Concurrent library correctness on the TSO memory model

Sebastian Burckhardt[1], Alexey Gotsman[2],
Madanlal Musuvathi[1], and Hongseok Yang[3]

[1] Microsoft Research
[2] IMDEA Software Institute
[3] University of Oxford

Abstract. Linearizability is a commonly accepted notion of correctness for libraries of concurrent algorithms. Unfortunately, it is only appropriate for sequentially consistent memory models, while the hardware and software platforms that algorithms run on provide weaker consistency guarantees. In this paper, we present the first definition of linearizability on a weak memory model, Total Store Order (TSO), implemented by x86 processors. We establish that our definition is a correct one in the following sense: while proving a property of a client of a concurrent library, we can soundly replace the library by its abstract implementation related to the original one by our generalisation of linearizability. This allows abstracting from the details of the library implementation while reasoning about the client. We have developed a tool for systematically testing concurrent libraries against our definition and applied it to several challenging algorithms.

1 Introduction

Concurrent software developers nowadays rely heavily on libraries of concurrency patterns and high-performance concurrent data structures, such as java.util.concurrent for Java and Intel's Threading Building Blocks for C++. The algorithms implemented by these libraries are very efficient, with the downside being that they are notoriously difficult to design and implement. More surprisingly, it is often difficult to understand even what it means for them to be correct! Correctness of concurrent libraries is commonly formalised by the notion of *linearizability* [10], which fixes a certain correspondence between the library and its abstract specification, the latter usually sequential, with methods implemented atomically. Unfortunately, the classical definition of linearizability is only appropriate for sequentially consistent (SC) memory models, in which accesses to shared memory occur in a global-time linear order. At the same time, most multiprocessors (x86 [14], Power [16], ARM [1]) and programming languages (Java [11], C++ [2]) provide weaker memory models that allow more efficient implementations at the expense of exhibiting counterintuitive behaviours in some cases.

In this paper, we present the first definition of linearizability on a weak memory model, Total Store Order (TSO), implemented by x86 processors [14] (Section 4). We show that our definition is a correct one in the sense that it validates what we call the Abstraction Theorem: while proving a property of a client of a concurrent library, we can soundly replace the library by its abstract implementation related to the original one by

our generalisation of linearizability (Theorem 4, Section 5). The abstract implementation is usually simpler than the original one, with commands executing at a coarser grain of atomicity. The Abstraction Theorem thus formalises the intuitive requirement for a good definition of linearizability, which is that the library should provide an illusion of such a simpler atomic implementation. It also has a practical value as a compositional verification technique: it allows abstracting from the details of the library implementation while reasoning about its client, despite subtle interactions between the two caused by the weak memory model. As a corollary of the Abstraction Theorem, we establish that the proposed notion of linearizability is compositional (Corollary 5, Section 5).

To demonstrate that our notion of linearizability is appropriate for practical concurrent algorithms, we have developed a tool for systematically testing such algorithms against the definition and applied it to several examples (Section 6). We have also proved the linearizability of one of the algorithms formally (Theorem 3, Section 4). The algorithms considered are challenging to reason about and to specify, as they sometimes exhibit behaviours not reproducible on a sequentially consistent memory model.

The TSO memory model. The most intuitive way to explain the TSO memory model is operationally (Section 2), using an abstract multiprocessor machine in which every CPU has a *store buffer*. The buffer holds write requests that were issued by the CPU, but have not yet been *flushed* into the shared memory. A command that would like to write to a location in memory stores the corresponding write request in the store buffer of the CPU executing it, thus avoiding the need to block the CPU while the write completes.The CPU may decide to flush a store buffer entry into the main memory at any time, subject to maintaining the FIFO ordering of the buffer: the oldest write will be flushed first. A command that would like to read from a location in memory returns the value stored in the newest entry for this location in the store buffer of the CPU executing it; if such an entry does not exist, it accesses the memory directly.

The behaviour of programs running on TSO can sometimes be counterintuitive. For example, consider two memory locations `x` and `y` initially holding 0. On standard x86 processors, if two CPUs respectively write 1 to `x` and `y` and then read from `y` and `x`, as in the following program, it is possible for both to read 0 in the same execution:

$$\begin{array}{c} \texttt{x = y = 0;} \\ \texttt{x = 1; b = y;} \quad \| \quad \texttt{y = 1; a = x;} \\ \{\texttt{a} = \texttt{b} = 0\} \end{array}$$

This outcome cannot happen on a sequentially consistent machine, where both reads and writes access the memory directly. On TSO, it happens when the reads from `y` and `x` occur before the writes to them have propagated from the store buffers of the corresponding CPUs to the main memory. To exclude such behaviours, TSO processors provide special instructions, called *memory barriers*, that force the store buffer of the corresponding CPU to be flushed completely before executing the next instruction. Adding memory barriers after the writes to `x` and `y` in the above program would make it produce only SC behaviours. However, barriers incur a performance penalty.

Technical challenges. The presence of store buffers leads to subtle interactions between a library and its client that make it challenging to define linearizability. Showing

linearizability requires us to provide, for every execution of the concrete library implementation, an execution of the abstract library interacting with the client in a similar way (in a certain technical sense). Interactions between the library and the client are usually defined in terms of *histories*, which, in the classical definition, are sequences of calls to and returns from the library, along with the values passed. In the case of TSO, however, this would not describe all interactions between the two components, since one of them can exhibit a side effect on the other via a store buffer. For example, a memory barrier inside a library method will flush entries written there by client as well as library code. More subtly, write commands in a library method can insert entries into the store buffer without ensuring that they get flushed by the time the method returns. For this reason, on TSO, the method return point does not characterise the time by which the effects of these writes will be visible to the client (see the seqlock example in Section 4). To define the notion of linearizability on TSO that validates the Abstraction Theorem and is compositional, we thus need histories to describe the information relevant to the client about how the library uses store buffers. The classical notion of linearizability [10], which is not aware of store buffers, cannot specify this.

Main ideas. Our main insight lies in identifying the additional information that we need to record in histories to get a definition of linearizability on TSO validating the Abstraction Theorem. Namely, the contents of a store buffer can be viewed as a sandwich consisting of blocks of entries inserted there by an invocation of a library method or a fragment of the client code between two such invocations. We show that the behaviour of the library with regards to the store buffer that can affect the client is completely described by the moments of time at which the first and the last elements of any given library layer in the sandwich get flushed. Roughly speaking, the time when a library layer starts to get flushed defines an assumption the library makes about the client: since store buffers are FIFO, the library requires the previous client layer in the buffer to be flushed completely before this. The time by which a library layer is flushed completely represents a guarantee the library provides to the client: this action enables the next client layer to be flushed starting from this point of time.

To specify this, we enrich histories with additional actions denoting the times when a layer of entries inserted by every library method invocation starts to get flushed and is flushed completely. Linearizability then requires preserving the order between some of these actions in a history of the concrete library implementation when providing a matching history of the abstract library implementation. As we show, this is sufficient to establish the Abstraction Theorem.

The proposed definition of linearizability on TSO requires a novel way of specifying libraries. In the classical definition, the specification of a library method often consists of one atomic action. Since on TSO writes can be delayed in the store buffer, such a specification according to our notion of linearizability is often given by two atomic actions: one that atomically writes entries into the store buffer, and one that flushes them into the memory, possibly after the method returns. The resulting specification captures the effects of using the store buffer visible to the client, yet is simpler than the implementation: it ensures that all the locations written to by a library method will be written to the memory atomically, albeit at some later time. We provide examples of such specifications in Section 4 and Appendix B.

2 TSO semantics

In this section, we present the operational semantics of the TSO memory model, following [14], along with our modifications to it needed to define linearizability.

Notation. We write A^+ and A^* for the sets of all nonempty, respectively, possibly empty finite sequences of elements of a set A. We denote the empty sequence with ε and the concatenation of sequences α_1 and α_2 with $\alpha_1\alpha_2$. When we deal with sequences of sequences, for clarity we sometimes put an element of a sequence that is itself a sequence into brackets $\langle\cdot\rangle$. For example, $\alpha_1 \langle\beta\rangle \alpha_2$ denotes a sequence containing another sequence β as one of its elements. We write $g[x : y]$ for the function that has the same value as g everywhere, except for x, where it has the value y. We write $\_$ for an expression whose value is irrelevant and implicitly existentially quantified. We denote the powerset of a set X with $\mathcal{P}(X)$, and the disjoint union of sets with $\uplus$.

Programming language. We consider a machine with n CPUs, indexed by $\mathsf{CPUid} = \{1, \ldots, n\}$ and a shared memory. The machine executes programs of the following form:

$$L \ ::= \ \{m = C_m \mid m \in M\} \qquad\qquad C(L) \ ::= \ \mathsf{let}\ L\ \mathsf{in}\ C_1 \parallel \ldots \parallel C_n$$

A program consists of a declaration of a library L, implementing a set of methods $M \subseteq \mathsf{Method}$, and its client, specifying a command C_t to be run by the (hardware) thread in each CPU t. For the above program we let $\mathsf{sig}(L) = M$. To simplify presentation, we assume that the program is stored separately from the memory.

It is technically convenient for us to abstract from a particular syntax of thread and method bodies C_t and C_m and represent them using *control-flow graphs*. Namely, assume a set of primitive commands PComm (defined below). A control-flow graph (CFG) over the set PComm is a tuple $(N, T, \mathsf{start}, \mathsf{end})$, consisting of the set of program positions N, the control-flow relation $T \subseteq N \times \mathsf{PComm} \times N$, and the initial and final positions $\mathsf{start}, \mathsf{end} \in N$. The edges of the CFG are annotated with primitive commands from PComm.

We represent a program $C(L)$ by a collection of CFGs: the client command C_t for a CPU t is represented by $(N_t, T_t, \mathsf{start}_t, \mathsf{end}_t)$, and the body C_m of a method m by $(N_m, T_m, \mathsf{start}_m, \mathsf{end}_m)$. We often view this collection of CFGs for $C(L)$ as a single graph consisting of the node set $N = \biguplus_{t=1}^{n} N_t \uplus \biguplus_{m \in \mathsf{sig}(L)} N_m$ and the edge set $T = \biguplus_{t=1}^{n} T_t \uplus \biguplus_{m \in \mathsf{sig}(L)} T_m$.

Machine configurations. The set of possible configurations Config of our machine is defined in Figure 1. The special configuration $\top$ results from the machine executing an illegal instruction, such as dereferencing a non-existent memory location. An ordinary configuration $(\mathsf{pc}, \theta, b, h, K) \in \mathsf{Config}$ consists of several components. The first one $\mathsf{pc} \in \mathsf{CPUid} \to \mathsf{Pos}$ gives the current instruction pointer of every CPU. When a CPU executes client code, its instruction pointer defines the program position of the client command being executed. Otherwise, it is given by a pair whose first component is the program position of the current library command, and the second one is the client position to return to when the library method finishes executing (one return position is sufficient, since, as explained below, we disallow nested method calls).

$$\mathsf{Loc} = \mathbb{N} \qquad \mathsf{Val} = \mathbb{Z} \qquad \mathsf{Heap} = \mathsf{Loc} \rightharpoonup_{fin} \mathsf{Val}$$
$$\mathsf{Pos} = N \uplus (N \times N) \qquad \mathsf{Reg} = \{\mathtt{r}_1, \ldots, \mathtt{r}_m\} \qquad \mathsf{RegBank} = \mathsf{Reg} \to \mathsf{Val}$$
$$\mathsf{Buff} = ((\mathsf{Loc} \times \mathsf{Val})^+ \cup \{\mathsf{lock}, \mathsf{call}, \mathsf{ret}\})^*$$
$$\mathsf{Config} = \{\top\} \cup ((\mathsf{CPUid} \to \mathsf{Pos}) \times (\mathsf{CPUid} \to \mathsf{RegBank}) \times (\mathsf{CPUid} \to \mathsf{Buff}) \times \mathsf{Heap} \times \mathcal{P}(\mathsf{CPUid}))$$

Fig. 1. The set of machine configurations

Each CPU in the machine has a set of registers Reg, whose values are defined by $\theta \in \mathsf{CPUid} \to \mathsf{RegBank}$. The machine memory $h \in \mathsf{Heap}$ is represented as a finite partial function from existing memory locations to the values they store. The component $K \in \mathcal{P}(\mathsf{CPUid})$ defines the set of *active* CPUs that can currently execute a command and is used to implement atomic execution of certain commands.

The component $b \in \mathsf{CPUid} \to \mathsf{Buff}$ describes the state of all store buffers in the machine, each represented by a sequence of write requests with newest coming first. The contents of store buffers in our configurations differ from those prescribed by the TSO memory model [14] in two ways.

First, in TSO every entry in a store buffer is represented by a single location-value pair, whereas we use a sequence of those. In our semantics, all the locations in such a sequence are written to the memory atomically. This functionality is not provided by the hardware; we use it for expressing the semantics of library specifications, which might include atomic blocks performing several writes (see the seqlock example in Section 4).

Second, to formulate linearizability, we need to maintain some auxiliary information about executions, recorded by call, ret and lock entries in a store buffer. The marker lock is used to implement atomic commands performing several writes to different locations in memory. The markers call and ret get added to the buffer upon a call to or a return from the library, respectively, and thus delimit entries added by library method invocations and client code. They are used to generate additional actions in histories of interactions between the client and the library needed to define linearizability on TSO. We note that, despite store buffers in our configurations including call and ret markers, the semantics we define below corresponds to the standard TSO one, in the sense that erasing the markers from store buffers in all configurations of a given execution yields a valid execution in the standard TSO semantics.

Primitive commands. The set of primitive commands is defined as follows:

$$\mathsf{PComm} = \mathsf{Local} \uplus \mathsf{Read} \uplus \mathsf{Write} \uplus \{m \mid m \in \mathsf{Method}\} \uplus \{\mathsf{lock}, \mathsf{unlock}, \mathsf{xlock}, \mathsf{xunlock}\}.$$

Here Local, Read and Write are unspecified sets of commands such that:

- commands in Local access only CPU registers;
- commands in Read read a single location in memory and write its contents into the register $\mathtt{r}_1$;
- commands in Write write to a single location in memory.

We also have library method calls and the commands lock and unlock that lock the machine, allowing several commands to be executed atomically, and unlock it. We assume that parameters and return values of methods are passed via CPU registers. If a

client needs to preserve register values when calling a library method, it can save them in memory before the call and restore them when the method returns. The xlock and xunlock commands act as lock and unlock, except they have a built-in memory barrier, flushing the store buffer of the CPU executing the command. We call a sequence of commands bracketed by lock and unlock, or xlock and xunlock, an *atomic block*.

For every command $c \in \mathsf{Local} \uplus \mathsf{Read} \uplus \mathsf{Write}$, we assume a transformer:

- $f_c : \mathsf{RegBank} \to \mathcal{P}(\mathsf{RegBank})$ for $c \in \mathsf{Local}$ defining how the command changes the registers of the CPU executing it;
- $f_c : \mathsf{RegBank} \to \mathcal{P}(\mathsf{Loc})$ for $c \in \mathsf{Read}$ defining the location read;
- $f_c : \mathsf{RegBank} \to \mathcal{P}(\mathsf{Loc} \times \mathsf{Val})$ for $c \in \mathsf{Write}$ defining the location and the value written.

Note that we allow the execution of primitive commands to be non-deterministic. As in this paper we are dealing with low-level programs, we do not assume a built-in allocator, and thus do not consider commands for memory (de)allocation as primitive.

We place certain restrictions on CFGs over the above set PComm. Namely, we assume that on any path in a CFG, (x)lock and (x)unlock commands alternate correctly. In particular, we disallow nested (x)lock instructions. We assume that every method called in the program is defined, and we disallow nested method calls as well as method calls inside atomic blocks.

Let E, F denote expressions over the set of registers Reg, and $[\![E]\!]r$ the result of evaluating the expression E in the register bank r. Then we can define sample primitive commands

$$\mathsf{havoc} \in \mathsf{Local}, \quad \mathsf{assume}(E) \in \mathsf{Local}, \quad \mathsf{read}(E) \in \mathsf{Read}, \quad \mathsf{write}(E, F) \in \mathsf{Write}$$

with the following semantics:

$$f_{\mathsf{havoc}}(r) = \mathsf{RegBank}; \qquad f_{\mathsf{assume}(E)}(r) = \{r\}, \text{ if } [\![E]\!]r \neq 0;$$
$$f_{\mathsf{read}(E)}(r) = \{[\![E]\!]r\}; \qquad f_{\mathsf{assume}(E)}(r) = \emptyset, \text{ if } [\![E]\!]r = 0;$$
$$f_{\mathsf{write}(E,F)}(r) = \{([\![E]\!]r, [\![F]\!]r)\}.$$

The read and write commands have the expected meaning. The havoc command assigns arbitrary values to all registers. The assume(E) command acts as a filter on states, choosing only those where E evaluates to non-zero values. Using assume(E), a conditional branch on the value of E can be implemented with the CFG edges $(v, \mathsf{assume}(E), v_1)$ and $(v, \mathsf{assume}(!E), v_2)$, where $!E$ denotes the C-style negation.

Given the above commands, a memory barrier can be implemented as "xlock; xunlock". We can also implement the well-known atomic compare-and-swap (CAS) operation. A CAS takes three arguments: a memory address `addr`, an expected value `v1` and a new value `v2`. It atomically reads the memory address and updates it with the new value when the address contains the expected value; otherwise, it does nothing. In our language, we define `CAS(addr, v1, v2)` as syntactic sugar for the control-flow graph representation of:

```
xlock;
if (*addr == v1) { *addr = v2; xunlock; return 1; }
else { xunlock; return 0; }
```

Actions and traces. Transitions in our operational semantics are labelled using *actions* of the form

$$\begin{array}{rl}\varphi \in \mathsf{Act} ::= & (t, \mathsf{read}(x,u)) \mid (t, \mathsf{write}(x,u)) \mid (t, \mathsf{flush}(x,u)) \mid (t, \mathsf{flush}(\mathsf{call})) \mid \\ & (t, \mathsf{flush}(\mathsf{ret})) \mid (t, \mathsf{lock}) \mid (t, \mathsf{unlock}) \mid (t, \mathsf{xlock}) \mid (t, \mathsf{xunlock}) \mid \\ & (t, \mathsf{call}\ m(r)) \mid (t, \mathsf{ret}\ m(r))\end{array}$$

where $t \in \mathsf{CPUid}$, $x \in \mathsf{Loc}$, $u \in \mathsf{Val}$, $m \in \mathsf{Method}$ and $r \in \mathsf{RegBank}$. Here $(t, \mathsf{write}(x,u))$ corresponds to enqueuing a pending write of u to the location x into the store buffer of CPU t, $(t, \mathsf{flush}(x,u))$ to flushing a pending write of u to the location x from the store buffer of t into the shared memory, $(t, \mathsf{flush}(\mathsf{call}))$ or $(t, \mathsf{flush}(\mathsf{ret}))$ to discarding a call or ret marker from the head of a store buffer. The last two actions record moments of time when entries in a store buffer written by a given library method invocation start to get flushed and are flushed completely, which are needed in the formulation of linearizability as we explained in Section 1. The rest of the actions have the expected meaning. Since parameters and return values of library methods are passed via CPU registers, we record their values in call and return actions.

