

Higher-order linearisability

Andrzej S. Murawski^{a,1}, Nikos Tzevelekos^{b,*,2}^a University of Oxford, United Kingdom^b Queen Mary University of London, United Kingdom

ARTICLE INFO

Article history:

Received 11 December 2017

Received in revised form 26 November 2018

Accepted 11 January 2019

Available online 21 January 2019

Keywords:

Linearisability

Concurrency

Higher-order computation

ABSTRACT

Linearisability is a central notion for verifying concurrent libraries: a library is proven correct if its operational history can be rearranged into a sequential one that satisfies a given specification. Until now, linearisability has been examined for libraries in which method arguments and method results were of ground type. In this paper we extend linearisability to the general higher-order setting, where methods of arbitrary type can be passed as arguments and returned as values, and establish its soundness.

© 2019 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Software libraries provide implementations of routines, often of specialised nature, to facilitate code reuse and modularity. To support the latter, they should follow specifications that describe the range of acceptable behaviours for correct and safe deployment. Adherence to specifications can be formalised using the classic notion of contextual approximation (refinement), which scrutinises the behaviour of code in any possible context. Unfortunately, the quantification makes it difficult to prove contextual approximation directly, which motivates research into sound techniques for establishing it.

In the concurrent setting, a notion that has been particularly influential is that of *linearisability* [1]. Linearisability requires that, for each history generated by a library, one should be able to find another history from the specification (a *linearisation*), which matches the former up to certain rearrangements of events. In the original formulation by Herlihy and Wing [1], these permutations were not allowed to disturb the order between library returns and client calls. Moreover, linearisations were required to be *sequential* traces, that is, sequences of method calls immediately followed by their returns.

In this paper we shall work with *open higher-order* libraries, which provide implementations of *public* methods and may themselves depend on *abstract* ones, to be supplied by parameter libraries. The classic notion of linearisability only applies to closed libraries (without abstract methods). Additionally, both method arguments and results had to be of *ground* type. The closedness limitation was recently lifted in [2,3], which distinguished between public (or *implemented*) and abstract methods (*callable*). Although [2] did not in principle exclude higher-order functions, those works focused on linearisability for the case where the allowable methods were restricted to first-order functions ($\text{int} \rightarrow \text{int}$). Herein, we give a systematic exposition of linearisability for general higher-order concurrent libraries, where methods can be of arbitrary higher-order types. In doing so, we also propose a corresponding notion of sequential history for higher-order library interactions.

* Corresponding author.

E-mail address: nikos.tzevelekos@qmul.ac.uk (N. Tzevelekos).¹ Supported by a Royal Society Leverhulme Trust Senior Research Fellowship.² Research supported by EPSRC (EP/P004172/1).




Fig. 1. A library $L : \Psi \rightarrow \Psi'$ in environment comprising a parameter library $L' : \emptyset \rightarrow \Psi, \Psi''$ and a client K of the form $\Psi', \Psi'' \vdash M_1 \parallel \dots \parallel M_N$.

We examine libraries L that can interact with their environments by means of public and abstract methods: a library L with abstract methods of types $\Psi = \theta_1, \dots, \theta_n$ and public methods $\Psi' = \theta'_1, \dots, \theta'_n$ is written as $L : \Psi \rightarrow \Psi'$. We shall work with arbitrary higher-order types generated from the ground types `unit` and `int`. Types in Ψ, Ψ' must always be function types, i.e. their order is at least 1.

A library L may be used in computations by placing it in a context that will keep on calling its public methods (via a client K) as well as providing implementations for the abstract ones (via a parameter library L'). The setting is depicted in Fig. 1. Note that, as the library L interacts with K and L' , they exchange functions between each other. Consequently, in addition to K making calls to public methods of L and L making calls to its abstract methods, K and L' may also issue calls to functions that were passed to them as arguments during higher-order interactions. Analogously, L may call functions that were communicated to it via library calls.

Our framework is operational in flavour and draws upon concurrent [4,5] and operational game semantics [6–8]. We shall model library use as a game between two participants: *Player* (P), corresponding to the library L , and *Opponent* (O), representing the environment (L', K) in which the library was deployed. Each call will be of the form $\text{call } m(v)$ with the corresponding return of the shape $\text{ret } m(v)$, where v is a value. As we work in a higher-order framework, v may contain functions, which can participate in subsequent calls and returns. Histories will be sequences of *moves*, which are calls and returns paired with thread identifiers. A history is sequential just if every move produced by O is immediately followed by a move by P in the same thread. In other words, the library immediately responds to each call or return delivered by the environment. In contrast to classic linearisability, the move by O and its response by P need not be a call/return pair, as the higher-order setting provides more possibilities (in particular, the P response may well be a call). Accordingly, linearisable higher-order histories can be seen as sequences of atomic segments (linearisation points), starting at environment moves and ending with corresponding library moves.

In the spirit of [3], we are going to consider two scenarios: one in which K and L' share an explicit communication channel (the general case) as well as a situation in which they can only communicate through the library (the encapsulated case). Further, we also handle the case in which extra closure assumptions can be made about the parameter library (the relational case), which can be useful for dealing with a variety of assumptions on the use of parameter libraries that may arise in practice. In each case, we present a candidate definition of linearisability and illustrate it with tailored examples. The suitability of each kind of linearisability is demonstrated by showing that it implies the relevant form of contextual approximation (refinement). We also examine compositionality of the proposed concepts. One of our examples will discuss the implementation of the flat-combining approach [9,3], adapted to higher-order types.

The paper is an extended version of [10] and contains complete proofs, fully elaborated examples and appendices with further technical material, e.g. on compositionality.

1.1. Example: a higher-order multiset library

Higher-order libraries are common in languages like ML, Java, Python, etc. As an illustrative example, we consider a library written in ML-like syntax which implements a multiset data structure with integer elements. For simplicity, we assume that its signature contains just two methods:

`count` : `int` \rightarrow `int`, `update` : (`int` \times (`int` \rightarrow `int`)) \rightarrow `int`.

The former method returns for each integer its multiplicity in the multiset – this is 0 if the integer is not a member of the multiset. On the other hand, `update` takes as an argument an integer i and a function g , and updates the multiplicity j of i in the multiset to $|g(j)|$ (we use the absolute value of $g(j)$ in order to meet the multiset requirement that element multiplicities not be negative; alternatively, we could have used exceptions to quarantine such client method behaviour). Methods with the same functionalities can be found in the multiset module of the `ocaml-containers` library [11]. While our example is simple, the same kind of analysis as below can be applied to more intricate examples such as `map` methods for integer-valued arrays, maps or multisets.

Example 1 (Multiset). Consider the concurrent multiset library L_{msset} in Fig. 2 on the LHS (the RHS will be discussed only later). It uses a private reference for storing the multiset's characteristic function and reads *optimistically*, without locking (cf. [12,13]). The `update` method in particular reads the current multiplicity of the given element i (via `count`) and computes its new multiplicity without acquiring a lock on the characteristic function. It only acquires a lock when it is ready to write

<pre> 1 public count, update; 2 Lock lock; 3 F := λx.0; 4 5 count = λi. (!F)i 6 update = λ(i, g). aux(i,g,count i) 7 8 aux = λ(i, g, j). 9 let y = g j in 10 lock.acquire(); 11 let f = !F in 12 if (j == (f i)) then (13 F := λx. if (x == i) then y 14 else (f x) ; 15 lock.release(); 16 y) 17 else (18 lock.release(); 19 aux(i,g,f i)) </pre>	<pre> 1 public count, update, reset; 2 abstract default; 3 Lock lock; 4 F := λx.0; 5 count = λi. (!F)i 6 update = λ(i, g). aux(i,g,count i) 7 8 ... 9 ... 10 ... 19 ... 20 reset = λi. 21 lock.acquire(); 22 let y = default i in 23 let f = !F in 24 F := λx. if (x == i) then y 25 else (f x); 26 lock.release(); 27 y </pre>
---	---

Fig. 2. Left: Multiset library L_{mset} with public methods $\text{count} : \text{int} \rightarrow \text{int}$ and $\text{update} : \text{int} \times (\text{int} \rightarrow \text{int}) \rightarrow \text{int}$. Right: Parameterised multiset library $L_{\text{mset}2}$ (lines 8–19 as in LHS) with public methods count , $\text{reset} : \text{int} \rightarrow \text{int}$, $\text{update} : \text{int} \times (\text{int} \rightarrow \text{int}) \rightarrow \text{int}$; abstract method $\text{default} : \text{int} \rightarrow \text{int}$.




Fig. 3. Example histories of L_{mset} . (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

the new value (line 10) in the hope that the value at i will still be the same and the update can proceed; if not, another attempt to update the value is made.

Let us look at some example executions of the library via their resulting histories, i.e. sequences of method calls and returns between the library and a client. In the topmost block (a) of history diagrams of Fig. 3, we see three such executions. Note that we do not record internal calls to count or aux , and use m and variants for method identifiers (names). We use the abbreviation cnt for count , and upd for update , and initially ignore the circled events for cnt . Each execution involves 2 threads.

In the first execution, the client calls $\text{update}(i, m)$ in the second thread, and subsequently calls $\text{count}(i)$ in the first thread. The code for update stipulates that first $\text{count}(i)$ be called internally, returning some multiplicity j for i , and then $m(j)$ should be called. As soon m returns a value j' , update sets the multiplicity of i to j' and itself returns j' . The last event in this history is a return of count in the first thread with the old value j . According to our proposed definition, this history will be linearisable to another, intuitively correct one: the last return can be moved to the circled position. At this point the notion of linearisability is used informally, but it will be made precise in the following sections. In the second execution, the last return of count in the first thread returns the updated value. In this case, we will be able to move $\text{call cnt}(i)$ to the circled position to obtain a linearisation, which is obviously correct. Finally, in the third execution we have a history that

will turn out non-linearisable to an intuitively correct history. Indeed, we should not be able to return the updated value in the first thread before m has returned it in the second one.

The two histories in block (b) in the same figure demonstrate the mechanism for updates. The first history will be linearisable to the second one. In the second history we see that both threads try to update the same element i , but the first one succeeds in it first and returns k on *update*. Then, the second thread realises that the value of i has been updated to k and calls m again, this time with argument k . An important feature of the second history is that it is *sequential*: each client event (call or return) is immediately followed by a library event.

Observe that the rearrangements discussed above involve either advancing a library action or postponing an environment action and that each action could be a call or a return. Definition 9 will capture this formally. For now, we note that this generalises the classic setting [1], where library method returns could be advanced and environment method calls deferred.

The next section will introduce histories along with the proposed notion of linearisability. In Section 3 we present the syntax for libraries and clients, and in Section 3.1 we define their semantics in terms of histories and co-histories respectively.

2. Higher-order linearisability

We examine higher-order libraries interacting with their context by means of abstract and public methods. In particular, we shall rely on types given by the grammar below. We let Meths stand for the set of *method names* and assume $\text{Meths} = \biguplus_{\theta, \theta'} \text{Meths}_{\theta, \theta'}$, where each set $\text{Meths}_{\theta, \theta'}$ contains names for methods of type $\theta \rightarrow \theta'$. Methods are ranged over by m (and variants). We let v range over computational *values*, which include a unit value, integers, methods and pairs of values.

$$\theta ::= \text{unit} \mid \text{int} \mid \theta \times \theta \mid \theta \rightarrow \theta \quad v ::= () \mid i \mid m \mid (v, v)$$

The framework of a higher-order library and its environment is depicted in Fig. 1. Given $\Psi, \Psi' \subseteq \text{Meths}$, a library L is said to have type $\Psi \rightarrow \Psi'$ if it defines public methods with names (and types) as in Ψ' , using abstract methods Ψ . The environment of L consists of a *client* K (which invokes public methods of Ψ'), and a *parameter library* L' (which provides code for the abstract methods Ψ). In general, K and L' may interact via a disjoint set of methods $\Psi'' \subseteq \text{Meths}$, to which L has no access.

In the rest of this paper, we shall implicitly assume that we work with a library L operating in an environment presented in Fig. 1. The client K will consist of a fixed number N of concurrent threads. Next we introduce a notion of history tailored to the setting and define how histories can be linearised.

2.1. Higher-order histories

The operational semantics of libraries will be given in terms of *histories*, which are sequences of method calls and returns, each decorated with a thread identifier t and a *polarity index* X , where $X \in \{O, P\}$, as shown below.

$$(t, \text{call } m(v))_X \quad (t, \text{ret } m(v))_X$$

We shall refer such decorated calls and returns as **moves**. Here, m is a method name and v is a value of a matching type. The index X specifies who produces the move: the library L (polarity P), or its environment (L', K) (polarity O). Using notation e.g. from [3], P corresponds to $!$, and O to $?$. We may be dropping the polarity of a move when it is not important or no confusion arises by doing so.

The choice of indices is motivated by the fact that the moves can be seen as defining a 2-player game between the library (L), which represents the *Proponent* player in the game (P), and its environment (L', K) that represents the *Opponent* (O). Finally, we let the **dual** polarity of X be X' , where $X \neq X'$.

Next we proceed to define histories. Their definition will rely on a more primitive concept of *prehistories*, which are sequences of O/P -indexed method calls and returns that respect a stack discipline.

Definition 2. *Prehistories* are sequences generated by one of the grammars:

$$\begin{aligned} \text{PreH}_O &::= \epsilon \mid \text{call } m(v)_O \text{ PreH}_P \text{ ret } m(v')_P \text{ PreH}_O \\ \text{PreH}_P &::= \epsilon \mid \text{call } m(v)_P \text{ PreH}_O \text{ ret } m(v')_O \text{ PreH}_P \end{aligned}$$

where, if $m \in \text{Meths}_{\theta, \theta'}$, the types of v, v' must match θ, θ' respectively. We let $\text{PreH} = \text{PreH}_O \cup \text{PreH}_P$.

Thus, prehistories from PreH_O start with an O -call, while those in PreH_P start with a P -call. In each case, the polarities inside a prehistory alternate between O and P , and the polarities of calls and matching returns are always dual (*returns dual to calls*).

Histories will be interleavings of prehistories tagged with thread identifiers (natural numbers), subject to a set of well-formedness constrains. In particular, a history h for library $L : \Psi \rightarrow \Psi'$ will have to begin with an O -move and satisfy the following conditions, to be formalised in Definition 3:

1. The name of any method called in h must come from Ψ or Ψ' , or be introduced earlier in h as a higher-order argument or result (*no methods out of thin air*). In addition:
 - if the method is from Ψ' , the call must be tagged with O (i.e. issued by K);
 - if the method is from Ψ , the call must be tagged with P (i.e. issued by L towards L');
 - for a call of method $m \notin \Psi \cup \Psi'$ to be valid, m must be introduced in an earlier move of dual polarity (*calls dual to introductions*).
2. Any method name appearing inside a call or return argument in h must be *fresh*, i.e. not used earlier (*introductions always fresh*).
 - This reflects the assumption that methods can be called and returned from, but not compared for identity equality. It is therefore a requirement towards the *completeness* of histories as a semantics for concurrent libraries. For example, this ensures that rules like η -equality are preserved in the semantics.
 - The condition serves the additional purpose of making the setting described in Fig. 1 robust, as it prevents method names in Ψ from being leaked to the client K . This ensures that encapsulation cannot be broken.

Given $h \in \text{PreH}$ and $t \in \mathbb{N}$, we write $t \times h$ for h in which each element is decorated with t :

$$t \times ((x_1)_{x_1} (x_2)_{x_2} \cdots (x_k)_{x_k}) = (t, x_1)_{x_1} (t, x_2)_{x_2} \cdots (t, x_k)_{x_k}.$$

We say that a move $(t, x)_X$ **introduces** a name $m \in \text{Meths}$ when $x \in \{\text{call } m'(v), \text{ret } m'(v)\}$ for some m', v such that v contains m .

Definition 3. Given Ψ, Ψ' , the set of **histories** over $\Psi \rightarrow \Psi'$, written $\mathcal{H}_{\Psi, \Psi'}$, is defined by

$$\mathcal{H}_{\Psi, \Psi'} = \bigcup_{N \geq 0} \bigcup_{h_1, \dots, h_N \in \text{PreH}_0} (1 \times h_1) \mid \cdots \mid (N \times h_N)$$

where $(1 \times h_1) \mid \cdots \mid (N \times h_N)$ is the set of all interleavings of $(1 \times h_1), \dots, (N \times h_N)$ satisfying:

1. For any $s_1 (t, \text{call } m(v))_X s_2 \in \mathcal{H}_{\Psi, \Psi'}$, either $m \in \Psi'$ and $X = O$, or $m \in \Psi$ and $X = P$, or there is a move $(t', x)_{X'}$ in s_1 that introduces m and $X \neq X'$.
2. For any $s_1 (t, x)_X s_2 \in \mathcal{H}_{\Psi, \Psi'}$ and any m , if m is introduced by x then m must not occur in s_1 .

Note that the definition supports scenarios in which a method sent as a parameter by one thread can be called by a different thread. This feature will be explored in Example 18.

A history $h \in \mathcal{H}_{\Psi, \Psi'}$ is called **sequential** if it is of the form

$$h = (t_1, x_1)_O (t_1, x'_1)_P \cdots (t_k, x_k)_O (t_k, x'_k)_P$$

for some t_i, x_i, x'_i . We write $\mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$ for the set of all sequential histories from $\mathcal{H}_{\Psi, \Psi'}$.

We shall range over $\mathcal{H}_{\Psi, \Psi'}$ using h, s (and variants). The subscripts Ψ, Ψ' will often be omitted. Given a history h , we shall write \bar{h} for the sequence of moves obtained from h by dualising all move polarities inside it. The set of **co-histories** over $\Psi \rightarrow \Psi'$ will be $\mathcal{H}_{\Psi, \Psi'}^{\text{co}} = \{\bar{h} \mid h \in \mathcal{H}_{\Psi, \Psi'}\}$.

While in this section histories will be extracted from example libraries informally, in Section 3.1 we give the formal semantics $\llbracket L \rrbracket$ of libraries. For each $L : \Psi \rightarrow \Psi'$, we shall have $\llbracket L \rrbracket \subseteq \mathcal{H}_{\Psi, \Psi'}$.

Remark 4. The notion of history introduced above extends the classic notion from [1] to higher-order types. It also extends the notion presented in [3]. The intuition behind the definition is that a history is a sequence of (well-bracketed) method calls and returns, called *moves*, each tagged with a thread identifier and a polarity, where polarities track the originators and recipients of moves. Moves may be calls or returns related to methods given in the library interface ($\Psi \rightarrow \Psi'$), or dynamically created methods that appear earlier inside the histories – recall that, in a higher-order setting, methods can be passed around as arguments to calls or be returned as results by other methods. On the other hand, a sequential history is one in which the operations performed by the library can be perceived as **atomic**, that is, each move produced by O is to be immediately followed by the library's response, which is a P move in the same thread.

Example 5 (Multiset spec). We now revisit our first example and provide a specification for it. Recall the multiset library L_{mset} from Fig. 2. Our verification goal will be to prove linearisability of L_{mset} to a specification $A_{\text{mset}} \subseteq \mathcal{H}_{\emptyset, \Psi}^{\text{seq}}$, where $\Psi = \{\text{count}, \text{update}\}$, which we define below. Intuitively, the specification stipulates that the multiset operations are functionally correct and only includes sequential histories. For example, the following histories are in the specification:

$(1, \text{call upd}(5, m))_O (1, \text{call } m(5))_P (1, \text{call cnt}(5))_O (1, \text{ret cnt}(0))_P$
 $(1, \text{ret } m(42))_O (1, \text{ret upd}(42))_P$
 $(1, \text{call upd}(5, m))_O (1, \text{call } m(5))_P (2, \text{call upd}(5, m'))_O (2, \text{call } m'(5))_P$
 $(1, \text{ret } m(42))_O (1, \text{ret upd}(42))_P (3, \text{call cnt}(5))_O (3, \text{ret cnt}(42))_P$
 $(2, \text{ret } m'(24))_O (2, \text{ret upd}(24))_P (1, \text{call cnt}(5))_O (1, \text{ret cnt}(24))_P$

while the next ones are not:

$(1, \text{call upd}(5, m))_O (1, \text{call } m(5))_P (1, \text{call cnt}(5))_O (1, \text{ret cnt}(42))_P \dots$
 $(1, \text{call upd}(5, m))_O (2, \text{call upd}(6, m'))_O \dots$
 $(1, \text{call upd}(5, m))_O (1, \text{call } m(5))_P (2, \text{call upd}(5, m'))_O (2, \text{call } m'(5))_P$
 $(1, \text{ret } m(42))_O (1, \text{ret upd}(42))_P (3, \text{call cnt}(5))_O (3, \text{ret cnt}(42))_P$
 $(2, \text{ret } m'(24))_O (2, \text{ret upd}(24))_P (1, \text{call cnt}(5))_O (1, \text{ret cnt}(42))_P$

A_{mset} will certify that L_{mset} correctly implements some integer multiset I whose elements change over time according to the moves in h . For a multiset I and natural numbers i, j , we write $I(i)$ for the multiplicity of i in I , and $I[i \mapsto j]$ for I with its multiplicity of i set to j . We shall stipulate that moves inside histories $h \in A_{\text{mset}}$ be annotatable with multisets I in such a way that the multiset is empty at the start of h (i.e. $I(i) = 0$ for all i) and:

- If I is changed between two consecutive moves in h then the second move is a P -move. In other words, the client cannot directly update the elements of I .
- Each call to *count* on argument i must be immediately followed by a return with value $I(i)$, and with I remaining unchanged.
- Each call to *update* on (i, m) must be followed by a call to m on $I(i)$, with I unchanged. Moreover, m must later return with some value j . Assuming at that point the multiset will have value J , if $I(i) = J(i)$ then the next move is a return of the original *update* call, with value j ; otherwise, a new call to m on $J(i)$ is produced, and so on.

We formally define the specification next.

Let $\mathcal{H}_{\emptyset, \Psi}^\circ$ contain sequences of moves from $\emptyset \rightarrow \Psi$ accompanied by a multiset (i.e. the sequences consist of elements of the form $(t, x, I)_X$). For each $s \in \mathcal{H}_{\emptyset, \Psi}^\circ$, we let $\pi_1(s)$ be the history extracted by projection, i.e. $\pi_1(s) \in \mathcal{H}_{\emptyset, \Psi}$. For each t , we let $s \upharpoonright t$ be the subsequence of s of elements with first component t . Writing \sqsubseteq_{pre} for the prefix relation, we define $A_{\text{mset}} = \{ \pi_1(s) \mid s \in A_{\text{mset}}^\circ \}$ where:

$$\begin{aligned}
 A_{\text{mset}}^\circ = \{ s \in \mathcal{H}_{\emptyset, \Psi}^\circ \mid & \pi_1(s) \in \mathcal{H}_{\emptyset, \Psi}^{\text{seq}} \wedge \forall t. s \upharpoonright t \in \mathcal{S} \wedge (s = (_, I)_O s' \implies \forall i. I(i) = 0) \\
 & \wedge \forall s' (_, I)_P (_, J)_O \sqsubseteq_{\text{pre}} s. I = J \}
 \end{aligned}$$

and, for each t , the set of t -indexed annotated histories \mathcal{S} is given by the following grammar:

$$\begin{aligned}
 \mathcal{S} & \rightarrow \epsilon \mid (t, \text{call cnt}(i), I)_O (t, \text{ret cnt}(I(i)), I)_P \mathcal{S} \\
 & \mid (t, \text{call upd}(i, m), I)_O \mathcal{M}_{I, J}^{i, j} (t, \text{ret upd}(|j|), J[i \mapsto |j|])_P \mathcal{S} \\
 \mathcal{M}_{I, J}^{i, j} & \rightarrow (t, \text{call } m(I(i)), I)_P \mathcal{S} (t, \text{ret } m(j), J)_O \quad \text{provided } J(i) = I(i) \\
 \mathcal{M}_{I, J}^{i, j} & \rightarrow (t, \text{call } m(I(i)), I)_P \mathcal{S} (t, \text{ret } m(j'), J')_O \mathcal{M}_{J', J}^{i, j} \quad \text{provided } J'(i) \neq I(i)
 \end{aligned}$$

By definition, all histories in A_{mset} are sequential. The elements of A_{mset}° carry along the multiset I that is being represented. The conditions on A_{mset}° stipulate that I is initially empty and that O cannot change the value of I , while the rest of the conditions above are imposed by the grammar for \mathcal{S} . With the notion of linearisability to be introduced next, we will be able to show that $\llbracket L_{\text{mset}} \rrbracket$ is indeed linearisable to A_{mset} .

Remark 6. In our framework (higher-order computation with state) specifications are necessarily close to implementations. For example, they need to preserve the exact number of calls/returns, because each of them could trigger a potential side effect. As in [1], specifications contain sequential histories.

