

Correctness of Concurrent Objects under Weak Memory Models

Graeme Smith Kirsten Winter
Robert J. Colvin

School of Information Technology and Electrical Engineering,
The University of Queensland, Australia

smith@itee.uq.edu.au kirsten@itee.uq.edu.au r.colvin@uq.edu.au

In this paper we develop a theory for correctness of concurrent objects under weak memory models. Central to our definitions is the concept of *observations* which determine when effects of operations become visible, and hence determine the semantics of objects, under a given memory model. The resulting notion of correctness, called *object refinement*, is generic as it is parameterised by the memory model under consideration. Our theory enforces the minimal constraints on the placing of observations and on the semantics of objects that underlie object refinement. Object refinement is suitable as a reference for correctness when proving new proof methods for objects under weak memory models to be sound and complete.

1 Introduction

Linearizability [15] is widely accepted as the standard correctness criterion for concurrent objects, i.e., objects designed to be accessed simultaneously by multiple threads [16]. Recent work [5, 13, 7, 20, 8, 10] has begun examining the applicability of linearizability in the context of weak memory models of modern multicore architectures [17, 18, 2, 12, 6]. These memory models improve hardware efficiency by limiting accesses to global memory. Individual threads operate on local copies of global variables, updates to the global memory being made by the hardware and largely out of the programmer’s control¹. This can cause threads executing on different cores to get out of sync with respect to the values of global variables.

For example, on the TSO (Total Store Order) architecture [18] a thread updating a global variable x stores the new value in a per-core FIFO buffer. Threads executing on that core will then read x from the buffer, rather than the global memory, until the new value is flushed from the buffer to global memory by the hardware. In the meantime, threads on other cores read the value of x from the global memory or from their own core’s buffer when it has a value for x . The situation is even more complex for weaker architectures, like ARM and Power [2]. In these architectures, values of global variables are flushed independently (and at different times) to different cores (this is referred to as *non-multi-copy atomicity*).

For example, in Figure 1 three different cores are depicted with one or two threads running on each (as P is also a thread) and referring to a global variable x . Assume at time t_0 thread T1 performs a write to x , which then gets flushed to core2 at time t_1 and later, at time t_2 to core3. The evaluation of x per core at these time instances is depicted in the table, which shows that for this scenario the value is consistent only at time t_2 (for scenarios other than the depicted one, there might not be a consistent state at all).

To evaluate proposed notions of linearizability for weak memory models we need to argue their soundness and completeness. That is, we need to prove that an object implementation linearizes to its

¹A programmer can add fences (or memory barriers) to code to force any pending updates to be written to memory. However, if used indiscriminately, fences cause the code to be less efficient.

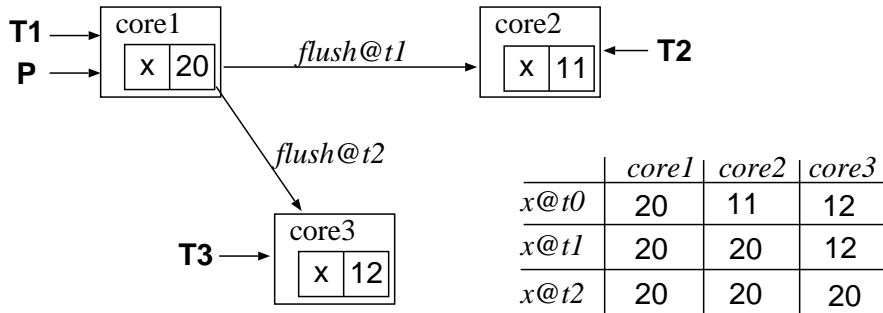


Figure 1: Scenario for non-multi-copy atomicity in ARM and Power architectures

specification if and only if it is *correct* with respect to its specification. However, this requires a notion of correctness as a reference point. What does it mean for an implementation to be correct when executed on a weak memory model?

Traditionally, trace refinement provides this notion of correctness for programs: an implementation is correct with respect to its specification if and only if each observable behaviour of the implementation can also be observed from the specification [4, 3, 1]. In the context of concurrent objects that are called by a *client program*, an object implementation is deemed correct if and only if the client program cannot differentiate between the object implementation and its abstract specification, as the observable behaviour is the same.