We call a (finite or infinite) sequence of actions a *trace* and adopt the standard notation: $\lambda(i)$ is the i-th action in the trace λ, $|\lambda|$ is the length of the trace λ ($|\lambda| = \omega$ if λ is infinite), and $\lambda|_t$ is the projection of λ to actions by CPU t.

Program semantics. The operational semantics of a program $C(L)$ is defined by the transition relation $\longrightarrow_{C(L)} : \mathsf{Config} \times \mathsf{Act}^* \times \mathsf{Config}$ in Figure 2. We remind the reader that T in the figure is the control-flow relation of $C(L)$. To handle transitions inside the library code, we lift it to program positions $N \uplus (N \times N)$ as follows:

$$\hat{T} = T \cup \{((v, v_0), c, (v', v_0)) \mid (v, c, v') \in T \wedge v_0 \in N\}.$$

The LOCAL rule handles the execution of commands that access registers only. These and other commands can only be executed by a CPU t if it is included into the set of active CPUs, represented by the last component of a configuration.

A write by a CPU to a location in memory does not happen immediately; instead, a pair of the location and the value to be written is added to the tail of the corresponding store buffer (WRITE). Recall that the newest entry comes first in the store buffer. When the location being written does not exist, the write command faults (WRITE-$\top$).

The READ rule uses $\mathsf{lookup}(\alpha, h, x)$ to find the value stored for the address x in the store buffer α of the CPU executing the command or the memory h:

$$\mathsf{lookup}(\alpha, h, x) = \begin{cases} u, & \text{if } \alpha = \alpha_1 \langle \beta_1 (x,u) \beta_2 \rangle \alpha_2 \text{ and} \\ & \alpha_1, \beta_1 \text{ do not contain entries for } x; \\ h(x), & \text{if } x \in \mathsf{dom}(h) \text{ and } \alpha \text{ does not contain entries for } x; \\ \top, & \text{otherwise.} \end{cases}$$

If there are entries for x in the store buffer, the read takes the value in the newest one; otherwise, it looks up the value in memory. If the location being read does not exist, lookup returns $\top$. According to READ, the value read is stored in the register $\mathtt{r}_1$.

$$\frac{t \in K \qquad (\rho, c, \rho') \in \hat{T} \qquad c \in \mathsf{Local} \qquad r' \in f_c(r)}{\mathsf{pc}[t:\rho], \theta[t:r], b, h, K \xrightarrow{\varepsilon}_{C(L)} \mathsf{pc}[t:\rho'], \theta[t:r'], b, h, K} \text{ LOCAL}$$

$$\frac{(\rho, c, \rho') \in \hat{T} \qquad c \in \mathsf{Write} \qquad (x, u) \in f_c(r) \qquad x \in \mathsf{dom}(h)}{\mathsf{pc}[t:\rho], \theta[t:r], b[t:\alpha], h, K \xrightarrow{(t,\mathsf{write}(x,u))}_{C(L)} \mathsf{pc}[t:\rho'], \theta[t:r], b[t:(x,u)\,\alpha], h, K} \text{ WRITE}$$

$$\frac{t \in K \qquad (\rho, c, \rho') \in \hat{T} \qquad c \in \mathsf{Write} \qquad (x, u) \in f_c(r) \qquad x \notin \mathsf{dom}(h)}{\mathsf{pc}[t:\rho], \theta[t:r], b, h, K \xrightarrow{\varepsilon}_{C(L)} \top} \text{ WRITE-}\top$$

$$\frac{t \in K \qquad (\rho, c, \rho') \in \hat{T} \qquad c \in \mathsf{Read} \qquad x \in f_c(r) \qquad u = \mathsf{lookup}(\alpha, h, x) \neq \top}{\mathsf{pc}[t:\rho], \theta[t:r], b[t:\alpha], h, K \xrightarrow{(t,\mathsf{read}(x,u))}_{C(L)} \mathsf{pc}[t:\rho'], \theta[t:r[\mathtt{r}_1:u]], b[t:\alpha], h, K} \text{ READ}$$

$$\frac{t \in K \qquad (\rho, c, \rho') \in \hat{T} \qquad c \in \mathsf{Read} \qquad x \in f_c(r) \qquad \mathsf{lookup}(\alpha, h, x) = \top}{\mathsf{pc}[t:\rho], \theta[t:r], b[t:\alpha], h, K \xrightarrow{\varepsilon}_{C(L)} \top} \text{ READ-}\top$$

$$\frac{(\rho, \mathsf{lock}, \rho') \in \hat{T}}{\mathsf{pc}[t:\rho], \theta, b[t:\alpha], h, \mathsf{CPUid} \xrightarrow{(t,\mathsf{lock})}_{C(L)} \mathsf{pc}[t:\rho'], \theta, b[t:\mathsf{lock}\,\alpha], h, \{t\}} \text{ LOCK}$$

$$\frac{(\rho, \mathsf{unlock}, \rho') \in \hat{T}}{\mathsf{pc}[t:\rho], \theta, b[t:(x_1,u_1)\ldots(x_l,u_l)\,\mathsf{lock}\,\alpha], h, \{t\} \xrightarrow{(t,\mathsf{unlock})}_{C(L)} \mathsf{pc}[t:\rho'], \theta, b[t:\langle(x_1,u_1)\ldots(x_l,u_l)\rangle\,\alpha], h, \mathsf{CPUid}} \text{ UNLOCK}$$

$$\frac{}{\mathsf{pc}, \theta, b[t:\alpha\,\langle(x_1,u_1)\ldots(x_l,u_l)\rangle], h, \mathsf{CPUid} \xrightarrow{(t,\mathsf{flush}(x_l,u_l))\ldots(t,\mathsf{flush}(x_1,u_1))}_{C(L)} \mathsf{pc}, \theta, b[t:\alpha], h[x_l:u_l]\ldots[x_1:u_1], \mathsf{CPUid}} \text{ FLUSH}$$

$$\frac{\beta \in \{\mathsf{call}, \mathsf{ret}\}}{\mathsf{pc}, \theta, b[t:\alpha\beta], h, \mathsf{CPUid} \xrightarrow{(t,\mathsf{flush}(\beta))}_{C(L)} \mathsf{pc}, \theta, b[t:\alpha], h, \mathsf{CPUid}} \text{ FLUSH-MARKER}$$

$$\frac{(\rho, \mathsf{xlock}, \rho') \in \hat{T}}{\mathsf{pc}[t:\rho], \theta, b[t:\varepsilon], h, \mathsf{CPUid} \xrightarrow{(t,\mathsf{xlock})}_{C(L)} \mathsf{pc}[t:\rho'], \theta, b[t:\varepsilon], h, \{t\}} \text{ XLOCK}$$

$$\frac{(\rho, \mathsf{xunlock}, \rho') \in \hat{T}}{\mathsf{pc}[t:\rho], \theta, b[t:(x_1,u_1)\ldots(x_l,u_l)], h, \{t\} \xrightarrow{(t,\mathsf{flush}(x_l,u_l))\ldots(t,\mathsf{flush}(x_1,u_1))(t,\mathsf{xunlock})}_{C(L)} \mathsf{pc}[t:\rho'], \theta, b[t:\varepsilon], h[x_l:u_l]\ldots[x_1:u_1], \mathsf{CPUid}} \text{ XUNLOCK}$$

$$\frac{(v, m, v') \in T}{\mathsf{pc}[t:v], \theta[t:r], b[t:\alpha], h, \mathsf{CPUid} \xrightarrow{(t,\mathsf{call}\ m(r))}_{C(L)} \mathsf{pc}[t:(\mathsf{start}_m, v')], \theta[t:r], b[t:\mathsf{call}\,\alpha], h, \mathsf{CPUid}} \text{ CALL}$$

$$\frac{}{\mathsf{pc}[t:(\mathsf{end}_m, v')], \theta[t:r], b[t:\alpha], h, \mathsf{CPUid} \xrightarrow{(t,\mathsf{ret}\ m(r))}_{C(L)} \mathsf{pc}[t:v'], \theta[t:r], b[t:\mathsf{ret}\,\alpha], h, \mathsf{CPUid}} \text{ RET}$$

Fig. 2. Operational TSO semantics

A CPU executing lock makes itself the only active CPU, preventing the others from executing commands[4] (LOCK). The commands executed within the corresponding atomic block, i.e., until the CPU calls unlock (UNLOCK) are thus not interleaved with commands of other CPUs. A lock command also adds a lock marker to the tail of the store buffer, thus delimiting the write requests issued within the atomic block. The corresponding unlock command then uses the lock marker to gather these write requests into a single buffer entry. Since we prohibit method calls inside atomic blocks, this entry does not contain call or ret markers.

A CPU may at any point decide to flush the entry at the head of the store buffer into memory (FLUSH). All the writes in the entry are flushed at the same time, thus ensuring that writes made in an atomic block take effect atomically. A CPU can also discard the marker at the head of the store buffer (FLUSH-MARKER). Although this does not modify the memory, we use the corresponding action, recorded in the transition relation, to formulate linearizability (Section 4). For technical reasons, it is convenient for us to prohibit flushes inside an atomic block delimited by lock and unlock. Thus, the FLUSH and FLUSH-MARKER require the set of active CPUs to be CPUid.

The xlock command (XLOCK) can only be executed when the store buffer is empty and thus forces the CPU to flush its store buffer beforehand using FLUSH and FLUSH-MARKER. For this reason, it does not need to insert a lock marker into the buffer: by the end of the atomic block the buffer will only contain writes issued inside it. The xunlock command flushes all these entries into the memory (XUNLOCK).

The rules CALL and RET handle calls to and returns from methods. Upon a method call, the return point is saved as a component in the new thread position, a call marker is added to the tail of the store buffer, and the method starts executing from the corresponding starting node of its CFG. Upon a return, the return point is read from the current program position, and a ret marker is added to the tail of the store buffer. Note that configurations in CALL and RET rules have CPUid as the set of active CPUs, since we prohibit method calls inside atomic blocks.

We note that the store buffers arising in executions of $C(L)$ as defined in Figure 2 are not arbitrary elements of Buff, but satisfy certain properties: e.g., call and ret markers in them alternate correctly, and they contain at most one lock marker. We formalise such properties in Appendix A.

Implementations of the TSO memory model usually guarantee that store buffers are fair, in the sense that, eventually, every write request in a buffer will be flushed into the memory. Our results can be extended to accommodate this constraint; however, we do not handle it in this paper so as not to obfuscate presentation.

A *computation* of $C(L)$ is a sequence of transitions using $\longrightarrow_{C(L)}$. For a computation τ, we let $\mathsf{trace}(\tau)$ be the trace obtained by concatenating all the annotations of transitions in τ. In the following, we assume that program properties of interest are linear-time properties over sets of program traces. We denote with $\xrightarrow{\lambda}{}^{*}_{C(L)}$ the reflexive and transitive closure of $\longrightarrow_{C(L)}$, where λ is obtained by concatenating the transition annotations.

[4] The semantics of TSO [14] locks only the memory bus in this case, which allows other CPUs to execute local commands affecting only their registers. For simplicity, we chose to disallow all commands.

Let $I \subseteq \mathsf{Heap}$ be the set of initial heaps that the program $C(L)$ expects to execute from. We define the set of its initial configurations as

$$\Sigma_0(I) = \{(\mathsf{pc}_0, \theta_0, b_0, h_0, \mathsf{CPUid}) \mid \forall t \in \mathsf{CPUid}.\, \mathsf{pc}_0(t) = \mathsf{start}_t \wedge b_0(t) = \varepsilon \wedge h_0 \in I\}.$$

We define the semantics $[\![C(L)]\!]I$ of $C(L)$ executing from I as the set of computations with initial configurations from $\Sigma_0(I)$. We say that the program $C(L)$ is *safe* for I, if it is not the case that $\sigma_0 \xrightarrow{\lambda}{}^*_{C(L)} \top$ for some λ and $\sigma_0 \in \Sigma_0(I)$. Informally, a program is safe when it accesses only allocated memory. Safety can be established using existing logics for reasoning about programs running on TSO [15, 19].

3 Library-local and client-local semantics

Consider a library L and a program $C(L)$ using this library:

$$L \;=\; \{m = C_m \mid m \in M\}, \qquad\qquad C(L) \;=\; \mathsf{let}\ L\ \mathsf{in}\ C_1 \parallel \ldots \parallel C_n.$$

To formulate the definition of linearizability and the Abstraction Theorem, we need to give a semantics to parts of $C(L)$: the library L considered in isolation from its client and the client C considered in isolation from the implementation of the library it uses. In this section, we specialise the semantics of programs in Section 2 to such *library-local* and *client-local* semantics describing all possible behaviours of the corresponding components.

Let us lift the operation of the disjoint union of heaps to sets of heaps pointwise:

$$\forall I_1, I_2 \subseteq \mathsf{Heap}.\, I_1 \circ I_2 = \{h_1 \uplus h_2 \mid h_1 \in I_1 \wedge h_2 \in I_2\}.$$

We assume that the set I of initial heaps of $C(L)$ satisfies $I = I_c \circ I_l$ for some $I_c, I_l \subseteq \mathsf{Heap}$ such that for any $h_c \in I_c$ and $h_l \in I_l$, $h_c \uplus h_l$ is defined. Here I_c and I_l are meant to represent parts of initial heaps used by the client C and the library L, respectively; the initial heaps of $C(L)$ are obtained as the $\circ$-combination of these.

Recall that n is the number of CPUs in our machine. To give a library-local semantics to L, we consider the program $\mathsf{MGC}(L) \;=\; \mathsf{let}\ L\ \mathsf{in}\ C_1^{\mathsf{mgc}} \parallel \ldots \parallel C_n^{\mathsf{mgc}}$, where C_t^{mgc} has the CFG

$$(\{v^t_{\mathsf{mgc}}\}, \{(v^t_{\mathsf{mgc}}, \mathsf{havoc}, v^t_{\mathsf{mgc}}), (v^t_{\mathsf{mgc}}, m, v^t_{\mathsf{mgc}}) \mid m \in \mathsf{sig}(L)\}, v^t_{\mathsf{mgc}}, v^t_{\mathsf{mgc}}).$$

The program $\mathsf{MGC}(L)$ is the *most general client* of the library L, whose hardware threads on every CPU repeatedly invoke library methods in any order and with any parameters possible. The latter are passed via registers, set arbitrarily by the havoc command. The set of computations $[\![\mathsf{MGC}(L)]\!]I_l$ thus includes all library behaviours under any possible client (this fact is formalised in Lemma 6, Section 5).

In practice, a library often tolerates only calls from clients adhering to a certain policy. For example, a spinlock implementation might expect client calls to `acquire` and `release` methods to alternate. We can take this into account by restricting the most general client appropriately. While libraries in our examples do rely on the client satisfying such constraints, to simplify presentation we do not formalise them here.

To define the client-local semantics of the client C, we consider the program

$$C_M(\cdot) \;=\; \mathsf{let}\ \{m = C_m^{\mathsf{stub}} \mid m \in M\}\ \mathsf{in}\ C_1 \parallel \ldots \parallel C_n$$

where the body C_m^{stub} of every method m has the CFG $(\{v_{\mathsf{start}}^m\}, \{(v_{\mathsf{start}}^m, \mathsf{havoc}, v_{\mathsf{end}}^m)\}, v_{\mathsf{start}}^m, v_{\mathsf{end}}^m)$. That is, every method in $C_M(\cdot)$ is implemented by a stub that returns immediately after having been called, scrambling all the registers. Since return values of library methods are stored in registers, the set of computations $[\![C_M(\cdot)]\!]I_c$ generates all executions of the client assuming any behaviour of the library it uses.

Note that both library-local and client-local semantics allow store buffer entries of the corresponding component to be flushed non-deterministically while the other component is running, since this is possible in the semantics of the whole program. Similarly, we add call and ret markers to the store buffer when calling a method stub in the client-local and library-local semantics.

We say that a client C, respectively, a library L is safe for I_c, respectively, I_l, if so is $C_M(\cdot)$, respectively, $\mathsf{MGC}(L)$ (see Section 2). As we have noted before, the safety of a library or a client can be established using logics for TSO [15, 19]. Note that in the client-local or the library-local semantics, the program runs on the state owned by the corresponding component and faults when accessing memory locations not belonging to it. Thus, the safety of the client and the library ensures that they cannot corrupt each other's state. We rely crucially on this in establishing the Abstraction Theorem for the notion of linearizability we propose. It can also be shown that, when the client C and the library L are safe, so is the complete program $C(L)$ (Lemma 6, Section 5).

4 Linearizability on TSO

When defining linearizability, we are not interested in internal steps recorded in library computations, but only in the interactions of the library with its client. We record such interactions using *histories*, which are traces including only actions from the following subset of Act:

$$\mathsf{HAct} \;::=\; (t, \mathsf{call}\ m(r)) \mid (t, \mathsf{ret}\ m(r)) \mid (t, \mathsf{flush}(\mathsf{call})) \mid (t, \mathsf{flush}(\mathsf{ret}))$$

where $t \in \mathsf{CPUid}$, $m \in \mathsf{Method}$, $r \in \mathsf{RegBank}$. Recall that here r records the values of registers of the CPU that calls a library method or returns from it, which serve as parameters or return values. We define the history $\mathsf{history}(\tau)$ corresponding to a computation τ of the program $C(L)$ by projecting $\mathsf{trace}(\tau)$ to actions from HAct.

In contrast to histories used in the classical definition of linearizability [10], ours include two new types of actions needed for defining linearizability on TSO: $(t, \mathsf{flush}(\mathsf{call}))$ and $(t, \mathsf{flush}(\mathsf{ret}))$, denoting times when the CPU t flushes a call or a ret marker from its store buffer. We first formulate our definition, and then explain the motivation behind it.

Definition 1. *The **linearizability relation** is a binary relation $\sqsubseteq$ on histories defined as follows: $H \sqsubseteq H'$ if $\forall t \in \mathsf{CPUid}.\, H|_t = H'|_t$ and there is a bijection $\pi\colon \{1, \ldots, |H|\} \to \{1, \ldots, |H'|\}$ such that $\forall i.\, H(i) = H'(\pi(i))$ and*

$$\begin{aligned}(i < j \wedge (H(i) = (\_, \mathsf{ret}\ \_) \vee H(i) = (\_, \mathsf{flush}(\mathsf{ret})))\\ \wedge (H'(j) = (\_, \mathsf{call}\ \_) \vee H'(j) = (\_, \mathsf{flush}(\mathsf{call})))) \Rightarrow \pi(i) < \pi(j).\end{aligned}$$

That is, a history H' linearizes a history H when it is a permutation of the latter preserving the order of certain types of actions. We lift the notion of linearizability to libraries using the library-local semantics of Section 3.

Definition 2. *For libraries L_1 and L_2 safe for I_l and such that $\mathsf{sig}(L_1) = \mathsf{sig}(L_2)$, we say that L_2 **linearizes** L_1, written $L_1 \sqsubseteq L_2$, if*

$$\forall H_1 \in \mathsf{history}([\![\mathsf{MGC}(L_1)]\!]I_l).\, \exists H_2 \in \mathsf{history}([\![\mathsf{MGC}(L_2)]\!]I_l).\, H_1 \sqsubseteq H_2.$$

Thus, L_2 linearizes L_1 if every behaviour of the latter under the most general client may be reproduced in a linearized form by the former.

