) using the Litmus tool [4].

In order to extend Memalloy to support the development of transactional memory models in this way, we augmented the form of executions as described in §3.1, and modified the litmus test generator as described in §3.2.

4.2 Synthesising and Running Conformance Tests

To build confidence in a model, we compare the behaviours it admits against those allowed by the architecture or language being modelled. It is vital that no behaviour allowed by the architecture/language is forbidden by the model, so we exhaustively generate all litmus tests (up to a bounded size) that our model forbids, and confirm using Litmus that none can be observed on existing hardware.

To achieve this, we extended Memalloy with a mode for exhaustively generating conformance tests for a given model M . The key to exhaustive generation is a suitable notion of *minimality*, without which we would obtain an infeasibly large number of tests. We closely follow Lustig et al. [40], and define execution minimality with respect to the following partial order between executions. Let $X \sqsubset Y$ hold when execution X can be obtained from execution Y by:

- (i) removing an event (plus any incident edges),
- (ii) removing a dependency edge (*addr, ctrl, data, rmw*), or
- (iii) downgrading an event (e.g. reducing an acquire-read to a plain read in ARMv8).

We then calculate the set $\text{min-inconsistent}(M) = \{X \in \mathbb{X} \mid X \notin \text{consistent}(M) \wedge \forall X' \sqsubset X. X' \in \text{consistent}(M)\}$.

Extending Memalloy to support the synthesis of transactional conformance tests requires minimality to take transactions into account. To do this, we arrange that $X \sqsubset Y$ also holds when X can be obtained from Y by:

- (v) making the first or last event in a transaction non-transactional (i.e. removing all of its incident *stxn* edges).

(We avoid the ‘middle’ of a transaction so as not to create non-contiguous transactions and hence ill-formed executions.)

Remark 4.1. While this is a slightly coarse notion of minimality – a more refined version would also allow a large transaction to be chopped into two smaller ones – it is cheap to implement in the constraint solver as it only requires quantification over a single event. As a result, Memalloy may generate some executions that appear non-minimal, but as we show in §5.3, this does not impede our ability to generate and run large batches of conformance tests.

Generating Allowed Tests Having generated the minimally-forbidden tests, the question naturally arises of whether we can generate the *maximally-allowed* tests too. Where the minimally-forbidden tests include just enough fences/dependencies/transactions to be forbidden (and failing to observe these tests empirically suggests that the model is not too strong), the maximally-allowed tests include just *not* enough (and observing them suggests that the model is not too weak). We found the maximally-allowed tests valuable for communicating with engineers about the detailed relaxations permitted by our models. However, in our experiments, allowed tests are less conclusive than forbidden ones, because where the observation of a forbidden test implies that the model is unsound, the non-observation of an allowed test may be caused by not performing enough runs, or by the machine under test being implemented conservatively.

Moreover, the notion of execution maximality is not as natural as minimality. For instance, an inconsistent execution is only considered minimally-inconsistent if *removing* any event makes it *consistent*, yet it is not sensible to deem a consistent execution maximally-consistent only when *adding* any event makes it *inconsistent* – such a condition is almost impossible to satisfy. Even with event addition/removal set aside, maximal-consistency tends to require executions to be littered with redundant fences and dependencies.

Therefore, we approximate the maximally-consistent executions as those obtained via a single $\sqsubset$ -step from a minimally-inconsistent execution. That is, we let $\text{max-consistent}(M) = \{X \in \mathbb{X} \mid \exists Y \in \text{min-inconsistent}(M). X \sqsubset Y\}$.

4.3 Checking Metatheoretical Properties

As explained at the start of this section, Memalloy is able to validate transformations between two memory models, providing they can be encoded as a π -relation between executions. In §8, we exploit this ability to check several TM-related transformations and compiler mappings.

In fact, Memalloy can also be used to check libraries under weak memory. Prior work has (manually) verified that stack, queue, and barrier libraries implement their specifications under weak memory models [8, 52]; here we show how checking these types of properties can be automated up to a bounded number of library and client events. We see this as a straightforward first-step towards a general verification effort. The idea, which we apply to a lock elision library in §8.3, is to treat the replacement of the library's specification with its implementation as a program transformation. To do this, we first extend executions with events that represent method calls. Second, we extend execution well-formedness so that illegal call sequences (such as popping from an empty stack) are rejected. Third, we strengthen the memory model's consistency predicate with axioms capturing the library's obligations (such as pops never returning data from later pushes). Finally, we constrain π so that it maps each method

$\text{acyclic}(po_{\text{loc}} \cup com)$	(COHERENCE)
$\text{empty}(rmw \cap (fr_e; co_e))$	(RMWISOL)
$\text{acyclic}(hb)$	(ORDER)
where $ppo = ((W \times W) \cup (R \times W) \cup (R \times R)) \cap po$	
$tfence = po \cap ((\neg stxn; stxn) \cup (stxn; \neg stxn))$	
$L = \text{domain}(rmw) \cup \text{range}(rmw)$	
$\text{implied} = [L]; po \cup po; [L] \cup tfence$	
$hb = mfence \cup ppo \cup \text{implied} \cup rf_e \cup fr \cup co$	
$\text{acyclic}(\text{stronglift}(com, stxn))$	(STRONGISOL)
$\text{acyclic}(\text{stronglift}(hb, stxn))$	(TXNORDER)

Figure 5. x86 consistency axioms [5], with our TM additions highlighted

call to an event sequence representing the implementation of that method.

5 Transactions in x86 and Power

Over the next three sections, we show how our methodology can be applied to four different targets. We begin with x86 and Power, which have both supported TM since 2013 [15, 32]. Intel's Transactional Synchronisation Extensions (TSX) provide XBEGIN, XEND, and XABORT instructions for starting, committing, and aborting transactions, while Power provides tbegin, tend, and tabort.

5.1 Background: the x86 and Power Memory Models

Both the x86 memory model [44] and the Power memory model [5, 47, 48] allow certain instructions to execute out of program order, with x86 allowing stores to be reordered with later loads and Power allowing many more relaxations. Both architectures provide fences (MFENCE in x86, and lwsync, sync, and isync in Power) to allow these relaxations to be controlled. The x86 architecture provides atomic RMWs via LOCK-prefixed instructions, while Power implements RMWs using *exclusive* instructions like those seen in Example 1.1. Moreover, x86 is *multicopy-atomic* [18], which means that writes are propagated to all other threads simultaneously. Power does not have this property, so its memory model includes explicit axioms to describe how writes propagate.

More formally, we extend executions with relations that connect events in program order that are separated by a fence event of a given type. For x86, we add an *mfence* relation, and for Power, we add *isync*, *lwsync*, and *sync*.