To capture what is observed by client programs under *sequentially consistent* (SC) architectures (i.e., those without a weak memory model), the notions of *observational refinement* [11] and *contextual refinement* [9] have been introduced. However, both of these definitions do not provide a notion of correctness under weak memory models, as they do not take into account that an event occurring on one core might be observable later on another. Instead events become observable immediately after their occurrence. For architectures with non-multi-copy atomicity, in which flushes occur out of write order, the definitions do not provide us with the right semantics of programs.

In this paper, we propose a general definition of correctness for concurrent objects running on *any* existing memory model. The aim is to provide a reference point for proving new notions of object correctness for weak memory models sound and complete. The result is a notion of object refinement which is parameterised by the memory model it refers to, and is therefore generic and can be instantiated for any weak memory model behaviour. Key to this result is a definition of the semantics of an object operating in the context of a calling client program under a weak memory model. In particular the semantics of an object's specification in the context of a client program under a weak memory model needs to be defined in such a way that its behaviour maintains the intention of the specification, namely that operations are atomic.

The paper is structured as follows. In Section 2 we introduce the basic concepts of our theory including that of *observations* which is key to our definitions. Based on these concepts, Section 3 formalises the semantics of programs under a given memory model. The semantics of a concurrent object in the context of a client, under a memory model, is elaborated in Section 4, distinguishing cases for the specification, which is atomic, and the implementation which includes non-atomic, and possibly non-terminating, operations. Section 5 ties these basics into the notion of *weak-memory trace refinement* which defines refinement under memory model M for a client program using an object and its specification, respectively. The definition is parameterised with a given memory model for which we assume a semantics is given.

Using weak-memory trace refinement we can then define our notion of *object refinement* under memory model M which delivers the notion of correctness for objects. Section 6, illustrates how our definition can be used to prove correctness of a case study. The paper concludes with a discussion of related work in Section 7.

2 Client programs and observation events

To investigate the behaviour of concurrent objects under weak memory models, and relate their implementation to their specification using refinement, we need to consider the calling context. Programs calling the operations of a concurrent object are referred to as *client programs*, or clients for short. A client program P is concurrent, spawning multiple threads T_i on multiple cores, and is affected by the memory model of the architecture it is running on. For some finite n , we have²

$$P \hat{=} T_1 \parallel T_2 \parallel \dots \parallel T_n$$

where P is a thread itself and might share a core with other threads.

Following other work on concurrent objects [15, 11, 9], the behaviour of a program is described in terms of *events* that occur. We allow events to be *program steps*, *operation events* or, as introduced in Section 2.1, *observation events*.

Program steps are steps, other than calling an object operation, performed by the client. These are assignments, conditional branch instructions (e.g., of if or while statements), other control instructions like various forms of fences, atomic read-write-modify instructions which atomically perform these three steps (e.g., the compare-and-swap construct CAS), and higher-level instructions which can, in many cases, be defined in terms of assignments and/or conditional branches. For example, a statement `await(z=1)` could be defined as `while(z ≠ 1) {}`.

Operation events abstract the effects of an operation call by a program. They include the *invocation* of the operation (i.e., when it is called) and the operation's *response* (i.e., when it returns). The operation events carry the operation's input and output values as parameters and thus reflect on the operation's externally visible behaviour. The internal behaviour of an object is elided.

All program steps and operations are deterministic; non-determinism in a program results from the interleaving of events on different threads varying between executions.

2.1 Observation events

Central to our definitions is the notion of an additional type of event called an *observation event*. Such an event denotes the point in an execution where a program step or the response of an operation can be deemed to have been observed by *all* threads T_i of client P . Such observed events are either

- program steps that write to global program variables, or
- responses of operations that write to object variables that are shared by threads accessing the object.

Note that generally operation responses as such are not observable, as the returned value (if any) is written to a thread-local variable only. A subsequent program step is required to write it to a global program variable [11, 19].

²For simplicity, we do not consider dynamically spawned threads.