Discussion. A good definition of linearizability has to allow replacing a library implementation with its specification while keeping client behaviours reproducible (as formalised by the Abstraction Theorem in Section 5). However, linearizability itself is defined between libraries considered in isolation from their clients. In Definition 2, this is achieved by considering executions of libraries under their most general clients (Section 3), which can only refer to store buffer entries inserted by write commands in library code. When a library is used by a client, the store buffer mixes entries inserted by the two components. As we noted in Section 1, in this case the library can affect the client via the store buffer, e.g., by executing a memory barrier or leaving an unflushed entry blocking newer client entries from being flushed. The $(\_, \mathsf{flush}(\mathsf{call}))$ and $(\_, \mathsf{flush}(\mathsf{ret}))$ actions in histories record the necessary information about library behaviour of this kind, as we now explain.

Recall the analogy from Section 1, where we viewed the contents of a store buffer as a sandwich consisting of blocks of entries inserted there by an invocation of a library method or a fragment of client computation between two such invocations. The call and ret markers delimit the layers in this sandwich. For example, at some point in an execution of $C(L)$, the store buffer of some CPU might have the following contents:

$$\mathsf{ret}\ (x_5, u_5)\ \mathsf{call}\ (x_4, u_4)\ \mathsf{ret}\ (x_3, u_3)\ (x_2, u_2)\ \mathsf{call}\ (x_1, u_1), \tag{1}$$

where the leftmost end contains the newest entry. From the call and ret markers, we can immediately conclude that the write to x_1 was inserted by the client before calling a library method, the writes to x_2 and x_3 were by the library method invocation, the write to x_4 was again by the client, and the write to x_5 was by the next method invocation on this CPU.

The most general client exercises the library methods under all possible input parameters, but does not perform writes by itself. For this reason, a store buffer in the most general client of a library never has entries between a call marker and an older ret marker (we formalise this in Appendix A). For example, a computation of the most general client of the library with the same library method invocations as in the one producing (1) might have the store buffer

$$\mathsf{ret}\ (x_5, u_5)\ \mathsf{call}\ \mathsf{ret}\ (x_3, u_3)\ (x_2, u_2)\ \mathsf{call}, \tag{2}$$

which contains only library entries from (1). Thus, when considering a library in isolation from its client in defining linearizability, the call and ret markers let us determine

the places in the store buffer where client entries might be located in a corresponding execution of a complete program.

Consider an execution of the most general client of a library in which the CPU flushes a library entry (e.g., (x_3, u_3) in (2)). Since store buffers are FIFO, in the corresponding execution of a particular client with the same library behaviour, this will *assume* that the client entries in the store buffer older than it have been flushed (e.g., (x_1, u_1) in (1)). Conversely, flushing a library entry (e.g., (x_3, u_3) in (1)) preceding a client one (e.g., (x_4, u_4) in (1)) will *guarantee* that the client entry can now be flushed. For the Abstraction Theorem to hold, in Definition 2 we need to make sure that the executions of the most general clients producing histories H_1 and H_2 make the same assumptions and give the same guarantees concerning times when client entries are flushed. This is the reason for including flushes of call and ret markers into histories. The position of a $(t, \mathsf{flush}(\mathsf{call}))$ action in a history produced by the most general client defines a moment of time by which, in a complete program, all older client writes in the store buffer of t *must* be flushed for the library to be able to flush the entries from the layer following the call marker. The position of a $(t, \mathsf{flush}(\mathsf{ret}))$ action defines a moment starting from which the client entries from the layer following the ret marker *may* be flushed. In our definition of linearizability, we require that the two histories considered have the same history actions describing how store buffers are modified during the execution. Hence, in two executions corresponding to the histories, libraries make the same assumptions and give the same guarantees concerning the use of store buffers.

Like the classical definition of linearizability, ours requires preserving the order between non-overlapping library method invocations; two invocations do not overlap in a history if the return of one precedes the call of the other. This is needed for the Abstraction Theorem to hold, since the client code executed in between two non-overlapping method invocations can notice their order. To handle TSO correctly, our definition also takes into account intervals during which all the writes of a library method invocation were being flushed: it requires preserving the order between two such non-overlapping intervals or non-overlapping interval of this kind and a library method invocation. This is expressed by preserving the order of $(\_, \mathsf{flush}(\mathsf{ret}))$ preceding $(\_, \mathsf{flush}(\mathsf{call}))$, $(\_, \mathsf{flush}(\mathsf{ret}))$ preceding $(\_, \mathsf{call}\_)$, and $(\_, \mathsf{ret}\_)$ preceding $(\_, \mathsf{flush}(\mathsf{call}))$. The requirement is again needed to validate the Abstraction Theorem.

We note that our definition of linearizability is flexible in the following sense: it puts restrictions on times when call and ret markers are flushed, but not on how many ordinary entries a given method invocation inserts into the store buffer. For example, this allows us to relate a library implementation writing to some part of the memory accessed only by a given CPU to its specification that does not write to any local state.

Example. Even though we formalise our results for programs represented by their CFGs, for readability in our examples we use a C-like language. Its programs can be translated to CFGs in the standard way. We assume that global variables are allocated at fixed addresses in memory, and local variables are stored in CPU registers.

Figure 3 presents a simplified version of a seqlock [3]—an efficient implementation of a readers-writer protocol based on version counters used in the Linux kernel. Two memory addresses `x1` and `x2` make up a conceptual register that a single hardware thread can write to, and any number of other threads can attempt to read from. A version

```
word x1 = 0, x2 = 0;
word c = 0;

write(in word d1, in word d2) {
  c++;
  x1 = d1; x2 = d2;
  c++;
}
```

```
read(out word d1, out word d2) {
  word c0;
  do {
    do { c0 = c; } while (c0 % 2);
    d1 = x1; d2 = x2;
  } while (c != c0);
}
```

Fig. 3. Seqlock implementation L_{seqlock}

```
word x1 = 0, x2 = 0;

write(in word d1, in word d2) { lock; x1 = d1; x2 = d2; unlock; }

read(out word d1, out word d2) { lock; d1 = x1; d2 = x2; unlock; }
```

Fig. 4. Seqlock specification $L^{\sharp}_{\mathsf{seqlock}}$. Here `nondet()` represents a non-deterministic choice.

number is stored at `c`. The writing thread maintains the invariant that the version number is odd during writing by incrementing it before the start of and after the finish of writing. A reader checks that the version number is even before attempting to read (otherwise it could see an inconsistent result by reading while `x1` and `x2` are being written). After reading, the reader checks that the version has not changed, thereby ensuring that no write has overlapped the read. Note that neither the `write` nor the `read` operation includes a memory barrier, which means that writes to `x1`, `x2` and `c` may not be visible to readers immediately.

We give a specification to seqlock using the abstract implementation in Figure 4. Instead of using a version counter, this implementation just locks the machine while reading from or writing to `x1` and `x2`. According to the semantics of Section 2, the writes to `x1` and `x2` performed by `write` are stored in a single entry of the corresponding store buffer and are written to the shared memory atomically. This specifies that the implementation of a seqlock indeed ensures the illusion of atomicity. However, we also need our specification to capture the effect of the library executing on a weak memory model—the fact that the writes to `x1` and `x2`, although executed atomically, may still be delayed due to the presence of store buffers. This is because the delay can be noticed by certain clients and can result in a non-SC behaviour. For example, using a seqlock, we can reproduce the example from Section 1 yielding non-SC behaviour as shown in Figure 4. To capture this, the specification of `write` ensures atomicity by a pair of lock and unlock commands, which do not flush the writes to the memory immediately.

Thus, we have two atomic actions associated with the abstract `write` method: one that writes to the store buffer and the other that flushes the writes to the memory, possibly after the method returns. This is different from the classical definition of linearizability on a sequentially consistent memory model [10], which requires methods in the specification to be implemented by one atomic action.

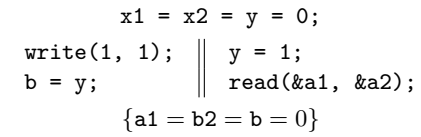

Fig. 5. A client of L_{seqlock} producing a non-SC behaviour

As the following theorem shows, the abstract implementation $L^{\sharp}_{\mathsf{seqlock}}$ in Figure 4 indeed linearizes the concrete one L_{seqlock} in Figure 3.

Theorem 3. $L_{\mathsf{seqlock}} \sqsubseteq L^{\sharp}_{\mathsf{seqlock}}$.

The proof is given in Appendix A; here we discuss it informally. The proof is similar to proofs of classical linearizability using linearization points [10], although here methods of the abstract implementation contain more than one atomic action. We consider the most general clients of the concrete and the abstract implementations of the library running alongside each other. For every execution of the client of the concrete library, we construct the corresponding execution of the client of the abstract one by firing transitions of the latter at certain times during the execution of the former.

For example, the abstract `read` method is executed when the corresponding concrete one reads `x2` for the last time. The code of the abstract `write` method is executed when the concrete one writes to `x2`. Finally, a store buffer entry containing writes to `x1` and `x2` by the abstract `write` method is flushed together with the second write to `c` by the corresponding concrete method invocation. To prove that this flush in the abstract implementation does not contradict the FIFO ordering of store buffers, we maintain an invariant relating the contents of the store buffers in the concrete and the abstract seqlock implementations.

Programs producing only SC behaviours. By this time, the reader may wonder whether it is always necessary to expose the behaviour of a library with respect to store buffers in its specification. After all, many programs running on TSO only produce SC behaviours, and there are ways of effectively checking this [13, 5, 6]. Therefore, a valid question is whether we can use the usual definition of linearizability for libraries producing only SC behaviours when they are used by clients also behaving SC. Unfortunately, in general the answer is no. This is because, even if the most general client of a library $\mathsf{MGC}(L)$ and its client $C_{\mathsf{sig}(L)}(\cdot)$ only produce SC behaviours when considered in isolation, this may not be the case for the complete program $C(L)$ due to interactions of the two components via the store buffer. For example, the most general client of a single seqlock produces only SC behaviours, as it satisfies the triangular race freedom criterion of [13]. However, Figure 4 shows that if we use a seqlock *together with* a client that also happens to be SC by itself, we can get non-SC behaviours. This is not surprising: a seqlock is meant to ensure the atomicity of writes to and reads from a pair of locations, but it is not meant to make these reads and writes strongly consistent. Thus, the classical definition of linearizability is not sufficient to specify libraries even when constraining separate components of a program to behave SC.

5 Abstraction Theorem

We now justify that the notion of linearizability proposed in Section 4 is a correct one by establishing the Abstraction Theorem that allows abstracting an implementation of a library with its specification while reasoning about its client.

For a computation τ of $C(L)$ obtained from the semantics of Section 2, we denote with $\mathsf{client}(\tau)$ the projection of its trace $\lambda = \mathsf{trace}(\tau)$ to actions relevant to the client, i.e., executed by the client code or corresponding to flushes of client entries in store buffers. Formally, we include an action φ such that $\lambda = \lambda'\varphi\lambda''$ into the projection if:

- φ is included into $\mathsf{history}(\tau)$; or
- φ is not a flush action and is outside an invocation of a library method, i.e., it is not the case that $\lambda|_t = \lambda_1\,(t, \mathsf{call}\ \_)\,\lambda_2\varphi\lambda_3$, where λ_2 does not contain a $(t, \mathsf{ret}\ \_)$ action; or
- φ corresponds to a flush of a client entry in a store buffer, i.e., it is not the case that $\lambda|_t = \lambda_1\,(t, \mathsf{flush}(\mathsf{call}))\,\lambda_2\varphi\lambda_3$, where λ_2 does not contain a $(t, \mathsf{flush}(\mathsf{ret}))$ action.

We lift client to sets of computations pointwise.

The Abstraction Theorem states that the behaviour of a client of a concurrent library will stay reproducible on TSO if we replace the library by its abstract implementation related to the original one by our definition of linearizability.

Theorem 4 (Abstraction). *Consider $C(L_1)$ and $C(L_2)$ such that C is safe for I_c, L_1 and L_2 are safe for I_l and $L_1 \sqsubseteq L_2$. Then $C(L_1)$ and $C(L_2)$ are safe for $I = I_c \circ I_l$ and $\mathsf{client}(\llbracket C(L_1) \rrbracket I) \subseteq \mathsf{client}(\llbracket C(L_2) \rrbracket I)$.*

We provide a proof outline below and give the complete proof in Appendix A. The requirement that the client C be safe in the theorem is required to replace one library implementation with another: it ensures that C cannot access the internals of the library implementation.

From Theorem 4 it follows that, while reasoning about a client $C(L_1)$ of a library L_1, we can soundly replace L_1 with a simpler library L_2 linearizing L_1: if a linear-time property over client actions holds over $C(L_2)$, it will also hold over $C(L_1)$. Note that the abstract implementation is usually simpler than the original one (in most cases implemented using atomic blocks, like the one in Figure 4), which eases the proof of the resulting program. Thus, the proposed notion of linearizability and Theorem 4 enable compositional reasoning about programs running on TSO: they allow decomposing the verification of a whole program into the verification of its constituent components. We give an example of using this technique in Section 6.

The following corollary of Theorem 4, proved in Appendix A, states that, like the classical notion of linearizability [10], ours is compositional: if several non-interacting libraries are linearizable, then so is their composition. Formally, consider libraries $L_1, \ldots, L_k$ with disjoint sets of declared methods and sets of initial heaps $I_1, \ldots, I_k$ such that

$$\forall \{i_1, \ldots, i_l\} \subseteq \{1, \ldots, k\}.\, \forall h_1 \in I_{i_1}, \ldots, h_l \in I_{i_l}.\, h_1 \uplus \ldots \uplus h_l \text{ is defined.}$$

We let the *composition* L of $L_1, \ldots, L_k$ be the library implementing all of their methods and having the set of initial heaps $I_1 \circ \ldots \circ I_k$.

Corollary 5 (Compositionality). *Consider libraries $L_1, \ldots, L_k$ and $L_1^\sharp, \ldots, L_k^\sharp$ such that L_j and $L_j^\sharp$ are safe for I_j, $j = 1..k$. Let L and $L^\sharp$ be the compositions of the respective sets of libraries. If $L_j \sqsubseteq L_j^\sharp$ for $j = 1..k$, then $L \sqsubseteq L^\sharp$.*

Proof outline for Theorem 4. The proof of Theorem 4 relies on the following lemmas, proved in Appendix A. The first lemma shows that a computation of $C(L)$ generates two computations in the client-local and library-local semantics with the same history.

Lemma 6 (Decomposition). *If $C_{\mathsf{sig}(L)}(\cdot)$ and $\mathsf{MGC}(L)$ are safe for I_c and I_l, respectively, then $C(L)$ is safe for $I_c \circ I_l$ and*

$$\forall \tau \in [\![C(L)]\!](I_c \circ I_l).\, \exists \eta \in [\![C_{\mathsf{sig}(L)}(\cdot)]\!] I_c.\, \exists \xi \in [\![\mathsf{MGC}(L)]\!] I_l. \\ \mathsf{history}(\eta) = \mathsf{history}(\xi) \wedge \mathsf{client}(\tau) = \mathsf{client}(\eta).$$

The following lemma presents the core of the transformation used to convert a computation of $C(L_1)$ into one of $C(L_2)$ in Theorem 4: it shows that a computation of a most general client can be transformed into another of its computations with a given history linearized by the history of the original one.

Lemma 7 (Rearrangement). *Consider a library L safe for I_l and histories H, H' such that $H \sqsubseteq H'$. Then*

$$\forall \tau' \in [\![\mathsf{MGC}(L)]\!] I_l.\, \mathsf{history}(\tau') = H' \Rightarrow \exists \tau \in [\![\mathsf{MGC}(L)]\!] I_l.\, \mathsf{history}(\tau) = H.$$

Finally, the following lemma states that any pair of client-local and library-local computations agreeing on the history can be combined into a valid computation of $C(L)$.

Lemma 8 (Composition). *If $C_{\mathsf{sig}(L)}(\cdot)$ and $\mathsf{MGC}(L)$ are safe for I_c and I_l, respectively, then*

$$\forall \eta \in [\![C_{\mathsf{sig}(L)}(\cdot)]\!] I_c.\, \forall \xi \in [\![\mathsf{MGC}(L)]\!] I_l.\, \mathsf{history}(\eta) = \mathsf{history}(\xi) \Rightarrow \\ \exists \tau \in [\![C(L)]\!](I_c \circ I_l).\, \mathsf{client}(\tau) = \mathsf{client}(\eta).$$

Most of the proof of the Decomposition Lemma (Lemma 6) deals with maintaining a splitting of the state of $C(L)$ into the parts owned by the client and the library, including store buffer entries. The resulting partial states then define the computations of $C_{\mathsf{sig}(L)}(\cdot)$ and $\mathsf{MGC}(L)$. Conversely, the Composition Lemma (Lemma 8) composes the states of $C_{\mathsf{sig}(L)}(\cdot)$ and $\mathsf{MGC}(L)$ into states of $C(L)$ to construct an execution of the latter. The proof of the Rearrangement Lemma (Lemma 7) transforms τ' into τ by repeatedly permuting transitions in the computation according to a certain strategy to make its history equal to H.

Proof of Theorem 4. Lemma 6 implies that $C(L)$ is safe. We now need to transform a computation $\tau_1 \in [\![C(L_1)]\!] I$ of $C(L_1)$ into a computation $\tau_2 \in [\![C(L_2)]\!] I$ with the same client trace projection: $\mathsf{client}(\tau_1) = \mathsf{client}(\tau_2)$. To this end, we use the semantics of Section 3, which defines the interpretation of L_1, L_2, $C_{\mathsf{sig}(L_1)}(\cdot)$ and their compositions. Namely, to transform τ_1 into τ_2, we first apply Lemma 6 to generate two computations from τ_1—a library-local computation $\xi_1 \in [\![\mathsf{MGC}(L_1)]\!] I_l$ and a client-local one

$\eta \in [\![C_{\mathsf{sig}(L_1)}(\cdot)]\!] I_c$—such that $\mathsf{client}(\tau_1) = \mathsf{client}(\eta)$ and $\mathsf{history}(\tau_1) = \mathsf{history}(\eta) = \mathsf{history}(\xi_1)$. Note that the computation η of C thus constructed excludes the internal library actions. Since $L_1 \sqsubseteq L_2$, for some computation $\xi_2 \in [\![\mathsf{MGC}(L_2)]\!] I_l$, we have $\mathsf{history}(\xi_1) \sqsubseteq \mathsf{history}(\xi_2)$. By Lemma 7, ξ_2 can be transformed into a computation $\xi_2' \in [\![\mathsf{MGC}(L_2)]\!] I_l$ such that $\mathsf{history}(\xi_2') = \mathsf{history}(\xi_1) = \mathsf{history}(\eta)$. We then use Lemma 8 to compose the library-local computation ξ_2' with the client-local one η into a computation $\tau_2 \in [\![C(L_2)]\!] I$ such that $\mathsf{client}(\tau_2) = \mathsf{client}(\eta) = \mathsf{client}(\tau_1)$. □

6 Checking linearizability on TSO

We have implemented a tool called LINTSO for systematically testing concurrent libraries for our notion of linearizability. Our intention in implementing the tool is twofold. First, the tool allows developers of concurrent libraries to find violations of linearizability quickly. The second (and more important) goal is to use the tool to perform a sanity check of our definition of linearizability by making sure that real-world algorithms that are commonly accepted as correct are linearizable with respect to it.