An x86 execution is consistent if it satisfies all of the axioms in Fig. 5 (ignoring the highlighted regions for now). The COHERENCE axiom forbids cycles in communication edges and program order among events on the same location; this guarantees programs that use only a single location to have SC semantics. Happens-before (*hb*) imposes the event-ordering constraints upon which all threads must agree, and ORDER ensures that hb^* is a partial order. The constraints

$\text{acyclic}(po_{\text{loc}} \cup com)$	(COHERENCE)
$\text{empty}(rmw \cap (fr_e; co_e))$	(RMWISOL)
$\text{acyclic}(hb)$	(ORDER)
where $ppo = (\text{preserved program order, elided})$	
$tfence = po \cap ((\neg stxn; stxn) \cup (stxn; \neg stxn))$	
$fence = sync \cup tfence \cup (lwsync \setminus (W \times R))$	
$ihb = ppo \cup fence$	
$thb = (rf_e \cup ((fr_e \cup co_e)^*; ihb))^*; (fr_e \cup co_e)^*; rf_e^?$	
$hb = (rf_e^?; ihb; rf_e^?) \cup \text{weaklift}(thb, stxn)$	
$\text{acyclic}(co \cup prop)$	(PROPAGATION)
where $efence = rf_e^?; fence; rf_e^?$	
$prop_1 = [W]; efence; hb^*; [W]$	
$prop_2 = com_e^*; efence^*; hb^*; (sync \cup tfence); hb^*$	
$tprop_1 = rf_e; stxn; [W]$	
$tprop_2 = stxn; rf_e$	
$prop = prop_1 \cup prop_2 \cup tprop_1 \cup tprop_2$	
$\text{irreflexive}(fr_e; prop; hb^*)$	(OBSERVATION)
$\text{acyclic}(\text{stronglift}(com, stxn))$	(STRONGISOL)
$\text{acyclic}(\text{stronglift}(hb, stxn))$	(TXNORDER)
$\text{empty}(rmw \cap tfence^*)$	(TXNCANCELSRMW)

Figure 6. Power consistency axioms [5], with our TM additions highlighted, and some details elided for brevity.

on hb arise from: fences placed by the programmer ($mfence$), fences created implicitly by LOCK'd operations ($implied$), the preserved fragment of the program order (ppo), inter-thread observations (rf_e) and communication edges (fr and co).

A Power execution is consistent if it satisfies all the axioms in Fig. 6 (again, ignoring the highlights). The first axiom not already seen is ORDER, which ensures that happens-before (hb) is acyclic. In contrast to x86, the happens-before relation in Power is formed from inter-thread observations (rf_e), the preserved fragment of the program order (ppo), and fences ($fence$). We elide the definition of ppo as it is complex and unchanged by our TM additions. The $prop$ relation governs how fences restrict “the order in which writes propagate” [5], and the PROPAGATION axiom ensures that this relation does not contradict the coherence order. OBSERVATION governs which writes a read can observe: if e_1 propagates before e_2 , then any read that happens after e_2 is prohibited from observing a write that precedes e_1 in coherence order.

5.2 Adding Transactions

To extend the x86 and Power memory models to support TM, we make the following amendments, each highlighted in Figs. 5 and 6.

Strong Isolation (x86 and Power) The Power manual says that transactions “appear atomic with respect to both transactional and non-transactional accesses performed by other

threads” [29, §5.1], and the TSX manual defines conflicts not just between transactions, but between a transaction and “another logical processor” (which is not required to be executing a transaction) [31, §16.2]. We interpret these statements to mean that x86 and Power transactions provide *strong* isolation, so we add our STRONGISOL axiom from §3.3.

Implicit Transaction Fences (x86 and Power) In both x86 and Power, fences are created at the boundaries of successful transactions. In x86, “a successfully committed [transaction] has the same ordering semantics as a LOCK prefixed instruction” [31, §16.3.6], and in Power, “[a] tbegin instruction that begins a successful transaction creates a [cumulative] memory barrier”, as does “a tend instruction that ends a successful transaction” [29, §1.8]. Hence, we define $tfence$ as the program-order edges that enter ($\neg stxn; stxn$) or exit ($stxn; \neg stxn$) a successful transaction, and add $tfence$ alongside the existing fence relations ($mfence$ and $sync$).

Transaction Atomicity (x86 and Power) We extend the prohibition on hb cycles among events to include cycles among transactions (TXNORDER). This essentially treats all the transaction’s events as one indivisible event, and is justified by the atomicity guarantee given to transactions, which in x86 is “that all memory operations [...] appear to have occurred instantaneously when viewed from other logical processors” [31, §16.2], and in Power is that each successful transaction “appears to execute as an atomic unit as viewed by other processors and mechanisms” [29, §1.8].

Barriers within Transactions (Power only) Each transaction contains an “integrated memory barrier”, which ensures that writes observed by a successful transaction are propagated to other threads before writes performed by the transaction itself [29, §1.8]. This behaviour is epitomised by the WRC-style execution below [15, Fig. 6],

which must be ruled out because the transaction’s write (c) has propagated to the third thread before a write (a) that the transaction observed. We capture this constraint by extending the $prop$ relation so that it connects any write observed by a transaction to any write within that transaction ($tprop_1$). In execution (1), this puts a $prop$ edge from a to c . The execution is thus forbidden by the existing OBSERVATION axiom.

Remark 5.1. The following executions are similar to (1), and like (1), they could not be observed empirically. However, the Power manual is ambiguous about whether they should be forbidden.

In particular, because the transactions are read-only, we cannot appeal to the integrated memory barrier. We have reported this ambiguity to IBM architects, and while we await a clarified specification, our model errs on the side of caution by permitting these executions.

Propagation of Transactional Writes (Power only) Although Power is not multicopy-atomic in general, *transactional* writes are multicopy-atomic; that is, the architecture will “propagate the transactional stores fully before committing the transaction” [15, §4.2]. This behaviour is epitomised by another **WRC**-style execution, in which the middle thread sees the transactional write to x before the right thread does.

To rule out such executions, it suffices to extend the *prop* relation with reads-from edges that exit a transaction (*tprop₂*), and then to invoke **OBSERVATION** again.

Read-modify-writes (Power only) In Power, when a store-exclusive is separated from its corresponding load-exclusive by “a state change from Transactional to Non-transactional or Non-transactional to Transactional”, the RMW operation will always fail [29, §1.8]. Therefore, the **TxNCANCELSRMW** axiom ensures that no consistent execution has an *rmw* edge crossing a transaction boundary.

Transaction Ordering (Power only) The Power manual states that “successful transactions are serialised in some order”, and that it is impossible for contradictions to this order to be observed [29, p. 824].

We capture this constraint by extending the *hb* relation to include a new *thb* relation between transactions. The *thb* relation imposes constraints on the order in which transactions can be serialised. By including it in *hb* and requiring *thb* to be a partial order, we guarantee the existence of a suitable transaction serialisation order, without having to construct this order explicitly.