Introducing observation events allows us to decouple the occurrence of an event and its observation which might not fall together under weak memory models. The semantics specific to the memory model under consideration determines the possible placings of observation events. For example, the operational semantics for TSO given in [17], or for ARM and Power (e.g., [6]), can produce all possible executions from which we can deduce the points of observations according to the following rules: Generally, an observation will occur when *all* threads can either

- (a) access a new value of a global program variable written by a program step,
- (b) access the values of *all* shared object variables written by an operation, *after* the response of the operation has occurred, or
- (c) when the response of a *covert* operation, whose steps or outcomes are not directly visible, becomes apparent to the client *P* because an observable step later in the program order became visible (or the client terminates if there is no later step).

Covert operations (as referred to in case (c)) are those which do not write to any shared object variables and either have no return value or have a return value which is not written to a global program variable ([11] does not consider these operations). The observation of the response of a covert operation will occur at the same time as the next observation that, due to the order enforced by the program, must have occurred later. Examples of cases (a) and (b) are simply the flushes of global program and shared object variables. If an operation writes to multiple shared object variables, the observation of the last write is considered. Note that observations of writes in an operation (as in case (b)) or covert operation responses (as in case (c)) can only occur after the response of the operation.

The effects of concurrency (interleaving the behaviour of threads) as well as the effects of weak memory models leads to a partial ordering of these events to which all possible executions must adhere.

2.2 Examples


Consider the example of a client program given in Figure 2 with two global program variables x and z , running on two threads T1 and T2. The threads call three operations of the same object: Operation A writes to a shared object variable and operations B and C read this object variable.

On an SC architecture (i.e., one without a weak memory model), writes to global variables occur instantaneously. Hence, the observation event for a program step occurs immediately after the program step, and that of an operation which writes to shared object variables immediately after its response. Note that we assume that the result of any operation call is implicitly stored in a local register (e.g., r_A for the operation call `object.A`).


The partial order of the events of the program of Figure 2 on SC is shown below, where inv_A denotes the invocation of `object.A`, etc.

T1		T2
$x := \text{object.A}();$ $z := 1;$ $\text{object.B}();$		$\text{await}(z=1);$ $\text{object.C}();$

Figure 2: Client program



On TSO, writes to global variables and shared object variables become available to threads on other cores when they are flushed. Furthermore, flushes occur in the same order as the writes occurred. Hence, the observation event of a program step occurs at the point when the global variable is flushed, and the observation of an operation that writes to shared object variables occurs at the final flush associated with the writes, or immediately after the response when the final flush occurs before the response. The partial order on events of the program of Figure 2 on TSO is



Note that obs_A occurs before $obs_{x=r_A}$ which occurs before $obs_{z=1}$ which occurs before obs_B due to the FIFO order of the per-core buffers on TSO.

3 Semantics of programs relative to the memory model

The semantics of a program P under memory model M is defined in terms of the set of *events* that can occur, and the *partial order* on the occurrence of those events. Note that the partial order, and in particular the placement of observation events, is determined by the memory model as outlined in Section 2.

The partial order of events is partly enforced by P through the program text; this is referred to as the *program order*. Additionally, it reflects the semantics of the memory model. For example, under TSO the order of writes to variables is determined by the program order, and the order of observations of such writes is the same as the order of the writes (due to the FIFO nature of the store buffer).

We formalise the semantics of programs as follows. Let T be the set of all thread ids, and $Call$ the set of all operation calls. An operation is then defined as a call by a particular thread.

$$Op \hat{=} T \times Call$$

Let PS denote the set of all program step events, and Val the set of all values (of input and output parameters) including a special element \perp meaning ‘no value’. The set of all events is defined as follows where each invocation is associated with an input, and each response and observation with an output.

$$Event \hat{=} step(T, PS) \mid obs(T, PS) \mid inv(Op, Val) \mid res(Op, Val) \mid obs(Op, Val)$$

In the remainder of this paper, we refer to $step(T, PS)$ and $obs(T, PS)$ as *program events*, and to $inv(Op, Val)$, $res(Op, Val)$, and $obs(Op, Val)$ as *object events*.