LINTSO is similar in spirit to the LINE-UP tool for checking linearizability on a sequentially consistent memory model [4]. It takes as input a concrete and an abstract implementation of a library (such as the ones in Figures 3 and 4) along with a (bounded) test harness that calls into the library. LINTSO then composes the input with an operational model of TSO such that sequentially consistent behaviors of the resulting program emulate TSO behaviors of the input. This allows LINTSO to use existing model checkers, such as CHESS [12], to systematically enumerate the behaviors of the harness and the library on TSO.

In a first phase, LINTSO exhaustively generates all histories of the input harness calling into the abstract version of the library. In a subsequent phase, LINTSO systematically enumerates the TSO behaviors of the harness and the concrete version of the library. For every such behavior, LINTSO uses the linearizability condition to check if the behavior is consistent with respect to some history observed in the first phase. Any violation is reported as an error.

If the enumeration in the second phase completes, then LINTSO guarantees that the abstract implementation linearizes the concrete one for the given harness. If the number of possible computations in this phase is too large, a subset of them can be considered by bounding the number of context switches [12]. Obviously, this does not provide a complete guarantee of linearizability, as only (possibly a subset of computations of) one of the infinitely many harnesses is considered.

In our experiments we considered the following concurrent algorithms that were identified as challenges in [13]:

- seqlock, the readers-writer lock we discussed in Section 4;
- simple spinlock, which does not provide fairness guarantees;
- ticketed spinlock, ensuring fairness using a variant of the Bakery algorithm;
- initialisation using double-checked locking.

We provide their code and specifications in Appendix B. The seqlock and the spinlock implementations are used in various versions of the Linux kernel [3]. The above algorithms are optimised for the TSO memory model and, when used in certain ways,

can exhibit behaviours that cannot be reproduced on a sequentially consistent memory model. In fact, the correctness of the spinlock implementations was a subject of debate among Linux developers [13].

In more detail, the simple and ticketed spinlocks do not execute a memory barrier after writing a value into the lock data structure saying that the lock is free. According to the semantics of TSO, this does not violate mutual exclusion: delaying the write in the store buffer can only lead to CPUs that want to acquire the lock waiting longer. As in the case of a seqlock (Figure 4), the specification of a spinlock captures the fact that the lock release can be delayed.

The initialisation using double-checked locking first checks if an object is initialised by reading a corresponding flag without acquiring the lock for the object. Since the read is not preceded by a memory barrier, on TSO this can cause it to return 'uninitialised' even after the object has been in fact initialised. This does not violate the correctness of the algorithm, since the flag is then re-checked with the lock held.

For simple harnesses of the above examples, consisting of up to 3 threads, each performing up to 3 operations, LINTSO performed the check in a matter of minutes. The specification histories were generated exhaustively, and the implementation histories for computations up to a maximum of two preemptions (the CHESS default). In all cases, the tool did not detect any errors. As a further sanity check, we introduced simple errors in the examples, e.g., by replacing xunlock with an unlock in the concrete version. LINTSO was able to find all of them.

We used Theorem 4 to modularise checking the linearizability of the intialisation using double-checked locking. Namely, Theorem 4 allowed us to consider the specification of the spinlock used in this example, instead of a particular implementation. This cut down the number of interleavings to be analysed and made the analysis more efficient. Additionally, it allowed us to prove the linearizability of the algorithm regardless of the particular spinlock implementation used (e.g., the simple or ticketed spinlock). This is just one example of using the Abstraction Theorem to verify concurrent programs compositionally.

7 Related work and conclusion

All the definitions of linearizability proposed for various settings so far [10, 7, 9, 8] have assumed a sequentially consistent memory model. This paper is the first to define a notion of linearizability on a weak memory model and show that it validates the Abstraction Theorem (Theorem 4). Our result is based on a novel insight about what information should be kept in histories to specify interactions between the library and the client due to the weak memory model. Even though in this paper we considered only one weak memory model—TSO, implemented by x86 processors [14]—our insights form a starting point for investigating weaker memory models, such as those of Power [16] and ARM [1] processors, and the C++ language [2].

Our work lays the foundation for future correctness proofs for implementations of concurrent algorithms in operating system kernels [3] and language run-times [2]. In particular, we hope that it should be possible to develop a logic for establishing the proposed notion of linearizability formally, based on existing logics for proving safety

properties on TSO [19, 15] and linearizability on sequentially consistent memory models [17, 18]. This should make proofs such as that of Theorem 3 easier to carry out.

We also intend to investigate definitions of linearizability on weak memory models in cases when the library and the client interact in more complicated ways. For example, in this paper we did not consider the transfer of data structure ownership between the library and the client, assuming that they communicate only by passing values of a primitive type. We believe that our approach to handling weak memory can be married with a previous generalisation of linearizability for ownership transfer on a sequentially consistent memory model [8].

Acknowledgements. We thank Scott Owens, Ian Wehrman and the anonymous reviewers for comments that helped to improve the paper. Yang was supported by EPSRC.

References

1. J. Alglave, A. Fox, S. Ishtiaq, M. O. Myreen, S. Sarkar, P. Sewell, and F. Zappa Nardelli. The semantics of Power and ARM multiprocessor machine code. In *DAMP*, 2009.
2. M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *POPL*, 2011.
3. D. Bovet and M. Cesati. *Understanding the Linux Kernel, 3rd ed.* O'Reilly, 2005.
4. S. Burckhardt, C. Dern, M. Musuvathi, and R. Tan. Line-up: A complete and automatic linearizability checker. In *PLDI*, 2010.
5. S. Burckhardt and M. Musuvathi. Effective program verification for relaxed memory models. In *CAV*, 2008.
6. E. Cohen and B. Schirmer. From total store order to sequential consistency: A practical reduction theorem. In *ITP*, 2010.
7. I. Filipović, P. O'Hearn, N. Rinetzky, and H. Yang. Abstraction for concurrent objects. In *ESOP*, 2009.
8. A. Gotsman and H. Yang. Linearizability with ownership transfer. Draft. Available from `www.software.imdea.org/~gotsman`, 2011.
9. A. Gotsman and H. Yang. Liveness-preserving atomicity abstraction. In *ICALP*, 2011.
10. M. P. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *TOPLAS*, 12, 1990.
11. J. Manson, W. Pugh, and S. V. Adve. The Java memory model. In *POPL*, 2005.
12. M. Musuvathi, S. Qadeer, T. Ball, G. Basler, P. A. Nainar, and I. Neamtiu. Finding and reproducing heisenbugs in concurrent programs. In *OSDI*, 2008.
13. S. Owens. Reasoning about the implementation of concurrency abstractions on x86-TSO. In *ECOOP*, 2010.
14. S. Owens, S. Sarkar, and P. Sewell. A better x86 memory model: x86-TSO. In *TPHOLs*, 2009.
15. T. Ridge. A rely-guarantee proof system for x86-TSO. In *VSTTE*, 2010.
16. S. Sarkar, P. Sewell, J. Alglave, L. Maranget, and D. Williams. Understanding POWER multiprocessors. In *PLDI*, 2011.
17. V. Vafeiadis. Modular fine-grained concurrency verification. PhD Thesis. Technical Report UCAM-CL-TR-726, University of Cambridge, 2008.
18. V. Vafeiadis. Automatically proving linearizability. In *CAV*, 2010.
19. I. Wehrman and J. Berdine. A proposal for weak-memory local reasoning. In *LOLA*, 2011.

A Proofs

A.1 Auxiliary definitions

Before proving the lemmas, we introduce some auxiliary notation.

We first define an operation analogous to client, but selecting actions relevant to the library. For a computation τ of $C(L)$ generated from the semantics of Section 2, we denote with $\mathsf{lib}(\tau)$ the projection of its trace $\lambda = \mathsf{trace}(\tau)$ to actions executed by the library code or corresponding to flushes of library entries in store buffers. Formally, we include an action φ such that $\lambda = \lambda'\varphi\lambda''$ into the projection if:

- φ is included into $\mathsf{history}(\tau)$; or
- φ not a flush action and is inside an invocation of a library method, i.e., $\lambda|_t = \lambda_1\,(t, \mathsf{call}\ \_)\,\lambda_2\varphi\lambda_3$, where λ_2 does not contain a $(t, \mathsf{ret}\ \_)$ action; or
- φ corresponds to a flush of a library entry in a store buffer, i.e., $\lambda|_t = \lambda_1\,(t, \mathsf{flush}(\mathsf{call}))\,\lambda_2\varphi\lambda_3$, where λ_2 does not contain a $(t, \mathsf{flush}(\mathsf{ret}))$ action.

For a finite computation τ_1 and a computation τ_2, their concatenation $\tau_1\tau_2$ is defined when the final configuration in τ_1 is the same as the initial one in τ_2. The concatenation glues the computations at this configuration.

We sometimes apply history, lib and client directly to traces, instead of computations.

In the following, we classify every transition in a computation of $C(L)$ distinct from FLUSH or FLUSH-MARKER as either a client or a library transition, depending on which part of the code of $C(L)$ executes the command.

We now define a partial operation $\circ : \mathsf{Config} \times \mathsf{Config} \rightharpoonup \mathsf{Config}$ that combines configurations in the local semantics of $C_{\mathsf{sig}(L)}(\cdot)$ and $\mathsf{MGC}(L)$ to yield a configuration of $C(L)$. We first define $\circ$ on the components of a configuration.

Let $\circ : \mathsf{Pos} \times \mathsf{Pos} \rightharpoonup \mathsf{Pos}$ combine CPU positions in the client-local and library-local semantics to obtain a position in the complete program: $v \circ v^t_{\mathsf{mgc}} = v$ for $v \in \biguplus_{t=1}^{n} N_t$; $(v^m_k, v) \circ (v', v^t_{\mathsf{mgc}}) = (v', v)$ for $k \in \{\mathsf{start}, \mathsf{end}\}$, $v \in \biguplus_{t=1}^{n} N_t$ and $v' \in N_m$; all other combinations are undefined. We lift $\circ$ to program counters pointwise.

For $r_1, r_2 \in \mathsf{RegBank}$ and $\rho \in \mathsf{Pos}$ we let

$$r_1 \circ_\rho r_2 = \begin{cases} r_1, & \text{if } \rho \in N; \\ r_2, & \text{if } \rho \in N \times N. \end{cases}$$

For $\theta_1, \theta_2 \in \mathsf{CPUid} \to \mathsf{RegBank}$ and $\mathsf{pc} \in \mathsf{CPUid} \to \mathsf{Pos}$ we then let

$$\forall t \in \mathsf{CPUid}.\ (\theta_1 \circ_{\mathsf{pc}} \theta_2)(t) = \theta_1(t) \circ_{\mathsf{pc}(t)} \theta_2(t).$$

The following simple proposition describes the structure of store buffers obtained in the operational semantics of $C(L)$, $C_{\mathsf{sig}(L)}$ and $\mathsf{MGC}(L)$. In the following, we often use it implicitly when we perform operations on store buffers.

Proposition 9. *Consider $\sigma_0 \in \Sigma_0(I)$ and let us define the following regular languages:*

$$\begin{aligned} \gamma &\in ((\mathsf{Loc} \times \mathsf{Val})^+)^* \\ \nu &\in (\mathsf{Loc} \times \mathsf{Val})^* \\ \delta &::= \gamma \mid \nu\,\mathsf{lock}\,\gamma \end{aligned}$$

– *If* $\sigma_0 \xrightarrow{\lambda}{}^*_{C(L)} \sigma$*, then every store buffer in* σ *belongs to the following regular language:*

$$(\delta \cup \delta\,\mathsf{call}\,\gamma)\,(\mathsf{ret}\,\gamma\,\mathsf{call}\,\gamma)^*\,(\varepsilon \cup \mathsf{ret}\,\gamma)$$

– *If* $\sigma_0 \xrightarrow{\lambda}{}^*_{C_{\mathsf{sig}(L)}(\cdot)} \sigma$*, then every store buffer in* σ *belongs to the following regular language:*

$$(\delta\ \cup \mathsf{call}\,\gamma)\,(\mathsf{ret}\,\mathsf{call}\,\gamma)^*\,(\varepsilon \cup \mathsf{ret})$$

– *If* $\sigma_0 \xrightarrow{\lambda}{}^*_{\mathsf{MGC}(L)} \sigma$*, then every store buffer in* σ *belongs to the following regular language:*

$$\delta \mid (\varepsilon \cup \delta\,\mathsf{call})\,(\mathsf{ret}\,\gamma\,\mathsf{call})^*\,(\varepsilon \cup \mathsf{ret}\,\gamma)$$

Consider a store buffer α containing at least one call or ret marker. In this case, we can classify every non-marker entry β in the buffer α as a client or a library entry, depending on its position relative to the markers in the buffer. Namely, β is a library entry if $\alpha = \alpha'\,\mathsf{ret}\,\alpha''\beta\alpha'''$, where α'' does not contain call, or $\alpha = \alpha'\beta\alpha''\,\mathsf{call}\,\alpha'''$, where α'' does not contain ret; otherwise, it is a client entry.

Consider two store buffers α_1 and α_2 from configurations in the client-local and library-local semantics, respectively. Then the former one does not have any entries in between a ret marker and an older call marker, and the latter—in between a call marker and an older ret marker. We define the combination of $\alpha_1 \circ_\rho \alpha_2$, parameterised by the current program position ρ. Assume that the projections of α_1 and α_2 to call and ret markers are equal; otherwise, their combination is undefined. If α_1 and α_2 do not contain any call or ret markers, then

$$\alpha_1 \circ_\rho \alpha_2 = \begin{cases} \alpha_1, & \text{if } \rho \in N \text{ and } \alpha_2 = \varepsilon; \\ \alpha_2, & \text{if } \rho \in N \times N \text{ and } \alpha_1 = \varepsilon; \\ \text{undefined}, & \text{otherwise.} \end{cases}$$

If α_1 and α_2 contain at least one call or ret marker, then there exists a unique sequence such that its projection to client entries (and call and ret markers) is α_1 and to library entries (and the markers) is α_2. We let $\alpha_1 \circ_\rho \alpha_2$ be equal to this sequence. We lift $\circ$ to vectors of store buffers as follows:

$$\forall t \in \mathsf{CPUid}.\, (b_1 \circ_{\mathsf{pc}} b_2)(t) = b_1(t) \circ_{\mathsf{pc}(t)} b_2(t).$$

We define $\circ$ on sets of active CPUs as follows: $\mathsf{CPUid} \circ \mathsf{CPUid} = \mathsf{CPUid}$, $\{t\} \circ \mathsf{CPUid} = \mathsf{CPUid} \circ \{t\} = \{t\}$; all other cases are undefined.

Finally, we lift $\circ$ to configurations as follows:

$$(\mathsf{pc}_1, \theta_1, b_1, h_1, K_1) \circ (\mathsf{pc}_2, \theta_2, b_2, h_2, K_2) = \\ (\mathsf{pc}_1 \circ \mathsf{pc}_2, \theta_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)} \theta_2, b_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)} b_2, h_1 \uplus h_2, K_1 \circ K_2).$$

A.2 Proof of Lemma 6

Let $M = \mathsf{sig}(L)$. Consider $\tau \in [\![C(L)]\!](I_c \circ I_l)$. We show that, if τ does not contain a $\top$ configuration, then there exist $\eta \in [\![C_M(\cdot)]\!]I_c$ and $\xi \in [\![\mathsf{MGC}(L)]\!]I_l$ such

that $\mathsf{client}(\eta) = \mathsf{client}(\tau)$ and $\mathsf{lib}(\xi) = \mathsf{lib}(\tau)$; then $\mathsf{history}(\eta) = \mathsf{history}(\xi)$. We also show that, if τ contains a $\top$ configuration, then either $C_M(\cdot)$ or $\mathsf{MGC}(L)$ is unsafe, contradicting the assumption of the theorem. The latter implies the safety of $C(L)$.

Let $\sigma^0 = (\mathsf{pc}_0, \theta_0, b_0, h_0, \mathsf{CPUid}) \in \Sigma_0(I_c \circ I_l)$ be the initial configuration of τ. Then for some $h_1^0 \in I_c$ and $h_2^0 \in I_l$ we have $h_0 = h_1^0 \uplus h_2^0$. Let

$$\sigma_1^0 = (\mathsf{pc}_0, \theta_0, b_0, h_1^0, \mathsf{CPUid}) \wedge \sigma_2^0 = (\mathsf{pc}_0', \theta_0, b_0, h_2^0, \mathsf{CPUid}),$$

where $\mathsf{pc}_0'(t) = v_{\mathsf{mgc}}^t$ for all $t \in \mathsf{CPUid}$. The computations η and ξ we construct start from configurations σ_1^0 and σ_2^0. Our construction first considers every finite prefix τ_1 of τ and builds client-local and library-local computations for τ_1. Then we construct the desired computations η and ξ as the limits of these sequences (our construction is such that this limit exists). The following claim lies at the core of our construction:

Consider a finite prefix of τ:

$$\sigma^0 \xrightarrow{\lambda_1}{}^*_{C(L)} \sigma,$$

where $\sigma \neq \top$, and configurations $\sigma_1, \sigma_2 \in \mathsf{Config} - \{\top\}$ such that $\sigma_1 \circ \sigma_2 = \sigma$,

$$\sigma_1^0 \xrightarrow{\mathsf{client}(\lambda_1)}{}^*_{C_M(\cdot)} \sigma_1 \ \wedge \ \sigma_2^0 \xrightarrow{\mathsf{lib}(\lambda_1)}{}^*_{\mathsf{MGC}(L)} \sigma_2,$$

and no CPU is at the program position from $\{(v_{\mathsf{end}}^m, \_) \mid m \in M\}$ in σ_1. If for some $\sigma' \in \mathsf{Config} - \{\top\}$

$$\sigma \xrightarrow{\lambda}_{C(L)} \sigma',$$

then for some $\sigma_1', \sigma_2' \in \mathsf{Config} - \{\top\}$ we have $\sigma' = \sigma_1' \circ \sigma_2'$ and

$$\sigma_1^0 \xrightarrow{\mathsf{client}(\lambda_1\lambda)}{}^*_{C_M(\cdot)} \sigma_1' \ \wedge \ \sigma_2^0 \xrightarrow{\mathsf{lib}(\lambda_1\lambda)}{}^*_{\mathsf{MGC}(L)} \sigma_2',$$

where no CPU is at the program position from $\{(v_{\mathsf{end}}^m, \_) \mid m \in M\}$ in σ_1'. Also, if

$$\sigma \xrightarrow{\lambda}_{C(L)} \top,$$

then

$$\sigma_1^0 \xrightarrow{\mathsf{client}(\lambda_1\lambda)}{}^*_{C_M(\cdot)} \top \ \vee \ \sigma_2^0 \xrightarrow{\mathsf{lib}(\lambda_1\lambda)}{}^*_{\mathsf{MGC}(L)} \top.$$

To prove the claim, we assume $\lambda_1, \lambda, \sigma, \sigma_1, \sigma_2, \sigma'$ satisfying the above assumptions. We consider several cases, depending on the rule of the operational semantics used to derive σ'.