The definition of the *thb* relation is a little convoluted, but the intuition is quite straightforward: it contains all non-empty chains of intra-thread happens-before edges (*ihb*) and inter-thread communication edges (*com_e*), except those that contain an *fr_e* or *co_e* edge followed by an *rf_e* edge that does not terminate the chain. The rationale for excluding *fr_e*; *rf_e* and *co_e*; *rf_e* chains is that these do not provide ordering in a non-multicopy-atomic architecture. That is, from

$$a \xrightarrow{fr_e} b \xrightarrow{rf_e} c \quad \text{or} \quad a \xrightarrow{co_e} b \xrightarrow{rf_e} c$$

we cannot deduce that a happens before c , because this behaviour can also be attributed to the write b being propagated to c 's thread before a 's thread.

Table 1. Testing our transactional x86 and Power models

Arch.	E	Synthesis time (s)	Forbid			Allow		
			T	S	¬S	T	S	¬S
x86	2	4	0	0	0	2	2	0
	3	22	4	0	4	24	23	1
	4	87	22	0	22	99	99	0
	5	260	42	0	42	249	244	5
	6	4402	133	0	133	895	832	63
	7	>7200	307	0	307	2457	1901	556
	Total (x86):		508	0	508	3726	3101	625
Power	2	13	2	0	2	7	7	0
	3	58	9	0	9	44	44	0
	4	318	60	0	60	184	175	9
	5	9552	353	0	353	1517	1330	187
	6	>7200	922	0	922	5043	4407	636
	Total (Power):		1346	0	1346	6795	5963	832

Cain et al. epitomise the transaction-ordering constraint using the **IRIW**-style execution reproduced below [15, Fig. 5].

The execution must be disallowed because different threads observe incompatible transaction orders: the second thread observes a before f , but the third observes f before a . Our model disallows this execution on the basis of a *thb* cycle between the two transactions.

We must be careful not to overgeneralise here, because a behaviour similar to (3) but with only *one* write transactional was observed during our empirical testing, and is duly allowed by our model.

5.3 Empirical Testing

Table 1 gives the results obtained using our testing strategy from §4.2. We use Memalloy to synthesise litmus tests that are forbidden by our transactional models but allowed under the non-transactional baselines (the **Forbid** set), up to a bounded number of events ($|E|$). We then derive the **Allow** sets by relaxing each test. We report synthesis times on a 4-core Haswell i7-4771 machine with 32GB RAM, using a timeout of 2 hours. For both sets we give the number of tests (T) found; we say this number is complete if synthesis did not reach timeout and non-exhaustive otherwise. We say a test is seen (S) if it is observed on any implementation, and not seen (¬S) otherwise. Each x86 test is run 1M times on four TSX implementations: a Haswell (i7-4771), a Broadwell-Mobile (i7-5650U), a Skylake (i7-6700), and a Kabylake (i7-7600U). Each Power test is run 10M times on an 80-core POWER8 (TN71-BP012) machine. When testing this machine, we use Litmus's *affinity* parameter [4], which places threads

Figure 7. The distribution of synthesis times for the 7-event x86 **Forbid** tests

incrementally across the logical processors to encourage **IRIW**-style behaviours.

We were able to generate the complete set of x86 **Forbid** executions that have up to 6 events, and the complete set of Power **Forbid** executions up to 5 events. Regarding these bounds: we remark that since our events only represent memory accesses and fences (not, for instance, starting or committing transactions), we can capture many interesting behaviours with relatively few events. For instance, these bounds are large enough to include all the executions discussed in this section.

Of our 508 x86 **Forbid** tests, 29% had one transaction, 44% had two, and 27% had three, and of the 1346 Power **Forbid** tests, 29% had one transaction, 54% had two, and 17% had three. No **Forbid** test was empirically observable on either architecture, which gives us confidence that our models are not too strong. Of the x86 **Allow** tests, 83% could be observed on at least one implementation, as could 88% of the Power **Allow** tests; this provides some evidence that our models are not excessively weak. Many of the unobserved Power **Allow** tests are based on the load-buffering (**LB**) shape, which has never actually been observed on a Power machine, even without transactions [6].

Increasing the timeout to 48 hours is sufficient to generate the complete set of x86 **Forbid** executions for 7 events. It takes 34 hours for Memalloy to find all 313 tests. Figure 7 shows how the percentage of executions found is affected by various caps on the synthesis time. We observe that many tests are found quickly: 98% of the tests are found within 2 hours (i.e., 6% of the total synthesis time), and all of the tests are found within 9 hours (the remaining synthesis time is used to confirm that there are no further tests). During the development process, we exploited this observation to obtain preliminary test results more rapidly.

6 Transactions in ARMv8

The ARMv8 memory model sits roughly between x86 and Power. Like x86, it is multicopy-atomic [45], but like Power, it permits several relaxations to the program order. Unwanted relaxations can be inhibited either using barriers (DMB, DMB LD, DMB ST, ISB) or using *release/acquire* instructions (LDAR, STLR) that act like one-way fences.

$\text{acyclic}(po_{\text{loc}} \cup com)$	(COHERENCE)
$\text{acyclic}(ob)$	(ORDER)
where $dob = (\text{order imposed by dependencies, elided})$	
$aob = (\text{order imposed by atomic RMWs, elided})$	
$bob = (\text{order imposed by barriers, elided})$	
$tfence = po \cap ((\neg stxn; stxn) \cup (stxn; \neg stxn))$	
$ob = com_e \cup dob \cup aob \cup bob \cup tfence$	
$\text{empty}(rmw \cap (fr_e; co_e))$	(RMWISOL)
$\text{acyclic}(\text{stronglift}(com, stxn))$	(STRONGISOL)
$\text{acyclic}(\text{stronglift}(ob, stxn))$	(TXNORDER)
$\text{empty}(rmw \cap tfence^*)$	(TXNCANCELSRMW)

Figure 8. ARMv8 consistency axioms [7, 21], with our TM additions highlighted, and some details elided for brevity.

Formally, ARMv8 executions are obtained by adding six extra fields: *Acq* and *Rel*, which are the sets of acquire and release events, and *dmb/dmbld/dmbst/isb*, which relate events in program order that are separated by barriers.

An ARMv8 execution is consistent if it satisfies all of the axioms in Fig. 8 (ignoring the highlighted regions). We have seen the COHERENCE and RMWISOL axioms already. The ordered-before relation (*ob*) plays the same role as happens-before in x86: it imposes the event-ordering constraints upon which all threads must agree, and must be free from cycles (ORDER). These constraints arise from communication (com_e), dependencies (*dob*), atomic RMWs (*aob*), and barriers (*bob*).

6.1 Adding Transactions

The ARMv8 architecture does not support TM, so the extensions proposed below (highlighted in Fig. 8) are unofficial. Nonetheless, the extensions we give are based upon a proposal currently being considered within ARM Research and upon extensive conversations with ARM architects.