A program P has a set of events, $events(P)$, such that for each invocation event in the events of P , $events(P)$ also contains the corresponding response and observation events. That is, each called operation can respond and be observed.

$$\begin{aligned} \forall op : Op; in : Val \bullet inv(op, in) \in events(P) \Rightarrow \\ \exists out : Val \bullet \{res(op, out), obs(op, out)\} \subseteq events(P) \end{aligned} \quad (1)$$

Since the observation of an operation can occur immediately after its response (or when the operation does not write to shared object variables, at the same time as the observation of a subsequent event), $<_{P_M}$

cannot enforce the observation of a later operation to occur before an earlier one. Of course, a later operation *can* be observed before an earlier operation (in particular under ARM and Power), but the point we are making here is that although the memory model might allow this changed order, it cannot enforce that the order of observations *must* be different to the order of corresponding events. Note that \langle_{P_M} contains only those constraints to the order that every possible behaviour adheres to.

$$\forall c, d : Op \times Val \bullet (res(c), inv(d)) \in \langle_{P_M} \Rightarrow (obs(d), obs(c)) \notin \langle_{P_M} \quad (2)$$

3.1 Traces

The semantics of a program P is a set of *finite* sequences of events, referred to as *traces*.³ For each trace t , each event is unique (similar events, e.g., calls to the same operation, may be annotated by their relative position in the trace), and an invocation of an operation always occurs before the associated response, which in turn occurs before the associated observation. Similarly, a program step always occurs before its observation, if any. In the following s_i denotes the i th element of a sequence s , and $\#s$ its length.

$$\begin{aligned} Trace \hat{=} \{ & t : seq Event \mid (\forall i, j \leq \#t \bullet i \neq j \Rightarrow t_i \neq t_j) \wedge \\ & \forall c : Op \times Val \bullet (\forall j \bullet t_j = res(c) \Rightarrow \exists i < j \bullet t_i = inv(c)) \wedge \\ & (\forall j \bullet t_j = obs(c) \Rightarrow \exists i < j \bullet t_i = res(c)) \wedge \\ & \forall s : T; p : PS \bullet (\forall j \bullet t_j = obs(s, p) \Rightarrow \exists i < j \bullet t_i = step(s, p)) \} \end{aligned}$$

The events of a trace and the order on these events are defined as follows. Note that the order \langle_t is a total order over the events in t , as a trace describes exactly one execution.

$$\begin{aligned} events(t) \hat{=} \{ & a : Event \mid \exists i \bullet t_i = a \} \\ \langle_t \hat{=} \{ & (a, b) : Event \times Event \mid \exists i, j \bullet i < j \wedge t_i = a \wedge t_j = b \} \end{aligned}$$

The semantics of program P on memory model M is then defined as the set of traces using only events from P and whose orders adhere to the constraints prescribed by the partial order \langle_{P_M} . We introduce the relation \Subset between orders which specifies whether an order is *allowed* by P on M .

$$\llbracket P \rrbracket_M \hat{=} \{ t : Trace \mid events(t) \subseteq events(P) \wedge \langle_{P_M} \Subset \langle_t \}$$

where $\langle_{P_M} \Subset \langle_t$ is defined as

$$\langle_{P_M} \Subset \langle_t \hat{=} \forall (a, b) : \langle_{P_M} \bullet b \in events(t) \Rightarrow (a, b) \in \langle_t$$

That is, for any event b that occurs in trace t , if this event is enforced to come after another event a by \langle_{P_M} , then event a must have also occurred in t before event b . Note that it is not suitable to use the simple subset relation here, i.e., $\langle_{P_M} \subseteq \langle_t$, since trace t will not, in general, include all events b that are restricted by \langle_{P_M} .

4 Semantics of objects under weak memory models

To define trace refinement between client programs using objects, we need to constrain the behaviour of the program to a particular object or collection of objects. If the program uses a collection of interacting

³Since we are interested in defining a notion of correctness that readily relates to linearizability, and hence only safety properties [14, 19], we do not consider infinite sequences of events in our semantics.

objects, we simply consider the collection of operations (and their events) provided by all objects in the collection. Below we consider a single object only.

The semantics of an object is given as a set of *histories*, where each history is a trace with only object events.

$$\text{History} \hat{=} \{t : \text{Trace} \mid t|_o = t\}$$

where $t|_o$ denotes the trace t restricted to only object events.