2.2. Three notions of linearisability

We present three notions of linearisability. First introduce a general notion that generalises classic linearisability [1] and parameterised linearisability [3]. We then develop two more specialised variants: a notion of encapsulated linearisability, following [3], that captures scenarios where the parameter library and the client cannot directly interact; and a relational notion whereby context behaviour (client and parameter library) is known to be relationally invariant.

2.2.1. General linearisability

We begin by introducing a class of reorderings on histories.

Definition 7. Let $\triangleleft_{PO} \subseteq \mathcal{H}_{\Psi, \Psi'} \times \mathcal{H}_{\Psi, \Psi'}$ be the smallest binary relation over $\mathcal{H}_{\Psi, \Psi'}$ satisfying, for any $t \neq t'$:

$$s_1(t', x')_{Z'}(t, x)_Z s_2 \triangleleft_{PO} s_1(t, x)_Z(t', x')_{Z'} s_2$$

whenever $Z = P$ or $Z' = O$.

Intuitively, two histories h_1, h_2 are related by \triangleleft_{PO} if the latter can be obtained from the former by swapping two adjacent moves from different threads in such a way that, after the swap, a P -move will occur earlier or an O -move will occur later. Note that the relation always applies to adjacent moves of the same polarity. On the other hand, we do *not* have $s_1(t, x)_P(t', x')_O s_2 \triangleleft_{PO} s_1(t', x')_O(t, x)_P s_2$.

Example 8. Let $\Psi = \{m : \text{int} \rightarrow \text{int}\}$ and $\Psi' = \{m' : \text{int} \rightarrow \text{int}\}$. Consider $h, h' \in \mathcal{H}_{\Psi, \Psi'}$ given below:

$$\begin{aligned} h &= (1, \text{call } m(1))_O (2, \text{call } m(5))_O (1, \text{call } m'(2))_P (1, \text{ret } m'(3))_O \\ &\quad (2, \text{call } m'(6))_P (2, \text{ret } m'(7))_O (2, \text{ret } m(8))_P (1, \text{ret } m(4))_P \\ h' &= (1, \text{call } m(1))_O (1, \text{call } m'(2))_P (1, \text{ret } m'(3))_O (1, \text{ret } m(4))_P \\ &\quad (2, \text{call } m(5))_O (2, \text{call } m'(6))_P (2, \text{ret } m'(7))_O (2, \text{ret } m(8))_P \end{aligned}$$

Note that $h \triangleleft_{PO}^* h'$ by permuting $(2, \text{call } m(5))_O$ rightwards and $(1, \text{ret } m(4))_P$ leftwards.

As another example, we can revisit the histories in Fig. 3. There, O -moves are coloured purple and P -moves are blue. In part (a) we can see that:

- the first history linearises to a sequential one by swapping a P -move of thread 1 to the left of two moves of thread 2,
- the second history linearises to a sequential one by swapping an O -move of thread 1 to the right of two moves of thread 2,
- the third history is already sequential and it cannot be linearised to a different one.

In part (b), on the other hand, the first history linearises to the second one by a series of swaps (left as exercise).

Analogously, one can consider the symmetric variant \triangleleft_{OP} of \triangleleft_{PO} , which will turn out useful in our soundness argument.

Definition 9 (General linearisability). Given $h_1, h_2 \in \mathcal{H}_{\Psi, \Psi'}$, we say that h_1 **is linearised by** h_2 , written $h_1 \triangleleft h_2$, if $h_1 \triangleleft_{PO}^* h_2$.

Given libraries $L, L' : \Psi \rightarrow \Psi'$ and a set of sequential histories $A \subseteq \mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$, we write $L \triangleleft A$, and say that L *can be linearised to* A , if for any $h \in \llbracket L \rrbracket$ there exists $h' \in A$ such that $h \triangleleft h'$. Moreover, we write $L \triangleleft L'$ if $L \triangleleft \llbracket L' \rrbracket \cap \mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$ (i.e. for all $h \in \llbracket L \rrbracket$ there is sequential $h' \in \llbracket L' \rrbracket$ such that $h \triangleleft h'$).

Remark 10. The classic notion of linearisability from [1] states that h linearises to h' just if the return/call order of h is preserved in h' (and h' is sequential), i.e. if a return move precedes a call move in h then so is the case in h' . Observing that, in [1], return and call moves coincide with P - and O -moves respectively, we can see that our higher-order notion of linearisability is a generalisation of the classic notion.

Our definition shows that the ownership of actions is the key determinant of what moves can be swapped rather than the call/return distinction, which was prominent in the classic case. It just so happens that, for $\Psi = \emptyset$ and $\Psi' = \{m' : \text{int} \rightarrow \text{int}\}$, the two coincide.

For further comparison, recall that the classic definition allowed for call/call, ret/ret and call/ret swaps, but ret/call was forbidden. According to our definition, what is allowed depends on polarity, so a call/call swap may well be illegal if the first call is a P -move and the second call is an O -move. Similarly, a ret/call swap is allowed as long as both actions belong to the same player or the return is an O -action and the call is a P -action. For instance, Example 8 involves the following kinds of swaps: call_O/call_P, call_O/ret_O, ret_P/ret_P, ret_O/ret_P, call_P/ret_P, call_O/ret_P.

Our emphasis on move ownership is motivated by Lemma 34, which will ultimately enable us to prove that, if $h \triangleleft h'$, then h' suffices to demonstrate the interactive potential of h . This intuition is formally captured in Theorem 35.

Remark 11. [3] defines linearisation using a “big-step” relation that applies a single permutation to the whole sequence. This contrasts with our definition as \triangleleft_{PO}^* , in which we combine multiple adjacent swaps. In Appendix A we show that the two definitions are equivalent.

2.2.2. Encapsulated linearisability

We next show that a more permissive notion of linearisability applies if the parameter library L' of Fig. 1 is encapsulated, that is, the client K can have no direct access to it (i.e. $\Psi'' = \emptyset$). To capture this scenario, we define a second polarity function on moves, which determines the **side** of the move:

- a move with side \mathcal{K} is played between the library L and the client K , while
- a move with side \mathcal{L} is played between the library L and the parameter library L' .

Formally, given a history $h \in \mathcal{H}_{\Psi, \Psi'}$, we define a **side** function on its moves by:

$$\text{side}((t, \text{call } m(v))) = \begin{cases} \mathcal{K} & \text{if } m \in \Psi' \\ \mathcal{L} & \text{if } m \in \Psi \\ \text{side}((t', x)) & \text{if } (t', x) \text{ introduces } m \end{cases}$$

$$\text{side}((t, \text{ret } m(v))) = \text{side}((t, \text{call } m(v')))$$

where, in the latter case, $(t, \text{call } m(v'))$ is the corresponding call of $(t, \text{ret } m(v))$. Thus, every move in h can be assigned a unique side polarity from $\{\mathcal{K}, \mathcal{L}\}$. For simplicity, we shall be tagging moves with a second index $Y \in \{\mathcal{K}, \mathcal{L}\}$ corresponding to their side polarity.

In this more restrictive nature of interaction, in which K and L' are separated, in addition to sequentiality in every thread we shall insist that a move made by the library in the \mathcal{L} or \mathcal{K} side must be followed by an O move from the *same* side.

Definition 12. We call a history $h \in \mathcal{H}_{\Psi, \Psi'}$ **encapsulated** if, for each thread t , we have that if

$$h = s_1(t, x)_{PY} s_2(t, x')_{OY'} s_3$$

and moves from t are absent from s_2 then $Y = Y'$. Moreover, if $L : \Psi \rightarrow \Psi'$, we set $\mathcal{H}_{\Psi, \Psi'}^{\text{enc}} = \{h \in \mathcal{H}_{\Psi, \Psi'} \mid h \text{ encapsulated}\}$ and $\llbracket L \rrbracket_{\text{enc}} = \llbracket L \rrbracket \cap \mathcal{H}_{\Psi, \Psi'}^{\text{enc}}$.

We define the corresponding linearisability notion as follows. First, let $\diamond \subseteq \mathcal{H}_{\Psi, \Psi'} \times \mathcal{H}_{\Psi, \Psi'}$ be the smallest binary relation on $\mathcal{H}_{\Psi, \Psi'}$ such that, for any X, X' , and any $Y, Y' \in \{\mathcal{K}, \mathcal{L}\}$ with $Y \neq Y'$ and $t \neq t'$:

$$s_1(t, m)_{XY}(t', m')_{X'Y'} s_2 \diamond s_1(t', m')_{X'Y'}(t, m)_{XY} s_2$$

Definition 13 (Encapsulated linearisability). Given $h_1, h_2 \in \mathcal{H}_{\Psi, \Psi'}^{\text{enc}}$, we say that h_1 is **enc-linearised** by h_2 , and write $h_1 \triangleleft_{\text{enc}} h_2$, if $h_1(\triangleleft_{PO} \cup \diamond)^* h_2$ and h_2 is sequential.

A library $L : \Psi \rightarrow \Psi'$ can be *enc-linearised* to A , written $L \triangleleft_{\text{enc}} A$, if $A \subseteq \mathcal{H}_{\Psi, \Psi'}^{\text{seq}} \cap \mathcal{H}_{\Psi, \Psi'}^{\text{enc}}$ and for any $h \in \llbracket L \rrbracket_{\text{enc}}$ there exists $h' \in A$ such that $h \triangleleft_{\text{enc}} h'$. We write $L \triangleleft_{\text{enc}} L'$ if $L \triangleleft_{\text{enc}} \llbracket L' \rrbracket_{\text{enc}} \cap \mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$.

Remark 14. Suppose $\Psi = \{m : \text{int} \rightarrow \text{int}\}$ and $\Psi' = \{m' : \text{int} \rightarrow \text{int}\}$. Histories from $\mathcal{H}_{\Psi, \Psi'}$ can contain the following actions:

$$\text{call } m'(i)_{OK}, \quad \text{ret } m(i)_{OL}, \quad \text{call } m(i)_{PK}, \quad \text{ret } m'(i)_{PK}.$$

Then, $(\triangleleft_{PO} \cup \diamond)^*$ prevents $(t, \text{call } m(i)_{PK})$ from being swapped with any subsequent $(t', \text{ret } m(i)_{OL})$, and similarly for $(t, \text{ret } m'(i)_{PK})$ and $(t', \text{call } m'(i)_{OK})$. Thus, our definition coincides in this example with Definition 3 of [3].

Remark 15. The encapsulated framework implies that the client and the parameter library are independent entities. Consequently, whenever their interaction with the library involves two adjacent moves $(t, m)_{XY}(t', m')_{X'Y'}$ with $t \neq t', X \neq X'$, permuting them will also generate a valid interaction. This justifies the extra freedom in rearranging moves in Definition 13. The soundness of this intuition is validated in Lemma 39 and Theorem 40.

Example 16 (Parameterised multiset). We revisit the multiset library of Example 1 and extend it with a public method *reset*, which performs multiplicity resets to default values using an abstract method *default* as the default-value function (again, we use absolute values to avoid negative multiplicities). The extended library is shown in the RHS of Fig. 2 and written


```

1  public run; ...;
2  Lock lock;
3  struct {fun, arg, wait, retv} requests[N];
4
5  run = λ (f,x).
6    requests[tid].fun := f;
7    requests[tid].arg := x;
8    requests[tid].wait := 1;
9    while (requests[tid].wait)
10     if (lock.tryacquire()) (
11       for (t = 0; t < N; t++)
12         if (requests[t].wait) (
13           requests[t].retv :=
14             requests[t].fun (requests[t].arg);
15           requests[t].wait := 0;
16         ); lock.release();
17     requests[tid].retv;

```

Fig. 4. Flat combination library L_{fc} .

$L_{mset2} : \Psi \rightarrow \Psi'$, with $\Psi = \{\text{default}\}$ and $\Psi' = \{\text{count}, \text{update}, \text{reset}\}$. In contrast to the *update* method of L_{mset} , *reset* is not optimistic: it retrieves the lock upon its call, and only releases it before return. In particular, the method *default* while it retains the lock.

Observe that, were *default* able to externally call *update*, we would reach a deadlock: *default* would be keeping the lock while waiting for the return of a method that requires the lock. On the other hand, if the library is encapsulated then the latter scenario is not possible. In such a case, L_{mset2} linearises to the specification A_{mset2} , defined next. Let $A_{mset2} = \{\pi_1(s) \mid s \in A_{mset2}^\circ\}$, where

$$A_{mset2}^\circ = \{s \in \mathcal{H}_{\Psi, \Psi'}^\circ \mid \pi_1(s) \in \mathcal{H}_{\Psi, \Psi'}^{\text{seq}} \cap \mathcal{H}_{\Psi, \Psi'}^{\text{enc}} \wedge \forall t. s \upharpoonright t \in \mathcal{S} \wedge (s = (_, I)_O \mathcal{S}' \implies \forall i. I(i) = 0) \\ \wedge \forall s'(_, I)_P(_, J)_O \sqsubseteq_{\text{pre}} s. I = J\}$$

and the set \mathcal{S} is now given by the grammar of Example 5 extended with the rule:

$$\mathcal{S} \rightarrow (t, \text{call reset}(i), I)_O \mathcal{K} (t, \text{call default}(i), I)_P \mathcal{L} (t, \text{ret default}(j), I)_O \mathcal{L} (t, \text{ret reset}(|j|), I')_P \mathcal{K} \mathcal{S}$$

with $I' = I[i \mapsto |j|]$. Our framework makes it possible to confirm that L_{mset2} enc-linearises to A_{mset2} .

2.2.3. Relational linearisability

We finally extend general linearisability to cater for situations where the client and the parameter library adhere to closure constraints expressed by relations \mathcal{R} on histories. Let Ψ, Ψ' be sets of abstract and public methods respectively. The closure relations we consider are closed under permutations of methods outside $\Psi \cup \Psi'$: if $h \mathcal{R} h'$ and π is a (type-preserving) permutation on $\text{Meths} \setminus (\Psi \cup \Psi')$ then $\pi(h) \mathcal{R} \pi(h')$. The requirement represents the fact that, apart from the method names from a library interface, the other method names are arbitrary and can be freely permuted without any observable effect. Thus, \mathcal{R} should not be distinguishing between such names.

Definition 17 (Relational linearisability). Let $\mathcal{R} \subseteq \mathcal{H}_{\Psi, \Psi'} \times \mathcal{H}_{\Psi, \Psi'}$ be closed under permutations of names in $\text{Meths} \setminus (\Psi \cup \Psi')$. Given $h_1, h_2 \in \mathcal{H}_{\Psi, \Psi'}$, we say that h_1 is \mathcal{R} -linearised by h_2 , and write $h_1 \triangleleft_{\mathcal{R}} h_2$, if $h_1 (\triangleleft_{PO} \cup \mathcal{R})^* h_2$ and h_2 is sequential. A library $L : \Psi \rightarrow \Psi'$ can be \mathcal{R} -linearised to A , written $L \triangleleft_{\mathcal{R}} A$, if $A \subseteq \mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$ and for any $h \in \llbracket L \rrbracket$ there exists $h' \in A$ such that $h \triangleleft_{\mathcal{R}} h'$. We write $L \triangleleft_{\mathcal{R}} L'$ if $L \triangleleft_{\mathcal{R}} \llbracket L' \rrbracket \cap \mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$.

Example 18. We consider a higher-order variant of an example from [3] that motivates relational linearisability. Flat combining [9] is a synchronisation paradigm that advocates the use of a single thread holding a global lock to process requests of all other threads. To facilitate this, threads share an array to which they write the details of their requests and wait either until they acquire a lock or their request has been processed by another thread. Once a thread acquires a lock, it executes all requests stored in the array and the outcomes are written to the array for access by the requesting threads.

Let $\Psi' = \{\text{run} \in \text{Meths}_{(\theta \rightarrow \theta') \times \theta, \theta'}\}$. The library $L_{fc} : \emptyset \rightarrow \Psi'$ in Fig. 4 is built following the flat combining approach and, on acquisition of the global lock, the winning thread acts as a combiner of all registered requests. Note that the requests will be attended to one after another (thus guaranteeing mutual exclusion) and only one lock acquisition will suffice to process one array of requests. Using our framework, one can show that L_{fc} can be \mathcal{R} -linearised to the specification given by the library L_{spec} defined by

$$\text{run} = \lambda (f, x). (\text{lock.acquire}(); \text{let result} = f(x) \text{ in } \text{lock.release}(); \text{result})$$

where each function call in L_{spec} is protected by a lock. Observe that we cannot hope for $L_{fc} \triangleleft_{\mathcal{R}} L_{\text{spec}}$, because clients may call library methods with functional arguments that recognise thread identity. Consequently, we can relate the two libraries

Libraries $L ::= B \mid \text{abstract } m; L \mid \text{public } m; L$ **Clients** $K ::= M \parallel \dots \parallel M$
Blocks $B ::= \epsilon \mid m = \lambda x.M; B \mid r := \lambda x.M; B \mid r := i; B$ **Values** $v ::= () \mid i \mid m \mid \langle v, v \rangle$
Terms $M ::= () \mid i \mid \text{tid} \mid x \mid m \mid M \oplus M \mid \langle M, M \rangle \mid \pi_1 M \mid \pi_2 M \mid \text{if } M \text{ then } M \text{ else } M \mid \lambda x^\theta.M \mid xM \mid mM \mid \text{let } x = M \text{ in } M \mid r := M \mid !r$

$$\begin{array}{c}
\frac{}{\Gamma \vdash () : \text{unit}} \quad \frac{}{\Gamma \vdash i : \text{int}} \quad \frac{}{\Gamma \vdash \text{tid} : \text{int}} \quad \frac{\Gamma(x) = \theta}{\Gamma \vdash x : \theta} \quad \frac{m \in \text{Meths}_{\theta, \theta'}}{\Gamma \vdash m : \theta \rightarrow \theta'} \quad \frac{\Gamma \vdash M : \text{int} \quad \Gamma \vdash M_0, M_1 : \theta}{\Gamma \vdash \text{if } M \text{ then } M_1 \text{ else } M_0 : \theta} \quad \frac{\Gamma \vdash M_1, M_2 : \text{int}}{\Gamma \vdash M_1 \oplus M_2 : \text{int}} \\
\frac{\Gamma \vdash M : \theta_1 \times \theta_2}{\Gamma \vdash \pi_i M : \theta_i \ (i = 1, 2)} \quad \frac{\Gamma \vdash M_i : \theta_i \ (i = 1, 2)}{\Gamma \vdash \langle M_1, M_2 \rangle : \theta_1 \times \theta_2} \quad \frac{\Gamma, x : \theta \vdash M : \theta'}{\Gamma \vdash \lambda x^\theta.M : \theta \rightarrow \theta'} \quad \frac{\Gamma(x) = \theta \rightarrow \theta' \quad \Gamma \vdash M : \theta}{\Gamma \vdash xM : \theta'} \quad \frac{m \in \text{Meths}_{\theta, \theta'} \quad \Gamma \vdash M : \theta}{\Gamma \vdash mM : \theta'} \\
\frac{\Gamma \vdash M : \theta \quad \Gamma, x : \theta \vdash N : \theta'}{\Gamma \vdash \text{let } x = M \text{ in } N : \theta'} \quad \frac{r \in \text{Refs}_{\text{int}} \quad \Gamma \vdash M : \text{int}}{\Gamma \vdash r := M : \text{unit}} \quad \frac{r \in \text{Refs}_{\theta, \theta'} \quad \Gamma \vdash M : \theta \rightarrow \theta'}{\Gamma \vdash r := M : \text{unit}} \quad \frac{r \in \text{Refs}_{\text{int}}}{\Gamma \vdash !r : \text{int}} \quad \frac{r \in \text{Refs}_{\theta, \theta'}}{\Gamma \vdash !r : \theta \rightarrow \theta'} \\
\frac{}{\vdash_B \epsilon : \emptyset} \quad \frac{m \in \text{Meths}_{\theta, \theta'} \quad x : \theta \vdash M : \theta' \quad \vdash_B B : \Psi}{\vdash_B m = \lambda x.M; B : \Psi \uplus \{m\}} \quad \frac{r \in \text{Refs}_{\theta, \theta'} \quad x : \theta \vdash M : \theta' \quad \vdash_B B : \Psi}{\vdash_B r := \lambda x.M; B : \Psi} \quad \frac{r \in \text{Refs}_{\text{int}} \quad \vdash_B B : \Psi}{\vdash_B r := i; B : \Psi} \\
\frac{\vdash_B B : \Psi}{\text{Meths}(B) \vdash_L B : \emptyset \rightarrow \Psi} \quad \frac{\Psi \uplus \{m\} \vdash_L L : \Psi' \rightarrow \Psi'' \quad m \in \Psi''}{\Psi \vdash_L \text{public } m; L : \Psi' \rightarrow \Psi''} \quad \frac{\Psi \uplus \{m\} \vdash_L L : \Psi' \rightarrow \Psi'' \quad m \notin \Psi''}{\Psi \vdash_L \text{abstract } m; L : \Psi' \uplus \{m\} \rightarrow \Psi''} \\
\frac{\vdash M_j : \text{unit} \ (j = 1, \dots, N)}{\Psi \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}} \quad \forall j. \text{Meths}(M_j) \subseteq \Psi
\end{array}$$

Fig. 5. Library syntax, and typing rules for terms (\vdash), blocks (\vdash_B), libraries (\vdash_L), clients (\vdash_K).

only if context behaviour is guaranteed to be independent of thread identifiers. This can be expressed through $\triangleleft_{\mathcal{R}}$, where $\mathcal{R} \subseteq \mathcal{H}_{\emptyset, \Psi'} \times \mathcal{H}_{\emptyset, \Psi'}$ is a relation capturing thread-blind client behaviour (see Subsection 3.2 for details).

3. Library syntax and semantics

We now look at the concrete syntax of libraries and clients. Libraries comprise collections of typed methods whose argument and result types adhere to the grammar: $\theta ::= \text{unit} \mid \text{int} \mid \theta \rightarrow \theta \mid \theta \times \theta$.

We shall use three disjoint enumerable sets of names, referred to as Vars, Meths and Refs, to name respectively variables, methods and references. x, f (and their decorated variants) will be used to range over Vars; m will range over Meths; and r over Refs. Methods and references are implicitly typed, i.e. $\text{Meths} = \biguplus_{\theta, \theta'} \text{Meths}_{\theta, \theta'}$ and $\text{Refs} = \text{Refs}_{\text{int}} \uplus \biguplus_{\theta, \theta'} \text{Refs}_{\theta, \theta'}$, where $\text{Meths}_{\theta, \theta'}$ contains names for methods of type $\theta \rightarrow \theta'$, Refs_{int} contains names of integer references and $\text{Refs}_{\theta, \theta'}$ contains names for references to methods of type $\theta \rightarrow \theta'$. We write \uplus for disjoint set union.

The syntax for libraries and clients is given in Fig. 5. Each library L begins with a series of method declarations (public or abstract) followed by a block B containing method implementations ($m = \lambda x.M$) and reference initialisations ($r := i$ or $r := \lambda x.M$). The typing rules ensure that each public method is implemented within the block, in contrast to abstract methods. Clients are parallel compositions of closed terms.

Terms M specify the shape of allowable method bodies. $()$ is the skip command, i ranges over integers, tid is the current thread identifier and \oplus represents standard arithmetic operations. Thanks to higher-order references, we can simulate divergence by $(!r)()$, where $r \in \text{Refs}_{\text{unit}, \text{unit}}$ is initialised with $\lambda x^{\text{unit}}.(!r)()$. Similarly, while $M \ N$ can be simulated by $(!r)()$ after $r := \lambda x^{\text{unit}}.\text{let } y = M \text{ in } (\text{if } y \text{ then } (N; (!r)()) \text{ else } ())$. We also use the standard derived syntax for sequential composition, i.e. $M; N$ stands for $\text{let } x = M \text{ in } N$, where x does not occur in N . For each term M , we write $\text{Meths}(M)$ for the set of method names occurring in M . We use the same notation for method names in blocks and libraries.

Remark 19. In Section 2 we used lock-related operations in our example libraries (*acquire*, *tryacquire*, *release*), on the understanding that they can be coded using shared memory. We assume that both *acquire* and *release* are blocking, while *tryacquire* is not. *tryacquire* makes an attempt to acquire the associated lock and returns 0 if the attempt was not successful or 1 otherwise. Similarly, the array of Example 18 in the sequel can be constructed using references.

For simplicity, we do not include private methods, yet the same effect could be achieved by storing them in higher-order references. As we explain in the next section, references present in library definitions are de facto private to the library. Note also that, according to our definition, sets of abstract and public methods are disjoint. However, given $m, m' \in \text{Refs}_{\theta, \theta'}$, one can define a “public abstract” method with: $\text{public } m; \text{abstract } m'; m = \lambda x^\theta.m'x$.