– A client transition using LOCAL, where $\lambda = \varepsilon$. In this case, for some t, v, c, v', pc, pc_1, pc_2, r, r', r'', θ, θ_1, θ_2, b, b_1, b_2, h, h_1, h_2, K, K_1 and K_2, we have $(v, c, v') \in T$, $t \in K$, $c \in \mathsf{Local}$, $r' \in f_c(r)$, $b(t)$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$, $\theta(t)$, $\theta_1(t)$, $\theta_2(t)$ are undefined and

$$\begin{array}{l}
\sigma = (\mathsf{pc}[t:v], \theta[t:r], b, h, K) \ \wedge \\
\sigma' = (\mathsf{pc}[t:v'], \theta[t:r'], b, h, K) \ \wedge \\
\sigma_1 = (\mathsf{pc}_1[t:v], \theta_1[t:r], b_1, h_1, K_1) \ \wedge \\
\sigma_2 = (\mathsf{pc}_2[t:v_{\mathsf{mgc}}^t], \theta_2[t:r''], b_2, h_2, K_2) \ \wedge \\
\mathsf{pc}_1 \circ \mathsf{pc}_2 = \mathsf{pc} \ \wedge \ \theta_1 \circ_{\mathsf{pc}} \theta_2 = \theta \ \wedge \ b_1 \circ_{\mathsf{pc}} b_2 = b \ \wedge \ h_1 \uplus h_2 = h \ \wedge \\
K = K_1 \circ K_2.
\end{array}$$

Let

$$\sigma_1' = (\mathsf{pc}_1[t : v'], \theta[t : r'], b_1, h_1, K_1)$$

and $\sigma_2' = \sigma_2$. Then $\sigma' = \sigma_1' \circ \sigma_2'$ and $\sigma_1 \xrightarrow{\lambda}_{C_M(\cdot)} \sigma_1'$.

– A client transition using WRITE, where $\lambda = (t, \mathsf{write}(x, u))$. In this case, for some v, c, v', pc, pc_1, pc_2, r, r', θ, θ_1, θ_2, α, α_1, α_2, b, b_1, b_2, h, h_1, h_2, K, K_1 and K_2, we have $(v, c, v') \in T$, $t \in K$, $c \in \mathsf{Write}$, $x \in \mathsf{dom}(h)$, $(x, u) \in f_c(r)$, $b(t)$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$, $\theta(t)$, $\theta_1(t)$, $\theta_2(t)$ are undefined and

$$\begin{array}{l}
\sigma = (\mathsf{pc}[t : v], \theta[t : r], b[t : \alpha], h, K) \ \wedge \\
\sigma' = (\mathsf{pc}[t : v'], \theta[t : r], b[t : (x, u)\,\alpha], h, K) \ \wedge \\
\sigma_1 = (\mathsf{pc}_1[t : v], \theta_1[t : r], b_1[t : \alpha_1], h_1, K_1) \ \wedge \\
\sigma_2 = (\mathsf{pc}_2[t : v^t_{\mathsf{mgc}}], \theta_2[t : r'], b_2[t : \alpha_2], h_2, K_2) \ \wedge \\
\mathsf{pc}_1 \circ \mathsf{pc}_2 = \mathsf{pc} \ \wedge \ \theta_1 \circ_{\mathsf{pc}} \theta_2 = \theta \ \wedge \ b_1 \circ_{\mathsf{pc}} b_2 = b \ \wedge \ h_1 \uplus h_2 = h \ \wedge \\
\alpha_1 \circ_v \alpha_2 = \alpha \ \wedge \ K_1 \circ K_2 = K.
\end{array}$$

Since $C_M(\cdot)$ is safe, we have $x \in \mathsf{dom}(h_1)$. Let

$$\sigma_1' = (\mathsf{pc}_1[t : v'], \theta[t : r], b_1[t : (x, u)\,\alpha_1], h_1, K_1)$$

and $\sigma_2' = \sigma_2$. Then $\sigma' = \sigma_1' \circ \sigma_2'$ and $\sigma_1 \xrightarrow{\lambda}_{C_M(\cdot)} \sigma_1'$.

– A client transition using LOCK, where $\lambda = (t, \mathsf{lock})$. Then for some v, v', pc, pc_1, pc_2, θ, θ_1, θ_2, α, α_1, α_2, b, b_1, b_2, h, h_1 and h_2, we have $(v, \mathsf{lock}, v') \in T$ $b(t)$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$ are undefined and

$$\begin{array}{l}
\sigma = (\mathsf{pc}[t : v], \theta, b[t : \alpha], h, \mathsf{CPUid}) \wedge \\
\sigma' = (\mathsf{pc}[t : v'], \theta, b[t : \mathsf{lock}\,\alpha], h, \{t\}) \wedge \\
\sigma_1 = (\mathsf{pc}_1[t : v], \theta_1, b_1[t : \alpha_1], h_1, \mathsf{CPUid}) \wedge \\
\sigma_2 = (\mathsf{pc}_2[t : v^t_{\mathsf{mgc}}], \theta_2, b_2[t : \alpha_2], h_2, \mathsf{CPUid}) \wedge \\
\mathsf{pc}_1 \circ \mathsf{pc}_2 = \mathsf{pc} \ \wedge \ \theta_1 \circ_{\mathsf{pc}[t:v]} \theta_2 = \theta \ \wedge \ b_1 \circ_{\mathsf{pc}} b_2 = b \ \wedge \ h_1 \uplus h_2 = h \ \wedge \\
\alpha_1 \circ_v \alpha_2 = \alpha.
\end{array}$$

Let

$$\sigma_1' = (\mathsf{pc}_1[t : v'], \theta_1, b_1[t : \mathsf{lock}\,\alpha_1], h_1, \{t\})$$

and $\sigma_2' = \sigma_2$. Then $\sigma' = \sigma_1' \circ \sigma_2'$ and $\sigma_1 \xrightarrow{\lambda}_{C_M(\cdot)} \sigma_1'$.

– A client transition using UNLOCK, where $\lambda = (t, \mathsf{unlock})$. Then for some v, v', pc, pc_1, pc_2, θ, θ_1, θ_2, α, α_1, α_2, b, b_1, b_2, h, h_1, h_2, K_1 and K_2, we have $(v, \mathsf{unlock}, v') \in T$, $b(t)$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$ are undefined and

$$\begin{array}{l}
\sigma = (\mathsf{pc}[t : v], \theta, b[t : (x_1, u_1) \ldots (x_l, u_l)\,\mathsf{lock}\,\alpha], h, \{t\}) \wedge \\
\sigma' = (\mathsf{pc}[t : v'], \theta, b[t : \langle (x_1, u_1) \ldots (x_l, u_l) \rangle\,\alpha], h, \mathsf{CPUid}) \wedge \\
\sigma_1 = (\mathsf{pc}_1[t : v], \theta_1, b_1[t : (x_1, u_1) \ldots (x_l, u_l)\,\mathsf{lock}\,\alpha_1], h_1, K_1) \wedge \\
\sigma_2 = (\mathsf{pc}_2[t : v^t_{\mathsf{mgc}}], \theta_2, b_2[t : \alpha_2], h_2, K_2) \wedge \\
\mathsf{pc}_1 \circ \mathsf{pc}_2 = \mathsf{pc} \ \wedge \ \theta_1 \circ_{\mathsf{pc}[t:v]} \theta_2 = \theta \ \wedge \ b_1 \circ_{\mathsf{pc}} b_2 = b \ \wedge \ h_1 \uplus h_2 = h \ \wedge \\
\alpha_1 \circ_v \alpha_2 = \alpha \ \wedge \ K_1 \circ K_2 = \{t\}.
\end{array}$$

The last equality implies that $K_1 = \{t\}$ and $K_2 = \mathsf{CPUid}$, or $K_1 = \mathsf{CPUid}$ and $K_2 = \{t\}$. Only the former holds. This is because, in our operational semantics, if

a CPU has a lock marker in its store buffer, then it is the only active CPU. In the case here, the lock marker is in the store buffer of t in the client-local computation, so $K_1 = \{t\}$. Let

$$\sigma_1' = (\mathsf{pc}_1[t : v'], \theta_1, b_1[t : \langle (x_1, u_1) \dots (x_l, u_l) \rangle\, \alpha_1], h_1, \mathsf{CPUid})$$

and $\sigma_2' = \sigma_2$. Then $\sigma' = \sigma_1' \circ \sigma_2'$ and $\sigma_1 \xrightarrow{\lambda}_{C_M(\cdot)} \sigma_1'$.

– A client transition using READ, where $\lambda = (t, \mathsf{read}(x, u))$. In this case, for some v, c, v', pc, pc_1, pc_2, r, r', θ, θ_1, θ_2, α, α_1, α_2, b, b_1, b_2, h, h_1, h_2, K, K_1 and K_2, we have $(v, c, v') \in T$, $t \in K$, $c \in \mathsf{Read}$, $x \in f_c(r)$, $u = \mathsf{lookup}(\alpha, h, x) \neq \top$, $b(t)$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$, $\theta(t)$, $\theta_1(t)$, $\theta_2(t)$ are undefined and

$$\begin{array}{l}
\sigma = (\mathsf{pc}[t : v], \theta[t : r], b[t : \alpha], h, K) \wedge \\
\sigma' = (\mathsf{pc}[t : v'], \theta[t : r[\mathtt{r}_1 : u]], b[t : \alpha], h, K) \wedge \\
\sigma_1 = (\mathsf{pc}_1[t : v], \theta_1[t : r], b_1[t : \alpha_1], h_1, K_1) \wedge \\
\sigma_2 = (\mathsf{pc}_2[t : v^t_{\mathsf{mgc}}], \theta_2[t : r'], b_2[t : \alpha_2], h_2, K_2) \wedge \\
\mathsf{pc}_1 \circ \mathsf{pc}_2 = \mathsf{pc} \ \wedge\ \theta_1 \circ_{\mathsf{pc}} \theta_2 = \theta \ \wedge\ b_1 \circ_{\mathsf{pc}} b_2 = b \ \wedge\ h_1 \uplus h_2 = h \ \wedge \\
\alpha_1 \circ_v \alpha_2 = \alpha \ \wedge\ K_1 \circ K_2 = K.
\end{array}$$

Since $C_M(\cdot)$ is safe, $\mathsf{lookup}(\alpha_1, h_1, x) \neq \top$. We now show that

$$\mathsf{lookup}(\alpha_1, h_1, x) = \mathsf{lookup}(\alpha, h, x) = u.$$

If α does not contain entries for x, then neither does α_1, so

$$\mathsf{lookup}(\alpha, h, x) = h(x) = h_1(x) = \mathsf{lookup}(\alpha_1, h_1, x).$$

Assume now that α contains entries for x, but the newest such entry is not in α_1. Then this entry was written by the library code. Note that in any reachable configuration in our semantics, a store buffer contains only write requests for cells in the domain of the current heap. Thus, the address x belongs to the library part of the heap throughout the computation. But then $h_1(x)$ is undefined and α cannot contain an entry for x, so $\mathsf{lookup}(\alpha_1, h_1, x) = \top$. The resulting contradiction shows that the newest entry for x in α is also the newest such entry in α_1 and, thus, $u = \mathsf{lookup}(\alpha_1, h_1, x)$. Let

$$\sigma_1' = (\mathsf{pc}_1[t : v'], \theta_1[t : r[\mathtt{r}_1 : u]], b_1[t : \alpha_1], h_1, K_1)$$

and $\sigma_2' = \sigma_2$. Then $\sigma' = \sigma_1' \circ \sigma_2'$ and $\sigma_1 \xrightarrow{\lambda}_{C_M(\cdot)} \sigma_1'$.

– A library transition using LOCAL, WRITE, LOCK, UNLOCK, READ. These cases are handled similarly to previous ones.

– A transition using FLUSH, where $\lambda = (t, \mathsf{flush}(x_l, u_l)) \dots (t, \mathsf{flush}(x_1, u_1))$ flushes a client entry, i.e., either the corresponding store buffer contains a call or ret marker and the entry is a client one, or the buffer does not contain the markers, but the program position for the corresponding CPU is inside the client code. Then for some pc, pc_1, pc_2, ρ, ρ_1, ρ_2, θ, θ_1, θ_2, α, α_1, α_2, b, b_1, b_2, h, h_1, h_2, K, K_1 and

K_2, we have $t \in K$, $b(t)$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$ are undefined and

$$\begin{array}{l}\sigma = (\mathsf{pc}[t:\rho], \theta, b[t:\alpha\,\langle(x_1,u_1)\ldots(x_l,u_l)\rangle], h, K) \ \wedge \\ \sigma' = (\mathsf{pc}[t:\rho], \theta, b[t:\alpha], h[x_l:u_l]\ldots[x_1:u_1], K) \ \wedge \\ \sigma_1 = (\mathsf{pc}_1[t:\rho_1], \theta_1, b_1[t:\alpha_1\,\langle(x_1,u_1)\ldots(x_l,u_l)\rangle], h_1, K_1) \ \wedge \\ \sigma_2 = (\mathsf{pc}_2[t:\rho_2], \theta_2, b_2[t:\alpha_2], h_2, K_2) \ \wedge \\ \mathsf{pc}_1 \circ \mathsf{pc}_2 = \mathsf{pc} \ \wedge \ \rho_1 \circ \rho_2 = \rho \ \wedge \ \theta_1 \circ_{\mathsf{pc}[t:\rho]} \theta_2 = \theta \ \wedge \ b_1 \circ_{\mathsf{pc}} b_2 = b \ \wedge \\ h_1 \uplus h_2 = h \ \wedge \ \alpha_1 \circ_\rho \alpha_2 = \alpha \ \wedge \ K_1 \circ K_2 = K.\end{array}$$

Let

$$\sigma_1' = (\mathsf{pc}_1[t:\rho_1], \theta_1, b_1[t:\alpha_1], h_1[x_l:u_l]\ldots[x_1:u_1], K_1)$$

and $\sigma_2' = \sigma_2$. There exist previous WRITE transitions in the client-local computation that put the write requests for $x_1, \ldots, x_l$ into the store buffer. Since $C_M(\cdot)$ is safe and the domain of the heap does not change during a computation, we get $x_1, \ldots, x_l \in \mathsf{dom}(h_1)$. Then $\sigma' = \sigma_1' \circ \sigma_2'$ and $\sigma_1 \xrightarrow{\lambda}_{C_M(\cdot)} \sigma_1'$.

– A transition using FLUSH, where $\lambda = (t, \mathsf{flush}(x_l, u_l))\ldots(t, \mathsf{flush}(x_1, u_1))$ flushes a library entry. This case is handled similarly to the previous one.

– A transition using FLUSH-MARKER, where $\lambda = (t, \mathsf{flush}(\beta))$ and $\beta \in \{\mathsf{call}, \mathsf{ret}\}$. Then for some pc, pc_1, pc_2, ρ, ρ_1, ρ_2, θ, θ_1, θ_2, α, α_1, α_2, b, b_1, b_2, h, h_1, h_2, K, K_1 and K_2, we have $t \in K$, $b(t)$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$ are undefined and

$$\begin{array}{l}\sigma = (\mathsf{pc}[t:\rho], \theta, b[t:\alpha\beta], h, K) \ \wedge \\ \sigma' = (\mathsf{pc}[t:\rho], \theta, b[t:\alpha], h, K) \ \wedge \\ \sigma_1 = (\mathsf{pc}_1[t:\rho_1], \theta_1, b_1[t:\alpha_1\beta], h_1, K_1) \ \wedge \\ \sigma_2 = (\mathsf{pc}_2[t:\rho_2], \theta_2, b_2[t:\alpha_2\beta], h_2, K_2) \ \wedge \\ \mathsf{pc}_1 \circ \mathsf{pc}_2 = \mathsf{pc} \ \wedge \ \theta_1 \circ_{\mathsf{pc}[t:\rho]} \theta_2 = \theta \ \wedge \ b_1 \circ_{\mathsf{pc}} b_2 = b \ \wedge \ h_1 \uplus h_2 = h \ \wedge \\ \alpha_1\beta \circ_\rho \alpha_2\beta = \alpha\beta \ \wedge \ \rho_1 \circ \rho_2 = \rho \ \wedge \ K_1 \circ K_2 = K.\end{array}$$

Let

$$\begin{array}{l}\sigma_1' = (\mathsf{pc}_1[t:\rho_1], \theta_1, b_1[t:\alpha_1], h_1, K_1) \ \wedge \\ \sigma_2' = (\mathsf{pc}_2[t:\rho_2], \theta_2, b_2[t:\alpha_2], h_2, K_2).\end{array}$$

Assume that α does not contain call or ret markers. Our semantics is such that, in this case $\alpha_1 \neq \varepsilon$ only if $\rho_1 \in N$ and $\alpha_2 \neq \varepsilon$ only if $\rho_2 \in N \times N$. Hence, $\alpha_1 \circ_\rho \alpha_2 = \alpha$, so that $\sigma' = \sigma_1' \circ \sigma_2'$. Besides, $\sigma_1 \xrightarrow{\lambda}_{C_M(\cdot)} \sigma_1'$ and $\sigma_2 \xrightarrow{\lambda}_{\mathsf{MGC}(L)} \sigma_2'$.

– A client transition using XLOCK, where $\lambda = (t, \mathsf{xlock})$. Then for some v, v', pc, pc_1, pc_2, θ, θ_1, θ_2, b, b_1, b_2, h, h_1 and h_2, we have $(v, \mathsf{xlock}, v') \in T$, $b(t)$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$ are undefined and

$$\begin{array}{l}\sigma = (\mathsf{pc}[t:v], \theta, b[t:\varepsilon], h, \mathsf{CPUid}) \wedge \\ \sigma' = (\mathsf{pc}[t:v'], \theta, b[t:\varepsilon], h, \{t\}) \wedge \\ \sigma_1 = (\mathsf{pc}_1[t:v], \theta_1, b_1[t:\varepsilon], h_1, \mathsf{CPUid}) \wedge \\ \sigma_2 = (\mathsf{pc}_2[t:v^t_{\mathsf{mgc}}], \theta_2, b_2[t:\varepsilon], h_2, \mathsf{CPUid}) \wedge \\ \mathsf{pc}_1 \circ \mathsf{pc}_2 = \mathsf{pc} \ \wedge \ \theta_1 \circ_{\mathsf{pc}[t:v]} \theta_2 = \theta \ \wedge \ b_1 \circ_{\mathsf{pc}} b_2 = b \ \wedge \ h_1 \uplus h_2 = h.\end{array}$$

Let

$$\sigma_1' = (\mathsf{pc}_1[t:v'], \theta_1, b_1[t:\varepsilon], h_1, \{t\})$$

and $\sigma_2' = \sigma_2$. Then $\sigma' = \sigma_1' \circ \sigma_2'$ and $\sigma_1 \xrightarrow{\lambda}_{C_M(\cdot)} \sigma_1'$.