4.1 Object implementation under weak memory models

An object implementation C has a set of object events, $\text{events}(C)$, and, on a particular memory model M , a prefix-closed set of histories made up of those events, $\llbracket C \rrbracket_M$. Observation events are not controlled by the object and hence can occur at any time after the associated response event that the memory model allows.

For any object implementation C , $P[C]$ denotes the object C operating in program P . It is only defined when all object events of P are events of C . The semantics of C operating in P on memory model M , $\llbracket P[C] \rrbracket_M$, is given as those traces of P on M whose object events correspond to a history of C on M .

$$\llbracket P[C] \rrbracket_M \hat{=} \{t : \llbracket P \rrbracket_M \mid \exists h : \llbracket C \rrbracket_M \bullet t|_o = h\}$$

4.2 Object specification under weak memory models

An object specification A similarly has a set of object events, $\text{events}(A)$, and a prefix-closed set of histories, $\llbracket A \rrbracket$. Since A represents a typical specification found in a software library, its set of histories is *independent* of the memory model (hence there is no subscript). Any weak memory model behaviour is absent from its histories due to its operations being *atomic*, i.e., they occur without interference from other operations.

To capture this in our semantics, the histories of A are restricted to those where only operations on one core are active at a time. For example, suppose the specification of a lock object, `lock`, has an operation `acquire` which waits until the value of a variable of the object, x , is 1 and sets it to 0, i.e., `acquire` is specified as `await(x=1); x=0`. Assuming x is initially 1, in the program of Figure 3 the intention would be that only one of y or z would be set to 1.

On SC, this intention is achieved when the invocation of `acquire` which occurs second does not happen until after the response of the `acquire` which occurs first. On TSO and assuming T1 and T2 are running on different cores, the intention is only achieved when the invocation of `acquire` which occurs second does not happen until after the flush of x from the `acquire` which occurs first. In both cases, the intention is met when the second occurrence of `acquire` is not invoked before the observation event of the first occurrence. In general, for any object specification to behave as intended, an operation invocation




Figure 3: Client program using lock

on one core does not occur before the observation event of a previously invoked operation on any other core. (The same does not need to hold for operations on the same core, as values are read locally.)

Let $core(a)$ denote the core on which an event a occurs. We assume the core can be determined from the thread of the event.

$$\begin{aligned} \forall h : \llbracket A \rrbracket \bullet \forall c : Op \times Val; i \bullet h_i = inv(c) \Rightarrow \\ \forall op : Op; in : Val; j < i \bullet h_j = inv(op, in) \wedge core(inv(op, in)) \neq core(inv(c)) \Rightarrow \\ \exists out : Val; k < i \bullet h_k = obs(op, out) \end{aligned} \quad (3)$$

Additionally, since operations in a specification are intended to always respond, for each history h of A with *pending* invocations, i.e., invocations for which there is no response, the history which extends h with the missing events is also in $\llbracket A \rrbracket$.

$$\begin{aligned} \forall h : \llbracket A \rrbracket; c : Op \times Val \bullet \exists i \bullet h_i = inv(c) \wedge (\nexists j \bullet h_j = res(c)) \Rightarrow \\ \exists h \wedge hr : \llbracket A \rrbracket; j \bullet (h \wedge hr)_j = res(c) \end{aligned} \quad (4)$$

This is not the case for object implementations which may have operations which never return (e.g., due to an infinite loop).

Provided all object events of a program P are events of a specification A , $P[A]$ denotes the program P operating with an object whose behaviour satisfies A . The semantics of P imposes restrictions from the memory model M on to the traces and restricts the placement of observation events. The semantics of $P[A]$ is given as follows.

$$\llbracket P[A] \rrbracket_M \hat{=} \{t : \llbracket P \rrbracket_M \mid \exists h : \llbracket A \rrbracket \bullet t|_o = h\}$$

5 Object refinement

Correctness of an object is defined from the client program's point of view. Such a program can only observe changes to *program variables*, i.e., variables that are not defined locally on a thread or as part on an object. Let $t|_{global}$ denote the *observable behaviour* of a trace t , i.e., the sequence of observation events of program steps which write to global program variables.

A program P using C on memory model M refines P using A on M when any observable behaviour of the former is a possible observable behaviour of the latter. We refer to this property as *weak-memory trace refinement*.