Terms are typed in environments $\Gamma = \{x_1 : \theta_1, \dots, x_n : \theta_n\}$. Method blocks are typed through judgements $\vdash_B B : \Psi$, where $\Psi \subseteq \text{Meths}$. The judgements collect the names of methods defined in a block as well as making sure that the definitions respect types and are not duplicated. Also, the initialisation statements must comply with types.

Finally, we type libraries using statements of the form $\Psi \vdash_L L : \Psi' \rightarrow \Psi''$, where $\Psi, \Psi', \Psi'' \subseteq \text{Meths}$ and $\Psi' \cap \Psi'' = \emptyset$. The judgement $\emptyset \vdash_L L : \Psi' \rightarrow \Psi''$ guarantees that any method occurring in L is present either in Ψ' or Ψ'' , that all methods

$(L) \longrightarrow_{\text{lib}} (L, \emptyset, S_{\text{init}})$	$(r := i; B, \mathcal{R}, S) \longrightarrow_{\text{lib}} (B, \mathcal{R}, S[r \mapsto i])$
$(\text{abstract } m; L, \mathcal{R}, S) \longrightarrow_{\text{lib}} (L, \mathcal{R}, S)$	$(m = \lambda x.M; B, \mathcal{R}, S) \longrightarrow_{\text{lib}} (B, \mathcal{R}_{**}, S)$
$(\text{public } m; L, \mathcal{R}, S) \longrightarrow_{\text{lib}} (L, \mathcal{R}, S)$	$(r := \lambda x.M; B, \mathcal{R}, S) \longrightarrow_{\text{lib}} (B, \mathcal{R}_{**}, S[r \mapsto m])$
<hr/>	
$(E[\text{tid}], \mathcal{R}, S) \rightarrow_t (E[t], \mathcal{R}, S)$	$(E[\text{if } i_* \text{ then } M_1 \text{ else } M_0], \mathcal{R}, S) \rightarrow_t (E[M_{j_*}], \mathcal{R}, S)$
$(E[i_1 \oplus i_2], \mathcal{R}, S) \rightarrow_t (E[i_{**}], \mathcal{R}, S)$	$(E[\pi_j(v_1, v_2)], \mathcal{R}, S) \rightarrow_t (E[v_j], \mathcal{R}, S)$
$(E[r], \mathcal{R}, S) \rightarrow_t (E[S(r)], \mathcal{R}, S)$	$(E[\text{let } x = v \text{ in } M], \mathcal{R}, S) \rightarrow_t (E[M\{v/x\}], \mathcal{R}, S)$
$(E[r := i], \mathcal{R}, S) \rightarrow_t (E[()], \mathcal{R}, S[r \mapsto i])$	$(E[r := \lambda x.M], \mathcal{R}, S) \rightarrow_t (E[()], \mathcal{R}_{**}, S[r \mapsto m])$
$(E[\lambda x.M], \mathcal{R}, S) \rightarrow_t (E[m], \mathcal{R}_{**}, S)$	$(E[mv], \mathcal{R}_{**}, S) \rightarrow_t (E[M\{v/x\}], \mathcal{R}_{**}, S)$
$E ::= \bullet \mid E \oplus M \mid i \oplus E \mid \text{if } E \text{ then } M \text{ else } M \mid \pi_j E \mid \langle E, M \rangle \mid \langle v, E \rangle \mid mE \mid \text{let } x = E \text{ in } M \mid r := E$	
<hr/>	
$\frac{(M, \mathcal{R}, S) \rightarrow_t (M', \mathcal{R}', S')}{(M_1 \parallel \dots \parallel M_{t-1} \parallel M \parallel M_{t+1} \parallel \dots \parallel M_N, \mathcal{R}, S) \Longrightarrow (M_1 \parallel \dots \parallel M_{t-1} \parallel M' \parallel M_{t+1} \parallel \dots \parallel M_N, \mathcal{R}', S')} (K_N)$	

Fig. 6. Evaluation rules for libraries ($\longrightarrow_{\text{lib}}$), terms (\rightarrow_t) and clients (\Longrightarrow). In the rules above we use the conditions/notations: $\mathcal{R}_{**} = \mathcal{R} \uplus (m \mapsto \lambda x.M)$, $i_{**} = i_1 \oplus i_2$, $\mathcal{R}_*(m) = \lambda x.M$, and $j_* = 0$ iff $i_* = 0$.

in Ψ' are declared as abstract and unimplemented, while all methods in Ψ'' are declared as public and defined. Thus, $\emptyset \vdash_L L : \Psi \rightarrow \Psi'$ is a library in which Ψ, Ψ' are the abstract and public methods respectively. In this case, we also write $L : \Psi \rightarrow \Psi'$.

3.1. Semantics

The semantics of our system is given in several stages. First, we define an operational semantics for sequential and concurrent terms that may draw methods from a repository. We then adapt it to capture interactions of concurrent clients with closed libraries (no abstract methods). This notion is then used to define contextual approximation for arbitrary libraries. Finally, we introduce a trace semantics of arbitrary libraries, which generates the histories on which our notions of linearisability are based.

3.1.1. Library-client evaluation

Libraries, terms and clients are evaluated in environments comprising:

- A method environment \mathcal{R} , called *own-method repository*, which is a finite partial map on Meths assigning to each m in its domain, with $m \in \text{Meths}_{\theta, \theta'}$, a term of the form $\lambda y.M$ (we omit type-superscripts from bound variables for economy).
- A finite partial map $S : \text{Refs} \rightarrow (\mathbb{Z} \cup \text{Meths})$, called *store*, which assigns to each r in its domain an integer (if $r \in \text{Refs}_{\text{int}}$) or name from $\text{Meths}_{\theta, \theta'}$ (if $r \in \text{Refs}_{\theta, \theta'}$).

The evaluation rules are presented in Fig. 6, where we also define *evaluation contexts* E .

Remark 20. We shall assume that reference names used in libraries are library-private, i.e. sets of reference names used in different libraries are assumed to be disjoint. Similarly, when libraries are being used by client code, this is done on the understanding that the references available to that code do not overlap with those used by libraries. Still, for simplicity, we shall rely on a single set Refs of references in our operational rules.

First we evaluate the library to create an initial repository and store. This is achieved by the first set of rules in Fig. 6, where we assume that S_{init} is empty. Thus, library evaluation produces a tuple $(\epsilon, \mathcal{R}_0, S_0)$ including a method repository and a store, which can be used as the initial repository and store for evaluating $M_1 \parallel \dots \parallel M_N$ using the (K_N) rule. We shall call the latter evaluation semantics for clients (denoted by \Longrightarrow) the *multi-threaded operational semantics*. The latter relies on closed-term reduction (\rightarrow_t), whose rules are given in the middle group, where t is the current thread index. Note that the rules for $E[\lambda x.M]$ in the middle group, along with those for $m = \lambda x.M$ and $r := \lambda x.M$ in the first group, involve the creation of a fresh method name m , which is used to put the function in the repository \mathcal{R} . Name creation is non-deterministic: any fresh m of the appropriate type can be chosen.

We define termination for clients linked with libraries that have no abstract methods. Recall our convention (Remark 20) that L and M_1, \dots, M_N must access disjoint parts of the store. Terms M_1, \dots, M_N can share reference names, though.

Definition 21. Let $L : \emptyset \rightarrow \Psi'$ and $\Psi' \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}$. We say that $M_1 \parallel \dots \parallel M_N$ *terminates with linked library* L if $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_0, S_0) \Longrightarrow^* ((), \dots, (), \mathcal{R}, S)$, for some \mathcal{R}, S , where $(L) \longrightarrow_{\text{lib}}^* (\epsilon, \mathcal{R}_0, S_0)$. We then write link L in $(M_1 \parallel \dots \parallel M_N) \Downarrow$.

We shall build a notion of contextual approximation of libraries on top of termination: one library will be said to approximate another if, whenever the former terminates when composed with any parameter library and client, so does the latter.

We will be considering the following notions for composing libraries. Let us denote a library L as $L = D; B$, where D contains all the (public/abstract) method declarations of L , and B is its method block. We write $\text{Refs}(L)$ for the set of references in L . Let $L_1 : \Psi_1 \rightarrow \Psi_2$ be of the form $D_1; B_1$. Given $L_2 : \Psi'_1 \rightarrow \Psi'_2 (= D_2; B_2)$ such that $\Psi_2 \cap \Psi'_2 = \text{Refs}(L_1) \cap \text{Refs}(L_2) = \emptyset$, $\Psi = \{m_1, \dots, m_n\} \subseteq \Psi_2$ and $L' : \emptyset \rightarrow \Psi_1, \Psi'$, we define the *union* of L_1 and L_2 , the Ψ -*hiding* of L_1 , and the *sequencing* of L' with L_1 respectively as:

$$L_1 \cup L_2 : (\Psi_1 \cup \Psi'_1) \setminus (\Psi_2 \cup \Psi'_2) \rightarrow \Psi_2 \cup \Psi'_2 = (D_1; B_1) \cup (D_2; B_2) = D'_1; D'_2; B_1; B_2$$

$$L_1 \setminus \Psi : \Psi_1 \rightarrow (\Psi_2 \setminus \Psi) = (D_1; B_1) \setminus \Psi = D'_1; B'_1 \{!r_1/m_1\} \dots \{!r_n/m_n\}$$

$$L'; L_1 : \emptyset \rightarrow \Psi_2, \Psi' = (L' \cup L_1) \setminus \Psi_1$$

where D'_1 is D_1 with any abstract m declaration removed for $m \in \Psi'_2$, dually for D'_2 ; and where D'_1 is D_1 without public m declarations for $m \in \Psi$ and each r_i is a fresh reference matching the type of m_i , and B'_1 is obtained from B_1 by replacing each $m_i = \lambda x.M$ by $r_i := \lambda x.M$. Thus, the union of libraries L_1 and L_2 corresponds to merging their code and removing any abstract declarations for methods defined in the union. On the other hand, the hiding of a public method simply renders it private via the use of references. Sequencing allows for the following notion.

Definition 22. Given $L_1, L_2 : \Psi \rightarrow \Psi'$, we say that L_1 **contextually approximates** L_2 , written $L_1 \sqsubseteq L_2$, if for all $L' : \emptyset \rightarrow \Psi, \Psi''$ and $\Psi'', \Psi'' \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}$, if $\text{link } L'; L_1 \text{ in } (M_1 \parallel \dots \parallel M_N) \Downarrow$ then $\text{link } L'; L_2 \text{ in } (M_1 \parallel \dots \parallel M_N) \Downarrow$. In this case, we also say that L_2 **contextually refines** L_1 .

Note that, according to this definition, the parameter library L' may communicate directly with the client terms through a common interface Ψ'' . We shall refer to this case as the *general* case. Later on, we shall also consider more restrictive testing scenarios in which this possibility of explicit communication is removed. Moreover, from the disjointness conditions in the definitions of sequencing and linking we have that L_i, L' and $M_1 \parallel \dots \parallel M_N$ access pairwise disjoint parts of the store.

Remark 23. Our ultimate goal will be to show that our notion of linearisability, written \triangleleft , provides a sound method for proving contextual approximation/refinement, written \sqsubseteq . Recall that in order to establish $L_1 \triangleleft L_2$, one has to exhibit a subset A_2 of *sequential* histories taken from $\llbracket L_2 \rrbracket$ such that L_1 is linearisable to A_2 , written $L_1 \triangleleft A_2$.

3.1.2. Trace semantics

Building on the earlier semantics, we next introduce a trace semantics of libraries in the spirit of game semantics [14]. As mentioned in Section 2, the behaviour of a library will be represented as an exchange of moves between two players called P and O , representing the library and its corresponding context respectively. The context consists of the client of the library as well as the parameter library, with an index on each move $(\mathcal{K}/\mathcal{L})$ specifying which of them is involved in the move.

In contrast to the semantics of the previous section, we handle scenarios in which methods need not be present in the repository \mathcal{R} . Calls to such undefined methods are represented by labelled transitions – calls to the context made on behalf of the library (P). The calls can later be responded to with labelled transitions corresponding to returns, made by the context (O). On the other hand, O is able to invoke methods in \mathcal{R} , which will also be represented through suitable labels. Because we work in a higher-order setting, calls and returns made by both players may involve methods as arguments or results. Such methods also become available for future calls: function arguments/results supplied by P are added to the repository and can later be invoked by O , while function arguments/results provided by O can be queried in the same way as abstract methods.

The trace semantics utilises configurations that carry more components than the previous semantics. We define two kinds of configurations:

$$O\text{-configurations } (\mathcal{E}, -, \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \quad \text{and} \quad P\text{-configurations } (\mathcal{E}, M, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$$

where the component \mathcal{E} is an *evaluation stack*, that is, a stack of the form $[X_1, X_2, \dots, X_n]$ with each X_i being either an evaluation context or a method name. On the other hand, $\mathcal{P} = (\mathcal{P}_{\mathcal{L}}, \mathcal{P}_{\mathcal{K}})$ with $\mathcal{P}_{\mathcal{L}}, \mathcal{P}_{\mathcal{K}} \subseteq \text{dom}(\mathcal{R})$ being sets of *public* method names, and $\mathcal{A} = (\mathcal{A}_{\mathcal{L}}, \mathcal{A}_{\mathcal{K}})$ is a pair of sets of *abstract* method names. \mathcal{P} will be used to record all the method names produced by P and passed to O : those passed to OK are stored in $\mathcal{P}_{\mathcal{K}}$, while those passed to OL are kept in $\mathcal{P}_{\mathcal{L}}$. Inside \mathcal{A} , the story is the opposite one: $\mathcal{A}_{\mathcal{K}} (\mathcal{A}_{\mathcal{L}})$ stores the method names produced by OK (resp. OL) and passed to P . Consequently, the sets of names stored in $\mathcal{P}_{\mathcal{L}}, \mathcal{P}_{\mathcal{K}}, \mathcal{A}_{\mathcal{L}}, \mathcal{A}_{\mathcal{K}}$ will always be disjoint.

- (INT) $(\mathcal{E}, M, \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \rightarrow_t (\mathcal{E}, M', \mathcal{R}', \mathcal{P}, \mathcal{A}, S')$, given that $(M, \mathcal{R}, S) \rightarrow_t (M', \mathcal{R}', S')$ and $\text{dom}(\mathcal{R}' \setminus \mathcal{R})$ consists of names that do not occur in \mathcal{E}, \mathcal{A} .
- (PCy) $(\mathcal{E}, E[mv], \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{call } m(v')_{PY}}_t (m :: E :: \mathcal{E}, -, \mathcal{R}', \mathcal{P}', \mathcal{A}, S)$, given $m \in \mathcal{A}_Y$ and (P).
- (OCy) $(\mathcal{E}, -, \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{call } m(v)_{OY}}_t (m :: \mathcal{E}, M\{v/x\}, \mathcal{R}, \mathcal{P}, \mathcal{A}', S)$, given $m \in \mathcal{P}_Y$, $\mathcal{R}(m) = \lambda x.M$ and (O).
- (PRy) $(m :: \mathcal{E}, v, \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{ret } m(v')_{PY}}_t (\mathcal{E}, -, \mathcal{R}', \mathcal{P}', \mathcal{A}, S)$, given $m \in \mathcal{P}_Y$ and (P).
- (ORy) $(m :: E :: \mathcal{E}, -, \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{ret } m(v)_{OY}}_t (\mathcal{E}, E[v], \mathcal{R}, \mathcal{P}, \mathcal{A}', S)$, given $m \in \mathcal{A}_Y$ and (O).
- (P) If v contains the names m_1, \dots, m_k then $v' = v\{m'_i/m_i \mid 1 \leq i \leq k\}$ with each m'_i being a fresh name. Moreover, $\mathcal{R}' = \mathcal{R} \uplus \{m'_i \mapsto \lambda x.m_i x \mid 1 \leq i \leq k\}$ and $\mathcal{P}' = \mathcal{P} \cup_Y \{m'_1, \dots, m'_k\}$.
- (O) If v contains names m_1, \dots, m_k then $m_i \in \phi(\mathcal{P}, \mathcal{A})$, for each i , and $\mathcal{A}' = \mathcal{A} \cup_Y \{m_1, \dots, m_k\}$.

Fig. 7. Trace semantics rules. The rule (INT) is for embedding internal rules. In the rule (PCy), the library (P) calls one of its abstract methods (either the original ones or those acquired via interaction), while in (PRy) it returns from such a call. The rules (OCy) and (ORy) are dual and represent actions of the context. In all of the rules, whenever we write $m(v)$ or $m(v')$, we assume that the type of v matches the argument type of m .

Given a pair \mathcal{P} as above and a set $Z \subseteq \text{Meths}$, we write $\mathcal{P} \cup_K Z$ for the pair $(\mathcal{P}_L, \mathcal{P}_K \cup Z)$. We define \cup_L in a similar manner, and extend it to pairs \mathcal{A} as well. Moreover, given \mathcal{P} and \mathcal{A} , we let $\phi(\mathcal{P}, \mathcal{A})$ be the set of *fresh* method names for \mathcal{P} , \mathcal{A} : $\phi(\mathcal{P}, \mathcal{A}) = \text{Meths} \setminus (\mathcal{P}_L \cup \mathcal{P}_K \cup \mathcal{A}_L \cup \mathcal{A}_K)$.

We give the rules generating the trace semantics in Fig. 7. Note that the rules are parameterised by P/O and Y , which together determine the polarity of the next move; C/R , which stands for the move being a call (C) or a return (R) respectively. The rules depict the intuition presented above. When in an O -configuration, the context may issue a call to a public method $m \in \mathcal{P}_Y$ and pass control to the library (rule (OCy)). Note that, when this occurs, the name m is added to the evaluation stack \mathcal{E} and a P -configuration is obtained. From there on, the library will compute internally using rule (INT), until: it either needs to evaluate an abstract method (i.e. some $m' \in \mathcal{A}_Y$), and hence issues a call via rule (PCy); or it completes its computation and returns the call (rule (PRy)). Calls to abstract methods, on the other hand, are met either by further calls to public methods (via (OCy)), or by returns (via (ORy)).

Finally, we extend the trace semantics to a concurrent setting where a fixed number of N -many threads run in parallel. Each thread has separate evaluation stack and term components, which we write as $\mathcal{C} = (\mathcal{E}, X)$ (where X is a term or “—”). Thus, a configuration now is of the following form:

$$N\text{-configuration } (\mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_N, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$$

where, for each i , $\mathcal{C}_i = (\mathcal{E}_i, X_i)$ and $(\mathcal{E}_i, X_i, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$ is a sequential configuration. We shall abuse notation a little and write $(\mathcal{C}_i, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$ for $(\mathcal{E}_i, X_i, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$. The concurrent traces are produced by the following two rules

$$\frac{(\mathcal{C}_i, \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \rightarrow_i (\mathcal{C}', \mathcal{R}, \mathcal{P}, \mathcal{A}, S')}{(\mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_N, \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \Longrightarrow (\mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_{i-1} \parallel \mathcal{C}' \parallel \mathcal{C}_{i+1} \parallel \dots \parallel \mathcal{C}_N, \mathcal{R}, \mathcal{P}, \mathcal{A}, S')} \quad (\text{PlNT})$$

$$\frac{(\mathcal{C}_i, \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{x_{XY}}_i (\mathcal{C}', \mathcal{R}, \mathcal{P}, \mathcal{A}, S')}{(\mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_N, \mathcal{R}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{(i,x)_{XY}} (\mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_{i-1} \parallel \mathcal{C}' \parallel \mathcal{C}_{i+1} \parallel \dots \parallel \mathcal{C}_N, \mathcal{R}, \mathcal{P}, \mathcal{A}, S')} \quad (\text{PEXT})$$

with the proviso that the names freshly produced internally in (PlNT) are fresh for the whole of $\vec{\mathcal{C}}$.

We can now define the trace semantics of a library L . We call a configuration component \mathcal{C}_i **final** if it is in one of the following forms, for O - and P -configurations respectively: $\mathcal{C}_i = ([], -)$ or $\mathcal{C}_i = ([], ())$. We call $(\vec{\mathcal{C}}, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$ final just if $\vec{\mathcal{C}} = \mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_N$ and each \mathcal{C}_i is final.

Definition 24. For each $L : \Psi \rightarrow \Psi'$, we define the N -trace semantics of L to be

$$\llbracket L \rrbracket_N = \{s \mid (\vec{\mathcal{C}}_0, \mathcal{R}_0, (\emptyset, \Psi'), (\Psi, \emptyset), S_0) \xRightarrow{s}^* \rho \wedge \rho \text{ final}\}$$

where $\vec{\mathcal{C}}_0 = ([], -) \parallel \dots \parallel ([], -)$ and $(L) \longrightarrow_{\text{lib}}^* (\epsilon, \mathcal{R}_0, S_0)$.

For economy, in the sequel we might be dropping the index N from $\llbracket L \rrbracket_N$. We conclude the presentation of the trace semantics by providing a semantics for library contexts.

Recall that in our setting (Fig. 1) a library $L : \Psi \rightarrow \Psi'$ is deployed in a context consisting of a parameter library $L' : \emptyset \rightarrow \Psi, \Psi''$ and a concurrent composition of client threads $\Psi', \Psi'' \vdash M_i : \text{unit}$ ($i = 1, \dots, N$). We shall write link $L'; -$ in $(M_1 \parallel \dots \parallel M_N)$, or simply C , to refer to such contexts.

Definition 25. Let $\Psi', \Psi'' \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}$ and $L' : \emptyset \rightarrow \Psi, \Psi''$. We define the semantics of the context formed by L' and M_1, \dots, M_N to be:

$$\llbracket \text{link } L'; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket = \{s \mid (\vec{\mathcal{C}}_0, \mathcal{R}_0, (\Psi, \emptyset), (\emptyset, \Psi'), S_0) \xRightarrow{s}^* \rho \wedge \rho \text{ final}\}$$

where $(L') \longrightarrow_{\text{lib}}^* (\epsilon, \mathcal{R}_0, S_0)$ and $\vec{\mathcal{C}}_0 = ([], M_1) \parallel \dots \parallel ([], M_N)$.

Lemma 26. For any $L : \Psi \rightarrow \Psi'$, $L' : \emptyset \rightarrow \Psi, \Psi''$ and $\Psi', \Psi'' \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}$ we have $\llbracket L \rrbracket_N \subseteq \mathcal{H}_{\Psi, \Psi'}$ and $\llbracket \text{link } L'; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket \subseteq \mathcal{H}_{\Psi, \Psi'}^{\text{co}}$.

3.2. Proofs of examples

With the definition of $\llbracket L \rrbracket$ in place, we can finally revisit the linearisability claims anticipated in Examples 1, 16 and 18.

Recall the multiset library L_{mset} and the specification A_{mset} of Example 1 and Fig. 2. We show that $L_{\text{mset}} \triangleleft A_{\text{mset}}$. More precisely, taking an arbitrary history $h \in \llbracket L_{\text{mset}} \rrbracket$ we show that h can be rearranged using \triangleleft_{PO}^* to match an element of A_{mset} . We achieve this by identifying, for each O -move $(t, x)_O$ and its following P -move $(t, x')_P$ in h , a *linearisation point* between them, i.e. a place in h to which $(t, x)_O$ can be moved right and to which $(t, x')_P$ can be moved left so that they become consecutive and, moreover, the resulting history is still produced by L_{mset} . After all these rearrangements, we obtain a sequential history \hat{h} such that $h \triangleleft \hat{h}$ and \hat{h} is also produced by L_{mset} . It then suffices to show that $\hat{h} \in A_{\text{mset}}$.

Lemma 27 (Multiset). L_{mset} linearises to A_{mset} .

Proof. Given some $h \in \llbracket L_{\text{mset}} \rrbracket$, let us assume that h has been generated by a sequence $\rho_1 \Rightarrow \rho_2 \Rightarrow \dots \Rightarrow \rho_k$ of atomic transitions and that the variable F of L_{mset} is instantiated with a reference r_F . We demonstrate the linearisation points for pairs of (O, P) moves in h , by case analysis on the moves (we drop \mathcal{K} indices from moves as they are ubiquitous). Line numbers below refer to the LHS of Fig. 2:

1. $h = \dots (t, \text{call cnt}(i))_O s (t, \text{ret cnt}(i'))_P \dots$. Here the linearisation point (LP) is the configuration ρ_j that dereferences r_F as per line 5 in L_{mset} (the $!F$ expression).
2. $h = \dots (t, \text{call upd}(i, m))_O s (t, \text{call m}(j))_P \dots$. The LP is the dereferencing of r_F in line 5 (called from within *update*).
3. $h = \dots (t, \text{ret m}(j'))_O s (t, \text{ret upd}(|j'|))_P \dots$. The LP is the update of r_F in line 13.
4. $h = \dots (t, \text{ret m}(j'))_O s (t, \text{call m}(j''))_P \dots$. The LP is the dereferencing of r_F in line 11.