– A client transition using XUNLOCK. This case is handled analogously to the cases of UNLOCK and FLUSH.
– A library transition using XLOCK or XUNLOCK. These cases are handled analogously to the previous ones.
– A transition using CALL, where $\lambda = (t, \mathsf{call}\ m(r))$. Then for some v, v', pc, pc_1, pc_2, r', θ, θ_1, θ_2, α, α_1, α_2, b, b_1, b_2, h, h_1 and h_2, we have $(v, m, v') \in T$, $b(t)$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$, $\theta(t)$, $\theta_1(t)$, $\theta_2(t)$ are undefined and

$$\begin{array}{l}\sigma = (\mathsf{pc}[t : v], \theta[t : r], b[t : \alpha], h, \mathsf{CPUid})\ \wedge \\ \sigma' = (\mathsf{pc}[t : (\mathsf{start}_m, v')], \theta[t : r], b[t : \mathsf{call}\,\alpha], h, \mathsf{CPUid})\ \wedge \\ \sigma_1 = (\mathsf{pc}_1[t : v], \theta_1[t : r], b_1[t : \alpha_1], h_1, \mathsf{CPUid})\ \wedge \\ \sigma_2 = (\mathsf{pc}_2[t : v^t_{\mathsf{mgc}}], \theta_2[t : r'], b_2[t : \alpha_2], h_2, \mathsf{CPUid})\ \wedge \\ \mathsf{pc}_1 \circ \mathsf{pc}_2 = \mathsf{pc}\ \wedge\ \theta_1 \circ_{\mathsf{pc}} \theta_2 = \theta\ \wedge\ b_1 \circ_{\mathsf{pc}} b_2 = b\ \wedge\ h_1 \uplus h_2 = h\ \wedge \\ \alpha_1 \circ_v \alpha_2 = \alpha.\end{array}$$

Let

$$\begin{array}{l}\sigma'_1 = (\mathsf{pc}_1[t : (v^m_{\mathsf{start}}, v')], \theta_1[t : r], b_1[t : \mathsf{call}\,\alpha_1], h_1, \mathsf{CPUid})\ \wedge \\ \sigma'_2 = (\mathsf{pc}_2[t : v^t_{\mathsf{mgc}}], \theta_2[t : r], b_2[t : \alpha_2], h_2, \mathsf{CPUid})\ \wedge \\ \sigma''_2 = (\mathsf{pc}_2[t : (\mathsf{start}_m, v^t_{\mathsf{mgc}})], \theta_2[t : r], b_2[t : \mathsf{call}\,\alpha_2], h_2, \mathsf{CPUid}).\end{array}$$

Then $\sigma' = \sigma'_1 \circ \sigma''_2$, $\sigma_1 \xrightarrow{\lambda}_{C_M(\cdot)} \sigma'_1$ and $\sigma_2 \xrightarrow{\varepsilon}_{\mathsf{MGC}(L)} \sigma'_2 \xrightarrow{\lambda}_{\mathsf{MGC}(L)} \sigma''_2$, where the first transition of $\mathsf{MGC}(L)$ results from executing havoc in the most general client.
– A client transition using RET. This case is similar to the previous one.
– A transition using WRITE-$\top$ or READ-$\top$. These two rules are applicable when the cell being written to or read from is not in the domain of the heap in σ. In this case, the cell is also not in the domain of the client-local heap in σ_1 or the library-local heap in σ_2. Thus, it is easy to show that either $\sigma_1 \xrightarrow{\varepsilon}_{C_M(\cdot)} \top$ or $\sigma_2 \xrightarrow{\varepsilon}_{\mathsf{MGC}(L)} \top$, depending on whether the faulting transition in τ is executed by the client or the library.

□

A.3 Proof of Lemma 7

Below we sometimes write $\sqsubseteq_\pi$ instead of $\sqsubseteq$ to make the bijection π used to establish the relation between histories explicit. Also, we say that the machine is locked in a configuration σ, if $\sigma = (\_, \_, \_, \_, \{t\})$ for some t.

In the following, we use the fact that, according to our semantics, the machine cannot be locked in a configuration from which it executes a transition producing a history action. This is because we prohibit method calls within atomic blocks and flushes within lock..unlock blocks, and because xlock..xunlock blocks start executing with an empty store buffer.

Consider a computation $\tau' \in [\![\mathsf{MGC}(L)]\!] I_l$. Assume histories H and H' such that $\mathsf{history}(\tau') = H'$ and $H \sqsubseteq H'$. We need to prove that there exists a computation $\tau \in [\![\mathsf{MGC}(L)]\!] I_l$ such that $\mathsf{history}(\tau) = H$. To this end, we define a (possibly infinite)

sequence of steps that transforms τ' into τ. The computation τ is constructed using a sequence of computations $\xi_k \in [\![\mathsf{MGC}(L)]\!] I_l$, defined for every finite prefix H_k of H of length k. Every computation ξ_k is such that for some prefix η_k of ξ_k we have $\mathsf{history}(\eta_k) = H_k$ and $H \sqsubseteq_\pi \mathsf{history}(\xi_k)$, where π is an identity on actions from H_k, i.e., the first k actions of H. Additionally, for all i, j such that $i < j$, η_i is a prefix of η_j. Hence, the sequence of computations η_k has a limit τ such that for every k, η_k is a prefix of τ and $\mathsf{history}(\tau) = H$. Then, as we show, $\tau \in [\![\mathsf{MGC}(L)]\!] I_l$.

To construct the sequence, we let $\xi_0 = \tau$ and let the prefix η_0 contain all transitions preceding the first history action in ξ_0. The computation ξ_{k+1} is constructed from the computation ξ_k by applying the following lemma for $\tau_1 = \eta_k$, $\tau_1\tau_2 = \xi_k$, $H'_1 = H_k$ and $H'_1 \psi H'_2 = H$.

Lemma 10. *Assume that* $\mathsf{MGC}(L)$ *is safe for* I_l*. Consider a history* $H'_1 \psi H'_2$ *and a computation* $\tau_1\tau_2 \in [\![\mathsf{MGC}(L)]\!] I_l$ *such that*

- *the machine is not locked in the final configuration of* τ_1*; and*
- *the following two properties hold:*

$$\mathsf{history}(\tau_1) = H'_1, \tag{3}$$
$$H'_1 \psi H'_2 \sqsubseteq_\pi \mathsf{history}(\tau_1\tau_2), \tag{4}$$

where π *is an identity on actions from* H'_1*.*

Then there exist computations τ'_2 *and* τ''_2 *such that*

- $\tau_1\tau'_2\tau''_2 \in [\![\mathsf{MGC}(L)]\!] I_l$*;*
- *the machine is not locked in the final configuration of* $\tau_1\tau'_2$*; and*
- *the following two properties hold:*

$$\mathsf{history}(\tau_1\tau'_2) = H'_1\psi, \tag{5}$$
$$H'_1 \psi H'_2 \sqsubseteq_{\pi'} \mathsf{history}(\tau_1\tau'_2\tau''_2) \tag{6}$$

where π' *is an identity on actions from* $H'_1\psi$*.*

To prove Lemma 10, we convert $\tau_1\tau_2$ into $\tau_1\tau'_2\tau''_2$ by applying a finite number of transformations that preserve its properties of interest, described by the lemmas below.

Lemma 11. *Assume*

$$\sigma_0 \xrightarrow{\lambda_0}{}^*_{\mathsf{MGC}(L)} \sigma_1 \xrightarrow{\lambda_1}{}^*_{\mathsf{MGC}(L)} \sigma_2 \xrightarrow{\lambda_2}_{\mathsf{MGC}(L)} \sigma_3,$$

where

- $\sigma_0 \in \Sigma_0(I_l)$*;*
- $\sigma_3 \neq \top$*;*
- *the machine is not locked in* σ_1*;*
- *the transition* λ_2 *is obtained using* LOCAL*,* CALL *or* FLUSH-MARKER *(the case of flushing* call*) for a CPU* t*;*
- *if the transition* λ_2 *is obtained using* LOCAL *or* CALL*, then CPU* t *executes only* FLUSH *transitions in the part* λ_1 *of the computation; and*

– *if the transition λ_2 is obtained using* FLUSH-MARKER *flushing* call*, then CPU t does not execute* CALL*,* RET*,* FLUSH *or* FLUSH-MARKER *transitions in the part λ_1 of the computation.*

Then there is a configuration σ_2' such that the machine is not locked in σ_2' and

$$\sigma_0 \xrightarrow{\lambda_0}{}^*_{\mathsf{MGC}(L)} \sigma_1 \xrightarrow{\lambda_2}_{\mathsf{MGC}(L)} \sigma_2' \xrightarrow{\lambda_1}{}^*_{\mathsf{MGC}(L)} \sigma_3.$$

Lemma 12. *Assume*

$$\sigma_0 \xrightarrow{\lambda_0}{}^*_{\mathsf{MGC}(L)} \sigma_1 \xrightarrow{\lambda_1}_{\mathsf{MGC}(L)} \sigma_2 \xrightarrow{\lambda_2}{}^*_{\mathsf{MGC}(L)} \sigma_3,$$

where

– *$\sigma_0 \in \Sigma_0(I_l)$;*
– *$\sigma_3 \neq \top$;*
– *the machine is not locked in σ_3;*
– *the transition λ_1 is obtained using* LOCAL*,* RET *or* FLUSH-MARKER *(the case of flushing* ret*) for a CPU t;*
– *if the transition λ_1 is obtained using* LOCAL *or* RET*, then CPU t executes only* FLUSH *transitions in the part λ_2 of the computation; and*
– *if the first transition λ_1 is obtained using* FLUSH-MARKER *flushing* ret*, then CPU t does not execute* CALL*,* RET*,* FLUSH *or* FLUSH-MARKER *transitions in the part λ_2 of the computation.*

Then there is a configuration σ_2' such that the machine is not locked in σ_2' and

$$\sigma_0 \xrightarrow{\lambda_0}{}^*_{\mathsf{MGC}(L)} \sigma_1 \xrightarrow{\lambda_2}{}^*_{\mathsf{MGC}(L)} \sigma_2' \xrightarrow{\lambda_1}_{\mathsf{MGC}(L)} \sigma_3.$$

Proof sketch for Lemmas 11 and 12. The proofs are mostly straightforward, with special care needed only when moving λ_2 or λ_1 over transitions affecting store buffers. The treatment of such cases relies on the fact that store buffers are FIFO. Thus, we can move a CALL transition earlier than a FLUSH transition by the same CPU in Lemma 11, and a RET transition later than FLUSH by the same CPU in Lemma 12. It also allows us to move a FLUSH-MARKER transition earlier or later than a WRITE transition by the same CPU in the two lemmas. In Lemma 11, we can also move a FLUSH-MARKER transition earlier than an XLOCK..XUNLOCK block of transitions by the same CPU. However, in Lemma 12, we cannot move a transition using FLUSH-MARKER flushing ret later than XLOCK or XUNLOCK transitions by the same CPU, because the latter include a memory barrier. We now show that in the case when λ_1 is obtained using FLUSH-MARKER flushing ret, a transition using XLOCK by CPU t cannot be present in the part λ_2 of the computation in Lemma 12 (since λ_2 is a history action, the machine is unlocked in σ_2 and thus a transition by t using XUNLOCK cannot be present in the part λ_2 either).

Assume the contrary. The XLOCK transition is executed by the library code, hence, in the computation there is a CALL transition preceding it such that there is no RET

transition between the CALL and the XLOCK transition. Since the part λ_2 of the computation does not contain CALL transitions by t, the CALL transition has to be in the part λ_0. Since the part λ_2 does not contain FLUSH or FLUSH-MARKER transitions, and the store buffer of t before the XLOCK transition has to be empty, the store buffer of t in σ_1 contains only a ret marker. Since there are no RET transitions between the CALL and the XLOCK transitions, in the configuration after the CALL transition, the store buffer of t already contains the ret marker, i.e., is of the form $\mathsf{call}\,\alpha_1\,\mathsf{ret}\,\alpha_2$. But this means that the buffer of t in σ_1 cannot contain only ret: if the ret marker has not been flushed at the end of λ_0, then neither has been the call marker. The contradiction shows that λ_2 cannot contain XLOCK transitions. □

Proof of Lemma 10. Let τ_3' be the part of the computation τ_2 preceding the transition producing ψ, and τ_4' be the part following it. Let CPU t be the one that executes the transition producing ψ. We first note that CPU t does not execute any transitions producing history actions in τ_3'. This is because such actions would precede ψ in $\mathsf{history}(\tau_1\tau_2)$, but by (3) and (4) would have to follow it in $H_1'\psi H_2'$, contradicting the fact that linearizability preserves the order of actions by the same CPU. We now have several cases, depending on the type of ψ.

- $\psi = (t, \mathsf{call}\ \_)$. By the definition of $\mathsf{MGC}(L)$, CPU t does not execute any transitions in τ_3' except FLUSH transitions and havoc transitions in the most general client. By Lemma 11, we can move the havoc transitions and the transition producing ψ to the beginning of τ_3', processing them left to right. Let τ_2' consist of the moved havoc transitions and the transition producing ψ, and τ_2'' consist of the transitions following them. Then, the machine is unlocked in the final configuration of τ_2'. Besides, any $(\_, \mathsf{ret}\ \_)$ or $(\_, \mathsf{flush}(\mathsf{ret}))$ action preceding ψ in $H_1'\psi H_2'$ has to be in H_1'. Since π is an identity on actions from H_1', this return action is in τ_1 and, hence, precedes ψ in the resulting computation. Thus, moving ψ to the left does not violate ordering constraints required by linearizability, and the computation satisfies (5) and (6).
- $\psi = (t, \mathsf{flush}(\mathsf{call}))$. The client in $\mathsf{MGC}(L)$ does not perform write commands, so, as we have observed before, store buffers in its computations do not contain entries between a call marker and an older ret marker. It follows that τ_3' does not contain any CALL, RET, FLUSH or FLUSH-MARKER transitions by t. Hence, by Lemma 11, we can move the transition producing ψ to the beginning of τ_3'. Let τ_2' be the transition of the resulting computation producing ψ, and τ_2'' consist of the transitions following it. Then the machine is unlocked in the final configuration of τ_2' and the computation satisfies (5) and (6).
- $\psi = (t, \mathsf{ret}\ \_)$ or $\psi = (t, \mathsf{flush}(\mathsf{ret}))$. Then τ_3' cannot contain an action $\varphi = (\_, \mathsf{call}\ \_)$ or $\varphi = (\_, \mathsf{flush}(\mathsf{call}))$, because in this case φ would precede ψ in $\mathsf{history}(\tau_1\tau_2)$. However, φ is in H_2' and, thus, ψ precedes φ in $H_1'\psi H_2'$, so this would contradict (4). Hence, there are no such actions in τ_3'. Moreover, by the definition of $\mathsf{MGC}(L)$, for any transition producing an action $\varphi = (t', \mathsf{ret}\ \_)$ in τ_3' there are no transitions by CPU t' in τ_3' following the transition φ other than FLUSH, FLUSH-MARKER flushing ret or havoc transitions. Since store buffers in computations of $\mathsf{MGC}(L)$ do not contain entries between a call marker and an older ret marker, for any transition producing $\varphi = (t', \mathsf{flush}(\mathsf{ret}))$ in τ_3' there are

no transitions by CPU t' in τ_3' following φ using FLUSH or FLUSH-MARKER. By Lemma 12, we can move all above-mentioned transitions producing actions $(t', \mathsf{ret}\ \_)$ or $(t', \mathsf{flush}(\mathsf{ret}))$ and havoc transitions in τ_3' to the position immediately preceding the transition producing ψ, processing them right to left. After this, we can them move them to the position immediately following ψ. Let $\tau_1\tau_2'$ be the prefix of the resulting computation up to and including ψ, and τ_2'' consist of the transitions following it. Then the machine is unlocked in the final configuration of τ_2'. Besides, since τ_3' does not contain actions of the form $(\_, \mathsf{call}\ \_)$ or $(\_, \mathsf{flush}(\mathsf{call}))$, any $(\_, \mathsf{ret}\ \_)$ or $(\_, \mathsf{flush}(\mathsf{ret}))$ action preceding another $(\_, \mathsf{call}\ \_)$ or $(\_, \mathsf{flush}(\mathsf{call}))$ action in the original computation will also precede it in the transformed one. Hence, the resulting computation satisfies (5) and (6).

□

A.4 Proof of Lemma 8

Let $M = \mathsf{sig}(L)$. Assume $\eta' \in [\![C_M(\cdot)]\!]I_c$ and $\xi' \in [\![\mathsf{MGC}(L)]\!]I_l$ such that $\mathsf{history}(\eta') = \mathsf{history}(\xi')$. Let η be the shortest prefix of η' such that $\mathsf{client}(\eta) = \mathsf{client}(\eta')$. Similarly, define ξ to be the shortest prefix of ξ' such that $\mathsf{history}(\xi) = \mathsf{history}(\xi')$. Then, every non-empty suffix of η or ξ contains transitions other than havoc transitions in method stubs or the threads of the most general client. Furthermore, no non-empty suffix of ξ continues with the lock held by some CPU and never released. Formally, this means that every non-empty suffix of ξ contains some configuration whose last component is CPUid, so that all CPUs can be scheduled. Finally, if ξ is infinite, so is $\mathsf{history}(\xi)$, and if ξ is finite, its last transition is in HAct and appears as the last element in $\mathsf{history}(\xi)$. We now construct the desired computation $\tau \in [\![C(L)]\!](I_c \circ I_l)$ from these preprocessed computations η and ξ. The properties of η and ξ noted above are used during the construction.

Let σ_0^1 and σ_0^2 be the initial configurations of the computations η and ξ; then $\sigma_0^1 \circ \sigma_0^2$ is defined by the assumption about initial heaps of clients and libraries in Section 3. We first build a series of finite computations

$$\tau_0, \tau_1, \tau_2, \ldots$$

such that for $i < j$, τ_i is a prefix of τ_j. Thus, the series has the limit, which is the desired computation τ.

The first element in the series is the empty computation consisting of the initial configuration $\sigma_0^1 \circ \sigma_0^2$ only. For the $(i{+}1)$-st element with $i > 0$, we assume that the i-th element τ_i has been constructed and satisfies the following property:

For some finite prefixes η_i of η and ξ_i of ξ of the form

$$\sigma_0^1 \xrightarrow{\mathsf{client}(\lambda_i)}{}^*_{C_M(\cdot)} \sigma_i^1 \ \wedge\ \sigma_0^2 \xrightarrow{\mathsf{lib}(\lambda_i)}{}^*_{\mathsf{MGC}(L)} \sigma_i^2,$$

we have

$$\mathsf{history}(\eta_i) = \mathsf{history}(\xi_i) \ \wedge\ \mathsf{client}(\tau_i) = \mathsf{client}(\eta_i) \ \wedge\ \mathsf{lib}(\tau_i) = \mathsf{lib}(\xi_i).$$

Furthermore, $\sigma_i^1 \circ \sigma_i^2$ is defined and τ_i is:

$$\sigma_0^1 \circ \sigma_0^2 \xrightarrow{\lambda_i}{}^*_{C(L)} \sigma_i^1 \circ \sigma_i^2.$$

We now define the $(i+1)$-th element τ_{i+1} that maintains this property. As we explained above, the computation τ_{i+1} is an extension of τ_i by one or more steps.