- STRONGISOL is a natural choice for hardware TM.
- As in x86 and Power, we place implicit fences (*tfence*) at the boundaries of successful transactions.
- We bring the TXNORDER axiom from x86 and Power to forbid *ob*-cycles among transactions.
- Like Power, ARMv8 has exclusive instructions, so it inherits the TXNCANCELSRMW axiom to ensure the failure of RMWs that straddle a transaction boundary.

6.2 Empirical Testing

ARM hardware does not support TM so we cannot test our model as we did for x86 and Power. However, we generated the **Forbid** and **Allow** suites anyway, and gave them to ARM architects. They were able to use these to reveal a bug (specifically, a violation of the TXNORDER axiom) in a register-transfer level (RTL) prototype implementation.

irreflexive ($hb; com^*$)	(HbCOM)
where $sw = (synchronises-with, elided)$	
$ecom = com \cup (co; rf)$	
$tsw = weaklift(ecom, stxn)$	
$hb = (sw \cup tsw \cup po)^+$	
empty ($rmw \cap (fr_e; co_e)$)	(RMWISOL)
acyclic ($po \cup rf$)	(NoTHINAIR)
acyclic (psc)	(SEQCST)
where $psc = (constraints\ on\ SC\ events, elided)$	
empty ($cnf \setminus Ato^2 \setminus (hb \cup hb^{-1})$)	(NoRACE)
where $cnf = ((W \times W) \cup (R \times W) \cup (W \times R)) \cap sloc \setminus id$	

Figure 9. C++ consistency and race-freedom axioms [38], with our TM additions highlighted, and some details elided.

7 Transactions in C++

We now turn our attention from hardware to software. TM is supported in C++ via an ISO technical specification that has been under development by the C++ TM Study Group since 2012 [34, 50]. In this section, we formalise how the proposed TM extensions interact with the existing C++ memory model, and detail a possible simplification to the specification.

C++ TM offers two main types of transactions: *relaxed transactions* (written `synchronized{...}`) can contain arbitrary code, but only promise weak isolation, while *atomic transactions* (written `atomic{...}`) promise strong isolation but cannot contain certain operations, such as atomic operations [34, §8.4.4]. Some atomic transactions can be aborted by the programmer, but we do not support these in this paper.

7.1 Background: the C++ Memory Model

Our presentation of the baseline C++ memory model follows Lahav et al. [38]. We choose to build on their formalisation because it incorporates a fix that allows correct compilation to Power – without this, we could not check the compilation of C++ transactions to Power transactions (§8.2).

C++ executions identify four additional subsets of events: *Ato* contains the events from atomic operations, while *Acq*, *Rel*, and *SC* contain events from atomic operations that use the acquire, release, and SC consistency modes [33, §29.3].

Unlike the architecture-level memory models, the C++ memory model defines *two* predicates on executions (Fig. 9). The first characterises the *consistent* candidate executions. If any consistent execution violates a second *race-freedom* predicate, then the program is completely undefined. Otherwise, the allowed executions are the consistent executions.

A C++ execution is consistent if it satisfies all of the consistency axioms given at the top of Fig. 9 (ignoring highlighted regions for now). The first, HbCOM, governs the happens-before relation, which is constructed from the program order and the *synchronises-with* relation (*sw*). Roughly speaking,

an *sw* edge is induced when an acquire read observes a release write; but it also handles fences and the ‘release sequence’ [9, 38]. The second axiom is standard for capturing the isolation of RMW operations. The NoTHINAIR axiom is Lahav et al.’s solution to C++’s ‘thin air’ problem [10]. Finally, SEQCST forbids certain cycles among SC events; we omit its definition as it does not interact with our TM extensions.

A consistent C++ execution is race-free if it satisfies the NoRACE axiom at the bottom of Fig. 9, which states that whenever two conflicting (*cnf*) events in different threads are not both atomic, they must be ordered by happens-before.

7.2 Adding Transactions

The specification for C++ TM makes two amendments to the C++ memory model: one for data races, and one for transactional synchronisation.

Transactions and Data Races The definition of a race is unchanged in the presence of TM. In particular, the program

```
atomic{ x=1; } || atomic_store(&x, 2);
```

is racy – which is perhaps contrary to the intuition that an atomic transaction with a single non-atomic store should be interchangeable with a non-transactional atomic store.

Remark 7.1. The specification also clarifies that although events in an unsuccessful transaction are unobservable, they can still participate in races. This implies that the program

```
atomic{ x=1; abort(); } || atomic_store(&x, 2);
```

must be considered racy. In our formalisation, transactions either succeed (giving rise to an *stxn*-class) or fail, giving rise to no events (cf. §3.1). This treatment correctly handles races involving unsuccessful transactions, because the race will be detected in the case where the transaction succeeds, but it cannot handle transactions that *never* succeed, such as the one above. Therefore, we leave the handling of `abort()` for future work.

Transactional Synchronisation The second amendment by the C++ TM extension defines when two transactions synchronise [34, §1.10]. An execution is deemed consistent only if there is a total order on transactions such that:

1. this order does not contradict happens-before, and
2. if transaction T_1 is ordered before conflicting transaction T_2 , then the end of T_1 synchronises with the start of T_2 .

We could incorporate these requirements into the formal model by extending executions with a transaction-ordering relation, *to*, that serialises all the *stxn*-classes in an order that does not contradict happens-before (point 1), and updating the *synchronises-with* relation to include events in conflicting transactions that are ordered by *to* (point 2).

However, this formulation is unsatisfying. It is awkward that *to* is used to *define* happens-before but is also forbidden to *contradict* happens-before. Moreover, having the consistency predicate involve quantification over all possible transaction serialisations makes simulation expensive [9].

Fortunately, we can formulate the C++ TM memory model without relying on a total order over transactions. The idea is that if two transactions are in conflict, then their order can be deduced from the existing *rf*, *co*, and *fr* edges, and if they are not, then there is no need to order them.

In more detail, and with reference to the highlighted parts of Fig. 9: observe that whenever two events in an execution conflict (*cnf*), they must be connected one way or the other by ‘extended communication’ (*ecom*), which is the communication relation extended with *co*; *rf* chains. That is, $cnf = ecom \cup ecom^{-1}$ [8]. We then say that transactions synchronise with (*tsw*) each other in *ecom* order, and we extend happens-before to include *tsw*.

By simply extending the definition of *hb* like this, we avoid the need for the *to* relation altogether, and we avoid adding any axioms to the consistency predicate. To make our proposal concrete, we provide in our companion material some text that the specification could incorporate (currently under review by the C++ TM Study Group).