Definition 1 *Weak-memory trace refinement*

$$P[A] \sqsubseteq_M P[C] \hat{=} \forall t : \llbracket P[C] \rrbracket_M \bullet \exists t' : \llbracket P[A] \rrbracket_M \bullet t'|_{global} = t|_{global}$$

An object implementation C refines an object specification A under weak memory model M if for all possible client programs P , P using C weak-memory trace refines P using A under memory model M . We refer to this property as *object refinement*.

Definition 2 *Object Refinement under memory model M*

$$A \sqsubseteq_M C \hat{=} \forall P \bullet P[A] \sqsubseteq_M P[C]$$

If $A \sqsubseteq_M C$ we say that C is correct with respect to A under memory model M .

6 Example application

6.1 Correctness on TSO

Consider a spinlock object with operations `acquire`, `release` and `tryAcquire` specified as follows.

<pre>acquire await(x=1); x=0</pre>	<pre>release x=1;</pre>	<pre>tryAcquire if (x=1) x=0; return 1 else return 0</pre>
--	---------------------------	--

A typical concurrent implementation which is correct on SC is

<pre>acquire while (true) { if (TAS(x, 1, 0)=1) return while (x=0) {} }</pre>	<pre>release x=1;</pre>	<pre>tryAcquire return TAS(x, 1, 0)</pre>
---	---------------------------	---

where $TAS(x,a,b)$ is the atomic hardware primitive test-and-set which, when x is a , sets x to b and returns 1, and otherwise returns 0. The TAS instruction has a built in fence to ensure any change it makes to x is immediately visible to all threads.

An early version of linearizability on TSO [7] proved that this implementation is also correct on TSO. However, using our definition of object refinement we can show that, in fact, it is not correct, and hence that the definition of linearizability in [7] is unsound.

Consider the client program in Figure 4, which uses a spinlock object `sl`, in which we assume that initially $x = 1$ and $z = 0$. One possible trace of this program is

$$\langle inv(T2, sl.acquire), \perp, res((T2, sl.acquire), \perp), obs((T2, sl.acquire), \perp), inv((T2, sl.release), \perp), res(T2, sl.release, \perp), step(T2, y = 0), step(T1, z = 1), obs(T1, z = 1), step(T3, await(z = 1)), inv((T3, sl.tryAcquire), \perp), res((T3, sl.tryAcquire), 0), obs((T3, sl.tryAcquire), 0), step(T3, w = 0), obs(T3, w = 0), obs((T2, obs.release), \perp), obs(T2, y = 0) \rangle$$

This trace corresponds to thread T2 acquiring and releasing the lock and reading the initial value of z , but not flushing the value written to x by the release operation until after the other two threads have run to completion. The observable behaviour of the trace is

$$\langle obs(T1, z = 1), obs(T3, w = 0), obs(T2, y = 0) \rangle$$

This is not an observable trace of the program running with an object satisfying the specification: if

T1	T2	T3
$z = 1;$	$sl.acquire();$	$await(z=1);$
	$sl.release();$	$w = sl.tryAcquire();$
	$y=z;$	

Figure 4: Program using spinlock

T1		T2
sl.acquire();		sl.acquire();
y= y + 1;		y=y + 1;
sl.release()		sl.release()

Figure 5: Another program using spinlock

$y=0$ then this step and hence `sl.release` on T2 must have occurred before $z=1$ on T1, and hence before `sl.tryAcquire` on T3. Hence, we do not have object refinement.

The spinlock implementation without the `tryAcquire` operation is, however, known to be correct on TSO [18]. Again we can show this using our definition.

The traces of the implementation can be derived from the operational semantics of TSO in [17]. These show that if an `acquire` has responded (and hence has been observed due to the fence in the TAS) then another `acquire` cannot respond until after a `release` on the same core has responded or a `release` on another core has been observed. This coincides with what can be observed from the abstract specification. Hence, object refinement holds.

6.2 Correctness on Power and ARM

On the Power and ARM architectures [2, 12, 6], writes to variables are, like on TSO, local to the core on which they occur and are made available to other cores by the hardware or the use of fence instructions in the program. However, they are not necessarily made available in FIFO order. A write to variable x may be made available globally after a write to variable y when the write to x occurs before the write to y in program order.