Let the following be the next transitions in the computations η and ξ (we consider the case when one of the computations has no next transition later):

$$\sigma_i^1 \xrightarrow{\varepsilon}{}^*_{C_M(\cdot)} \sigma_1 \xrightarrow{\lambda'}_{C_M(\cdot)} \sigma_1' \ \wedge \ \sigma_i^2 \xrightarrow{\varepsilon}{}^*_{\mathsf{MGC}(L)} \sigma_2 \xrightarrow{\lambda''}_{\mathsf{MGC}(L)} \sigma_2', \tag{7}$$

where $\sigma_1, \sigma_2, \sigma_1', \sigma_2' \in \mathsf{Config} - \{\top\}$, the computation transforming σ_i^1 into σ_1 is the maximal prefix of the first computation consisting only of havoc transitions in method stubs, and the computation transforming σ_i^2 into σ_2 is the maximal prefix of the second computation consisting only of havoc transitions in the client code of $\mathsf{MGC}(L)$. The non-havoc transitions exist in both cases, because of the preprocessing step described above. It is easy to show that $\sigma_i^1 \circ \sigma_i^2 = \sigma_1 \circ \sigma_2$, so that

$$\sigma_0^1 \circ \sigma_0^2 \xrightarrow{\lambda_i}{}^*_{C(L)} \sigma_1 \circ \sigma_2.$$

To construct τ_{i+1}, we make a case-split on the rules of the operational semantics used to obtain the non-havoc transitions λ' and λ''. We use one of these two transitions to extend τ_i to τ_{i+1}. The construction described below defines τ_{i+1} with a degree of non-determinism: the computation of $C(L)$ can sometimes be extended either using the transition from $C_M(\cdot)$ or the one from $\mathsf{MGC}(L)$. All possible results produce a valid computation of $C(L)$, and, as we show below, $\mathsf{client}(\tau) = \mathsf{client}(\eta)$.

Below we use symbols K_1 and K_2 to denote the last components of σ_1 and σ_2, respectively.

– CALL in η and CALL in ξ such that $\lambda' = \lambda'' = (t, \mathsf{call}\ m(r))$. Then for some v, v', pc_1, pc_2, θ_1, θ_2, b_1, b_2, α_1, α_2, h_1 and h_2, we have $(v, m, v') \in T$, $b_1(t)$, $b_2(t)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$, $\theta_1(t)$, $\theta_2(t)$ are undefined and

$$\begin{aligned}
&\sigma_1 = (\mathsf{pc}_1[t : v], \theta_1[t : r], b_1[t : \alpha_1], h_1, \mathsf{CPUid}) \ \wedge \\
&\sigma_2 = (\mathsf{pc}_2[t : v^t_{\mathsf{mgc}}], \theta_2[t : r], b_2[t : \alpha_2], h_2, \mathsf{CPUid}) \ \wedge \\
&\sigma_1' = (\mathsf{pc}_1[t : (v^m_{\mathsf{start}}, v')], \theta_1[t : r], b_1[t : \mathsf{call}\,\alpha_1], h_1, \mathsf{CPUid}) \ \wedge \\
&\sigma_2' = (\mathsf{pc}_2[t : (\mathsf{start}_m, v^t_{\mathsf{mgc}})], \theta_2[t : r], b_2[t : \mathsf{call}\,\alpha_2], h_2, \mathsf{CPUid}) \ \wedge \\
&\sigma_1 \circ \sigma_2 = ((\mathsf{pc}_1 \circ \mathsf{pc}_2)[t : v], (\theta_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)} \theta_2)[t : r], \\
&\qquad\qquad (b_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)} b_2)[t : \alpha_1 \circ_v \alpha_2], h_1 \uplus h_2, \mathsf{CPUid}).
\end{aligned}$$

Hence,

$$\begin{aligned}
\sigma_1' \circ \sigma_2' = ((\mathsf{pc}_1 \circ \mathsf{pc}_2)[t : (\mathsf{start}_m, v')], (\theta_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)} \theta_2)[t : r], \\
(b_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)} b_2)[t : \mathsf{call}\,(\alpha_1 \circ_v \alpha_2)], h_1 \uplus h_2, \mathsf{CPUid}).
\end{aligned}$$

Then, $\sigma_1 \circ \sigma_2 \xrightarrow{\lambda'}_{C(L)} \sigma_1' \circ \sigma_2'$. Thus, in this case the desired τ_{i+1} is obtained by extending τ_i with this transition.

– RET in η and RET in ξ such that $\lambda' = \lambda''$. This case is handled similarly to the previous one.
– The transition in ξ is generated by LOCAL, LOCK, UNLOCK, WRITE or READ, and we have $K_2 \subseteq K_1$. Note that by the definition of $\circ$ on K, at least one of K_1 and K_2 is CPUid. Our condition thus ensures that $K_1 = \mathsf{CPUid}$. We consider only the case of the LOCAL rule; the others are analogous. In this case for some $k \in \{\mathsf{start}, \mathsf{end}\}$, v, v', v_1, θ_1, θ_2, r, r', r'', pc_1, pc_2, b_1, b_2, h_1 and h_2, we have $(v, c, v') \in T$, $c \in \mathsf{Local}$, $r' \in f_c(r)$, $\mathsf{pc}_1(t)$, $\mathsf{pc}_2(t)$, $\theta_1(t)$, $\theta_2(t)$ are undefined and

$$\begin{array}{l}\sigma_1 = (\mathsf{pc}_1[t : (v_k^m, v_1)], \theta_1[t : r''], b_1, h_1, \mathsf{CPUid}) \ \wedge \\ \sigma_2 = (\mathsf{pc}_2[t : (v, v_{\mathsf{mgc}}^t)], \theta_2[t : r], b_2, h_2, K_2) \ \wedge \\ \sigma_2' = (\mathsf{pc}_2[t : (v', v_{\mathsf{mgc}}^t)], \theta_2[t : r'], b_2, h_2, K_2) \ \wedge \\ \sigma_1 \circ \sigma_2 = ((\mathsf{pc}_1 \circ \mathsf{pc}_2)[t : (v, v_1)], (\theta_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)} \theta_2)[t : r], \\ \qquad\qquad b_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)[t:(v,v_1)]} b_2, h_1 \uplus h_2, K_2).\end{array}$$

Hence,

$$\sigma_1 \circ \sigma_2' = ((\mathsf{pc}_1 \circ \mathsf{pc}_2)[t : (v', v_1)], (\theta_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)} \theta_2)[t : r'], \\ b_1 \circ_{(\mathsf{pc}_1 \circ \mathsf{pc}_2)[t:(v',v_1)]} b_2, h_1 \uplus h_2, K_2).$$

Then, $\sigma_1 \circ \sigma_2 \xrightarrow{\lambda''}_{C(L)} \sigma_1 \circ \sigma_2'$. The desired τ_{i+1} is obtained by extending τ_i with this transition.
– The transition in η is generated by LOCAL, LOCK, UNLOCK, WRITE or READ, and we have $K_1 \subseteq K_2$. This case is similar to the previous one.
– The transition in ξ is generated by FLUSH, and we have $K_2 \subseteq K_1$. In this case, by the definition of $\circ$, K_1 should be CPUid. Let $\lambda'' = (t, \mathsf{flush}(\beta))$ for $\beta \notin \{\mathsf{call}, \mathsf{ret}\}$. Let α_1 and $\alpha_2 \langle\beta\rangle$ be the store buffers of t in σ_1 and σ_2, respectively. We show that the FLUSH rule is enabled in $\sigma_1 \circ \sigma_2$. Assume this is not the case. Then the store buffer of t in $\sigma_1 \circ \sigma_2$ is $\alpha_1 \circ_\rho (\alpha_2 \langle\beta\rangle) = \alpha' \langle\beta\rangle \alpha''$, where $\alpha'' \neq \varepsilon$ and ρ is the program position of the CPU executing the FLUSH transition in σ_2. But this contradicts the definition of $\circ$: for $\alpha_1 \circ_\rho (\alpha_2 \langle\beta\rangle)$ to be defined, the library buffer would have to contain a ret marker to separate the library entry β and client entries in α''. Then, we can obtain ξ_{i+1} and τ_{i+1} by extending ξ_i and τ_i with the FLUSH transition by the CPU t.
– The transition in η is generated by FLUSH, and we have $K_1 \subseteq K_2$. This case is similar to the previous one.
– The transitions in η and ξ are generated by FLUSH-MARKER flushing the same marker. Note that in this case, $K_1 = K_2 = \mathsf{CPUid}$ by our operational semantics. Let $\alpha_1\beta$ and $\alpha_2\beta$ be store buffers of t in σ_1 and σ_2, respectively, where β is the marker to be flushed. Then β is the oldest entry in $\alpha_1\beta \circ_\_ \alpha_2\beta$. Hence, we can extend τ_i, η_i and ξ_i by the FLUSH-MARKER transition.
– The transition in ξ is generated by XLOCK, and we have $K_2 \subseteq K_1$. By the definition of $\circ$, $K_1 = \mathsf{CPUid}$ in this case. Let t be the CPU executing the XLOCK transition. Then the store buffer of t in σ_2 is empty. Let α_1 be the store buffer of t in σ_1, then its store buffer in $\sigma_1 \circ \sigma_2$ is also α_1. The program position of t in σ_2 is inside the library code; thus, by the definition of $\circ$ on store buffers, $\alpha_1 = \varepsilon$. Hence,

the store buffer of t is empty in $\sigma_1 \circ \sigma_2$, and so XLOCK is enabled there. We can then obtain ξ_{i+1} and τ_{i+1} by extending ξ_i and τ_i with the XLOCK transition.
- The transition in η is generated by XLOCK, and we have $K_1 \subseteq K_2$. The construction in this case is analogous to the previous one.
- The transition in ξ is generated by XUNLOCK, and we have $K_2 \subseteq K_1$. Let t be the CPU executing the XUNLOCK transition. Then the store buffer of t in σ_2 is of the form $(x_1, u_1) \dots (x_l, u_l)\,\mathsf{lock}$. The program position of t in σ_2 is inside the library code; thus, by the definition of $\circ$ on store buffers, the store buffer of t in σ_1 has to be empty. Hence, the buffer of t in $\sigma_1 \circ \sigma_2$ is the same as its buffer in σ_2, and we can obtain ξ_{i+1} and τ_{i+1} by extending ξ_i and τ_i with the XLOCK transition.
- The transition in η is generated by XUNLOCK, and we have $K_1 \subseteq K_2$. This case is similar to the previous one.

One of the above cases is always applicable. When both K_1 and K_2 are CPUid, the uncovered cases are those when both transitions in η and ξ are obtained using CALL, RET or FLUSH-MARKER and the actions produced are different. However, this case is impossible, since $\mathsf{history}(\eta) = \mathsf{history}(\xi)$ and $\mathsf{history}(\eta_i) = \mathsf{history}(\xi_i)$. When one of K_1 and K_2 is not CPUid, our construction covers all the cases: the transition of the local computation whose current configuration has the machine locked is always enabled in the corresponding configuration of the global computation.

Consider now the case when only one transition in (7) exists. Due to our preprocessing step, ξ is either empty or ends with a history action. Besides, the above construction consumes the latter together with the corresponding transition in η. Hence, the only transition in (7) has to be one from η, and the last configuration in ξ does not have the machine locked. This implies that the transition in η is enabled in the global configuration, and we can obtain η_{i+1} and τ_{i+1} by extending η_i and τ_i with this transition.

Thus, we have shown how to construct τ_i for all possible cases, sometimes nondeterministically. If ξ and η are both finite, our construction consumes both computations completely. Additionally, the construction consumes HAct transitions from η and ξ together. Thus, if ξ and η both have infinite histories, then our construction consumes also both computations. As we argued above, if one trace is shorter than the other one, then ξ has to be the shorter one. Thus, the only remaining case is when η is infinite and ξ is finite and either is empty or ends with a history action. Since history transitions in η and ξ are consumed together, ξ is consumed completely. As we argued above, after this, all transitions in η are reproducible in the global configuration, so η is consumed completely as well. Thus, in the limit we will exhaust all transitions from η and ξ, and the limits of η_i and ξ_i will be η and ξ. From this fact, it follows that $\mathsf{client}(\tau) = \mathsf{client}(\eta)$, as desired. □

A.5 Proof of Corollary 5

To simplify presentation, in our development we have assumed that any method called in a program by the client belongs to the library. The proof of Theorem 4 can be simply generalised to omit this requirement. Namely, we assume a setting where the client is allowed to have its private methods (as before, nested method calls are disallowed). We still insert call and ret markers into the store buffer when calling such methods; however, now we annotate them with the method called. The set of actions stays the same.

Annotating call and ret markers allows us to define generalisations client_M, lib_M and $\mathsf{history}_M$ of operations on computations defined earlier: the new operations interpret the given set of methods M as constituting a library. Theorem 4 still holds in this setting: the effect of non-library calls and returns on the store buffer is analogous to the one of write commands. We now prove Corollary 5 using this generalisation.

The idea of the proof is simple: to show $L \sqsubseteq L^\sharp$ we linearize libraries $L_1, \ldots, L_k$ one by one using Theorem 4. Let $I = I_1 \circ \ldots \circ I_k$. The program $\mathsf{MGC}(L) = \mathsf{MGC}(L_1, \ldots, L_k)$ can be viewed as consisting of the library L_1 and its client including the implementations of methods in $L_2, \ldots, L_k$. Since $L_1 \sqsubseteq L_1^\sharp$, by Theorem 4, we can linearize L_1, obtaining

$$\mathsf{client}_{\mathsf{sig}(L_1)}([\![\mathsf{MGC}(L_1, L_2, L_3, \ldots, L_k)]\!]I) \subseteq \mathsf{client}_{\mathsf{sig}(L_1)}([\![\mathsf{MGC}(L_1^\sharp, L_2, L_3, \ldots, L_k)]\!]I). \quad (8)$$

The resulting program $\mathsf{MGC}(L_1^\sharp, L_2, L_3, \ldots, L_k)$ can be viewed as consisting of the library implementing the methods from $L_1^\sharp$ and L_2 and the client including the methods from $L_3, \ldots, L_k$. Additionally in the program $\mathsf{MGC}(L_1^\sharp, L_2)$ we can view L_2 as the library and $L_1^\sharp$ as the client. Since $L_2 \sqsubseteq L_2^\sharp$, by Theorem 4, we get

$$\mathsf{client}_{\mathsf{sig}(L_2)}([\![\mathsf{MGC}(L_1^\sharp, L_2)]\!](I_1 \circ I_2)) \subseteq \mathsf{client}_{\mathsf{sig}(L_2)}([\![\mathsf{MGC}(L_1^\sharp, L_2^\sharp)]\!](I_1 \circ I_2)),$$

which implies

$$\mathsf{history}_{\mathsf{sig}(L_1,L_2)}([\![\mathsf{MGC}(L_1^\sharp, L_2)]\!](I_1 \circ I_2)) \subseteq \mathsf{history}_{\mathsf{sig}(L_1,L_2)}([\![\mathsf{MGC}(L_1^\sharp, L_2^\sharp)]\!](I_1 \circ I_2)).$$

Hence, $(L_1^\sharp, L_2) \sqsubseteq (L_1^\sharp, L_2^\sharp)$, where $(L_1^\sharp, L_2)$ and $(L_1^\sharp, L_2^\sharp)$ are the compositions of libraries $L_1^\sharp$ and L_2, respectively, $L_1^\sharp$ and $L_2^\sharp$. Applying Theorem 4 with this linearization to the program $\mathsf{MGC}(L_1^\sharp, L_2, L_3, \ldots, L_k)$, we obtain

$$\mathsf{client}_{\mathsf{sig}(L_1,L_2)}([\![\mathsf{MGC}(L_1^\sharp, L_2, L_3, \ldots, L_k)]\!]I) \subseteq \mathsf{client}_{\mathsf{sig}(L_1,L_2)}([\![\mathsf{MGC}(L_1^\sharp, L_2^\sharp, L_3, \ldots, L_k)]\!]I). \quad (9)$$

From (8), we get

$$\mathsf{client}_{\mathsf{sig}(L_1,L_2)}([\![\mathsf{MGC}(L_1, L_2, L_3, \ldots, L_k)]\!]I) \subseteq \mathsf{client}_{\mathsf{sig}(L_1,L_2)}([\![\mathsf{MGC}(L_1^\sharp, L_2, L_3, \ldots, L_k)]\!]I).$$

From this and (9), we get

$$\mathsf{client}_{\mathsf{sig}(L_1,L_2)}([\![\mathsf{MGC}(L_1, L_2, L_3, \ldots, L_k)]\!]I) \subseteq \mathsf{client}_{\mathsf{sig}(L_1,L_2)}([\![\mathsf{MGC}(L_1^\sharp, L_2^\sharp, L_3, \ldots, L_k)]\!]I).$$

```
word x1 = 0, x2 = 0;
word c = 0;

write(in word d1, in word d2) {
v0: c++;
v1: x1 = d1;
v2: x2 = d2;
v3: c++;
v4:
}
```

```
read(out word d1, out word d2) {
word c0;
v5 : c0 = c;
v6 : if (nondet()) goto v9;
v7 : assume(c0 % 2);
v8 : goto v5:
v9 : assume(!(c0 % 2));
v10: d1 = x1;
v11: d2 = x2;
v12: if (nondet()) goto v15;
v13: assume(c != c0);
v14: goto v5:
v15: assume(c == c0);
v16:
}
```

Fig. 6. Seqlock implementation L_{seqlock} with code labels

Repeatedly applying Theorem 4 as above to linearize $L_3, \ldots, L_k$, we get

$$\mathsf{client}_{\mathsf{sig}(L_1,L_2,L_3,\ldots,L_k)}([\![\mathsf{MGC}(L_1, L_2, L_3, \ldots, L_k)]\!]I) \subseteq \mathsf{client}_{\mathsf{sig}(L_1,L_2,L_3,\ldots,L_k)}([\![\mathsf{MGC}(L_1^\sharp, L_2^\sharp, L_3^\sharp, \ldots, L_k^\sharp)]\!]I),$$

which implies $L \sqsubseteq L^\sharp$.

A.6 Proof of Theorem 3

To ease the exposition, Figures 6 and 7 show the concrete and abstract implementations of a seqlock from Figures 3 and 4 including code labels and with loops and conditionals translated into non-deterministic branching and assume commands. The result is close to the CFG representation of programs we defined in Section 2.