Strong Isolation for Atomic Transactions The semantics described thus far provides the desired weak-isolating behaviour for *relaxed* transactions; that is, the WEAKISOL axiom follows from the other C++ consistency axioms [8]. However, *atomic* transactions must be strongly isolated. In fact, atomic transactions enjoy strong isolation simply by being forbidden to contain atomic operations. The idea is that for a non-transactional event to observe or affect a transaction’s intermediate state, it must conflict with an event in that transaction. If this event cannot be atomic, there must be a race. Thus, for race-free programs, atomic transactions are guaranteed to be strongly isolated.

To formalise this property, we extend C++ executions with an $stxn_{at}$ relation that identifies a subset of transactions as atomic. It satisfies $stxn_{at} \subseteq stxn$ and $(stxn_{at}; stxn) \subseteq stxn_{at}$. We then prove the following theorem.

Theorem 7.2 (Strong isolation for atomic transactions). If NoRACE holds, and atomic transactions contain no atomic operations (i.e., $\text{domain}(stxn_{at}) \cap Ato = \emptyset$), then

$$\text{acyclic}(\text{stronglift}(com, stxn_{at})).$$

Proof sketch. A cycle in $\text{stronglift}(com, stxn_{at})$ is either a *com*-cycle or an *r*-cycle, where $r = stxn_{at}; (com \setminus stxn_{at})^+; stxn_{at}$. From NoRACE, we have $com \setminus Ato^2 \subseteq hb$. Using this and the expansion $com^+ = ecom \cup (fr; rf)$ we can obtain $r \subseteq hb$. To finish the proof, note that execution well-formedness forbids *com*-cycles, and that *r*-cycles are forbidden too because they are also *hb*-cycles, which violate HbCom.

Table 2. Summary of our metatheoretical results. Timings are for a machine with four 16-core Opteron processors and 128GB RAM, using the Plingeling solver [11]. A means the property holds up to the given number of events, a means a counterexample was found, and $\ominus$ indicates a timeout.

Property	§	Target	Events	Time	C'ex?
Monotonicity	8.1	x86	6	20m
		Power	2	<1s
		ARMv8	2	<1s
		C++	6	91h
Compilation	8.2	C++/x86	6	14h
		C++/Power	6	16h
		C++/ARMv8	6	20h
Lock elision	8.3	x86	8	>48h	$\ominus$
		Power	9	>48h	$\ominus$
		ARMv8	7	63s
		ARMv8 (fixed)	8	>48h	$\ominus$

A Transactional SC-DRF Guarantee A central property of the C++ memory model is its SC-DRF guarantee [2, 13]: all race-free C++ programs that avoid non-SC atomic operations enjoy SC semantics. This guarantee can be lifted to a transactional setting [19, 50]: all race-free C++ programs that avoid relaxed transactions and non-SC atomic operations enjoy TSC semantics (cf. §3.4). This is formalised in the following theorem, which we prove in our companion material.

Theorem 7.3 (Transactional SC-DRF guarantee). If a C++ consistent execution has

- no relaxed transactions (i.e. $stxn = stxn_{at}$),
- no non-SC atomics (i.e. $Ato = SC$), and
- no data races (i.e. NoRACE holds),

then it is consistent under TSC.

8 Metatheory

We now study several metatheoretical properties of our proposed models. For instance, one straightforward but important property, which follows immediately from the model definitions, is that our TM models give the same semantics to transaction-free programs as the original models [8]. In this section, we use Memalloy to check some more interesting properties of our models, as summarised in Tab. 2.

8.1 Monotonicity

We check that adding *stxn*-edges can never make an inconsistent execution consistent. This implies that all of the following program transformations are sound: introducing a transaction (e.g., $\bullet \cdots \bullet \rightarrow \bullet \bullet$), enlarging a transaction (e.g., $\bullet \bullet \cdots \bullet \rightarrow \bullet \bullet \bullet$), and coalescing two consecutive transactions (e.g., $\bullet \bullet \rightarrow \bullet \bullet$).

Memalloy confirmed that the transactional x86 and C++ models enjoy this monotonicity property for all executions with up to 6 events. For Power and ARMv8, it found the following counterexample:

The left execution is inconsistent in both models because of the TXNCANCELSRMW axiom: that a store-exclusive separated from its corresponding load-exclusive by a transaction boundary always fails. The right execution, however, is consistent. This counterexample implies that techniques that involve transaction coalescing [17, 53] must be applied with care in the presence of RMWs.

8.2 Mapping C++ Transactions to Hardware

We check that it is sound to compile C++ transactions to x86, Power, and ARMv8 transactions. A realistic compiler would be more complex – perhaps including fallback options for when hardware transactions fail – but our direct mapping is nonetheless instructive for comparing the guarantees given to transactions in software and in hardware.

Specifically, we use Memalloy to search for a pair of executions, X and Y , such that X is an inconsistent C++ execution, Y is a consistent x86/Power/ARMv8 execution, and X is related to Y via the relevant compilation mapping, encoded in the relation π . Such a pair would demonstrate that the compilation mapping is invalid. Wickerson et al. [55] have encoded non-transactional compilation mappings; we only need to extend them to handle transactions, which we do by requiring π to preserve all stxn -edges:

$$\text{stxn}_Y = \pi^{-1}; \text{stxn}_X; \pi.$$

Memalloy confirmed that compilation to x86, Power, and ARMv8 is sound for all C++ executions with up to 6 events.

8.3 Checking Lock Elision

We now check the soundness of lock elision in x86, Power, and ARMv8 using the technique proposed in §4.3.

First, we extend executions with four new event types:

- L , the $\text{lock}()$ calls that will be implemented by acquiring the lock in the ordinary fashion,
- U , the corresponding $\text{unlock}()$ calls,
- L^t , the $\text{lock}()$ calls that will be transactionalised, and
- U^t , the corresponding $\text{unlock}()$ calls.

When generating candidate executions from programs, we assume that each $\text{lock}()$ / $\text{unlock}()$ pair gives rise to a L - U pair or a L^t - U^t pair. (Distinguishing these two modes at the execution level eases the definition of the mapping relation.) We obtain from these lock/unlock events a derived scr relation that forms an equivalence class among all the events in the same CR. Similarly, scr^t is a subrelation of scr that comprises just those CRs that are to be transactionalised.

Table 3. Key constraints on π for defining lock elision

Source event, e	Target event(s), $\pi(e)$			
	x86	Power	ARMv8	ARMv8 (fixed)
L	$R \downarrow \text{ctrl}$	$R \downarrow \text{rmw, ctrl}$	$R, \text{Acq} \downarrow \text{rmw, ctrl}$	$R, \text{Acq} \downarrow \text{rmw, ctrl}$
	$R \downarrow \text{rmw}$	$W \downarrow \text{ctrl}$	W	$W \downarrow \text{po}$
	W	isync		dmb
U	W	W	W, Rel	W, Rel
L^t	R	R	R	R
U^t	-	-	-	-

Moreover:

$$\text{sloc}_Y = I^2 \cup ((\neg I)^2 \cap (\pi^{-1}; \text{sloc}_X; \pi)) \quad (\text{LOCKVAR})$$

$$\text{where } I = \pi(L \cup U \cup L^t \cup U^t)$$

$$\text{scr}^t \setminus (\neg U^t)^2 = \pi; \text{stxn}_Y; \pi^{-1} \quad (\text{TXNINTRO})$$

$$\text{empty}([L]; \pi; \text{rf}; \pi^{-1}; [L^t]) \quad (\text{TXNREADSLOCKFREE})$$

Second, we extend execution well-formedness so that every L event must be followed by a U event without an intervening L^t or U^t , and so on.