Each of the linearisation points above specifies a PO -rearrangement of moves. For instance, for $h = s_0 (t, \text{call cnt}(i))_O s (t, \text{ret cnt}(i'))_P s'$, let $s = s_1 s_2$ where $s_0 (t, \text{call cnt}(i))_O s_1$ is the prefix of h produced by $\rho_1 \Rightarrow \rho_2 \Rightarrow \dots \Rightarrow \rho_j$. The rearrangement of h is then $\hat{h} = s_0 s_1 (t, \text{call cnt}(i))_O (t, \text{ret cnt}(i'))_P s_2 s'$. We thus obtain $h \triangleleft_{PO}^* \hat{h}$.

The selection of linearisation points is such that it guarantees that $\hat{h} \in \llbracket L_{\text{mset}} \rrbracket$. E.g. in case 1, the transitions occurring in thread t between $(t, \text{call cnt}(i))_O$ and configuration ρ_j do not access r_F . Hence, we can postpone them and fire them in sequence just before ρ_j . After ρ_{j+1} and until $(t, \text{ret cnt}(i'))_P$ there is again no access of r_F in t and we can thus bring forward the corresponding transitions just after ρ_{j+1} . Similar reasoning applies to case 2. In case 4, we reason similarly but also take into account that rendering the acquisition of the lock by t atomic is sound (i.e. the semantics can produce the rearranged history). Case 3 is similar, but we also use the fact that the access to r_F in lines 10–15 is inside the lock, and hence postponing dereferencing (line 11) to occur in sequence before update (line 13) is sound.

Now, any transition sequence α which produces \hat{h} (in $\llbracket L_{\text{mult}} \rrbracket$) can be used to derive an annotated history $h^\circ \in A_{\text{mult}}^\circ$, by attaching to each move in \hat{h} the multiset represented in the configuration that produces the move (ρ produces the move x if $\rho \xrightarrow{x} \rho'$ in α). By projection we obtain $\hat{h} \in A_{\text{mult}}$. \square

Lemma 28 (Parameterised multiset). L_{mset2} enc-linearises to A_{mset2} .

Proof. Again, we identify linearisation points, this time for given $h \in \llbracket L_{\text{mult2}} \rrbracket_{\text{enc}}$. For cases 1–4 as above we reason as in Lemma 27. For *reset* we have the following case:

$$h = s (t, \text{call reset}(i))_{OK} s_1 (t, \text{call default}(j))_{PL} s_2 (t, \text{ret default}(j'))_{OL} s_3 (t, \text{ret reset}(|j'|))_{PK} \dots$$

Here, we need a linearisation point for all four moves above. We pick this to be the point corresponding to the update of the multiset reference F on lines 24–25 (Fig. 2, RHS). We now transform h to \hat{h} so that the four moves become consecutive, in two steps:

- Let us write s_3 as $s_3 = s_3^1 s_3^2$, where the split is at the linearisation point. Since the lock is constantly held by thread t in $s_2 s_3^1$, there can be no calls or returns to *default* in $s_2 s_3^1$. Hence, all moves in $s_2 s_3^1$ are in component \mathcal{K} and can be transposed with the \mathcal{L} -moves above, using \diamond^* , to obtain $h' = s (t, \text{call reset}(i))_{OK} s_1 s_2 s_3^1 (t, \text{call default}(j))_{PL} (t, \text{ret default}(j'))_{OL} s_3^2 (t, \text{ret reset}(|j'|))_{PK} \dots$
- Next, by PO -rearrangement we obtain $\hat{h} = s s_1 s_2 s_3^1 (t, \text{call reset}(i))_{OK} (t, \text{call default}(j))_{PL} (t, \text{ret default}(j'))_{OL} (t, \text{ret reset}(|j'|))_{PK} s_3^2 \dots$. Thus, $h(\triangleleft_{PO} \cup \diamond^*) \hat{h}$.

To prove that $\hat{h} \in A_{\text{mult2}}$ we work as in Lemma 27, i.e. via showing that $\hat{h} \in \llbracket L_{\text{mult2}} \rrbracket_{\text{enc}}$. For the latter, we rely on the fact that the linearisation point was taken at the reference update point (so that any dereferencings from other threads are preserved), and that the dereferencings of lines 22 and 23 are within the same lock as the update. \square

For our last example, recall the flat combination library $L_{\text{fc}} : \emptyset \rightarrow \Psi'$ of Example 18, and Fig. 4, along with its specification library $L_{\text{spec}} : \emptyset \rightarrow \Psi'$, where $\Psi' = \{\text{run} \in \text{Meths}_{(\theta \rightarrow \theta') \times \theta, \theta'}\}$.

Remark 29. It is worth observing that in the higher-order setting a client thread may try to call *run*, even though the previous call to *run* by the same thread did not complete yet. This scenario happens, for example, when the first call to *run* passes a functional argument to the library that itself calls *run*. Observe that in this case both L_{fc} and L_{spec} will deadlock. Consequently, non-trivial histories (all calls are matched by returns) arise only if each client thread uses *run* serially, i.e. without nesting.

Let $\mathcal{R} = <^*$, where $< \subseteq \mathcal{H}_{\emptyset, \Psi'} \times \mathcal{H}_{\emptyset, \Psi'}$ is the smallest relation such that (for economy we omit methods from calls/returns):

- $s_1(t, \text{call})_P s_2(t, \text{ret})_O s_3 < s_1(t', \text{call})_P s_2(t', \text{ret})_O s_3$
- $s_1(t, \text{call})_P s_2(t, \text{call})_O s_3(t, \text{ret})_P s_4(t, \text{ret})_O s_5 < s_1(t', \text{call})_P s_2(t', \text{call})_O s_3(t', \text{ret})_P s_4(t', \text{ret})_P s_5$

for any s_1, s_2, s_3, s_4, s_5 such that s_2, s_4 do not contain any *t*-moves.

Intuitively, $<$ is about piecewise delegation of client computations to other existing threads subject to forming a correct history. Because the results do not change, this condition corresponds to thread-blind client behaviour.

Lemma 30 (Flat combining). L_{fc} \mathcal{R} -linearises to L_{spec} .

Proof. Observe that histories from $\llbracket L_{\text{spec}} \rrbracket$ feature threads built from segments of one of the three forms:

- $(t, \text{call run}(f, x))_O (t, \text{call } f(x'))_P \dots (t, \text{ret } f(v))_O (t, \text{ret run}(v'))_P$, or
- $(t', \text{call } w(v))_O (t', \text{call } w'(v'))_P$, where w is a name introduced in an earlier move $(t'', x)_P$ and w' is a corresponding name introduced by the move preceding $(t'', x)_P$ in t'' , or
- $(t', \text{ret } w'(v'))_O (t', \text{ret } w(v'''))_P$ such that a segment $(t', \text{call } w(v))_O (t', \text{call } w'(v'))_P$ already occurred earlier.

The first shape represents interaction of the client with the library: a call to *run* followed by a call to f , possibly some intermediate computation (using calls/returns to higher-order values that have been introduced in the trace), and a return of f followed by a return of *run*. The value introduced in the last return may well be a function, which – along with method names introduced earlier – provides method names that can be used in calls and returns later. As these methods are related to concrete functions, our trace semantics interprets them in a symbolic manner: each call is forwarded to the move preceding the one in which it was introduced. Note that threads can exchange higher-order values, so we need to allow for scenarios in which the three kinds of interaction are located in different threads.

We shall refer to moves in the second and third kind of segments as *inspection moves* and write ϕ to refer to sequences built exclusively from such sequences. Note that \dots in the first kind of block also stand for a segment of inspection moves in t .

Let us write \mathcal{X} for the subset of $\llbracket L_{\text{spec}} \rrbracket$ containing (sequential) plays of the form:

$$\begin{aligned} & (t_0, \text{call run}(f_0, x_0))(t_0, \text{call } f_0(x'_0))\phi_0(t_0, \text{ret } f_0(v_0))(t_0, \text{ret run}(v'_0))\phi_1 \\ & (t_1, \text{call run}(f_1, x_1))(t_1, \text{call } f_1(x'_1))\phi_2(t_1, \text{ret } f_1(v_1))(t_1, \text{ret run}(v'_1))\phi_3 \\ & \dots (t_k, \text{call run}(f_k, x_k))(t_k, \text{call } f_k(x'_k))\phi_{2k}(t_k, \text{ret } f_k(v_k))(t_k, \text{ret run}(v'_k))\phi_{2k+1}, \end{aligned}$$

where ϕ_{2j}, ϕ_{2j+1} may also contain inspection moves not in t_j . We take \mathcal{X} to be our linearisation target (specification).

Consider $h_1 \in \llbracket L_{\text{fc}} \rrbracket$. Threads in h_1 are built from blocks of shapes:

$$\begin{aligned} & (t, \text{call run}(f, x))_O ((t, \text{call } f_j(x'_j))_P \phi_j(t, \text{ret } f_j(v_j))_O)^* (t, \text{ret run}(v'))_P \\ & \text{or } (t', \text{call } w(v))_O (t', \text{call } w'(v'))_P \text{ or } (t', \text{ret } w'(v'''))_O (t', \text{ret } w(v'''))_P. \end{aligned}$$

In the first case, the j 's are meant to represent possibly different values used in each iteration. In the second kind of block, w needs to be introduced earlier by some $(t'', x)_P$ move and w' is then a name introduced by the preceding move. For the third kind, an earlier calling sequence of the second kind must exist in the same thread.

Observe that each segment $S_j = (t, \text{call } f_j(x'_j))_P \phi_j(t, \text{ret } f_j(v_j))_O$ in t must be preceded (in h_1) by a matching public call $(t', \text{call run}(f_j, x_j))_O$ followed by a corresponding return $(t', \text{ret run}(v_j))_P$, where t' need not be equal to t . We can obtain

the requisite h (for $\triangleleft_{\mathcal{R}}$) by changing t to t' in the whole of S_j for each S_j . Note that run-moves are not affected and we get $h_1 \mathcal{R}^* h$.

Note that, due to locking and sequentiality of loops, the segments S_j must be disjoint in h_1 , although they may be interleaved with inspection moves from other threads. We shall show how to obtain $h_2 \in \mathcal{X}$ with $h \triangleleft_{pO}^* h_2$:

- First the call to run associated with each S_j should be moved right to immediately precede the renamed S_j . Next the corresponding return of run should be move left to follow S_j .
- Subsequently, inspection moves need to be rearranged to yield a sequential play. This can be done by permuting inspection moves by O to the left through other O actions from different threads until a P -move is encountered and moving the corresponding inspection move P left to immediately follow the O move.

Then we have $h \triangleleft_{pO}^* h_2$ and, hence, $h_1(\triangleleft_{pO} \cup \mathcal{R})^* h_2$. \square

4. Soundness

To conclude, we clarify in what sense all the notions of linearisability are sound. Recall the general notion of contextual approximation (refinement) from Definition 22. In the encapsulated case libraries are being tested by clients that do not communicate with the parameter library explicitly. The corresponding definition of contextual approximation is defined below.

Definition 31 (*Encapsulated \sqsubseteq*). Given libraries $L_1, L_2 : \Psi \rightarrow \Psi'$, we write $L_1 \sqsubseteq_{\text{enc}} L_2$ when, for all $L' : \emptyset \rightarrow \Psi$ and $\Psi' \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}$, if $\text{link } L'; L_1$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$ then $\text{link } L'; L_2$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$.

For relational linearisability, we need yet another notion that will link \mathcal{R} to contextual testing.

Definition 32. Let $\mathcal{R} \subseteq \mathcal{H}_{\Psi, \Psi'} \times \mathcal{H}_{\Psi, \Psi'}$ be a set closed under permutation of names in $\text{Meths} \setminus (\Psi \cup \Psi')$. We say that a context formed by L' and M_1, \dots, M_N is \mathcal{R} -closed if, for any $h \in \llbracket \text{link } L'; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket$, $\bar{h} \mathcal{R} \bar{h}'$ implies $h' \in \llbracket \text{link } L'; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket$. Given $L_1, L_2 : \Psi \rightarrow \Psi'$, we write $L_1 \sqsubseteq_{\mathcal{R}} L_2$ if, for all \mathcal{R} -closed contexts formed from L', M_1, \dots, M_N , whenever $\text{link } L'; L_1$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$ then we also have $\text{link } L'; L_2$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$.

In what follows, we shall aim to establish three correctness results:

- $L_1 \triangleleft L_2$ implies $L_1 \sqsubseteq L_2$,
- $L_1 \triangleleft_{\text{enc}} L_2$ implies $L_1 \sqsubseteq_{\text{enc}} L_2$, and
- $L_1 \triangleleft_{\mathcal{R}} L_2$ implies $L_1 \sqsubseteq_{\mathcal{R}} L_2$.

Finally, we note that linearisability is compatible with library composition. \triangleleft is closed under union with libraries that use disjoint stores, while $\triangleleft_{\text{enc}}$ is closed under a form of sequencing that respects encapsulations (Appendix E).

4.1. Correctness

In this section we prove that the linearisability notions we introduce are correct: linearisability implies contextual approximation. The approach is based on showing that, in each case, the semantics of contexts is saturated relatively to conditions that are dual to linearisability. Hence, linearising histories does not alter the observable behaviour of a library. We start by presenting two compositionality theorems on the trace semantics, which will be used for relating library and context semantics.

4.2. Compositionality

The semantics we defined is compositional in the following ways:

- To compute the semantics of a library L inside a context C , it suffices to compose the semantics of C with that of L , for a suitable notion of context-library composition ($\llbracket C \rrbracket \otimes \llbracket L \rrbracket$).
- To compute the semantics of a union library $L_1 \cup L_2$, we can compose the semantics of L_1 and L_2 , for a suitable notion of library-library composition ($\llbracket L_1 \rrbracket \otimes \llbracket L_2 \rrbracket$).

The above are proven using bisimulation techniques for connecting syntactic and semantic compositions, and are presented in Appendix C and Appendix D respectively.

The latter correspondence is used in Appendix E for proving that linearisability is a congruence for library composition. From the former correspondence we obtain the following result, which we shall use for showing correctness.

```

1  public run;
2  Lock lock;
3  r := 0;
4
5  run = λ ().
6    lock.acquire();
7    r := !r + 1;
8    if (!r = 1) then lock.release();
9    while (!r < 2) do ();

```

Fig. 8. A library without a sequential history.

Theorem 33. Let $L : \Psi \rightarrow \Psi'$, $L' : \emptyset \rightarrow \Psi$, Ψ'' and $\Psi', \Psi'' \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}$, with L , L' and $M_1; \dots; M_N$ accessing pairwise disjoint parts of the store. Then:

$$\text{link } L' ; L \text{ in } (M_1 \parallel \dots \parallel M_N) \Downarrow \iff \exists h \in \llbracket L \rrbracket_N. \bar{h} \in \llbracket \text{link } L' ; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket$$

4.3. General linearisability

Recall the general notion of linearisability defined in Section 2.2, which is based on move-reorderings inside histories.

In Definitions 24 and 25 we have defined the trace semantics of libraries and contexts. The semantics turns out to be closed under \triangleleft_{OP}^* .

Lemma 34 (Saturation). Let $X = \llbracket L \rrbracket$ (Definition 24) or $X = \llbracket \text{link } L' ; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket$ (Definition 25). Then if $h \in X$ and $h \triangleleft_{OP}^* h'$ then $h' \in X$.

Proof. Recall that the same labelled transition system underpins the definition of X in either case. We make several observations about the single-threaded part of that system.

- The store is examined and modified only during ϵ -transitions.
- The only transition possible after a P -move is an O -move. In particular, it is never the case that a P -move is separated from the following O -move by an ϵ -transition.

Let us now consider the multi-threaded system and $t \neq t'$.

- Suppose $\rho \xrightarrow{(t', m')_P} \rho_1 \xrightarrow{\epsilon^*} \rho_2 \xrightarrow{(t, m)} \rho_3$. Then the $(t', m')_P$ -transition can be delayed inside t' until after (t, m) , i.e. $\rho \xrightarrow{\epsilon^*} \rho'_1 \xrightarrow{(t, m)} \rho'_2 \xrightarrow{(t', m')_P} \rho_3$ for some ρ'_1, ρ'_2 . This is possible because the $((t', m')_P$ -labelled) transition does not access or modify the store, and none of the ϵ -transitions distinguished above can be in t' , thanks to our earlier observations about the behaviour of the single-threaded system.
- Analogously, suppose $\rho \xrightarrow{(t', m')_O} \rho_1 \xrightarrow{\epsilon^*} \rho_2 \xrightarrow{(t, m)_O} \rho_3$. Then the $(t, m)_O$ -transition can be brought forward, i.e. $\rho \xrightarrow{(t, m)_O} \rho'_1 \xrightarrow{(t', m')_O} \rho'_2 \xrightarrow{\epsilon^*} \rho_3$, because it does not access or modify the store and the preceding ϵ -transitions cannot be from t . \square

This, along with the fact that

$$h_1 \triangleleft_{XX'} h_2 \iff h_2 \triangleleft_{X'X} h_1 \iff \bar{h}_1 \triangleleft_{X'X} \bar{h}_2$$

lead us to the notion of linearisability defined in Definition 9.

We now prove the main theorem of this subsection.

Theorem 35. $L_1 \triangleleft L_2$ implies $L_1 \sqsubseteq L_2$.

Proof. Consider C such that $C[L_1] \Downarrow$. We need to show $C[L_2] \Downarrow$. Because $C[L_1] \Downarrow$, Theorem 33 implies that there exists $h_1 \in \llbracket L_1 \rrbracket$ such that $\bar{h}_1 \in \llbracket C \rrbracket$. Because $L_1 \triangleleft L_2$, there exists $h_2 \in \llbracket L_2 \rrbracket$ with $h_1 \triangleleft_{PO}^* h_2$. Note that $\bar{h}_1 \triangleleft_{OP}^* \bar{h}_2$. By Lemma 34, $\bar{h}_2 \in \llbracket C \rrbracket$. Because $h_2 \in \llbracket L_2 \rrbracket$ and $\bar{h}_2 \in \llbracket C \rrbracket$, using Theorem 33 we can conclude $C[L_2] \Downarrow$. \square

Remark 36. A natural question to ask is whether the converse of Theorem 35 is true. The answer is negative and can be traced back to the fact that \triangleleft is defined using sequential histories: in order to establish $L_1 \triangleleft L_2$ (for $L_1, L_2 : \Psi \rightarrow \Psi'$) one needs to identify a subset $A_2 \subseteq \llbracket L_2 \rrbracket \cap \mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$ (i.e. consisting of *sequential* histories only) such that $L_1 \triangleleft A_2$.

Unfortunately, some libraries generate only non-sequential histories. We present an example of such a library, call it L , in Fig. 8. Because of locks, the library from Fig. 8 will only allow two threads to complete a computation. Additionally, the first thread (i.e. the one that will increment r to 1) must wait until a second thread increments the internal counter r to 2.

Observe that if L does not generate any sequential histories then we vacuously have $L \sqsubseteq L$, but cannot have $L \triangleleft L$. We conjecture that a completeness result would be possible if we allowed for non-sequential specs in the definition of \triangleleft .

4.4. Encapsulated linearisability

In this case libraries are being tested by clients that do not communicate with the parameter library explicitly. Recall from Definition 31 that, given libraries $L_1, L_2 : \Psi \rightarrow \Psi'$, we write $L_1 \sqsubseteq_{\text{enc}} L_2$ when, for all $L' : \emptyset \rightarrow \Psi$ and $\Psi' \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}$, if $\text{link } L'; L_1$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$ then $\text{link } L'; L_2$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$.

We call contexts of the above kind *encapsulated*, because the parameter library L' can no longer communicate directly with the client, unlike in Definition 22, where they shared methods in Ψ'' . Consequently, $\llbracket \text{link } L'; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket$ can be decomposed via parallel composition into two components, whose labels correspond to \mathcal{L} (parameter library) and \mathcal{K} (client) respectively.

Lemma 37 (Decomposition). Suppose $L' : \emptyset \rightarrow \Psi$ and $\Psi' \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}$, where $\Psi \cap \Psi' = \emptyset$. Then, setting $C' \equiv \text{link } \emptyset; - \text{ in } (M_1 \parallel \dots \parallel M_N)$, we have:

$$\llbracket \text{link } L'; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket = \{ h \in \mathcal{H}_{\Psi, \Psi'}^{\text{co}} \mid (h \upharpoonright \mathcal{L}) \in \llbracket L' \rrbracket, (h \upharpoonright \mathcal{K}) \in \llbracket C' \rrbracket \}.$$

Remark 38. Consider parameter library $L' : \emptyset \rightarrow \{m\}$ and client $\{m'\} \vdash_K M : \text{unit}$ with $m, m' \in \text{Meths}_{\text{unit} \rightarrow (\text{unit} \rightarrow \text{unit})}$, and suppose we insert in their context a “copycat” library L which implements m' as $m' = \lambda x.mx$. Then the following scenario may seem to contradict encapsulation:

- M calls $m'()$;
- L calls $m()$;
- L' returns with $m(m'')$ to L ;
- and finally L copycats this return to M .

However, by definition the latter copycat is done by L returning $m'(m''')$ to M , for some *fresh* name m''' , and recording internally that $m''' \mapsto \lambda x.m''x$. Hence, no methods of L' can leak to M and encapsulation holds.

Because of the above decomposition, the context semantics satisfies a stronger closure property than that already specified in Lemma 34, which in turn leads to the notion of encapsulated linearisability of Definition 13. The latter is defined in term of the symmetric reordering relation \diamond , which allows for swaps (in either direction) between moves from different threads if they are tagged with \mathcal{K} and \mathcal{L} respectively.

Moreover, we can show the following:

Lemma 39 (Encapsulated saturation). Consider $X = \llbracket \text{link } L'; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket$ (Definition 25). Then:

- If $h \in X$ and $h \triangleleft_{OP} \cup \diamond^* h'$ then $h' \in X$.
- Let $s_1(t, x)_{OY} s_2(t, x')_{PY'} s_3 \in X$ be such that no move in s_2 comes from thread t . Then $Y = Y'$, i.e. inside a thread only O can switch between \mathcal{K} and \mathcal{L} .

Proof. For the first claim, closure under \triangleleft_{OP} (resp. \diamond) follows from Lemma 34. (resp. Lemma 37).

Suppose $h = s_1(t, x)_{OY} s_2(t, x')_{PY'} s_3$ violates the second claim and $(t, x), (t, x')$ is the earliest such violation in h , i.e. no violations occur in s_1 . Observe that then h restricted to moves of the form $(t, z)_{XY'}$ would not be alternating, which contradicts the fact that $h \upharpoonright Y'$ is a history (Lemma 37). \square

Due to Theorem 33, the above property of contexts means that, in order to study termination in the encapsulated case, one can safely restrict attention to library traces satisfying a dual property to the one above, i.e. to elements of $\llbracket L \rrbracket_{\text{enc}}$. Note that $\llbracket L \rrbracket_{\text{enc}}$ can be obtained directly from our labelled transition system by restricting its single-threaded part to reflect the switching condition. Observe that Theorem 33 will still hold for $\llbracket L \rrbracket_{\text{enc}}$ (instead of $\llbracket L \rrbracket$), because we have preserved all the histories that are compatible with context histories. We are ready to prove correctness of encapsulated linearisability.

Theorem 40. $L_1 \triangleleft_{\text{enc}} L_2$ implies $L_1 \sqsubseteq_{\text{enc}} L_2$.

Proof. Similarly to Theorem 35, except we invoke Lemma 39 instead of Lemma 34. \square

4.5. Relational linearisability

Finally, we examine relational linearisability (Definition 17).

Theorem 41. $L_1 \triangleleft_{\mathcal{R}} L_2$ implies $L_1 \sqsubseteq_{\mathcal{R}} L_2$.

Proof. Consider \mathcal{R} -closed C such that $C[L_1] \Downarrow$. We need to show $C[L_2] \Downarrow$. Because $C[L_1] \Downarrow$, Theorem 33 implies that there exists $h_1 \in \llbracket L_1 \rrbracket$ such that $\overline{h_1} \in \llbracket C \rrbracket$. Because $L_1 \triangleleft_{\mathcal{R}} L_2$, there exists $h_2 \in \llbracket L_2 \rrbracket$ such that $h_1 (\triangleleft_{PO} \cup \mathcal{R})^* h_2$. Because C is \mathcal{R} -closed by definition and closed under \triangleleft_{OP} by Lemma 34, we have $\overline{h_2} \in \llbracket C \rrbracket$. Because $h_2 \in \llbracket L_2 \rrbracket$ and $\overline{h_2} \in \llbracket C \rrbracket$, we can conclude $C[L_2] \Downarrow$. \square

5. Related and future work

Linearisability has been consistently used as a correctness criterion for concurrent algorithms on a variety of data structures [15], and has inspired a variety of proof methods [16]. An explicit connection between linearisability and refinement was made in [17], where it was shown that, in base-type settings, linearisability and refinement coincide. Similar results have been proved in [18–20,3]. Our contributions are notions of linearisability that serve as correctness criteria for libraries with methods of arbitrary order and have a similar relationship to refinement. The next natural target is to investigate proof methods for establishing linearisability of higher-order concurrent libraries. The examples proved herein are only an initial step in that direction.