We consider the most general clients $\mathsf{MGC}(L_{\mathsf{seqlock}})$ of the concrete and $\mathsf{MGC}(L^\sharp_{\mathsf{seqlock}})$ of the abstract libraries, where only one CPU is a writer and all the others are readers. We assume that in both programs, CPU 0 runs the code of the writer. Thus, each most general client is described by two kinds of control flow graphs, one for the writer and another one for readers. Let v^t_{rc} and v^t_{ra}, $t \in \mathsf{CPUid} - \{0\}$, be the start nodes of the CFGs for readers in $\mathsf{MGC}(L_{\mathsf{seqlock}})$ and $\mathsf{MGC}(L^\sharp_{\mathsf{seqlock}})$. Also, let v_{wc} and v_{wa} be the start nodes of the CFGs for writers in these most general clients. Let $I = \{[\mathtt{x1} : 0, \mathtt{x2} : 0, \mathsf{c} : 0]\}$. Consider $\tau \in [\![\mathsf{MGC}(L_{\mathsf{seqlock}})]\!]I$, starting from a configuration $\sigma_0 = (\mathsf{pc}_0, \theta_0, b_0, h_0, \mathsf{CPUid})$, such that

$$(\forall t \in \mathsf{CPUid} - \{0\}.\, \mathsf{pc}_0(t) = v^t_{\mathsf{rc}}) \wedge \mathsf{pc}_0(0) = v_{\mathsf{wc}} \wedge (\forall t \in \mathsf{CPUid}.\, b_0(t) = \varepsilon) \wedge h_0 = [\mathtt{x1} : 0, \mathtt{x2} : 0, \mathsf{c} : 0].$$

```
word x1 = 0, x2 = 0;
word c;

write(in word d1, in word d2) {
w0: lock; x1 = d1; x2 = d2; unlock;
w1:
}

read(out word d1, out word d2) {
w2: lock; d1 = x1; d2 = x2; unlock;
w3:
}
```

Fig. 7. Seqlock specification $L^\sharp_{\mathsf{seqlock}}$ with code labels

We construct a corresponding execution $\tau' \in [\![\mathsf{MGC}(L^\sharp_{\mathsf{seqlock}})]\!] I$ starting from $\sigma'_0 = (\mathsf{pc}'_0, \theta_0, b_0, h_0, \mathsf{CPUid})$, where

$$(\forall t \in \mathsf{CPUid} - \{0\}.\, \mathsf{pc}'_0(t) = v^t_{\mathsf{ra}}) \wedge \mathsf{pc}'_0(0) = v_{\mathsf{wa}}.$$

This is done by induction on the length of τ: for every transition in τ, we construct zero or more transitions of $\mathsf{MGC}(L^\sharp_{\mathsf{seqlock}})$. During this construction, we maintain an invariant $\sim \subseteq \mathsf{Config}_1 \times (\mathsf{CPUid} \to \mathbb{B}) \times \mathsf{Config}_2$ relating the states of the two libraries. Here Config_1 and Config_2 are sets of configurations arising in computations of $\mathsf{MGC}(L_{\mathsf{seqlock}})$ and $\mathsf{MGC}(L^\sharp_{\mathsf{seqlock}})$. The extra parameter $p \in \mathsf{CPUid} \to \mathbb{B}$ is a prophecy variable, stating a fact about the future computation of $\mathsf{MGC}(L_{\mathsf{seqlock}})$: if in $\mathsf{MGC}(L_{\mathsf{seqlock}})$ CPU t is executing the `read` method, then $p(t)$ is true when the method is performing the last iteration of the outer loop of `read`.

We define the relation $\sim$ as follows:

$$(\mathsf{pc}_1, \theta_1, b_1, h_1, K_1) \sim_p (\mathsf{pc}_2, \theta_2, b_2, h_2, K_2) \Leftrightarrow$$
$$\mathsf{pc}_1 \sim_p \mathsf{pc}_2 \wedge \theta_1 \sim_{p,\mathsf{pc}_1,\mathsf{pc}_2,h_1} \theta_2 \wedge b_1 \sim_{\theta_1(0),\mathsf{pc}_1(0),h_1} b_2 \wedge h_1 \sim h_2 \wedge K_1 \sim K_2,$$

where the relations on parts of configurations are defined below.

We let $\mathsf{pc}_1 \sim_p \mathsf{pc}_2 \Leftrightarrow \forall t \in \mathsf{CPUid}.\, \mathsf{pc}_1(t) \sim_{t,p(t)} \mathsf{pc}_2(t)$, where for all $t \in \mathsf{CPUid} - \{0\}$ we have $v^t_{\mathsf{rc}} \sim_{t,p(t)} v^t_{\mathsf{ra}}$, $v_{\mathsf{wc}} \sim_{t,p(t)} v_{\mathsf{wa}}$ and

$$\begin{array}{ll}
(v, v_{\mathsf{wc}}) \sim_{0,p(t)} (\mathtt{w0}, v_{\mathsf{wa}},), & \text{if } v \in \{\mathtt{v0}, \ldots, \mathtt{v3}\};\\
(\mathtt{v4}, v_{\mathsf{wc}}) \sim_{0,p(t)} (\mathtt{w1}, v_{\mathsf{wa}}); & \\
(v, v^t_{\mathsf{rc}}) \sim_{t,p(t)} (\mathtt{w2}, v^t_{\mathsf{ra}}), & \text{if } v \in \{\mathtt{v5}, \ldots, \mathtt{v11}, \mathtt{v13}, \mathtt{v14}\};\\
(v, v^t_{\mathsf{rc}}) \sim_{t,p(t)} (\mathtt{w3}, v^t_{\mathsf{ra}}), & \text{if } v \in \{\mathtt{v15}, \mathtt{v16}\};\\
(\mathtt{v12}, v^t_{\mathsf{rc}}) \sim_{t,p(t)} (\mathtt{w2}, v^t_{\mathsf{ra}}), & \text{if } \neg p(t);\\
(\mathtt{v12}, v^t_{\mathsf{rc}}) \sim_{t,p(t)} (\mathtt{w3}, v^t_{\mathsf{ra}}), & \text{if } p(t).
\end{array}$$

To define $\sim$ on store buffers, we first define the following function

$$g : \mathsf{RegBank} \times \mathsf{Pos} \times \mathsf{Heap} \times ((\mathsf{Loc} \times \mathsf{Val})^+ \cup \{\mathsf{call}, \mathsf{ret}\})^* \rightharpoonup ((\mathsf{Loc} \times \mathsf{Val})^+ \cup \{\mathsf{call}, \mathsf{ret}\})^*$$

$$
\begin{aligned}
g(r,\rho,h_1,\varepsilon) &= \varepsilon,\\
&\quad h_1(\mathtt{c}) \text{ is even, } \rho \in \{(\mathtt{v4}, v_{\mathsf{wc}}), (\mathtt{v0}, v_{\mathsf{wc}}), v_{\mathsf{wc}}\}\\
g(r,\rho,h_1,\varepsilon) &= \varepsilon,\\
&\quad h_1(\mathtt{c}) \text{ is odd, } \rho \in \{(\mathtt{v1}, v_{\mathsf{wc}}), (\mathtt{v2}, v_{\mathsf{wc}}), (\mathtt{v3}, v_{\mathsf{wc}})\},\\
(\rho \in \{(\mathtt{v2}, v_{\mathsf{wc}}), (\mathtt{v3}, v_{\mathsf{wc}})\} &\Rightarrow h_1(\mathtt{x1}) = r(\mathtt{d1})),\ \rho = (\mathtt{v3}, v_{\mathsf{wc}}) \Rightarrow h_2(\mathtt{x2}) = r(\mathtt{d2})\\
g(r,\rho,h_1,\mathsf{call}\,\alpha) &= \mathsf{call}\, g(r,\rho,h_1,\alpha),\\
g(r,\rho,h_1,\mathsf{ret}\,\alpha) &= \mathsf{ret}\, g(r,\rho,h_1,\alpha),\\
g(r,\rho,h_1,(\mathtt{c},c)\,\alpha) &= g(r,\rho,h_1,\alpha),\\
&\quad c \text{ is odd, } h_1(\mathtt{c}) = c-1,\ \rho = (\mathtt{v1}, v_{\mathsf{wc}})\\
g(r,\rho,h_1,(\mathtt{x1}, r(\mathtt{d1}))(\mathtt{c},c)\,\alpha) &= g(r,\rho,h_1,\alpha),\\
&\quad c \text{ is odd, } h_1(\mathtt{c}) = c-1,\ \rho = (\mathtt{v2}, v_{\mathsf{wc}})\\
g(r,\rho,h_1,(\mathtt{x2}, r(\mathtt{d2}))(\mathtt{x1}, r(\mathtt{d1}))(\mathtt{c},c)\,\alpha) &= g(r,\rho,h_1,\alpha),\\
&\quad c \text{ is odd, } h_1(\mathtt{c}) = c-1,\ \rho = (\mathtt{v3}, v_{\mathsf{wc}})\\
g(r,\rho,h_1,(\mathtt{c},c+1)(\mathtt{x2},x_2)(\mathtt{x1},x_1)(\mathtt{c},c)\,\alpha) &= \langle(\mathtt{x2},x_2)(\mathtt{x1},x_1)\rangle\, g(r,\rho,h_1,\alpha), \quad c \text{ is odd}\\
g(r,\rho,h_1,(\mathtt{c},c')(\mathtt{x2},x_2)(\mathtt{x1},x_1)) &= \langle(\mathtt{x2},x_2)(\mathtt{x1},x_1)\rangle,\\
&\quad c' \text{ is even, } h_1(\mathtt{c}) = c'-1\\
g(r,\rho,h_1,(\mathtt{c},c')(\mathtt{x2},x_2)) &= \langle(\mathtt{x2},x_2)(\mathtt{x1},h_1(\mathtt{x1}))\rangle,\\
&\quad c' \text{ is even, } h_1(\mathtt{c}) = c'-1\\
g(r,\rho,h_1,(\mathtt{x2}, r(\mathtt{d2}))(\mathtt{x1}, r(\mathtt{d1}))) &= \varepsilon,\\
&\quad h_1(\mathtt{c}) \text{ is odd, } \rho = (\mathtt{v3}, v_{\mathsf{wc}})\\
g(r,\rho,h_1,(\mathtt{c},c')) &= \langle(\mathtt{x2},h_1(\mathtt{x2}))(\mathtt{x1},h_1(\mathtt{x1}))\rangle,\\
&\quad c' \text{ is even, } h_1(\mathtt{c}) = c'-1\\
g(r,\rho,h_1,(\mathtt{x2}, r(\mathtt{d2}))) &= \varepsilon,\\
&\quad h_1(\mathtt{c}) \text{ is odd, } h_1(\mathtt{x1}) = r(\mathtt{d1}),\ \rho = (\mathtt{v3}, v_{\mathsf{wc}})\\
g(r,\rho,h_1,(\mathtt{x1}, r(\mathtt{d1}))) &= \varepsilon,\\
&\quad h_1(\mathtt{c}) \text{ is odd, } \rho = (\mathtt{v2}, v_{\mathsf{wc}})
\end{aligned}
$$

Fig. 8. An auxiliary function on store buffers

converting the store buffers of the writer CPU in $\mathsf{MGC}(L_{\mathsf{seqlock}})$ to the corresponding one in $\mathsf{MGC}(L^{\sharp}_{\mathsf{seqlock}})$; see Figure 8. We then let

$$
\begin{aligned}
b_1 \sim_{r,\rho,h_1} b_2 \Leftrightarrow\ & b_2(0) = g(r,\rho,h_1,b_1(0)) \wedge (\forall t \in \mathsf{CPUid} - \{0\}.\, b_1(t) = b_2(t) = \varepsilon) \wedge\\
& (\exists \alpha, c'.\, b_1(0) = (\mathtt{c},c')\,\alpha \wedge (c' \text{ is even}) \Rightarrow \rho = (\mathtt{v4}, v_{\mathsf{wc}})) \wedge\\
& (\exists \alpha.\, b_1(0) = \mathsf{call}\,\alpha \Rightarrow \rho = (\mathtt{v0}, v_{\mathsf{wc}})) \wedge\\
& (\exists \alpha.\, b_1(0) = \mathsf{ret}\,\alpha \Rightarrow \rho = v_{\mathsf{wc}}) \wedge\\
& (\exists \alpha, \beta, c.\, b_1(0) = \beta\,(\mathtt{c},c)\,\alpha \wedge (\alpha \text{ does not contain entries for } \mathtt{c}) \Rightarrow h_1(\mathtt{c}) = c-1).
\end{aligned}
$$

We let

$$
\begin{aligned}
\theta_1 \sim_{p,\mathsf{pc}_1,\mathsf{pc}_2,h_1} \theta_2 \Leftrightarrow\ & \theta_1(0) = \theta_2(0) \wedge\\
& \forall t \in \mathsf{CPUid} - \{0\}.\, (\mathsf{pc}_2(t) = \mathtt{w3} \Rightarrow \exists c.\, \theta_1(t) = \theta_2(t)[\mathtt{c0} : c]) \wedge\\
& (\exists v.\, \mathsf{pc}_1(t) \in (v, v^t_{\mathsf{rc}}) \wedge v \in \{\mathtt{v6}, \ldots, \mathtt{v10}\} \wedge p(t) \Rightarrow\\
& \theta_1(t)(\mathtt{c0}) = h_1(\mathtt{c}) \wedge (c \text{ is even})) \wedge\\
& (\mathsf{pc}_1(t) = (\mathtt{v11}, v^t_{\mathsf{rc}}) \wedge p(t) \Rightarrow\\
& \theta_1(t)(\mathtt{d1}) = h_1(\mathtt{x1}) \wedge \theta_1(t)(\mathtt{c0}) = h_1(\mathtt{c}) \wedge (c \text{ is even})).
\end{aligned}
$$

We define $h_1 \sim h_2$ as follows:

$$[\mathtt{x1} : x_1, \mathtt{x2} : x_2, \mathsf{c} : c] \sim [\mathtt{x1} : x_1, \mathtt{x2} : x_2, \mathsf{c} : \_], \quad \text{if } c \text{ is even;}$$
$$[\mathtt{x1} : x_1, \mathtt{x2} : x_2, \mathsf{c} : c] \sim [\mathtt{x1} : \_, \mathtt{x2} : \_, \mathsf{c} : \_], \quad \text{if } c \text{ is odd.}$$

Finally, we let $K_1 \sim K_2 \Leftrightarrow K_1 = K_2 = \mathsf{CPUid}$.

The following lemma describes the induction step in constructing τ'.

Lemma 13. *Consider a transition λ_1 inside τ:*

$$\sigma_0 \xrightarrow{\lambda_0}{}^*_{\mathsf{MGC}(L_{\mathsf{seqlock}})} \sigma_1 \xrightarrow{\lambda_1}_{\mathsf{MGC}(L_{\mathsf{seqlock}})} \sigma_2 \ldots$$

Assume the computation

$$\sigma_0' \xrightarrow{\lambda_0'}{}^*_{\mathsf{MGC}(L^\sharp_{\mathsf{seqlock}})} \sigma_1'$$

such that $\sigma_0 \sim_{\lambda t.\mathsf{false}} \sigma_0'$, $\sigma_1 \sim_{p_1} \sigma_1'$, $\mathsf{history}(\lambda_0) = \mathsf{history}(\lambda_0')$ and $p_1(t)$ is true if CPU t is at a program position (v, v^t_{rc}) for $v \in \{\mathtt{v6}, \ldots, \mathtt{v12}\}$ in σ_1, and t either reaches $(\mathtt{v16}, v^t_{\mathsf{rc}})$ in the rest of the computation τ without going via $(\mathtt{v13}, v^t_{\mathsf{rc}})$ before that, or the computation ends without t going via $(\mathtt{v13}, v^t_{\mathsf{rc}})$.

Then there exists $\sigma_2' \in \mathsf{Config} - \{\top\}$ and a computation

$$\sigma_1' \xrightarrow{\lambda_1'}{}^*_{\mathsf{MGC}(L^\sharp_{\mathsf{seqlock}})} \sigma_2'$$

such that $\sigma_2 \sim_{p_2} \sigma_2'$, $\mathsf{history}(\lambda_0\lambda_1) = \mathsf{history}(\lambda_0'\lambda_1')$ and $p_2(t)$ is true if CPU t is at the program position (v, v^t_{rc}) for $v \in \{\mathtt{v6}, \ldots, \mathtt{v12}\}$ in σ_2 and t reaches $(\mathtt{v16}, v^t_{\mathsf{rc}})$ in the rest of the computation τ without going via $(\mathtt{v13}, v^t_{\mathsf{rc}})$ before that, or the computation ends without t going via $(\mathtt{v13}, v^t_{\mathsf{rc}})$.

From the lemma we immediately get that there exists $\tau' \in [\![\mathsf{MGC}(L^\sharp_{\mathsf{seqlock}})]\!]I$ such that $\mathsf{history}(\tau) = \mathsf{history}(\tau')$, which implies $L_{\mathsf{seqlock}} \sqsubseteq L^\sharp_{\mathsf{seqlock}}$.

Proof sketch for Lemma 13. For each transition of the seqlock implementation, we show how the abstract implementation can make zero or more steps such that the $\sim$ relation between the end configurations is preserved. We now list the transitions of the concrete implementation for which the abstract implementation makes one or more steps.

- A concrete havoc transition is matched by an abstract havoc transition assigning the same values to registers.
- A concrete CALL, RET or FLUSH-MARKER transition is matched by the same abstract transition.
- The command at v3 is matched by the sequence of commands at w0.
- The command by CPU t at v11 is matched by the sequence of commands at w2 when $p(t)$ is true.
- A FLUSH of an even value to c from a concrete store buffer is matched by flushing the oldest entry in the corresponding abstract store buffer.

It is easy to check that the $\sim$ relation ensures that all the above transitions (plus the transitions of the concrete system that do not cause the abstract implementation to make a step) are executable and preserve the $\sim$ relation. □

B Additional examples

In the following examples we assume that integers are unbounded (or, equivalently, consider only executions where an overflow does not occur).

1. Spinlock.

```
word x = 1;

acquire() {
  while (1) {
    xlock;
    x--;
    if (x >= 0) {
      xunlock;
      return;
    }
    xunlock;
    while (x <= 0) ;
  }
}

release() { x = 1; }
```

Note that `release` writes `1` to `x` without executing a memory barrier. On TSO this can result in an additional delay before the write releasing the lock becomes visible to another CPU trying to acquire it. The abstract implementation is as follows:

```
word x = 1;

acquire() {
  xlock;
  assume(x == 1);
  x = 0;
  xunlock;
}

release() { x = 1; }
```

Here the write to `x` in the `release` can be delayed in the store buffer as well. It is easy to see that the resulting specification still ensures mutual exclusion.

2. Ticketed spinlock.

```
word x = 1, y = 1;

acquire() {
  word ticket;
```

```
  xlock;
  ticket = y++;
  xunlock;
  while (x != ticket) ;
}

release () { x++; }
```

Unlike the previous spinlock implementation, this one ensures fairness using a variant of the Bakery algorithm. We give it the same specification as the previous implementation.

3. Initialisation using double-checked locking. The function `ensureinit` initialises an object. We assume that several copies of the function can be run concurrently, while the initialisation is meant to be performed only once.

```
word x = UNINITIALISED;

ensureinit() {
  if (x == INITIALISED) return;
  acquire();
  if (x == UNINITIALISED) {
    // Initialise the object...
    x = INITIALISED;
  }
  release();
}
```

The implementation assumes a lock with the specification given in Example 1 above. The first read of `x` in `ensureinit` can read `UNINITIALISED` when there is a pending write of `INITIALISED` to `x` by another CPU. For this reason, the implementation can exhibit non-SC behaviours [13]. Its abstract implementation is as follows:

```
word x = UNINITIALISED;

ensureinit() {
  word flag = nondet();
  if (flag) xlock; else lock;
  if (x == UNINITIALISED) {
    // Initialise the object...
    x = INITIALISED;
  }
  if (flag) xunlock; else unlock;
}
```

The abstract implementation non-deterministically flushes the store buffer, which can happen in the concrete implementation as a side effect of `acquire`.