Third, the consistency predicates from Figs. 5, 6, and 8 are extended with the following axiom that forces the serialisability of CRs.

$$\text{acyclic}(\text{weaklift}(\text{po} \cup \text{com}, \text{scr})) \quad (\text{CROORDER})$$

Finally, we define a mapping π from the events of an ‘abstract’ execution X to those of a ‘concrete’ execution Y , that captures the implementation of lock elision. Table 3 sketches the main constraints we impose on π so that it captures lock elision for x86, Power, and ARMv8. It preserves all the execution structure except for lock/unlock events. The LOCKVAR constraint imposes that all the reads/writes in Y that are introduced by the mapping (call these I) access the same location (i.e., the lock variable) and that this location is not accessed elsewhere in Y . The TXNINTRO constraint imposes that events in the same transactionalised CR in X become events in the same transaction in Y . The L and U events are mapped to sequences of events that represent the recommended spinlock implementation for each architecture. Each L event maps to a successful RMW on the lock variable, which in ARMv8 is an acquire-RMW [7, §K.9.3.1], in Power is followed by a control dependency³ and an isync [29, §B.2.1.1], and in x86 is preceded by an additional read (the ‘test-and-test-and-set’ idiom) [31, §8.10.6.1]. Each U event maps to a write on the lock variable, which in ARMv8 is a release-write [7, §K.9.3.2], and in Power is

³In Power, ctrl edges can begin at a store-exclusive [47].

Figure 10. A pair of executions that demonstrates lock elision being unsound in ARMv8

preceded by a sync [29, §B.2.2.1]. Each L^t event maps to a read of the lock variable. This read does not observe a write from an L event (TxnReadsLockFree), to ensure that it sees the lock as free. Finally, U^t events vanish (because we do not have explicit events for beginning/ending transactions).

Figure 10 shows a pair of ARMv8 executions, X (left) and Y (right), and a π relation (dotted arrows), that satisfy all of the constraints above. From this example, which was automatically generated using Memalloy in 63 seconds, we manually constructed the pair of litmus tests shown in Example 1.1. It thus demonstrates that lock elision is unsound in ARMv8. This example is actually one of several found by Memalloy; we provide another example in our companion material.

We also used Memalloy to check lock elision in x86 and Power, and again in ARMv8 after applying the fix proposed in §1.1 (appending a DMB to the lock() implementation). Given that each architecture implements L events with a different number of primitive events (Tab. 3), we ensured that the event count was large enough in each case to allow examples similar to Fig. 10 to be found. We were unable to find bugs in any of these cases, but Memalloy timed out before it could verify their absence. As such, we cannot claim lock elision in x86 and Power to be *verified*, but the timeout provides a high degree of confidence that these designs are bug-free up to the given bounds because, as Tab. 2 shows, when counterexamples exist they tend to be found quickly.

9 Related Work

In concurrent but independent work, Dongol et al. [23] have also proposed adding TM to the x86, Power, and ARMv8 memory models. Like us, Dongol et al. build their axioms by lifting relations from events to transactions. However, their models are significantly weaker than ours, because they capture only the *atomicity* of transactions, not the *ordering* of transactions. Because of this, their Power model is too weak to validate the natural compiler mapping from C++. This is demonstrated by the following execution, which is forbidden by C++ (owing to an *hb* cycle), but allowed by their Power model (though not actually observable on hardware).

Moreover, unlike our work, Dongol et al.'s models have not been empirically validated – and nor have earlier models that combine TM and weak memory [20, 41]. Nonetheless, our models being stronger than Dongol et al.'s implies that our endeavours are complementary: our experiments validate their models, and their proofs carry over to our models.

Cerone et al. [16] have studied the weak consistency guarantees provided by transactions in database systems. A key difference is that for Cerone et al., weak behaviours are attributed to weakly consistent transactions, but in our work, weak behaviours are attributed to weakly consistent non-transactional events surrounding strongly consistent transactions. Nonetheless, similar axiomatisations can be used in both settings, and similar weak behaviours can manifest.

Our models follow the axiomatic style, but there also exist operational memory models for x86 [44], Power [48], ARMv8 [45], and C++ [43]. It would be interesting to consider how these could be extended to handle TM.

Other architectures and languages that could be targeted by our methodology include RISC-V, which plans to incorporate TM in the future [54], and Java. Indeed, Grossman et al. [25] and Shpeisman et al. [51] identify several tricky corner cases that arise when attempting to extend Java's weak memory model to handle transactions, and our methodology can be seen as a way to automate the generation of these.

Regarding the analysis of programs that *provide* TM, an automatic tool for testing (software) TM implementations above a weak memory model has been developed by Manovit et al. [42]. Like us, they use automatically-generated litmus tests to probe the implementations, but where our test suites are constructed to be exhaustive and to contain only 'interesting' tests, their tests are randomly generated. Regarding the analysis of programs that *use* TM, we note that the formulation of the C++ memory model by Lahav et al. [38] leads to an efficient model checker for multithreaded C++ Kokologianakis et al. [37]. Since our C++ TM model builds on Lahav et al.'s model, it may be possible to get a model checker for C++ TM similarly.

Regarding tooling for axiomatic memory models in general: our methodology builds on tools due to Wickerson et al. [55] and Lustig et al. [40], both of which build on Alloy [35]. Related tools include Diy [3], which generates litmus tests by enumerating relaxations of SC. Compared to Diy, Memalloy is more easily extensible with constructs such as transactions, and only generates the tests needed to validate a model. MemSynth [14] can synthesise memory models from a corpus of litmus tests and their expected outcomes, though it does not currently handle software models or control dependencies.

10 Conclusion

We have extended axiomatic memory models for x86, Power, ARMv8, and C++ to support transactions. Using our extensions to Memalloy, we synthesised meaningful sets of litmus tests that precisely capture the subtle interactions between weak memory and transactions. These tests allowed us to validate our new models by running them on available hardware, discussing them with architects, and checking them against technical manuals. We also used Memalloy to check several metatheoretical properties of our models, including the validity of program transformations and compiler mappings, and the correctness – or lack thereof – of lock elision.