At the conceptual level, [17] proposed that the verification goal behind linearisability is observational refinement. In this vein, [21] utilised logical relations as a direct method for proving refinement in a higher-order concurrent setting, while [22] introduced a program logic that builds on logical relations. On the other hand, proving conformance to a history specification has been addressed in [23] by supplying history-aware interpretations to off-the-shelf Hoare logics for concurrency. Other logic-based approaches for concurrent higher-order libraries, which do not use linearisability, include Higher-Order and Impredicative Concurrent Abstract Predicates [24,25].

Acknowledgements

We thank the authors of [3] for bringing the higher-order linearisability problem to our attention. We would also like to thank Kasper Svendsen and Radha Jagadeesan for constructive comments. This work was partially funded by the Engineering and Physical Sciences Research Council (EP/P004172/1) and a Royal Society Leverhulme Trust Senior Research Fellowship.

Appendix A. Big-step vs small-step reorderings

[3] defines linearisation in the general case using a “big-step” relation that applies a single permutation to the whole sequence. This contrasts with our definition as \triangleleft_{PO}^* , in which we combine multiple adjacent swaps. We show that the two definitions are equivalent.

Definition 42 ([3]). Let $h_1, h_2 \in \mathcal{H}_{\psi, \psi'}$ of equal length. We write $h_1 \triangleleft_{PO}^{\text{big}} h_2$ if there is a permutation $\pi : \{1, \dots, |h_1|\} \rightarrow \{1, \dots, |h_2|\}$ such that, writing $h_i(j)$ for the j -th element of h_i : for all j , we have $h_1(j) = h_2(\pi(j))$ and, for all $i < j$:

$$((\exists t. h_1(i) = (t, -) \wedge h_1(j) = (t, -)) \vee (\exists t_1, t_2. h_1(i) = (t_1, -)_P \wedge h_1(j) = (t_2, -)_O)) \implies h_2(i) < h_2(j)$$

In other words, h_2 is obtained from h_1 by permuting moves in such a way that their order in threads is preserved and whenever a O -move occurred after an P -move in h_1 , the same must apply to their permuted copies in h_2 .

Lemma 43. $\triangleleft_{PO}^{\text{big}} = \triangleleft_{PO}^*$.

Proof. It is obvious that $\triangleleft_{PO}^* \subseteq \triangleleft_{PO}^{\text{big}}$, so it suffices to show the converse.

Suppose $h_1 \triangleleft_{PO}^{\text{big}} h_2$. Consider the set $X_{h_1, h_2} = \{h \mid h_1 \triangleleft_{PO}^* h, h \triangleleft_{PO}^{\text{big}} h_2\}$. Note that X_{h_1, h_2} is not empty, because $h_1 \in X_{h_1, h_2}$. For two histories h', h'' , define $\delta(h', h'')$ to be the length of the longest common prefix of h' and h'' . Let $N = \max_h \{\delta(h, h_2) \mid h \in X_{h_1, h_2}\}$. Note that $N \leq |h_1| = |h_2|$.

- If $N = |h_2|$ then we are done, because $N = |h_2|$ implies $h_2 \in X_{h_1, h_2}$ and, thus, $h_1 \triangleleft_{PO}^* h_2$.
- Suppose $N < |h_2|$ and consider h such that $N = \delta(h, h_2)$. We are going to arrive at a contradiction by exhibiting $h' \in X_{h_1, h_2}$ such that $\delta(h', h_2) > N$.

Because $N = \delta(h, h_2)$ and $N < |h_2|$, we have

$$\begin{aligned} h_2 &= a_1 \cdots a_N(t, m)u \\ h &= a_1 \cdots a_N(t_1, m_1) \cdots (t_k, m_k)(t, m)u', \end{aligned}$$

where $t_i \neq t$, because order in threads must be preserved. Consider

$$h' = a_1 \cdots a_N(t, m)(t_1, m_1) \cdots (t_k, m_k)u'.$$

Clearly $\delta(h', h_2) > N$ so, for a contradiction, it suffices to show that $h' \in X_{h_1, h_2}$. Note that because $h \prec_{PO}^{\text{big}} h_2$, we must also have $h' \prec_{PO}^{\text{big}} h_2$, because the new PO dependencies in h' (wrt h) caused by moving (t, m) forward are consistent with h_2 . Hence, we only need to show that $h \prec_{PO}^* h'$. Let us distinguish two cases.

- If (t, m) is a P -move then, clearly, $h \prec_{PO}^* h'$ (P -move moves forward).
- If (t, m) is an O -move then, because $h \prec_{PO}^{\text{big}} h_2$, all of the (t_i, m_i) actions must be O -moves (otherwise their position wrt (t, m) would have to be preserved in h_2 and it isn't). Hence, $h \prec_{PO}^* h'$, as required. \square

Appendix B. Auxiliary lemmas about histories

Recall the notions of history and history complementation (Definition 3). We next define a dual notion of history that is used for assigning semantics to contexts.

Definition 44. The set of **co-histories** over $\Psi \rightarrow \Psi'$ is: $\mathcal{H}_{\Psi, \Psi'}^{\text{co}} = \{\bar{h} \mid h \in \mathcal{H}_{\Psi, \Psi'}\}$.

We shall range over $\mathcal{H}_{\Psi, \Psi'}^{\text{co}}$ again using variables h, s . We can show the following.

Lemma 45.

- For all $h \in \mathcal{H}_{\Psi, \Psi'}$ we have $h \upharpoonright \mathcal{L} \in \mathcal{H}_{\emptyset, \Psi}^{\text{co}}$ and $h \upharpoonright \mathcal{K} \in \mathcal{H}_{\emptyset, \Psi'}$.
- For all $h \in \mathcal{H}_{\Psi, \Psi'}^{\text{co}}$ we have $h \upharpoonright \mathcal{L} \in \mathcal{H}_{\emptyset, \Psi}$ and $h \upharpoonright \mathcal{K} \in \mathcal{H}_{\emptyset, \Psi'}^{\text{co}}$.

Lemma 46. For any $L : \Psi \rightarrow \Psi'$, $L' : \emptyset \rightarrow \Psi, \Psi''$ and $\Psi', \Psi'' \vdash_K M_1 \parallel \cdots \parallel M_N : \text{unit}$ we have $\llbracket L \rrbracket_N \subseteq \mathcal{H}_{\Psi, \Psi'}$ and $\llbracket \text{link } L'; - \text{ in } (M_1 \parallel \cdots \parallel M_N) \rrbracket \subseteq \mathcal{H}_{\Psi, \Psi'}^{\text{co}}$.

Proof. The relevant sequences of moves are clearly alternating and well-bracketed, when projected on single threads, because the LTS is bipartite (O - and P -configurations) and separate evaluation stacks control the evolution in each thread. Other conditions for histories follow from the partitioning of names into $\mathcal{A}_{\mathcal{K}}, \mathcal{A}_{\mathcal{L}}, \mathcal{P}_{\mathcal{K}}, \mathcal{P}_{\mathcal{L}}$ and suitable initialisation: Ψ, Ψ' are inserted into $\mathcal{A}_{\mathcal{L}}, \mathcal{P}_{\mathcal{K}}$ respectively (for $\llbracket L \rrbracket$) and into $\mathcal{P}_{\mathcal{L}}, \mathcal{A}_{\mathcal{K}}$ for $\llbracket C \rrbracket$. \square

Appendix C. Trace compositionality

In this section we demonstrate how the semantics of a library inside a context can be drawn by composing the semantics of the library and that of the context. The result played a crucial role in our arguments about linearisability and contextual refinement in Section 4.1.

Let us divide (reachable) evaluation stacks into two classes: L -stacks, which can be produced in the trace semantics of a library; and C -stacks, which appear in traces of a context.

$$\begin{aligned} \mathcal{E}_L &::= [] \mid m :: \mathcal{E}'_L & \mathcal{E}_C &::= [] \mid m :: \mathcal{E}'_C \\ \mathcal{E}'_L &::= m :: \mathcal{E}_L & \mathcal{E}'_C &::= m :: \mathcal{E}_C \end{aligned}$$

From the trace semantics definition we have that N -configurations in the semantics of a library feature evaluation stacks of the forms \mathcal{E}_L (in O -configurations) and \mathcal{E}'_L (in P -configurations): these we will call **L -stacks**. On the other hand, those produced from a context utilise **C -stacks** which are of the forms \mathcal{E}_C (in P -configurations) and \mathcal{E}'_C (in O -configurations).

From here on, when we write \mathcal{E} we will mean an L -stack or a C -stack. Moreover, we will call an N -configuration ρ an **L -configuration** (or a **C -configuration**), if $\rho = (\vec{C}, \dots)$ and, for each i , $C_i = (\mathcal{E}_i, \dots)$ with \mathcal{E}_i an L -stack (resp. a C -stack).

Let ρ, ρ' be N -configurations and suppose $\rho = (\vec{C}, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$ is a C -configuration and $\rho' = (\vec{C}', \mathcal{R}', \mathcal{P}', \mathcal{A}', S')$ an L -configuration. We say that ρ and ρ' are **compatible**, written $\rho \succ \rho'$, if S and S' have disjoint domains and, for each i :

- $C_i = (\mathcal{E}_C, M)$ and $C'_i = (\mathcal{E}_L, -)$, or $C_i = (\mathcal{E}'_C, -)$ and $C'_i = (\mathcal{E}'_L, M)$.

- If the public and abstract names of \mathcal{C}_i are $(\mathcal{P}_{\mathcal{L}}, \mathcal{P}_{\mathcal{K}})$ and $(\mathcal{A}_{\mathcal{L}}, \mathcal{A}_{\mathcal{K}})$ respectively, and those of \mathcal{C}'_i are $(\mathcal{P}'_{\mathcal{L}}, \mathcal{P}'_{\mathcal{K}})$ and $(\mathcal{A}'_{\mathcal{L}}, \mathcal{A}'_{\mathcal{K}})$, then $\mathcal{P}_{\mathcal{L}} = \mathcal{A}'_{\mathcal{L}}$, $\mathcal{P}_{\mathcal{K}} = \mathcal{A}'_{\mathcal{K}}$, $\mathcal{A}_{\mathcal{L}} = \mathcal{P}'_{\mathcal{L}}$ and $\mathcal{A}_{\mathcal{K}} = \mathcal{P}'_{\mathcal{K}}$.
- The private names of ρ (i.e. those in $\text{dom}(\mathcal{R}) \setminus \mathcal{P}_{\mathcal{L}} \setminus \mathcal{P}_{\mathcal{K}}$) do not appear in ρ' , and dually for the private names of ρ' .
- If $\mathcal{C}_i = (\mathcal{E}, \dots)$ and $\mathcal{C}'_i = (\mathcal{E}', \dots)$ then \mathcal{E} and \mathcal{E}' are in turn compatible, that is:

- either $\mathcal{E} = m :: E :: \mathcal{E}_1$, $\mathcal{E}' = m :: \mathcal{E}'_1$ and $\mathcal{E}_1, \mathcal{E}'_1$ are compatible,
- or $\mathcal{E} = m :: \mathcal{E}_1$, $\mathcal{E}' = m :: E :: \mathcal{E}'_1$ and $\mathcal{E}_1, \mathcal{E}'_1$ are compatible,

or $\mathcal{E} = \mathcal{E}' = []$.

Note, in particular, that if $\rho \simeq \rho'$ then ρ must be a context configuration, and ρ' a library configuration.

We next define a trace semantics on compositions of compatible such N -configurations. We use the symbol \odot for configuration composition: we call this **external composition**, to distinguish it from the composition of ρ and ρ' we can obtain by merging their components, which we will examine later.

$$\begin{array}{c}
\frac{\rho_1 \Rightarrow \rho'_1}{\rho_1 \odot \rho_2 \longrightarrow \rho'_1 \odot \rho_2} \text{INT}_1 \quad \frac{\rho_1 \xrightarrow{(t, \text{call } m(v))} \rho'_1 \quad \rho_2 \xrightarrow{(t, \text{call } m(v))} \rho'_2}{\rho_1 \odot \rho_2 \longrightarrow \rho'_1 \odot \rho'_2} \text{CALL} \\
\frac{\rho_2 \Rightarrow \rho'_2}{\rho_1 \odot \rho_2 \longrightarrow \rho_1 \odot \rho'_2} \text{INT}_2 \quad \frac{\rho_1 \xrightarrow{(t, \text{ret } m(v))} \rho'_1 \quad \rho_2 \xrightarrow{(t, \text{ret } m(v))} \rho'_2}{\rho_1 \odot \rho_2 \longrightarrow \rho'_1 \odot \rho'_2} \text{RETN}
\end{array}$$

The INT rules above have side-conditions imposing that the resulting pairs of configurations are still compatible. Concretely, this means that the names created fresh in internal transitions do not match the names already present in the configurations of the other component. Note that external composition is not symmetric, due to the context/library distinction we mentioned.

Our next target is to show a correspondence between the above-defined semantic composition and the semantics obtained by (syntactically) merging compatible configurations. This will demonstrate that composing the semantics of two components is equivalent to first syntactically composing them and then evaluating the result. In order to obtain this correspondence, we need to make the semantics of syntactically composed configurations more verbose: in external composition methods belong either to the context or the library, and when e.g. the client wants to evaluate mm' , with m a library method, the call is made explicit and, more importantly, m' is replaced by a fresh method name. On the other hand, when we compose syntactically such a call will be done internally, and without refreshing m' .

To counter-balance the above mismatch, we extend the syntax of terms and evaluation contexts, and the operational semantics of closed terms as follows. The semantics will now involve quadruples of the form:

$$(E[M], \mathcal{R}_1, \mathcal{R}_2, S) \text{ written also } (E[M], \vec{\mathcal{R}}, S)$$

where the two repositories correspond to context and library methods respectively, so in particular $\text{dom}(\mathcal{R}_1) \cap \text{dom}(\mathcal{R}_2) = \emptyset$. Moreover, inside $E[M]$ we tag method names and lambda-abstractions with indices 1 and 2 to record which of the two components (context or library) is enclosing them: the tag 1 is used for the context, and 2 for the library. Thus e.g. a name m^1 signals an occurrence of method m inside the context. Tagged methods are passed around and stored as ordinary methods, but their behaviour changes when they are applied. Moreover, we extend (tagged) evaluation contexts by explicitly marking return points of methods:

$$E ::= \bullet \mid \dots \mid \text{let } x = E \text{ in } M \mid mE \mid r := E \mid \langle m^i \rangle E$$

In particular, $E[M]$ may not necessarily be a (tagged) term, due to the return annotations. The new reduction rules are as follows (we omit indices when they are not used in the rules).

$$\begin{array}{l}
(E[i_1 \oplus i_2], \vec{\mathcal{R}}, S) \rightarrow'_t (E[i], \vec{\mathcal{R}}, S') \quad (i = i_1 \oplus i_2) \\
(E[\text{tid}], \vec{\mathcal{R}}, S) \rightarrow'_t (E[t], \vec{\mathcal{R}}, S') \\
(E[\pi_j \langle v_1, v_2 \rangle], \vec{\mathcal{R}}, S) \rightarrow'_t (E[v_j], \vec{\mathcal{R}}, S') \\
(E[\text{if } i \text{ then } M_0 \text{ else } M_1], \vec{\mathcal{R}}, S) \rightarrow'_t (E[M_j], \vec{\mathcal{R}}, S) \quad (j = (i > 0)) \\
(E[\lambda^i x.M], \vec{\mathcal{R}}, S) \rightarrow'_t (E[m^i], \vec{\mathcal{R}} \uplus_i (m \mapsto \lambda x.M), S) \\
(E[m^i v], \vec{\mathcal{R}}, S) \rightarrow'_t (E[M\{v/x\}^i], \vec{\mathcal{R}}, S) \quad \text{if } \mathcal{R}_i(m) = \lambda x.M \\
(E[m^i v], \vec{\mathcal{R}}, S) \rightarrow'_t (E[\langle m^i \rangle M\{v'/x\}^{3-i}], \vec{\mathcal{R}}', S) \quad \text{if } \mathcal{R}_{3-i}(m) = \lambda x.M \text{ with} \\
\text{Meths}(v) = \{m_1, \dots, m_k\}, v' = v\{m'_j/m_j \mid 1 \leq j \leq k\}, \vec{\mathcal{R}}' = \vec{\mathcal{R}} \uplus_i \{m'_j \mapsto \lambda y.m_j y \mid 1 \leq j \leq k\}
\end{array}$$

$$\begin{aligned}
(E[\langle m^i \rangle v], \vec{\mathcal{R}}, S) &\rightarrow'_t (E[v'^i], \vec{\mathcal{R}} \uplus_{3-i} \{m'_j \mapsto \lambda y. m_j y\}, S) \text{ with } m_j, m'_j \text{ and } v' \text{ as above} \\
(E[\text{let } x = v \text{ in } M], \vec{\mathcal{R}}, S) &\rightarrow'_t (E[M\{v/x\}], \vec{\mathcal{R}}, S) \\
(E[!r], \vec{\mathcal{R}}, S) &\rightarrow'_t (E[S(r)], \vec{\mathcal{R}}, S) \\
(E[r := i], \vec{\mathcal{R}}, S) &\rightarrow'_t (E, \vec{\mathcal{R}}, S[r \mapsto i]) \\
(E[r := m^i], \vec{\mathcal{R}}, S) &\rightarrow'_t (E, \vec{\mathcal{R}}, S[r \mapsto m^i])
\end{aligned}$$

Above we write M^i for the term M with all its methods and lambdas tagged (or re-tagged) with i . Moreover, we use the convention e.g. $\vec{\mathcal{R}} \uplus_1 (m \mapsto \lambda x. M) = (\mathcal{R}_1 \uplus (m \mapsto \lambda x. M), \mathcal{R}_2)$. Note that the repositories need not contain tags as, whenever a method is looked up, we subsequently tag its body explicitly.

Thus, the computationally observable difference of the new semantics is in the rule for reducing $E[m^i v]$ when m is not in the domain of \mathcal{R}_i : this corresponds precisely to the case where e.g. a library method is called by the context with another method as argument. A similar behaviour is exposed when such a method is returning. However, this novelty merely adds fresh method names by η -expansions and does not affect the termination of the reduction.

Defining parallel reduction \Rightarrow' in an analogous way to \Rightarrow , we can show the following. We let a quadruple $(M_1 \parallel \dots \parallel M_N, \mathcal{R}, S)$ be *final* if $M_i = ()$ for all i , and we write $(M_1 \parallel \dots \parallel M_N, \mathcal{R}, S) \Downarrow$ if $(M_1 \parallel \dots \parallel M_N, \mathcal{R}, S)$ can reduce to some final quadruple; these notions are defined for $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S)$ in the same manner.

Lemma 47. *For any legal $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S)$, we have that $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S) \Downarrow$ iff $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1 \cup \mathcal{R}_2, S) \Downarrow$.*

We now proceed to syntactic composition of N -configurations. Given a pair $\rho_1 \succ \rho_2$, we define a single quadruple corresponding to their syntactic composition, called their **internal composition**, as follows. Let $\rho_1 = (\vec{\mathcal{C}}, \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1)$ and $\rho_2 = (\vec{\mathcal{C}}', \mathcal{R}_2, \mathcal{P}_2, \mathcal{A}_2, S_2)$ and, for each i , $\mathcal{C}_i = (\mathcal{E}_i, X_i)$ and $\mathcal{C}'_i = (\mathcal{E}'_i, X'_i)$, with $\{X_i, X'_i\} = \{M_i, -\}$, and we let $k_i = 1$ just if $X_i = M_i$. We let the internal composition of ρ_1 and ρ_2 be the quadruple:

$$\rho_1 \bowtie \rho_2 = ((\mathcal{E}_1 \bowtie \mathcal{E}'_1)[M_1^{k_1}] \parallel \dots \parallel (\mathcal{E}_N \bowtie \mathcal{E}'_N)[M_N^{k_N}], \mathcal{R}_1, \mathcal{R}_2, S_1 \uplus S_2)$$

where compatible evaluation stacks $\mathcal{E}, \mathcal{E}'$ are composed into a single evaluation context $\mathcal{E} \bowtie \mathcal{E}'$, as follows:

$$\begin{aligned}
(m :: E :: \mathcal{E}) \bowtie (m :: \mathcal{E}') &= (\mathcal{E} \bowtie \mathcal{E}')[E[\langle m \rangle \bullet]^1] \\
(m :: \mathcal{E}') \bowtie (m :: E :: \mathcal{E}) &= (\mathcal{E} \bowtie \mathcal{E}')[E[\langle m \rangle \bullet]^2]
\end{aligned}$$

and $[] \bowtie [] = \bullet$. Unfolding the above, we have that, for example:

$$\begin{aligned}
[m_k, E_k, m_{k-1}, m_{k-2}, E_{k-2}, \dots, m_1, E_1] \\
\bowtie [m_k, m_{k-1}, E_{k-1}, m_{k-2}, \dots, m_1] &= E_1^1[\langle m_1 \rangle E_2^2[\dots E_k^{k'}[\langle m_k \rangle \bullet] \dots]]
\end{aligned}$$

where $k' = 2 - (k \bmod 2)$.

We proceed to fleshing out the correspondence. We observe that an L -configuration ρ can be the final configuration of a trace just if all its components are O -configurations with empty evaluation stacks. On the other hand, for C -configurations, we need to reach P -configurations with terms $()$. Thus, we call an N -configuration ρ *final* if $\rho = (\vec{\mathcal{C}}, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$ and either $\mathcal{C}_i = ([], -)$ for all i , or $\mathcal{C}_i = ([], ())$ for all i .

Let us write $(S_1, \hookrightarrow_1, \mathcal{F}_1)$ for the transition system induced from external composition, and $(S_2, \hookrightarrow_2, \mathcal{F}_2)$ be the transition system derived from internal composition:

- $S_1 = \{\rho \odot \rho' \mid \rho \succ \rho'\}$, $\mathcal{F}_1 = \{\rho \odot \rho' \in S_1 \mid \rho, \rho' \text{ final}\}$, and \hookrightarrow_1 the transition relation \longrightarrow defined previously.
- $S_2 = \{(M_1 \parallel \dots \parallel M_N, \mathcal{R}, S) \mid (M_1 \parallel \dots \parallel M_N, \mathcal{R}_1 \uplus \mathcal{R}_2, S) \text{ valid}\}$, $\mathcal{F}_2 = \{x \in S_2 \mid x \text{ final}\}$, and \hookrightarrow_2 the transition relation \Rightarrow' defined above.

A relation $R \subseteq S_1 \times S_2$ is called a *bisimulation* if, for all $(x_1, x_2) \in R$:

- $x_1 \in \mathcal{F}_1$ iff $x_2 \in \mathcal{F}_2$,
- if $x_1 \hookrightarrow_1 x'_1$ then $x_2 \hookrightarrow_2 x'_2$ and $(x'_1, x'_2) \in R$,
- if $x_2 \hookrightarrow_2 x'_2$ then $x_1 \hookrightarrow_1 x'_1$ and $(x'_1, x'_2) \in R$.

Given $(x_1, x_2) \in S_1 \times S_2$, we say that x_1 and x_2 are *bisimilar*, written $x_1 \sim x_2$, if $(x_1, x_2) \in R$ for some bisimulation R .

Lemma 48. *Let $\rho \succ \rho'$ be compatible N -configurations. Then, $(\rho \odot \rho') \sim (\rho \bowtie \rho')$.*

Recall we write \bar{h} for the O/P complement of the history h . We can now prove Theorem 33, which states that the behaviour of a library L inside a context C can be deduced by composing the semantics of L and C .