It is easy to show that the spinlock implementation of Section 6.1, even without the `tryAcquire` operation, is not correct on such an architecture using our definition of object refinement. For example, consider the client program in Figure 5 for which we assume that initially $x = 1$ and $y = 0$.

Following the operational semantics of ARM and Power given in [6], one possible trace of this program is⁴

$$\langle \text{inv}((T1, \text{acquire}), \perp), \text{res}((T1, \text{acquire}), \perp), \text{obs}((T1, \text{acquire}), \perp), \text{step}(T1, y = 1), \\ \text{inv}((T1, \text{release}), \perp), \text{res}((T1, \text{release}), \perp), \text{obs}((T1, \text{release}), \perp), \text{inv}((T2, \text{acquire}), \perp), \\ \text{res}((T2, \text{acquire}), \perp), \text{obs}((T2, \text{acquire}), \perp), \text{step}(T2, y = 1), \text{inv}((T2, \text{release}), \perp), \text{res}((T2, \text{release}), \perp), \\ \text{obs}((T2, \text{release}), \perp), \text{obs}(T1, y = 1), \text{obs}(T2, y = 1) \rangle$$

This trace corresponds to the response of T1's `release` operation being observed before its update to y . This allows T2's `acquire` to occur followed by its update of y before T1's new value of y is observable by T2. Hence, both threads update y to 1.

Since the observable behaviour $\langle \text{obs}(T1, y = 1), \text{obs}(T2, y = 1) \rangle$ of the above trace is not possible using the specification, the implementation is not correct on Power or ARM: object refinement does not hold.

⁴In the operational semantics of [6] the placement of observations can be derived from the model of the “storage subsystem” which keeps track of which updates to global variables have been seen by which threads.

7 Conclusion

In this paper, we have defined object refinement which provides a reference point for correctness of concurrent objects on any weak memory model. This allows notions of linearizability on weak memory models to be proved sound and complete.

Observational refinement [11] and contextual refinement [9] are existing notions of correctness for client/object systems on SC architectures which coincide with our notion. At the fundamental level the meaning of a program is found in how it modifies the observable state. Where we differ with those frameworks is in how interactions between the client and object are treated. In particular, to handle operations taking effect outside of their invocation/response behaviour, we introduce the notion of observation events to abstract when a change to a global variable is visible.

Doherty and Derrick [8] and Dongol et al. [10] address linearizability for weak memory models. Both papers provide arguments that their definitions of linearizability are sound. At the level of client/object systems their notions of correctness are similarly based on changes to observable variables.

In Doherty and Derrick's framework the client's perspective using an object is not constrained to the atomic behaviour of an object's specification. This does not affect their results since the client programs are restricted. In particular, programs like in Figure 4 which have a race between the release and tryAcquire operations are not allowed. However, our goal is to provide a general notion of correctness for *any* client program.

Dongol et al. [10] integrate the semantics of allowed reorderings of invocation and response events with the reordering semantics framework of Alglave et al. [2]. The latter is a well accepted semantics which can be specialised to TSO as well as ARM and Power memory models. In their framework, traces of clients using specifications do not include constraints that the specifications place on the order of operations, e.g., that successive acquire operations are separated by a release operation. These constraints (captured by the specification order so) are added separately. A consequence of this construction is that correctness requires linearizability on their notion of traces and an extra condition (referred to as HB-SATISFACTION) to account for the specification order. In contrast to other work in this area, the definition of linearizability is *standard* linearizability as originally defined by Herlihy and Wing [15].

Our framework is closer to that original work on linearizability. Specifically, the allowable order of operations is included in the traces of the client program calling the specification. Therefore, unlike Dongol et al., we do not require an extra condition to handle specification order. Using our results from this paper, we would like to investigate whether we can define a notion of linearizability that can be proven sound and complete for weak memory models.

Acknowledgement This work was supported by Australian Research Council Discovery Grant DP160102457.