Acknowledgements

We are grateful to Stephan Diestelhorst, Matt Horsnell, and Grigorios Magklis for extensive discussions of TM and the ARM architecture, to Nizamudheen Ahmed and Vishwanath HV for RTL testing, and to Peter Sewell for letting us access his Power machine. We thank the following people for their insightful comments on various drafts of this work: Mark Batty, Andrea Cerone, George Constantinides, Stephen Dolan, Alastair Donaldson, Brijesh Dongol, Hugues Evrard, Shaked Flur, Graham Hazel, Radha Jagadeesan, Jan Kończak, Dominic Mulligan, Christopher Pulte, Alastair Reid, James Riely, the anonymous reviewers, and our shepherd, Julian Dolby. This work was supported by an Imperial College Research Fellowship and the EPSRC (EP/K034448/1).

References

- [1] Allon Adir, Dave Goodman, Daniel Hershovich, Oz Hershkovitz, Bryan Hickerson, Karen Holtz, Wisam Kadry, Anatoly Koyfman, John Ludden, Charles Meissner, Amir Nahir, Randall R. Pratt, Mike Schliffli, Brett St Onge, Brian Thompston, Elena Tsanko, and Avi Ziv. 2014. Verification of Transactional Memory in POWER8. In *Design Automation Conference (DAC)*. <https://doi.org/10.1145/2593069.2593241>
- [2] Sarita V. Adve and Mark D. Hill. 1990. Weak Ordering - A New Definition. In *Int. Symp. on Computer Architecture (ISCA)*. <https://doi.org/10.1145/325096.325100>
- [3] Jade Alglave, Luc Maranget, Susmit Sarkar, and Peter Sewell. 2010. Fences in Weak Memory Models. In *Computer Aided Verification (CAV)*. https://doi.org/10.1007/978-3-642-14295-6_25
- [4] Jade Alglave, Luc Maranget, Susmit Sarkar, and Peter Sewell. 2011. Litmus: Running Tests Against Hardware. In *Int. Conf. on Tools and Algorithms for Construction and Analysis of Systems (TACAS)*. https://doi.org/10.1007/978-3-642-19835-9_5
- [5] Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding cats: modelling, simulation, testing, and data-mining for weak memory. *ACM Trans. on Programming Languages and Systems (TOPLAS)* 36, 2 (2014). <https://doi.org/10.1145/2627752>
- [6] Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding cats: modelling, simulation, testing, and data-mining for weak memory (online companion material). (2014). <http://moscova.inria.fr/~maranget/cats/model-power/all.html#sec4>
- [7] ARM. 2017. *ARMv8 Architecture Reference Manual*. https://static.docs.arm.com/ddi0487/b/DDI0487B_a_armv8_arm.pdf
- [8] Mark Batty, Mike Dodds, and Alexey Gotsman. 2013. Library Abstraction for C/C++ Concurrency. In *ACM Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/2429069.2429099>
- [9] Mark Batty, Alastair F. Donaldson, and John Wickerson. 2016. Overhauling SC atomics in C11 and OpenCL. In *ACM Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/2914770.2837637>
- [10] Mark Batty, Kayvan Memarian, Kyndylan Nienhuis, Jean Pichon-Pharabod, and Peter Sewell. 2015. The Problem of Programming Language Concurrency Semantics. In *Europ. Symp. on Programming (ESOP)*. https://doi.org/10.1007/978-3-662-46669-8_12
- [11] Armin Biere. 2010. *Lingeling, Plingeling, PicoSAT and PrecoSAT at SAT Race 2010*. Technical Report 10/1. Institute for Formal Models and Verification, Johannes Kepler University. <http://fmv.jku.at/papers/Biere-FMV-TR-10-1.pdf>
- [12] Colin Blundell, E. C. Lewis, and Milo M. K. Martin. 2006. Subtleties of Transactional Memory Atomicity Semantics. *IEEE Computer Architecture Letters* 5, 2 (2006). <https://doi.org/10.1109/L-CA.2006.18>
- [13] Hans-J. Boehm and Sarita V. Adve. 2008. Foundations of the C++ Concurrency Memory Model. In *ACM Conf. on Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/1379022.1375591>
- [14] James Bornholt and Emina Torlak. 2017. Synthesizing Memory Models from Framework Sketches and Litmus Tests. In *ACM Conf. on Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/3062341.3062353>
- [15] Harold W. Cain, Brad Frey, Derek Williams, Maged M. Michael, Cathy May, and Hung Le. 2013. Robust Architectural Support for Transactional Memory in the Power Architecture. In *Int. Symp. on Computer Architecture (ISCA)*. <https://doi.org/10.1145/2485922.2485942>
- [16] Andrea Cerone, Giovanni Bernardi, and Alexey Gotsman. 2015. A Framework for Transactional Consistency Models with Atomic Visibility. In *Int. Conf. on Concurrency Theory (CONCUR)*. <https://doi.org/10.4230/LIPIcs.CONCUR.2015.58>
- [17] JaeWoong Chung, Michael Dalton, Hari Kannan, and Christos Kozyrakis. 2008. Thread-Safe Dynamic Binary Translation using Transactional Memory. In *Int. Symp. on High Performance Computer Architecture (HPCA)*. <https://doi.org/10.1109/HPCA.2008.4658646>
- [18] William W. Collier. 1992. *Reasoning about Parallel Architectures*. Prentice Hall.
- [19] Luke Dalessandro and Michael L. Scott. 2009. Strong Isolation is a Weak Idea. In *ACM Workshop on Transactional Computing (TRANSACT)*. http://transact09.cs.washington.edu/33_paper.pdf
- [20] Luke Dalessandro, Michael L. Scott, and Michael F. Spear. 2010. Transactions as the Foundation of a Memory Consistency Model. In *Int. Conf. on Distributed Computing (DISC)*. https://doi.org/10.1007/978-3-642-15763-9_4
- [21] Will Deacon. 2016. The ARMv8 Application Level Memory Model. <https://github.com/herd/herdtools7/blob/master/herd/libdir/aarch64.cat>. (2016).
- [22] Dave Dice, Yossi Lev, Mark Moir, and Dan Nussbaum. 2009. Early Experience with a Commercial Hardware Transactional Memory Implementation. In *Int. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. <https://doi.org/10.1145/2528521.1508263>
- [23] Brijesh Dongol, Radha Jagadeesan, and James Riely. 2018. Transactions in Relaxed Memory Architectures. In *ACM Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/3158106>
- [24] Shaked Flur, Kathryn E. Gray, Christopher Pulte, Susmit Sarkar, Ali Sezgin, Luc Maranget, Will Deacon, and Peter Sewell. 2016. Modelling the ARMv8 Architecture, Operationally: Concurrency and ISA. In *ACM Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/2837614.2837615>
- [25] Dan Grossman, Jeremy Manson, and William Pugh. 2006. What Do High-Level Memory Models Mean for Transactions?. In *ACM Workshop on Memory Systems Performance & Correctness (MSPC)*. <https://doi.org/10.1145/1178597.1178609>