Theorem 33. *Let $L : \Psi \rightarrow \Psi'$, $L' : 1 \rightarrow \Psi$, Ψ_1 and $\Psi' \vdash M_1, \dots, M_N : \text{unit}$, with L , L' and $M_1; \dots; M_N$ accessing pairwise disjoint parts of the store. Then, $\text{link } L'; L \text{ in } (M_1 \parallel \dots \parallel M_N) \Downarrow$ iff there is $h \in \llbracket L \rrbracket_N$ such that $\bar{h} \in \llbracket \text{link } L'; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket$.*

Proof. Let C be the context $\text{link } L'; - \text{ in } (M_1 \parallel \dots \parallel M_N)$, and suppose $(L) \rightarrow_{\text{lib}}^* (\epsilon, \mathcal{R}_0, S_0)$ and $(L') \rightarrow_{\text{lib}}^* (\epsilon, \mathcal{R}'_0, S'_0)$ with $\text{dom}(\mathcal{R}_0) \cap \text{dom}(\mathcal{R}'_0) = \text{dom}(S_0) \cap \text{dom}(S'_0) = \emptyset$. We set:

$$\rho_0 = ((\square, -) \parallel \dots \parallel (\square, -), \mathcal{R}_0, (\emptyset, \Psi'), (\Psi, \emptyset), S_0)$$

$$\rho'_0 = ((\square, M_1) \parallel \dots \parallel (\square, M_N), \mathcal{R}'_0, (\Psi, \emptyset), (\emptyset, \Psi'), S'_0)$$

We pick these as the initial N -configurations for $\llbracket L \rrbracket_N$ and $\llbracket C \rrbracket$ respectively. Moreover, we have that $(L'; L) \rightarrow_{\text{lib}}^* (\epsilon, \mathcal{R}''_0, S''_0)$ where $\mathcal{R}''_0 = \{(m, (\mathcal{R}_0 \uplus \mathcal{R}'_0)(m) \{ \bar{r}/\bar{m} \}) \mid m \in \text{dom}(\mathcal{R}_0 \uplus \mathcal{R}'_0)\}$ and $S''_0 = (S_0 \uplus S'_0) \{ \bar{r}/\bar{m} \} \uplus_s \{(r_i, m_i) \mid i = 1, \dots, n\}$, assuming $\Psi = \{m_1, \dots, m_n\}$ and r_1, \dots, r_n are fresh references of corresponding types. Hence, the initial triple for $\llbracket C[L] \rrbracket$ is taken to be $\phi_0 = ((\square, M_1) \parallel \dots \parallel (\square, M_N), \mathcal{R}''_0, S''_0)$. On the other hand, $\rho'_0 \bowtie \rho_0 = ((\square, M_1) \parallel \dots \parallel (\square, M_N), \mathcal{R}'_0, \mathcal{R}_0, S_0 \uplus S'_0)$ and, using also Lemma 47, we have that $\phi_0 \Downarrow$ iff $\rho'_0 \bowtie \rho_0 \Downarrow$.

Then, for the forward direction of the claim, from $\phi_0 \Downarrow$ we obtain that $\rho'_0 \bowtie \rho_0 \Downarrow$. From the previous lemma, we have that so does $\rho'_0 \circ \rho_0$. From the latter reduction we obtain the required common history. Conversely, suppose $h \in \llbracket L \rrbracket_N$ and $\bar{h} \in \llbracket C \rrbracket$. WLOG, assume that $\text{Meths}(h) \cap (\text{dom}(\mathcal{R}_0) \cup \text{dom}(\mathcal{R}'_0)) \subseteq \Psi \cup \Psi_1 \cup \Psi'$ (we can appropriately alpha-covert \mathcal{R}_0 and \mathcal{R}'_0 for this). Then, ρ_0 and ρ'_0 both produce h , with opposite polarities. By definition of the external composite reduction, we then have that $\rho'_0 \circ \rho_0$ reduces to some final state. By the previous lemma, we have that $\rho'_0 \bowtie \rho_0$ reduces to some final quadruple, which in turn implies that $\phi_0 \Downarrow$, i.e. $\text{link } L'; L \text{ in } (M_1 \parallel \dots \parallel M_N) \Downarrow$. \square

We conclude this section with the proofs of the last two lemmata used.

C.1. Proof of Lemma 47

We purpose to show that, for any legal $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S)$, $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S) \Downarrow$ iff $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1 \cup \mathcal{R}_2, S) \Downarrow$.

We prove something stronger. For any repository \mathcal{R} whose entries are of the form $(m, \lambda x.m'x)$, we define a directed graph $\mathcal{G}(\mathcal{R})$ where vertices are all methods appearing in \mathcal{R} , and (m, m') is a (directed) edge just if $\mathcal{R}(m) = \lambda x.m'x$. In such a case, we call \mathcal{R} an **expansion class** if $\mathcal{G}(\mathcal{R})$ is acyclic and all its vertices have at most one outgoing edge. Moreover, given an expansion class \mathcal{R} , we define the method-for-method substitution $\{\mathcal{R}\}$ that assigns to each vertex m of $\mathcal{G}(\mathcal{R})$ the (unique) leaf m' such that there is a directed path from m to m' in $\mathcal{G}(\mathcal{R})$. Let us write $\mathcal{L}(\mathcal{R})$ for the set of leaves of $\mathcal{G}(\mathcal{R})$. For any quadruple $\phi = (E_1[M_1] \parallel \dots \parallel E_N[M_N], \mathcal{R}_1, \mathcal{R}_2, S)$ and expansion class $\mathcal{R} \subseteq \mathcal{R}_1 \cup \mathcal{R}_2$, we define the triple:

$$\phi^{\# \mathcal{R}} = (E_1[M_1] \parallel \dots \parallel E_N[M_N], \mathcal{R}_1 \cup \mathcal{R}_2, S) \{\mathcal{R}\} = (E_1[M_1] \{\mathcal{R}\} \parallel \dots \parallel E_N[M_N] \{\mathcal{R}\}, (\mathcal{R}_1 \cup \mathcal{R}_2) \{\mathcal{R}\}, S \{\mathcal{R}\})$$

where $\mathcal{R}' \{\mathcal{R}\} = \{(m, \mathcal{R}'(m) \{\mathcal{R}\}) \mid m \in \text{dom}(\mathcal{R}' \setminus \mathcal{R}) \cup \mathcal{L}(\mathcal{R})\}$, $S \{\mathcal{R}\} = (S \upharpoonright \text{Refs}_{\text{int}}) \cup \{(r, S(r) \{\mathcal{R}\}) \mid r \in \text{dom}(S) \setminus \text{Refs}_{\text{int}}\}$, and $E[M]$ is the term obtained from $E[M]$ by removing all tagging.

We next define a notion of indexed bisimulation between the transition systems produced from quadruples and triples respectively. Given an expansion class \mathcal{R} , a relation $R_{\mathcal{R}}$ between quadruples and triples is called an \mathcal{R} -bisimulation if, whenever $\phi_1 R_{\mathcal{R}} \phi_2$:

- ϕ_1 final implies ϕ_2 final
- ϕ_2 final implies $\phi_2 \Downarrow$
- $\phi_1 \Rightarrow' \phi'_1$ implies $\phi_2 \Rightarrow' \phi'_2$ and $\phi'_1 R_{\mathcal{R}} \phi'_2$ for some expansion class $\mathcal{R}' \supseteq \mathcal{R}$
- $\phi_2 \Rightarrow' \phi'_2$ implies $\phi_1 \Rightarrow'^* \phi'_1$ and $\phi'_1 R_{\mathcal{R}} \phi'_2$ for some expansion class $\mathcal{R}' \supseteq \mathcal{R}$.

Thus, Lemma 47 directly follows from the next result.

Lemma 49. *For all expansion classes \mathcal{R} , the relation*

$$R_{\mathcal{R}} = \{(\phi, \phi^{\# \mathcal{R}}) \mid \phi = (E_1[M_1] \parallel \dots \parallel E_N[M_N], \vec{\mathcal{R}}, S) \text{ legal} \wedge \mathcal{R} \subseteq \mathcal{R}_1 \cup \mathcal{R}_2\}$$

is a bisimulation.

Proof. Suppose $\phi R_{\mathcal{R}} \phi^{\# \mathcal{R}}$. We note that finality conditions are satisfied: if ϕ is final then so is $\phi^{\# \mathcal{R}}$; while if $\phi^{\# \mathcal{R}}$ is final then all its contexts are from the grammar:

$$E' ::= \bullet \mid \langle m^i \rangle E'$$

so $\phi \Downarrow$ by acyclicity of $\mathcal{G}(\mathcal{R})$.

Suppose now $\phi \Rightarrow \phi'$, say due to $(E_1[M_1], \mathcal{R}_1, \mathcal{R}_2, S) \rightarrow_1' (E'_1[M'_1], \mathcal{R}'_1, \mathcal{R}'_2, S')$. In case the reduction is not a function call or return, then it can be clearly simulated by $\phi^{\# \mathcal{R}}$. Otherwise, suppose:

- $(E_1[m^i v], \vec{\mathcal{R}}, S) \rightarrow_1' (E_1[M\{v/x\}^i], \vec{\mathcal{R}}, S)$. If $m \notin \text{dom}(\mathcal{R})$ then, writing \mathcal{R}_{12} for $\mathcal{R}_1 \cup \mathcal{R}_2$, the above can be simulated by $(E_1[m v], \mathcal{R}_{12}, S)\{\mathcal{R}\} \rightarrow_1 (E_1[M\{v/x\}], \mathcal{R}_{12}, S)\{\mathcal{R}\}$. If, on the other hand, $m \in \text{dom}(\mathcal{R})$, suppose $\mathcal{R}_i(m) = \lambda x.m'x$, then $M = m'x$ and $m\{\mathcal{R}\} = m'\{\mathcal{R}\}$ so we have:

$$E_1[M\{v/x\}^i]\{\mathcal{R}\} = E_1[(m'v)^i]\{\mathcal{R}\} = E_1[(mv)^i]\{\mathcal{R}\}$$

and $E_1[(mv)^i] = E_1[m^i v]$ by the way the semantics was defined, so $\phi^{\# \mathcal{R}} = \phi^{\# \mathcal{R}'}$.

- $(E_1[m^i v], \vec{\mathcal{R}}, S) \rightarrow_1' (E_1[(m^i) M\{v'/x\}^{3-i}], \vec{\mathcal{R}}', S)$, with $\mathcal{R}_{3-i}(m) = \lambda x.M$, $\text{Meths}(v) = \{m_1, \dots, m_k\}$, $v' = \{\vec{m}'/\vec{m}\}$ and $\vec{\mathcal{R}}' = \vec{\mathcal{R}} \uplus_i \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\}$. Let $\mathcal{R}' = \mathcal{R} \uplus \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\} \subseteq \mathcal{R}'_1 \cup \mathcal{R}'_2$. If $m \notin \text{dom}(\mathcal{R})$ then $(E_1[m v], \mathcal{R}_{12}, S)\{\mathcal{R}\} \rightarrow_1 (E_1[M\{v/x\}], \mathcal{R}_{12}, S)\{\mathcal{R}\}$, and we have:

$$E_1[(m^i) M\{v'/x\}^{3-i}]\{\mathcal{R}'\} = E_1[M\{v'/x\}]\{\mathcal{R}'\} = E_1[M\{v/x\}]\{\mathcal{R}\}$$

Moreover, $\mathcal{R}_{12}\{\mathcal{R}\} = (\mathcal{R}'_1 \cup \mathcal{R}'_2)\{\mathcal{R}'\}$ and $S\{\mathcal{R}\} = S\{\mathcal{R}'\}$, so $\phi^{\# \mathcal{R}} = (E_1[M\{v/x\}], \mathcal{R}_{12}, S)\{\mathcal{R}\}$.

On the other hand, if $\mathcal{R}(m) = \lambda x.m''x$ then:

$$E[(m^i) M\{v'/x\}^{3-i}]\{\mathcal{R}'\} = E[m''v']\{\mathcal{R}'\} = E[m''v]\{\mathcal{R}\} = E[mv]\{\mathcal{R}\}$$

so $\phi^{\# \mathcal{R}} = \phi^{\# \mathcal{R}'}$.

- Finally, the cases for method-return reductions are treated similarly as above.

Suppose now $\phi^{\# \mathcal{R}} \Rightarrow \phi'$, where recall that we write ϕ as $(E_1[M_1] \parallel \dots \parallel E_N[M_N], \vec{\mathcal{R}}, S)$. We show by induction on $\text{size}_{\mathcal{R}}(E_1[M_1], \dots, E_N[M_N])$ that $\phi \Rightarrow \phi''$ and $\phi' R_{\mathcal{R}} \phi''$ for some $\mathcal{R}' \supseteq \mathcal{R}$. The size-function we use measures the length of $\mathcal{G}(\mathcal{R})$ -paths that appear inside its arguments:

$$\text{size}_{\mathcal{R}}(E_1[M_1], \dots, E_N[M_N]) = \text{size}_{\mathcal{R}}(E_1[M_1]) + \dots + \text{size}_{\mathcal{R}}(E_N[M_N])$$

$$\text{size}_{\mathcal{R}}(E[M]) = \sum_{m \in X_1} 2|m|_{\mathcal{R}} + \sum_{m \in X_2} 1$$

where X_1 is the multiset containing all occurrences of methods $m \in \text{dom}(\mathcal{R})$ inside $E[M]$ in call position (e.g. mM'), and X_2 contains all occurrences of methods $m \in \text{dom}(\mathcal{R})$ inside $E[M]$ in return position (i.e. $\langle m^i \rangle \dots$). We write $|m|_{\mathcal{R}}$ for the length of the unique directed path from m to a leaf in $\mathcal{G}(\mathcal{R})$. The fact that X_1, X_2 are multisets reflects that we count all occurrences of m in call/return positions. Suppose WLOG that the reduction to ϕ' is due to some $(E_1[M_1], \mathcal{R}_{12}, S)\{\mathcal{R}\} \rightarrow_1 (E'_1[M'], \mathcal{R}', S')$. If the reduction happens inside $M_1\{\mathcal{R}\}$ (this case also encompasses the base case of the induction) then the only case we need to examine is that of the reduction being a method call. In such a case, suppose we have $E_1[M_1]\{\mathcal{R}\} = E[mv]$, $E' = E$, $M' = M\{v/x\}$ and $\mathcal{R}_{12}\{\mathcal{R}\}(m) = \lambda x.M$. Then, $E_1[M_1] = \tilde{E}[\tilde{m}^i \tilde{v}]$ for some $\tilde{E}, \tilde{m}, \tilde{v}$ such that $\tilde{m}\{\mathcal{R}\} = m$, $\tilde{v}\{\mathcal{R}\} = v$ and $\tilde{E}\{\mathcal{R}\} = E$. If $m \neq \tilde{m}$ then, supposing $\mathcal{R}(\tilde{m}) = \lambda x.\tilde{m}'x$ we have the following cases:

- $(\tilde{E}[\tilde{m}^i \tilde{v}], \vec{\mathcal{R}}, S) \rightarrow_1' (\tilde{E}[\tilde{m}'^i \tilde{v}], \vec{\mathcal{R}}, S) = \phi''_1$
- $(\tilde{E}[\tilde{m}^i \tilde{v}], \vec{\mathcal{R}}, S) \rightarrow_1' (\tilde{E}[(\tilde{m}^i) (\tilde{m}'v')^{3-i}], \vec{\mathcal{R}}', S) = \phi''_1$, with $\vec{\mathcal{R}}' = \vec{\mathcal{R}} \uplus_{3-i} \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\}$, etc.

Let ϕ'' be the extension of ϕ''_1 to an N -quadruple by using the remaining $E_i[M_i]$'s of ϕ , so that $\phi \Rightarrow \phi''$. In the first case above we have that $\phi^{\# \mathcal{R}} = \phi$, and in the latter that $\phi^{\# \mathcal{R}'} = \phi$ (with $\mathcal{R}' = \mathcal{R} \uplus \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\}$), and we appeal to the IH.

Suppose now that $\tilde{m} = m$ and $\mathcal{R}_{12}(m) = \lambda x.\tilde{M}$. Then, one of the following is the case:

- $(\tilde{E}[\tilde{m}^i \tilde{v}], \vec{\mathcal{R}}, S) \rightarrow_1' (\tilde{E}[\tilde{M}\{\tilde{v}/x\}^i], \vec{\mathcal{R}}, S) = \phi''_1$
- $(\tilde{E}[\tilde{m}^i \tilde{v}], \vec{\mathcal{R}}, S) \rightarrow_1' (\tilde{E}[(\tilde{m}^i) \tilde{M}\{v'/x\}^{3-i}], \vec{\mathcal{R}}', S) = \phi''_1$, with $\vec{\mathcal{R}}' = \vec{\mathcal{R}} \uplus_{3-i} \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\}$, etc.

Extending ϕ''_1 to ϕ'' as above, in the former case we then have that $\phi^{\# \mathcal{R}} = \phi'$, and in the latter that $\phi^{\# \mathcal{R}'} = \phi'$, as required.

Finally, let us suppose that M_1 is some value v . Then, we can write E_1 as $E_1 = E_2[E']$, with E' coming from the grammar $E' ::= \bullet \mid \langle m^i \rangle E'$ and E_2 not being of the form $E''[\langle m^i \rangle \bullet]$. Observe that $\underline{E}_1 = \underline{E}_2$. If $E' = \bullet$ then by a case analysis on E_1 we can see that $\phi^{\# \mathcal{R}}$ can simulate the reduction. Otherwise, $(E_2[E'[v]], \vec{\mathcal{R}}, S) \rightarrow_1' (E_2[E'[v^i]], \vec{\mathcal{R}}', S)$ whereby $E' = E''[\langle m^i \rangle \bullet]$ and $\vec{\mathcal{R}}' = \vec{\mathcal{R}} \uplus_{3-i} \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\}$, etc. We have that

$$\phi_1'' = (E_2[E''[v'^i]], \vec{\mathcal{R}}', S)\{\mathcal{R}'\} = (E_2[E'[v]], \vec{\mathcal{R}}, S)\{\mathcal{R}\}$$

and hence, extending ϕ_1'' to ϕ'' , we have $\phi'' \# \mathcal{R}' = \phi \# \mathcal{R}$. We can now appeal to the IH. \square

C.2. Proof of Lemma 48

Let $\rho \asymp \rho'$ be compatible N -configurations. Then, $(\rho \odot \rho') \sim (\rho \bowtie \rho')$.

We prove that the relation $R = \{(\rho_1 \odot \rho_2, \rho_1 \bowtie \rho_2) \mid \rho_1 \asymp \rho_2\}$ is a bisimulation. Let us suppose that $(\rho_1 \odot \rho_2, \rho_1 \bowtie \rho_2) \in R$.

- Suppose $\rho_1 \odot \rho_2 \hookrightarrow_1 \rho_1' \odot \rho_2'$. If the transition is due to (INT1) then $\rho_2 = \rho_2'$ and we can see that $\rho_1 \bowtie \rho_2 \Longrightarrow' \rho_1' \bowtie \rho_2$. Similarly if the transition is due to (INT2). Suppose now we used instead (CALL), e.g. $\rho_1 \xrightarrow{(1, \text{call } m(v))} \rho_1'$ and $\rho_2 \xrightarrow{(1, \text{call } m(v))} \rho_2'$, and let us consider the case where $v \in \text{Meths}$ (the other case is simpler). Then, assuming $\rho_1 = (C_1^1 \parallel \dots \parallel \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1)$ and $\rho_2 = (C_2^2 \parallel \dots \parallel \mathcal{R}_2, \mathcal{P}_2, \mathcal{A}_2, S_2)$, we have that either of the following scenarios holds, for some $x \in \{\mathcal{K}, \mathcal{L}\}$: $C_1^1 = (\mathcal{E}_1, E[mm'])$, $C_2^2 = (\mathcal{E}_2, -)$ and

$$\begin{aligned} (\mathcal{E}_1, E[mm'], \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1) &\xrightarrow{\text{call } m(v)}_1 (m :: E :: \mathcal{E}_1, \mathcal{R}_1 \uplus (v \mapsto \lambda x.m'x), \mathcal{P}_1 \cup_x \{v\}, \mathcal{A}_1, S_1) \\ (\mathcal{E}_2, -, \mathcal{R}_2, \mathcal{P}_2, \mathcal{A}_2, S_2) &\xrightarrow{\text{call } m(v)}_1 (m :: \mathcal{E}_2, M\{v/x\}, \mathcal{R}_2, \mathcal{P}_1, \mathcal{A}_1 \cup_x \{v\}, S_2) \end{aligned}$$

or its dual, where ρ_2 contains the code initiating the call. Focusing WLOG in the former case and setting $S = S_1 \uplus S_2$:

$$\begin{aligned} \rho_1 \bowtie \rho_2 &= ((\mathcal{E}_1 \bowtie \mathcal{E}_2)[E[m^1m']] \parallel \dots \parallel \mathcal{R}_1, \mathcal{R}_2, S) \\ &\hookrightarrow_2 ((\mathcal{E}_1 \bowtie \mathcal{E}_2)[E[\langle m^1 \rangle M\{v/x\}^2]] \parallel \dots \parallel \mathcal{R}_1', \mathcal{R}_2, S) = \rho_1' \bowtie \rho_2' \quad (\mathcal{R}_1' = \mathcal{R}_1 \uplus (v \mapsto \lambda x.m'x)) \end{aligned}$$

The case for (RETN) is treated similarly.

- Suppose $\rho_1 \bowtie \rho_2 = (E[M_1] \parallel M_2 \parallel \dots \parallel M_N, \vec{\mathcal{R}}, S) \hookrightarrow_2 (E[M_1'] \parallel M_2 \parallel \dots \parallel M_N, \vec{\mathcal{R}}', S')$ and let $\rho_1 = ((\mathcal{E}_1, M_1') \parallel \dots \parallel \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1)$ and $\rho_2 = ((\mathcal{E}_2, -) \parallel \dots \parallel \mathcal{R}_2, \mathcal{P}_2, \mathcal{A}_2, S_2)$, where $(\mathcal{E}_1 \bowtie \mathcal{E}_2)[M_1'] = E[M_1]$. If the redex M_1 is not of the forms $M_1 = m^1 v$ or $M_1 = \langle m^1 \rangle v$, with $m \in \text{dom}(\mathcal{R}_2)$, then the reduction can clearly be simulated by $\rho_1 \odot \rho_2$ (internally, by ρ_1). Otherwise, similarly as above, the reduction can be simulated by a mutual call/return of m .

Finally, it is clear that $\rho_1 \odot \rho_2$ is final iff $\rho_1 \bowtie \rho_2$ is final. \square

Appendix D. Library compositionality

This compositionality result will allow us to compose histories of component libraries in order to obtain those of their composite library. Let $L_1 : \Psi_1 \rightarrow \Psi_2$ and $L_2 : \Psi_1' \rightarrow \Psi_2'$. The semantic composition will be guided by two sets of names Π, P . Π contains method names that are shared between by the respective libraries and their context. Thus $\Pi \supseteq \Psi_1 \cup \Psi_1' \cup \Psi_2 \cup \Psi_2'$. The names in P , on the other hand, will be used for private communication between L_1 and L_2 . Consequently, $\Pi \cap P$ consists of names that can be used both for internal communication between L_1 and L_2 , and for contextual interactions, i.e. $\Pi \cap P = (\Psi_1 \cup \Psi_1') \cap (\Psi_2 \cup \Psi_2')$.

Given $h_i \in \llbracket L_i \rrbracket$ ($i = 1, 2$), we define the *composition* of h_1 and h_2 , written $h_1 \bowtie_{\Pi, P}^\sigma h_2$, as a partial operation depending on Π, P and an additional parameter $\sigma \in \{0, 1, 2\}^*$ which we call a *scheduler*. It is given inductively as follows. We let $\epsilon \bowtie_{\Pi, P}^\epsilon \epsilon = \epsilon$ and:

$$\begin{aligned} (t, \text{call } m(v))_{S_1} \bowtie_{\Pi, P}^{0\sigma} (t, \text{call } m(v))_{S_2} &= s_1 \bowtie_{\Pi, P'}^\sigma s_2 \\ (t, \text{ret } m(v))_{S_1} \bowtie_{\Pi, P}^{0\sigma} (t, \text{ret } m(v))_{S_2} &= s_1 \bowtie_{\Pi, P'}^\sigma s_2 \\ (t, \text{call } m(v))_{PY S_1} \bowtie_{\Pi, P}^{1\sigma} s_2 &= (t, \text{call } m(v))_{PY} (s_1 \bowtie_{\Pi', P}^\sigma s_2) \\ (t, \text{ret } m(v))_{PY S_1} \bowtie_{\Pi, P}^{1\sigma} s_2 &= (t, \text{ret } m(v))_{PY} (s_1 \bowtie_{\Pi', P}^\sigma s_2) \\ (t, \text{call } m(v))_{OY S_1} \bowtie_{\Pi, P}^{1\sigma} s_2 &= (t, \text{call } m(v))_{OY} (s_1 \bowtie_{\Pi', P}^\sigma s_2) \\ (t, \text{ret } m(v))_{OY S_1} \bowtie_{\Pi, P}^{1\sigma} s_2 &= (t, \text{ret } m(v))_{OY} (s_1 \bowtie_{\Pi', P}^\sigma s_2) \end{aligned}$$

along with the dual rules for the last four cases (i.e. where we schedule 2 in each case). Note that the definition uses sequences of moves that are suffixes of histories (such as s_i). The above equations are subject to the following side conditions:

- $\text{Meths}(v) \cap (\Pi \cup P) = \emptyset$, $\Pi' = \Pi \uplus \text{Meths}(v)$ and $P' = P \uplus \text{Meths}(v)$;
- $m \in P$ in the 0-scheduling cases;
- $m \in \Pi$ in the 1-scheduling cases and, also, $m \in \Pi \setminus P$ in the third case (the P -call);
- in the 1-scheduling cases, we also require that the leftmost move with thread index t in s_2 is not a P -move.