References

- [1] M. Abadi & L. Lamport (1991): *The existence of refinement mappings*. *Theoretical Computer Science* 82(2), pp. 253–284.
- [2] J. Alglave, L. Maranget & M. Tautschnig (2014): *Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory*. *ACM Trans. Program. Lang. Syst.* 36(2), pp. 7:1–7:74.
- [3] R.-J.R. Back (1990): *Refinement calculus, part II: Parallel and reactive programs*. In: *Stepwise Refinement of Distributed Systems Models, Formalisms, Correctness*, Springer, pp. 67–93.

- [4] R.-J.R. Back & J. von Wright (1994): *Trace refinement of action systems*. In: *CONCUR '94*, LNCS 836, Springer, pp. 367–384.
- [5] S. Burckhardt, A. Gotsman, M. Musuvathi & H. Yang (2012): *Concurrent Library Correctness on the TSO Memory Model*. In H. Seidl, editor: *ESOP 2012*, LNCS 7211, Springer, pp. 87–107.
- [6] R.J. Colvin & G. Smith (2018): *A wide-spectrum language for verification of programs on weak memory models*. In: *FM 2018*. To appear.
- [7] J. Derrick, G. Smith & B. Dongol (2014): *Verifying linearizability on TSO architectures*. In E. Albert & E. Sekerinski, editors: *iFM 2014*, LNCS 8739, Springer, pp. 341–356.
- [8] S. Doherty & J. Derrick (2016): *Linearizability and Causality*. In: *SEFM 2016*, LNCS 9763, Springer, pp. 45–60.
- [9] B. Dongol & L. Groves (2016): *Contextual Trace Refinement for Concurrent Objects: Safety and Progress*. In K. Ogata, M. Lawford & S. Liu, editors: *ICFEM 2016*, Springer, pp. 261–278.
- [10] B. Dongol, R. Jagadeesan, J. Riely & A. Armstrong (2018): *On abstraction and compositionality for weak-memory linearisability*. In I. Dillig & J. Palsberg, editors: *VMCAI'18*, LNCS 10747, Springer, pp. 183–204.
- [11] I. Filipović, P. W. O'Hearn, N. Rinetzky & H. Yang (2010): *Abstraction for concurrent objects*. *Theoretical Computer Science* 411(51-52), pp. 4379 – 4398.
- [12] S. Flur, K.E. Gray, C. Pulte, S. Sarkar, A. Sezgin, L. Maranget, W. Deacon & P. Sewell (2016): *Modelling the ARMv8 Architecture, Operationally: Concurrency and ISA*. In R. Bodik & R. Majumdar, editors: *POPL 2016*, ACM, pp. 608–621.
- [13] A. Gotsman, M. Musuvathi & H. Yang (2012): *Show No Weakness: Sequentially Consistent Specifications of TSO Libraries*. In M. Aguilera, editor: *DISC 2012*, LNCS 7611, Springer, pp. 31–45.
- [14] A. Gotsman & H. Yang (2011): *Liveness-preserving atomicity abstraction*. In: *ICALP 2011*, LNCS 6756, Springer, pp. 453–465.
- [15] M. Herlihy & J. M. Wing (1990): *Linearizability: A Correctness Condition for Concurrent Objects*. *ACM Trans. Program. Lang. Syst.* 12(3), pp. 463–492.
- [16] M. Moir & N. Shavit (2007): *Concurrent data structures*. *Handbook of Data Structures and Applications*, pp. 47:1–47:30.
- [17] S. Owens, S. Sarkar & P. Sewell (2009): *A Better x86 Memory Model: x86-TSO*. In S. Berghofer, T. Nipkow, C. Urban & M. Wenzel, editors: *TPHOLs 2009*, LNCS 5674, Springer, pp. 391–407.
- [18] P. Sewell, S. Sarkar, S. Owens, F.Z. Nardelli & M.O. Myreen (2010): *x86-TSO: a rigorous and usable programmer's model for x86 multiprocessors*. *Commun. ACM* 53(7), pp. 89–97.
- [19] G. Smith & K. Winter (2017): *Relating trace refinement and linearizability*. *Formal Aspects of Computing* 29(6), pp. 935–950.
- [20] O. Travkin, A. Mütze & H. Wehrheim (2013): *SPIN as a Linearizability Checker under Weak Memory Models*. In V. Bertacco & A. Legay, editors: *HVC2013*, LNCS 8244, Springer, pp. 311–326.