- [26] Mark Hachman. 2014. Intel finds specialized TSX enterprise bug on Haswell, Broadwell CPUs. *PCWorld* (August 2014). <http://www.pcworld.com/article/2464880>
- [27] Tim Harris, James Larus, and Ravi Rajwar. 2010. *Transactional Memory* (2nd ed.). Morgan & Claypool. <https://doi.org/10.2200/S00272ED1V01Y201006CAC011>
- [28] Maurice Herlihy and J. Eliot B. Moss. 1993. Transactional Memory: Architectural Support for Lock-Free Data Structures. In *Int. Symp. on Computer Architecture (ISCA)*. <https://doi.org/10.1145/173682.165164>
- [29] IBM. 2015. *Power ISA (Version 3.0)*.
- [30] Intel. 2017. 6th Generation Intel Processor Family: Specification Update. (June 2017). <https://www3.intel.com/content/dam/www/public/us/en/documents/specification-updates/desktop-6th-gen-core-family-spec-update.pdf>
- [31] Intel. 2017. *Intel 64 and IA-32 Architectures: Software Developer's Manual*. <https://software.intel.com/en-us/articles/intel-sdm>
- [32] Intel Developer Zone. 2012. Transactional Synchronization in Haswell. (February 2012). <https://software.intel.com/en-us/blogs/2012/02/07/transactional-synchronization-in-haswell>
- [33] ISO/IEC. 2011. *Programming languages – C++*. International standard 14882:2011.
- [34] ISO/IEC. 2015. *Technical Specification for C++ Extensions for Transactional Memory*. Draft technical specification. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2015/n4514.pdf>
- [35] Daniel Jackson. 2012. *Software Abstractions – Logic, Language, and Analysis* (revised ed.). MIT Press.
- [36] Eric H. Jensen, Gary W. Hagensen, and Jeffrey M. Broughton. 1987. *A New Approach to Exclusive Data Access in Shared Memory Multiprocessors*. Technical Report 97663. Lawrence Livermore National Laboratory. <https://e-reports-ext.llnl.gov/pdf/212157.pdf>
- [37] Michalis Kokologianakis, Ori Lahav, Konstantinos Sagonas, and Viktor Vafeiadis. 2018. Effective Stateless Model Checking for C/C++ Concurrency. In *ACM Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/3158105>
- [38] Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. Repairing Sequential Consistency in C/C++11. In *ACM Conf. on Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/3062341.3062352>
- [39] Leslie Lamport. 1979. How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs. *IEEE Trans. Comput.* C-28, 9 (1979). <https://doi.org/10.1109/TC.1979.1675439>
- [40] Daniel Lustig, Andrew Wright, Alexandros Papakonstantinou, and Olivier Giroux. 2017. Automated Synthesis of Comprehensive Memory Model Litmus Test Suites. In *Int. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. <https://doi.org/10.1145/3037697.3037723>
- [41] Jan-Willem Maessen and Arvind. 2007. Store Atomicity for Transactional Memory. *Electronic Notes in Theoretical Computer Science* 174, 9 (2007). <https://doi.org/10.1016/j.entcs.2007.04.009>
- [42] Chaiyasit Manovit, Sudheendra Hangal, Hassan Chafi, Austen McDonald, Christos Kozyrakis, and Kunle Olukotun. 2006. Testing Implementations of Transactional Memory. In *Int. Conf. on Parallel Architectures and Compilation Techniques (PACT)*. <https://doi.org/10.1145/1152154.1152177>
- [43] Kyndylan Nienhuis, Kayvan Memarian, and Peter Sewell. 2016. An Operational Semantics for C/C++11 Concurrency. In *ACM Int. Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/3022671.2983997>
- [44] Scott Owens, Susmit Sarkar, and Peter Sewell. 2009. A Better x86 Memory Model: x86-TSO. In *Theorem Proving in Higher Order Logics (TPHOLs)*. https://doi.org/10.1007/978-3-642-03359-9_27
- [45] Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. 2018. Simplifying ARM Concurrency: Multicopy-atomic Axiomatic and Operational Models for ARMv8. In *ACM Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/3158107>
- [46] Ravi Rajwar and James R. Goodman. 2001. Speculative Lock Elision: Enabling Highly Concurrent Multithreaded Execution. In *Int. Symp. on Microarchitecture (MICRO)*. <https://doi.org/10.1109/MICRO.2001.991127>
- [47] Susmit Sarkar, Kayvan Memarian, Scott Owens, Mark Batty, Peter Sewell, Luc Maranget, Jade Alglave, and Derek Williams. 2012. Synchronising C/C++ and POWER. In *ACM Conf. on Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/2254064.2254102>
- [48] Susmit Sarkar, Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams. 2011. Understanding POWER Multiprocessors. In *ACM Conf. on Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/1993498.1993520>
- [49] Dennis Shasha and Marc Snir. 1988. Efficient and Correct Execution of Parallel Programs that Share Memory. *ACM Trans. on Programming Languages and Systems (TOPLAS)* 10, 2 (1988). <https://doi.org/10.1145/42190.42277>
- [50] Tatiana Shpeisman, Ali-Reza Adl-Tabatabai, Robert Geva, Yang Ni, and Adam Welc. 2009. Towards Transactional Memory Semantics for C++. In *Symp. on Parallelism in Algorithms and Architectures (SPAA)*. <https://doi.org/10.1145/1583991.1584012>
- [51] Tatiana Shpeisman, Vijay Menon, Ali-Reza Adl-Tabatabai, Steven Balensiefer, Dan Grossman, Richard L. Hudson, Katherine F. Moore, and Bratin Saha. 2007. Enforcing Isolation and Ordering in STM. In *ACM Conf. on Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/1273442.1250744>
- [52] Tyler Sorensen, Alastair F. Donaldson, Mark Batty, Ganesh Gopalakrishnan, and Zvonimir Rakamarić. 2016. Portable Inter-workgroup Barrier Synchronisation for GPUs. In *ACM Int. Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/2983990.2984032>
- [53] Srdan Stipić, Vesna Smiljković, Osman Unsal, Adrián Cristal, and Mateo Valero. 2013. Profile-Guided Transaction Coalescing—Lowering Transactional Overheads by Merging Transactions. *ACM Transactions on Architecture and Code Optimization* 10, 4 (2013). <https://doi.org/10.1145/2541228.2555306>
- [54] Andrew Waterman and Krste Asanović (Eds.). 2017. *The RISC-V Instruction Set Manual, Volume I: User-Level ISA (version 2.2)*. RISC-V Foundation. <https://riscv.org/specifications/>
- [55] John Wickerson, Mark Batty, Tyler Sorensen, and George A. Constantinides. 2017. Automatically Comparing Memory Consistency Models. In *ACM Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/3009837.3009838>
- [56] Michael Wong. 2014. Transactional Language Constructs for C++. In *C++ Conference (CppCon)*. <http://bit.ly/2tWk4uz>