History composition is a partial function: if the conditions above are not met, or h_1, h_2, σ are not of the appropriate form, then the composition is undefined. The above conditions ensure that the composed histories are indeed compatible and can be produced by composing actual libraries. For instance, the last condition corresponds to determinacy of threads: there can only be at most one component starting with a P -move in each thread t . We then have the following correspondence:

Theorem 50. Let $L_1 : \Psi_1 \rightarrow \Psi_2$ and $L_2 : \Psi'_1 \rightarrow \Psi'_2$ be libraries accessing disjoint parts of the store. Then,

$$\llbracket L_1 \cup L_2 \rrbracket_N = \{ h \in \mathcal{H} \mid \exists \sigma, h_1 \in \llbracket L_1 \rrbracket_N, h_2 \in \llbracket L_2 \rrbracket_N. h = h_1 \bowtie_{\Pi_0, P_0}^\sigma h_2 \}$$

with $\Pi_0 = \Psi_1 \cup \Psi_2 \cup \Psi'_1 \cup \Psi'_2$ and $P_0 = (\Psi_1 \cup \Psi'_1) \cap (\Psi_2 \cup \Psi'_2)$.

The rest of this section is devoted in proving the Theorem.

Recall that we examine library composition in the sense of union of libraries. This scenario is more general than the one of Appendix C as, during composition via union, the calls and returns of each of the component libraries may be caught by the other library or passed as a call/return to the outer context. Thus, the setting of this section comprises given libraries $L_1 : \Psi_1 \rightarrow \Psi_2$ and $L_2 : \Psi'_1 \rightarrow \Psi'_2$, such that $\Psi_2 \cap \Psi'_2 = \emptyset$, and relating their semantics to that of their union $L_1 \cup L_2 : (\Psi_1 \cup \Psi'_1) \setminus (\Psi_2 \cup \Psi'_2) \rightarrow \Psi_2 \cup \Psi'_2$.

Given configurations for L_1 and L_2 , in order to be able to reduce them together we need to determine which of their methods can be used for communication between them, and which for interacting with the external context, which represents player O in the game. We will therefore employ a set of method names, denoted by Π and variants, to register those methods used for interaction with the external context. Another piece of information we need to know is in which component in the composition was the last call played, or whether it was an internal call instead. This is important so that, when O (or P) has the choice to return to both components, in the same thread, we know which one was last to call and therefore has precedence. We use for this purpose sequences $w = (w_1, \dots, w_N)$ where, for each i , $w_i \in \{0, 1, 2\}^*$. Thus, if e.g. $w_1 = 2w'_1$, this would mean that, in thread 1, the last call to O , was done from the second component; if, on the other hand, $w_1 = 0w'_1$ then the last call in thread 1 was an internal one between the two components. Given such a w and some $j \in \{0, 1, 2\}$, for each index t , we write $j +_t w$ for $w[t \mapsto (jw_t)]$.

Let us fix libraries $L_1 : \Psi_1 \rightarrow \Psi_2$ and $L_2 : \Psi'_1 \rightarrow \Psi'_2$. Let ρ_1, ρ_2 be N -configurations, and in particular L -configurations, and suppose that $\rho_1 = (\vec{C}, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$ and $\rho_2 = (\vec{C}', \mathcal{R}', \mathcal{P}', \mathcal{A}', S')$. Moreover, let $\Psi_1 \cup \Psi_2 \cup \Psi'_1 \cup \Psi'_2 \subseteq \Pi$. We say that ρ_1 and ρ_2 are (w, Π) -**compatible**, written $\rho_1 \asymp_w^\Pi \rho_2$, if S, S' have disjoint domains and, for each i :

- $\mathcal{C}_i = (\mathcal{E}'_i, M)$ and $\mathcal{C}'_i = (\mathcal{E}_i, -)$, or $\mathcal{C}_i = (\mathcal{E}_i, -)$ and $\mathcal{C}'_i = (\mathcal{E}'_i, M)$, or $\mathcal{C}_i = (\mathcal{E}_{L1}, -)$ and $\mathcal{C}'_i = (\mathcal{E}_{L2}, -)$.
- We have $\Psi_1 \subseteq \mathcal{A}_i$, $\Psi_2 \subseteq \mathcal{P}_K$, $\Psi'_1 \subseteq \mathcal{A}'_L$, $\Psi'_2 \subseteq \mathcal{P}'_K$ and, setting

$$P = (\mathcal{P}_K \cap \mathcal{A}'_L) \uplus (\mathcal{P}_L \cap \mathcal{A}'_K) \uplus (\mathcal{P}'_K \cap \mathcal{A}_L) \uplus (\mathcal{P}'_L \cap \mathcal{A}_K)$$

we also have:

- $(\mathcal{P}_L \uplus \mathcal{P}_K \uplus \mathcal{A}_i \uplus \mathcal{A}_K) \cap (\mathcal{P}'_L \uplus \mathcal{P}'_K \uplus \mathcal{A}'_i \uplus \mathcal{A}'_K) = P \uplus (\Psi_1 \cap \Psi'_1)$,
- $\Pi \cap P = (\Psi_2 \cup \Psi'_2) \cap (\Psi_1 \cup \Psi'_1)$,
- $\Pi \cup P = \mathcal{P}_L \cup \mathcal{P}_K \cup \mathcal{P}'_L \cup \mathcal{P}'_K \cup \mathcal{A}_L \cup \mathcal{A}_K \cup \mathcal{A}'_L \cup \mathcal{A}'_K$.

- The private names of \mathcal{R} do not appear in ρ_2 , and dually for the private names of \mathcal{R}' .
- If $\mathcal{C}_i = (\mathcal{E}, \dots)$ and $\mathcal{C}'_i = (\mathcal{E}', \dots)$ then \mathcal{E} and \mathcal{E}' are w_i -compatible, that is, either $\mathcal{E} = \mathcal{E}' = []$ or:

- $\mathcal{E} = m :: \mathcal{E}_1$ and $\mathcal{E}' \in \mathcal{E}_L$, with $m \in \Pi$, $w_i = 1u$ and $\mathcal{E}_1, \mathcal{E}'$ are u -compatible,
- or $\mathcal{E} = m :: \mathcal{E}_1$ and $\mathcal{E}' = m :: \mathcal{E} :: \mathcal{E}_2$, with $m \in P$, $w_i = 0u$ and $\mathcal{E}_1, \mathcal{E}_2$ are u -compatible,
- or $\mathcal{E} = m :: \mathcal{E} :: \mathcal{E}_1$ and $\mathcal{E}' \in \mathcal{E}_L$, with $m \in \Pi \setminus P$, $w_i = 1u$ and $\mathcal{E}_1, \mathcal{E}'$ are u -compatible,

or the dual of one of the three conditions above holds.

Given $\rho_1 \asymp_w^\Pi \rho_2$, we let their external composition be denoted as $\rho_1 \otimes_w^\Pi \rho_2$ (and note that now the notation is symmetric for ρ_1 and ρ_2) and define the semantics for external composition by these rules:

$$\begin{array}{c} \frac{\rho_1 \Rightarrow \rho'_1}{\rho_1 \otimes_w^\Pi \rho_2 \longrightarrow \rho'_1 \otimes_w^\Pi \rho_2} \text{INT}_1 \\ \frac{\rho_1 \xrightarrow{(t, \text{call } m(v))} \rho'_1 \quad \rho_2 \xrightarrow{(t, \text{call } m(v))} \rho'_2 \quad m \in P}{\rho_1 \otimes_w^\Pi \rho_2 \longrightarrow \rho'_1 \otimes_{\Pi}^{0+_t w} \rho'_2} \text{CALL} \quad \frac{\rho_1 \xrightarrow{(t, \text{ret } m(v))} \rho'_1 \quad \rho_2 \xrightarrow{(t, \text{ret } m(v))} \rho'_2 \quad m \in P}{\rho_1 \otimes_{\Pi}^{0+_t w} \rho_2 \longrightarrow \rho'_1 \otimes_w^\Pi \rho'_2} \text{RETN} \end{array}$$

$$\begin{array}{c}
\frac{\rho_1 \xrightarrow{(t, \text{call } m(v))_{PY}} \rho'_1 \quad m \in \Pi \setminus P}{\rho_1 \otimes_{\Pi}^W \rho_2 \xrightarrow{(t, \text{call } m(v))_{PY}} \rho'_1 \otimes_{\Pi'}^{1+t} \rho_2} \text{PCALL}_1 \\
\frac{\rho_1 \xrightarrow{(t, \text{call } m(v))_{OY}} \rho'_1 \quad m \in \Pi}{\rho_1 \otimes_{\Pi}^W \rho_2 \xrightarrow{(t, \text{call } m(v))_{OY}} \rho'_1 \otimes_{\Pi'}^{1+t} \rho_2} \text{OCALL}_1 \\
\frac{\rho_1 \xrightarrow{(t, \text{ret } m(v))_{PY}} \rho'_1 \quad m \in \Pi}{\rho_1 \otimes_{\Pi}^{1+t} \rho_2 \xrightarrow{(t, \text{ret } m(v))_{PY}} \rho'_1 \otimes_{\Pi'}^W \rho_2} \text{PRETN}_1 \\
\frac{\rho_1 \xrightarrow{(t, \text{ret } m(v))_{OY}} \rho'_1 \quad m \in \Pi \setminus P}{\rho_1 \otimes_{\Pi}^{1+t} \rho_2 \xrightarrow{(t, \text{ret } m(v))_{OY}} \rho'_1 \otimes_{\Pi'}^W \rho_2} \text{ORETN}_1
\end{array}$$

along with their dual counterparts (INT₂, XCALL₂, XRET₂). The internal rules above have the same side-conditions on name privacy as before. Moreover, in (XRET₂)_i and (XCALL₂)_i, for X = O, P, we let $\Pi' = \Pi \uplus_t \text{Meths}(v)$ and impose that the t -th component of ρ_{3-i} be an O -configuration and $\text{Meths}(v) \cap \text{Meths}(\rho_{3-i}) = \emptyset$.

We can now show the following.

Lemma 51. *Let $\rho_1 \succ_{\Pi}^W \rho_2$ and suppose $\rho_1 \otimes_{\Pi}^W \rho_2 \xrightarrow{s}^* \rho'_1 \otimes_{\Pi'}^W \rho'_2$ for some sequence s of moves. Then, $\rho'_1 \succ_{\Pi'}^W \rho'_2$.*

We next juxtapose the semantics of external composition to that obtained by internally composing the libraries and then deriving the multi-threaded semantics of the result. As before, we call the latter form *internal composition*. The traces we obtain are produced from a transition relation, written \Longrightarrow' , between configurations of the form $(\mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_N, \mathcal{R}_1, \mathcal{R}_2, \mathcal{P}, \mathcal{A}, S)$, also written $(\vec{\mathcal{C}}, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S)$. In particular, in each $\mathcal{C}_i = (\mathcal{E}_i, X_i)$ with $X_i = E_i[M_i]$ or $X_i = -$, E_i is selected from the extended evaluation contexts and \mathcal{E}_i is an *extended L-stack*, that is, of either of the following two forms:

$$\mathcal{E}_{\text{ext}} ::= [] \mid m^i :: E :: \mathcal{E}'_{\text{ext}} \quad \mathcal{E}'_{\text{ext}} ::= m :: \mathcal{E}_{\text{ext}}$$

where E is again from the extended evaluation contexts.

First, given u -compatible evaluation stacks $\mathcal{E}, \mathcal{E}'$, we construct a pair $\mathcal{E} \bowtie^u \mathcal{E}'$ consisting of an extended evaluation context and an extended L -stack, as follows. Given $\mathcal{E} \bowtie^u \mathcal{E}' = (E', \mathcal{E}'')$:

$$\begin{aligned}
(m :: E :: \mathcal{E}) \bowtie^{0u} (m :: \mathcal{E}') &= (E'[E[\langle m \rangle \bullet]^1], \mathcal{E}'') \\
(m :: \mathcal{E}) \bowtie^{0u} (m :: E :: \mathcal{E}') &= (E'[E[\langle m \rangle \bullet]^2], \mathcal{E}'') \\
(m :: \mathcal{E}) \bowtie^{1u} \mathcal{E}' &= \mathcal{E} \bowtie^{2u} (m :: \mathcal{E}') = (\bullet, m :: E' :: \mathcal{E}'') \\
(m :: E :: \mathcal{E}) \bowtie^{1u} \mathcal{E}' &= \mathcal{E} \bowtie^{2u} (m :: E :: \mathcal{E}') = (\bullet, m :: E'[E] :: \mathcal{E}'') \text{ if } \mathcal{E}' \in \mathcal{E}_L
\end{aligned}$$

and $[] \bowtie^e [] = (\bullet, [])$.

For each pair $\rho_1 \succ_{\Pi}^W \rho_2$, we define a configuration corresponding to their syntactic composition as follows. Let $\rho_1 = (\mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_N, \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1)$ and $\rho_2 = (\mathcal{C}'_1 \parallel \dots \parallel \mathcal{C}'_N, \mathcal{R}_2, \mathcal{P}_2, \mathcal{A}_2, S_2)$ and, for each i , $\mathcal{C}_i = (\mathcal{E}_i, X_i)$ and $\mathcal{C}'_i = (\mathcal{E}'_i, X'_i)$. If $\mathcal{E}_i \bowtie^u \mathcal{E}'_i = (E_i, \mathcal{E}''_i)$, we set:

$$\mathcal{C}_i \bowtie^u \mathcal{C}'_i = \begin{cases} (\mathcal{E}''_i, E_i[M^1]) & \text{if } X_i = M \text{ and } X'_i = - \\ (\mathcal{E}''_i, E_i[M^2]) & \text{if } X_i = - \text{ and } X'_i = M \\ (\mathcal{E}''_i, -) & \text{if } X_i = X'_i = - \end{cases}$$

We then let the internal composition of ρ_1 and ρ_2 be:

$$\rho_1 \bowtie_{\Pi}^W \rho_2 = (\mathcal{C}_1 \bowtie^{W1} \mathcal{C}'_1 \parallel \dots \parallel \mathcal{C}_N \bowtie^{WN} \mathcal{C}'_N, \mathcal{R}_1, \mathcal{R}_2, \mathcal{P}', \mathcal{A}', S_1 \uplus S_2)$$

where we set $\mathcal{P}' = ((\mathcal{P}_{1\mathcal{L}} \uplus \mathcal{P}_{2\mathcal{L}}) \cap \Pi, (\mathcal{P}_{1\mathcal{K}} \uplus \mathcal{P}_{2\mathcal{K}}) \cap \Pi)$ and $\mathcal{A}' = ((\mathcal{A}_{1\mathcal{L}} \cup \mathcal{A}_{2\mathcal{L}}) \cap (\Pi \setminus P), (\mathcal{A}_{1\mathcal{K}} \uplus \mathcal{A}_{2\mathcal{K}}) \cap \Pi)$.

Now, as expected, the definition of \Longrightarrow' builds upon \rightarrow'_t . The definition of the latter is given by the following rules.

$$\begin{aligned}
&\frac{(E[M], \vec{\mathcal{R}}, S) \rightarrow'_t (E'[M'], \vec{\mathcal{R}}', S')}{(\mathcal{E}, E[M], \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \rightarrow'_t (\mathcal{E}, E'[M'], \vec{\mathcal{R}}', \mathcal{P}, \mathcal{A}, S')} \text{ (INT')} \\
&(\mathcal{E}, E[m^i v], \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{call } m(v')_{PY}}'_t (m^i :: E :: \mathcal{E}, -, \vec{\mathcal{R}}', \mathcal{P}', \mathcal{A}, S) \text{ (PCY')} \\
&(m :: \mathcal{E}, v, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{ret } m(v')_{PY}}'_t (\mathcal{E}, -, \vec{\mathcal{R}}', \mathcal{P}', \mathcal{A}, S) \text{ (PRY')} \\
&(\mathcal{E}, -, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{call } m(v)_{OY}}'_t (m :: \mathcal{E}, M\{v/x\}^i, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}', S) \text{ (OCY')} \\
&(m^i :: E :: \mathcal{E}, -, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{ret } m(v)_{OY}}'_t (\mathcal{E}, E[v^i], \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}', S) \text{ (ORY')}
\end{aligned}$$

The side-conditions are similar to those for the relation \rightarrow_t between ordinary configurations, with the following exceptions: in (PCY'), if $\text{Meths}(v) = \{m_1, \dots, m_k\}$ then $v' = v\{m'_j/m_j \mid 1 \leq j \leq k\}$, for fresh m'_j 's, and $\vec{\mathcal{R}}' = \vec{\mathcal{R}} \uplus_i \{m'_j \mapsto \lambda x.m_j x\}$; and in (PRY'), if $m \in \text{dom}(\mathcal{R}_i)$ then $\vec{\mathcal{R}}' = \vec{\mathcal{R}} \uplus_i \{m'_j \mapsto \lambda x.m_j x\}$, etc. Moreover, in (OCY') we have that $m \in \text{dom}(\mathcal{R}_i)$. Finally, we let

$$(\vec{\mathcal{C}}, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{(t,x)_{XY}}' (\vec{\mathcal{C}}[t \mapsto C'], \vec{\mathcal{R}}', \mathcal{P}', \mathcal{A}', S')$$

just if $(\mathcal{C}_t, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{XY} (\mathcal{C}', \vec{\mathcal{R}}', \mathcal{P}', \mathcal{A}', S')$.

We next relate the transition systems induced by external (via \otimes) and internal composition (via \mathbb{M}). Let us write $(\mathcal{S}_1, \hookrightarrow_1, \mathcal{F}_1)$ for the transition system induced by external composition of compatible N -configurations (so \hookrightarrow_1 is \rightarrow), and $(\mathcal{S}_2, \hookrightarrow_2, \mathcal{F}_2)$ be the one for internal composition (so \hookrightarrow_2 is \Rightarrow). Finality of extended N -configurations $(\mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_N, \vec{\mathcal{R}}, \dots)$ is defined as expected: all \mathcal{C}_i 's must be $([], -)$. A relation $R \subseteq \mathcal{S}_1 \times \mathcal{S}_2$ is called a *bisimulation* if, for all $(x_1, x_2) \in R$:

- $x_1 \in \mathcal{F}_1$ iff $x_2 \in \mathcal{F}_2$,
- if $x_1 \hookrightarrow_1 x'_1$ then $x_2 \hookrightarrow_2 x'_2$ and $(x'_1, x'_2) \in R$,
- if $x_1 \xrightarrow{(t,x)_{XY}}_1 x'_1$ then $x_2 \xrightarrow{(t,x)_{XY}}_2 x'_2$ and $(x'_1, x'_2) \in R$,
- if $x_2 \hookrightarrow_2 x'_2$ then $x_1 \hookrightarrow_1 x'_1$ and $(x'_1, x'_2) \in R$,
- if $x_2 \xrightarrow{(t,x)_{XY}}_2 x'_2$ then $x_1 \xrightarrow{(t,x)_{XY}}_1 x'_1$ and $(x'_1, x'_2) \in R$.

Again, we say that x_1 and x_2 are *bisimilar*, and write $x_1 \sim x_2$, if there exists a bisimulation R such that $(x_1, x_2) \in R$.

Lemma 52. Let $\rho \asymp_{\Pi}^W \rho'$ be compatible N -configurations. Then, $(\rho \otimes_{\Pi}^W \rho') \sim (\rho \mathbb{M}_{\Pi}^W \rho')$.

Proof. We prove that the relation $R = \{(\rho_1 \otimes_{\Pi}^W \rho_2, \rho_1 \mathbb{M}_{\Pi}^W \rho_2) \mid \rho_1 \asymp_{\Pi}^W \rho_2\}$ is a bisimulation. Let us suppose that $(\rho_1 \otimes_{\Pi}^W \rho_2, \rho_1 \mathbb{M}_{\Pi}^W \rho_2) \in R$.

- Let $\rho_1 \otimes_{\Pi}^W \rho_2 \xrightarrow{(t,x)} \rho'_1 \otimes_{\Pi'}^{W'} \rho'_2$ with the transition being due to (XCALL₁), e.g. $\rho_1 \xrightarrow{(1, \text{call } m(v))} \rho'_1$ and $\rho'_2 = \rho_2$, $w' = 1 +_1 w$ and $\Pi' = \Pi \uplus_1 \text{Meths}(v)$, $\text{Meths}(v) = \{m'_1, \dots, m'_j\}$, and recall that $\text{Meths}(v) \cap \text{Meths}(\rho_2) = \emptyset$. Then, assuming $\rho_1 = (\mathcal{C}_1^1 \parallel \dots, \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1)$, we have that one of the following holds, for some $x \in \{\mathcal{K}, \mathcal{L}\}$:

$$\begin{aligned} \mathcal{C}_1^1 &= (\mathcal{E}_1, E[mv']) \text{ and } (\mathcal{C}_1^1, \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1) \\ &\xrightarrow{\text{call } m(v)}_1 (m :: E :: \mathcal{E}_1, -, \mathcal{R}_1 \uplus \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\}, \mathcal{P}_1 \cup_x \text{Meths}(v), \mathcal{A}_1, S_1) \\ \mathcal{C}_1^1 &= (\mathcal{E}_1, -) \text{ and } (\mathcal{C}_1^1, \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1) \xrightarrow{\text{call } m(v)}_1 (m^1 :: \mathcal{E}_1, mv, \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1 \cup_x \text{Meths}(v), S_1) \end{aligned}$$

In the former case, if $\rho_2 = ((\mathcal{E}_2, -) \parallel \dots, \mathcal{R}_2, \mathcal{P}_2, \mathcal{A}_2, S_2)$ with $\mathcal{E}_1 \mathbb{M}^{w_1} \mathcal{E}_2 = (E', \mathcal{E})$, we get:

$$\begin{aligned} \rho_1 \mathbb{M}_{\Pi}^W \rho_2 &= ((\mathcal{E}, E'[E[mv']^1]) \parallel \dots, \mathcal{R}_1, \mathcal{R}_2, \mathcal{P}, \mathcal{A}, S) \\ &\xrightarrow{(1, \text{call } m(v))}' ((m^1 :: E'[E^1] :: \mathcal{E}, -) \parallel \dots, \mathcal{R}_1 \uplus \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\}, \mathcal{R}_2, \mathcal{P}', \mathcal{A}, S) \end{aligned}$$

with \mathcal{P}, \mathcal{A} as in the definition of composition and $\mathcal{P}' = \mathcal{P} \cup_x \text{Meths}(v)$, and the latter N -configuration equals $\rho'_1 \mathbb{M}_{\Pi'}^{W'} \rho'_2$. The other case is treated in the same manner, and we work similarly for (RET_N).

- On the other hand, if the transition is due to (CALL) or (RET_N) then we work as in the proof of Lemma 48.
- Suppose $\rho_1 \mathbb{M}_{\Pi}^W \rho_2 = (\mathcal{C}_1 \parallel \dots, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{(1, \text{call } m(v))}' (\mathcal{C}'_1 \parallel \dots, \vec{\mathcal{R}}', \mathcal{P}', \mathcal{A}', S)$. Then, assuming WLOG that $v \in \text{Meths}$, one of the following must be the case, for some $x \in \{\mathcal{K}, \mathcal{L}\}$ and $i \in \{1, 2\}$:

$$\begin{aligned} \mathcal{C}_1 &= (\mathcal{E}, E[m^i v']) \text{ and } (\mathcal{C}_1, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{call } m(v)}'_i \\ &\quad (m^i :: E :: \mathcal{E}, \vec{\mathcal{R}} \uplus_i \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\}, \mathcal{P} \cup_x \text{Meths}(v), \mathcal{A}, S) \\ \mathcal{C}_1(\mathcal{E}, -) &\text{ and } (\mathcal{C}_1, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{\text{call } m(v)}'_i (m :: \mathcal{E}, M\{v/x\}^i, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A} \cup_x \text{Meths}(v), S) \end{aligned}$$

We only examine the former case, as the latter one is similar, and suppose that $i = 1$. Taking $\rho_j = (\mathcal{C}_1^j \parallel \dots, \mathcal{R}_j, \mathcal{P}_j, \mathcal{A}_j, S_j)$, for $j = 1, 2$, we have that $(\mathcal{C}_1^1, \mathcal{C}_2^2) = ((\mathcal{E}_1, E'[mv']), (\mathcal{E}_2, -))$, for some $E, \mathcal{E}_1, \mathcal{E}_2$ such that $\mathcal{E}_1 \mathbb{M}^{w_1} \mathcal{E}_2 = (E'', \mathcal{E})$ and $E = E''[E'^1]$. Moreover, taking $\mathcal{R}'_1 = \mathcal{R}_1 \uplus \{m'_j \mapsto \lambda x.m_j x \mid 1 \leq j \leq k\}$, $\mathcal{P}'_1 = \mathcal{P}_1 \uplus_x \{v\}$, $w' = 1 +_1 w$ and $\Pi' = \Pi \uplus \text{Meths}(v)$ (note $\text{Meths}(v) = \{m'_1, \dots, m'_k\}$),

$$\rho_1 \otimes_{\Pi}^W \rho_2 \xrightarrow{(1, \text{call } m(v))} ((m :: E' :: \mathcal{E}_1, -) \parallel \dots, \mathcal{R}'_1, \mathcal{P}'_1, \mathcal{A}_1, S_1) \otimes_{\Pi'}^{W'} \rho_2 = \rho'_1 \otimes_{\Pi'}^{W'} \rho'_2$$

and $\rho'_1 \bowtie_{\Pi}^{w'} \rho_2 = (C'_1 \parallel \dots \parallel \vec{\mathcal{R}}', \mathcal{P}', \mathcal{A}', S)$ as required. The case for return transitions is similar.

- On the other hand, if the transition out of $\rho_1 \bowtie_{\Pi}^w \rho_2$ does not have a label then we work as in the proof of Lemma 48.

Moreover, by definition of syntactic composition, $\rho_1 \otimes_{\Pi}^w \rho_2$ is final iff $\rho_1 \bowtie_{\Pi}^w \rho_2$ is. \square

Given an N -configuration ρ and a history h , let us write $\rho \Downarrow h$ if $\rho \xRightarrow{h} \rho'$ for some final configuration ρ' . Similarly if ρ is of the form $(\vec{C}, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$. We have the following connections in history productions. The next lemma is proven in a similar fashion as Lemma 47.

Lemma 53. *For any legal $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, \mathcal{P}, \mathcal{A}, S)$ and history h , we have that $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, \mathcal{P}, \mathcal{A}, S) \Downarrow h$ iff $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1 \cup \mathcal{R}_2, \mathcal{P}, \mathcal{A}, S) \Downarrow h$.*

Lemma 54. *For any compatible N -configurations $\rho_1 \asymp_{\Pi}^w \rho_2$ and history h , $(\rho_1 \otimes_{\Pi}^w \rho_2) \Downarrow h$ iff:*

$$\exists h_1, h_2, \sigma. \rho_1 \Downarrow h_1 \wedge \rho_2 \Downarrow h_2 \wedge h = h_1 \bowtie_{\Pi, P}^{\sigma} h_2$$

where P is computed from ρ_1, ρ_2 and Π as before.

Proof. We show that, for any compatible N -configurations $\rho_1 \asymp_{\Pi}^w \rho_2$ and history suffix s , $(\rho_1 \otimes_{\Pi}^w \rho_2) \Downarrow s$ iff:

$$\exists s_1, s_2, \sigma. \rho_1 \Downarrow s_1 \wedge \rho_2 \Downarrow s_2 \wedge s = s_1 \bowtie_{\Pi, P}^{\sigma} s_2$$

where P is computed from ρ_1, ρ_2 and Π as in the beginning of this section.

The left-to-right direction follows from straightforward induction on the length of the reduction that produces s . For the right-to-left direction, we do induction on the length of σ . If $\sigma = \epsilon$ then $s_1 = s_2 = s = \epsilon$. Otherwise, we do a case analysis on the first element of σ . We only look at the most interesting subcase, namely of $\sigma = 0\sigma'$. Then, for some $m \in P$:

$$s_1 = (t, \text{call } m(v))s'_1 \quad s_2 = (t, \text{call } m(v))s'_2$$

By $\rho_i \Downarrow s_i$ and $\rho_1 \asymp_{\Pi}^w \rho_2$ we have that $\rho_1 \otimes_{\Pi}^w \rho_2 \longrightarrow \rho'_1 \otimes_{\Pi}^{w'} \rho'_2$, where $w' = 0 +_t w$ and $\rho'_1 \asymp_{\Pi}^{w'} \rho'_2$. Also, $\rho'_i \Downarrow s'_i$ and $s = s'_1 \bowtie_{\Pi, P'}^{\sigma'} s'_2$ so, by IH, $(\rho'_1 \otimes_{\Pi}^{w'} \rho'_2) \Downarrow s$. \square

We can now prove the correspondence between the traces of component libraries and those of their union.

Theorem 50. *Let $L_1 : \Psi_1 \rightarrow \Psi_2$ and $L_2 : \Psi'_1 \rightarrow \Psi'_2$ be libraries accessing disjoint parts of the store. Then,*

$$\llbracket L_1 \cup L_2 \rrbracket_N = \{h \in \mathcal{H} \mid \exists \sigma, h_1 \in \llbracket L_1 \rrbracket_N, h_2 \in \llbracket L_2 \rrbracket_N. h = h_1 \bowtie_{\Pi_0, P_0}^{\sigma} h_2\}$$

with $\Pi_0 = \Psi_1 \cup \Psi_2 \cup \Psi'_1 \cup \Psi'_2$ and $P_0 = (\Psi_1 \cup \Psi'_1) \cap (\Psi_2 \cup \Psi'_2)$.

Proof. Let us suppose $(L_i) \longrightarrow_{\text{lib}}^* (\epsilon, \mathcal{R}_i, S_i)$, for $i = 1, 2$, with $\text{dom}(\mathcal{R}_1) \cap \text{dom}(\mathcal{R}_2) = \text{dom}(S_1) \cap \text{dom}(S_2) = \emptyset$. We set:

$$\rho_1 = ((\square, -) \parallel \dots \parallel (\square, -), \mathcal{R}_1, (\emptyset, \Psi_2), (\Psi_1, \emptyset), S_1)$$

$$\rho_2 = ((\square, -) \parallel \dots \parallel (\square, -), \mathcal{R}_2, (\emptyset, \Psi'_2), (\Psi'_1, \emptyset), S_2)$$

We pick these as the initial configurations for $\llbracket L_1 \rrbracket_N$ and $\llbracket L_2 \rrbracket_N$ respectively. Then, $(L_1 \cup L_2) \longrightarrow_{\text{lib}}^* (\epsilon, \mathcal{R}_0, S_0)$ where $\mathcal{R}_0 = \mathcal{R}_1 \uplus \mathcal{R}_2$ and $S_0 = S_1 \uplus S_2$, and we take

$$\rho_0 = ((\square, -) \parallel \dots \parallel (\square, -), \mathcal{R}_0, (\emptyset, \Psi_2 \cup \Psi'_2), ((\Psi_1 \cup \Psi'_1) \setminus P_0, \emptyset), S_0)$$

as the initial N -configuration for $\llbracket L_1 \cup L_2 \rrbracket_N$. On the other hand, we have $\rho_1 \bowtie_{\Pi_0}^{\epsilon} \rho_2 = ((\square, -) \parallel \dots \parallel (\square, -), \mathcal{R}_1, \mathcal{R}_2, (\emptyset, \Psi_2 \cup \Psi'_2), ((\Psi_1 \cup \Psi'_1) \setminus P_0, S_0))$. From Lemma 53, we have that $\rho_0 \Downarrow h$ iff $\rho_1 \bowtie_{\Pi_0}^{\epsilon} \rho_2 \Downarrow h$, for all h .

Pick a history h . For the forward direction of the claim, $\rho_0 \Downarrow h$ implies $\rho_1 \bowtie_{\Pi_0}^{\epsilon} \rho_2 \Downarrow h$ which, from Lemma 52, implies $\rho_1 \otimes_{\Pi_0}^{\epsilon} \rho_2 \Downarrow h$. We now use Lemma 54 to obtain h_1, h_2, σ such that $\rho_i \Downarrow h_i$ and $h = h_1 \bowtie_{\Pi_0, P_0}^{\sigma} h_2$. Conversely, suppose that $h_i \in \llbracket L_i \rrbracket_N$ and $h = h_1 \bowtie_{\Pi_0, P_0}^{\sigma} h_2$. WLOG assume that $(\text{Meths}(h_1) \cup \text{Meths}(h_2)) \cap (\text{dom}(\mathcal{R}_1) \cup \text{dom}(\mathcal{R}_2)) \subseteq \Pi_0$ (or we appropriately alpha-covert \mathcal{R}_1 and \mathcal{R}_2). Then, $\rho_i \Downarrow h_i$, for $i = 1, 2$, and therefore $\rho_1 \otimes_{\Pi_0}^{\epsilon} \rho_2 \Downarrow h$ by Lemma 54. By Lemma 52 we have that $\rho_1 \bowtie_{\Pi_0}^{\epsilon} \rho_2 \Downarrow h$, which in turn implies that $\rho_0 \Downarrow h$, i.e. $h \in \llbracket L_1 \cup L_2 \rrbracket_N$. \square

Appendix E. Composition congruence

Theorem 55. *If $L_1 \triangleleft L_2$ then, for suitably typed L accessing disjoint part of the store than L_1 and L_2 , we have $L \cup L_1 \triangleleft L \cup L_2$.*

Proof. Assume $L_1 \triangleleft L_2$ and suppose $h_1 \in \llbracket L \cup L_1 \rrbracket$. By Theorem 50, $h_1 = h' \bowtie_{\Pi, P}^{\sigma} h'_1$, where $h' \in \llbracket L \rrbracket$ and $h'_1 \in \llbracket L_1 \rrbracket$. Because $L_1 \triangleleft L_2$, there exists $h'_2 \in \llbracket L_2 \rrbracket$ such that $h'_1 \triangleleft h'_2$, i.e. $h'_1 \triangleleft_{P, O}^* h'_2$. Note that some of the rearrangements necessary to transform h'_1 into h'_2 may concern actions shared by h'_1 and h' ; their polarity will then be different in h' . Let h'' be obtained by applying such rearrangements to h' . We claim that $h' \triangleleft_{O, P}^* h''$. Indeed, suppose that $(t', x')(t, x)_P$ are consecutive in h'_1 , but swapped in order to obtain h'_2 , and $(t, x)_P$ appears in h' as $(t, x)_O$. Now, the move (t', x') either appears in h_1 , or it appears in h' and gets hidden in h_1 . In every case, let s contain the moves of h' that are after (t', x') in the composition to h_1 , and before $(t, x)_O$. We have that $s(t, x)_O$ is a subsequence of h' and $h' \triangleleft_{O, P}^* h''$ holds just if s contains no moves from t . But, if s contained moves from t then the rightmost one such would be some $(t, y)_P$. Moreover, in the composition towards h_1 , the move would be scheduled with 1. The latter would break the conditions for trace composition as, at that point, the corresponding subsequence of h'_1 has as leftmost move in t the P-move $(t, x)_P$. We can show similarly that $h' \triangleleft_{O, P}^* h''$ holds in the case that the permutation in h'_1 is on consecutive moves $(t, x)_O(t', x')$. Finally, the rearrangements in h'_1 that do not affect moves shared with h' can be treated in a simpler way: e.g. in the case of $(t', x')(t, x)_P$ consecutive in h'_1 and swapped in h'_2 , if $(t, x)_P$ does not appear in h' then we can check that h' cannot contain any t -moves between (t', x') and (t, x) as the conditions for trace composition impose that only O is expected to play in that part of h' (and any t -move would swap this polarity).

Now, since $h' \in \llbracket L \rrbracket$, Lemma 34 implies $h'' \in \llbracket L \rrbracket$. Take h_2 to be $h'' \bowtie_{\Pi, P}^{\sigma'} h'_2$, where σ' is obtained from σ following these move rearrangements. We then have $h_2 \in \llbracket L \cup L_2 \rrbracket$. Moreover, $h_1 \triangleleft h_2$ thanks to $h'_1 \triangleleft h'_2$. Hence, $h_2 \in \llbracket L \cup L_2 \rrbracket$ and $h_1 \triangleleft h_2$. Thus, $L \cup L_1 \triangleleft L \cup L_2$. \square

We next examine the behaviour of $\triangleleft_{\text{enc}}$ with respect to library composition. In contrast to general linearisability, we need to restrict composition for it to be compatible with encapsulation.

Remark 56. The general case of union does not conform with encapsulation in the sense that encapsulated testing of $L \cup L_i$ ($i = 1, 2$) according to Definition 31 may subject L_i to unencapsulated testing. For example, because method names of L and L_i are allowed to overlap, methods in L may call public methods from L_i as well as implementing abstract methods from L_i . This amounts to L playing the role of both \mathcal{K} and \mathcal{L} , which in addition can communicate with each other, as both are inside L .

Even if we make L and L_i non-interacting (i.e. without common abstract/public methods), if higher-order parameters are still involved, the encapsulated tests of $L \cup L_i$ can violate the encapsulation hypothesis for L_i . For instance, consider the methods $m_2, m'_1, m'_2 \in \text{Meths}_{\text{unit}, \text{unit}}$ and $m_1 \in \text{Meths}_{(\text{unit} \rightarrow \text{unit}), \text{unit}}$, and libraries $L_1, L_2 : \{m_1\} \rightarrow \{m_2\}$ and $L : \{m'_1\} \rightarrow \{m'_2\}$, as well as the unions $L \cup L_i : \{m_1, m_2\} \rightarrow \{m'_1, m'_2\}$. A possible trace in $\llbracket L \cup L_i \rrbracket_{\text{enc}}$ is this one:

$$\begin{aligned} h_i &= (1, \text{call } m_2())_{O\mathcal{K}} (1, \text{call } m_1(v))_{P\mathcal{L}} (1, \text{ret } m_1())_{O\mathcal{L}} \\ &\quad (1, \text{ret } m_2())_{P\mathcal{K}} (1, \text{call } m'_2())_{O\mathcal{K}} (1, \text{call } m'_1())_{P\mathcal{L}} (1, \text{call } v())_{O\mathcal{L}} \end{aligned}$$

which decomposes as $h_i = h' \bowtie_{\Pi, \emptyset}^{\sigma} h'_i$, with $\Pi = \{m_1, m_2, m'_1, m'_2\}$, $\sigma = 2222112$, $h' = (1, \text{call } m'_2())_{O\mathcal{K}} (1, \text{call } m'_1())_{P\mathcal{L}}$ and:

$$h'_i = (1, \text{call } m_2())_{O\mathcal{K}} (1, \text{call } m_1(v))_{P\mathcal{L}} (1, \text{ret } m_1())_{O\mathcal{L}} (1, \text{ret } m_2())_{P\mathcal{K}} (1, \text{call } v())_{O\mathcal{L}}$$

We now see that $h'_i \notin \llbracket L_i \rrbracket_{\text{enc}}$ as in the last move O is changing component from \mathcal{K} to \mathcal{L} .

We therefore look at compositionality for two specific cases: encapsulated sequencing (e.g. of $L : \Psi \rightarrow \Psi'$ with $L' : \Psi' \rightarrow \Psi''$) and disjoint union for first-order methods. Given $L : \Psi_1 \rightarrow \Psi_2$ and $L' : \Psi'_1 \rightarrow \Psi'_2$, we define their *disjoint union* $L \uplus L' = L \cup L' : (\Psi_1 \cup \Psi'_1) \rightarrow (\Psi_2 \cup \Psi'_2)$ under the assumption that $(\Psi_1 \cup \Psi_2) \cap (\Psi'_1 \cup \Psi'_2) = \emptyset$.

Theorem 57. *Let $L_1, L_2 : \Psi_1 \rightarrow \Psi_2$ and $L : \Psi'_1 \rightarrow \Psi'_2$. If $L_1 \triangleleft_{\text{enc}} L_2$ then:*

- assuming $\Psi'_2 = \Psi_1$, we have $L; L_1 \triangleleft_{\text{enc}} L; L_2$ and $L_1; L \triangleleft_{\text{enc}} L_2; L$;
- if $\Psi_1, \Psi_2, \Psi'_1, \Psi'_2$ are first-order then $L \uplus L_1 \triangleleft_{\text{enc}} L \uplus L_2$.

Proof. Let us consider the first sequencing case (the second one is dual), and assume that $L_1, L_2 : \Psi \rightarrow \Psi'$ and $L : \Psi'' \rightarrow \Psi$. Assume $L_1 \triangleleft_{\text{enc}} L_2$ and suppose $h_1 \in \llbracket L; L_1 \rrbracket_{\text{enc}}$. By Theorem 50, $h_1 = h' \bowtie_{\Pi, P}^{\sigma} h'_1$, where $h' \in \llbracket L \rrbracket$, $h'_1 \in \llbracket L_1 \rrbracket$ and method calls from Ψ are always scheduled with 0. The fact that O cannot switch between \mathcal{L}/\mathcal{K} components in (threads of) h_1 implies that the same holds for h', h'_1 , hence $h' \in \llbracket L \rrbracket_{\text{enc}}$ and $h'_1 \in \llbracket L_1 \rrbracket_{\text{enc}}$. Because $L_1 \triangleleft_{\text{enc}} L_2$, there exists $h'_2 \in \llbracket L_2 \rrbracket_{\text{enc}}$ such that $h'_1 \triangleleft h'_2$, i.e. $h'_1 \triangleleft_{P, O \cup \diamond}^* h'_2$. As before, some of the rearrangements necessary to transform h'_1 into h'_2 may concern actions shared by h'_1 and h' ; we need to check that these can lead to compatible $h'' \in \llbracket L \rrbracket_{\text{enc}}$. Let h'' be obtained by applying such

rearrangements to h' . We claim that $h' \triangleleft_{OP}^* h''$. The transpositions covered by \triangleleft_{PO} are treated as in Lemma 55. Suppose now that $(t', x')_{PK}(t, x)_{OL}$ are consecutive in h'_1 but swapped in order to obtain h'_2 , and $(t, x)_{OL}$ appears in h' as $(t, x)_{PK}$. Now, the move (t', x') cannot appear in h' as it is in L_1 's \mathcal{K} -component (L is the \mathcal{L} -component of L_1). Let s contain the moves of h' that are after (t', x') in the composition to h_1 , and before $(t, x)_{PK}$. We claim that s contains no moves from t , so h' can be directly composed with h'_2 as far as this transposition is concerned. Indeed, if s contained moves from t then, taking into account the encapsulation conditions, the leftmost one such would be some $(t, y)_{OK}$. But the \mathcal{K} -component of L is L_1 , which contradicts the fact that the moves we consider are consecutive in h'_1 . Hence, taking h_2 to be $h'' \wedge_{\Pi, P}^{\sigma'} h'_2$, where σ' is obtained from σ following the \triangleleft_{PO} move rearrangements, we have $h_2 \in \llbracket L; L_2 \rrbracket_{\text{enc}}$ and $h_1 \triangleleft_{\text{enc}} h_2$. Thus, $L; L_1 \triangleleft_{\text{enc}} L; L_2$.

The case of $L \uplus L_1 \triangleleft_{\text{enc}} L \uplus L_2$ is treated in a similar fashion. In this case, because of disjointness, the moves transposed in h'_1 do not have any counterparts in h' . Again, we consider consecutive moves $(t', x')_{PK}(t, x)_{OL}$ in h'_1 that are swapped in order to obtain h'_2 . Let s contain the moves of h' that are after (t', x') in the composition to h_1 , and before (t, x) . As Ψ_1, Ψ'_1 is first-order, $(t, x)_{OL}$ must be a return move and the t -move preceding it in h_1 must be the corresponding call. The latter is a move in h'_1 , which therefore implies that there can be no moves from t in s . Similarly for the other transposition case. \square

References

- [1] M. Herlihy, J.M. Wing, Linearizability: a correctness condition for concurrent objects, *ACM Trans. Program. Lang. Syst.* 12 (3) (1990) 463–492.
- [2] R. Jagadeesan, G. Petri, C. Pitcher, J. Riely, Quarantining weakness – compositional reasoning under relaxed memory models (extended abstract), in: *ESOP 2013*, 2013, pp. 492–511.
- [3] A. Cerone, A. Gotsman, H. Yang, Parameterised linearisability, in: *Proceedings of ICALP'14*, in: *Lecture Notes in Computer Science*, vol. 8573, Springer, 2014, pp. 98–109.
- [4] J. Laird, A game semantics of idealized CSP, in: *Proceedings of MFPS'01*, in: *Electron. Notes Theor. Comput. Sci.*, vol. 45, Elsevier, 2001, pp. 1–26.
- [5] D.R. Ghica, A.S. Murawski, Angelic semantics of fine-grained concurrency, in: *Proceedings of FOSSACS*, in: *Lecture Notes in Computer Science*, vol. 2987, Springer-Verlag, 2004, pp. 211–225.
- [6] A. Jeffrey, J. Rathke, A fully abstract may testing semantics for concurrent objects, *Theor. Comput. Sci.* 338 (1–3) (2005) 17–63.
- [7] J. Laird, A fully abstract trace semantics for general references, in: *Proceedings of ICALP*, in: *Lecture Notes in Computer Science*, vol. 4596, Springer, 2007, pp. 667–679.
- [8] D.R. Ghica, N. Tzevelekos, A system-level game semantics, *Electron. Notes Theor. Comput. Sci.* 286 (2012) 191–211.
- [9] D. Hendler, I. Incze, N. Shavit, M. Tzafir, Flat combining and the synchronization-parallelism tradeoff, in: *Proceedings of SPAA 2010*, 2010, pp. 355–364.
- [10] A.S. Murawski, N. Tzevelekos, Higher-order linearizability, in: *Proceedings of CONCUR*, in: *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 85, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, Chapter 34, 18 pp.
- [11] <http://c-cube.github.io/ocaml-containers/0.21/CCMultiSet.S.html>.
- [12] S. Heller, M. Herlihy, V. Luchangco, M. Moir, W.N. Scherer III, N. Shavit, A lazy concurrent list-based set algorithm, in: *OPDIS*, 2005, pp. 3–16.
- [13] P.W. O'Hearn, N. Rinetzky, M.T. Vechev, E. Yahav, G. Yorsh, Verifying linearizability with hindsight, in: *PODC*, 2010, pp. 85–94.
- [14] S. Abramsky, G. McCusker, Game semantics, in: H. Schwichtenberg, U. Berger (Eds.), *Logic and Computation*, Springer-Verlag, 1998, *Proceedings of the 1997 Marktoberdorf Summer School*.
- [15] M. Moir, N. Shavit, Concurrent data structures, in: *Handbook of Data Structures and Applications*, Chapman and Hall/CRC, 2004.
- [16] B. Dongol, J. Derrick, Verifying linearisability: a comparative survey, *ACM Comput. Surv.* 48 (2) (2015) 19.
- [17] I. Filipovic, P.W. O'Hearn, N. Rinetzky, H. Yang, Abstraction for concurrent objects, *Theor. Comput. Sci.* 411 (51–52) (2010) 4379–4398.
- [18] J. Derrick, G. Schellhorn, H. Wehrheim, Mechanically verified proof obligations for linearizability, *ACM Trans. Program. Lang. Syst.* 33 (1) (2011) 4.
- [19] A. Gotsman, H. Yang, Liveness-preserving atomicity abstraction, in: *Automata, Languages and Programming – 38th International Colloquium, ICALP 2011. Proceedings, Part II*, 2011, pp. 453–465.
- [20] H. Liang, X. Feng, Modular verification of linearizability with non-fixed linearization points, in: *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13. Proceedings*, 2013, pp. 459–470.
- [21] A.J. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, D. Dreyer, Logical relations for fine-grained concurrency, in: *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13*, 2013, pp. 343–356.
- [22] A. Turon, D. Dreyer, L. Birkedal, Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency, in: *ACM SIGPLAN International Conference on Functional Programming, ICFP'13*, 2013, pp. 377–390.
- [23] I. Sergey, A. Nanevski, A. Banerjee, Specifying and verifying concurrent algorithms with histories and subjectivity, in: *Programming Languages and Systems – 24th European Symposium on Programming, ESOP 2015. Proceedings*, 2015, pp. 333–358.
- [24] K. Svendsen, L. Birkedal, Impredicative concurrent abstract predicates, in: *Programming Languages and Systems – 23rd European Symposium on Programming, ESOP 2014. Proceedings*, 2014, pp. 149–168.
- [25] K. Svendsen, L. Birkedal, M.J. Parkinson, Joins: a case study in modular specification of a concurrent reentrant higher-order library, in: *ECOOP 2013 – Object-Oriented Programming – 27th European Conference. Proceedings*, 2013, pp. 327–351.