

Higher-Order Linearisability

Andrzej S. Murawski^a, Nikos Tzevelekos^{b,1}

^a*University of Oxford*

^b*Queen Mary University of London*

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.

Keywords: Linearisability, Concurrency, Higher-Order Computation

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 focussed on linearisability for the case where the allowable methods were restricted to first-order

¹Research supported by EPSRC (EP/P004172/1).




Figure 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 $\Psi', \Psi'' \vdash M_1 \parallel \dots \parallel M_N$.

26 functions ($\text{int} \rightarrow \text{int}$). Herein, we give a systematic exposition of linearisability for general
 27 higher-order concurrent libraries, where methods can be of arbitrary higher-order
 28 types. In doing so, we also propose a corresponding notion of sequential history for
 29 higher-order library interactions.

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

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

43 Our framework is operational in flavour and draws upon concurrent [4, 5] and
 44 operational game semantics [6, 7, 8]. We shall model library use as a game between two
 45 participants: *Player* (P), corresponding to the library L , and *Opponent* (O), representing
 46 the environment (L', K) in which the library was deployed. Each call will be of the
 47 form $\text{call } m(v)$ with the corresponding return of the shape $\text{ret } m(v)$, where v is a
 48 value. As we work in a higher-order framework, v may contain functions, which can
 49 participate in subsequent calls and returns. Histories will be sequences of *moves*, which
 50 are calls and returns paired with thread identifiers. A history is sequential just if every
 51 move produced by O is immediately followed by a move by P in the same thread. In
 52 other words, the library immediately responds to each call or return delivered by the
 53 environment. In contrast to classic linearisability, the move by O and its response by P
 54 need not be a call/return pair, as the higher-order setting provides more possibilities (in
 55 particular, the P response may well be a call). Accordingly, linearisable higher-order
 56 histories can be seen as sequences of atomic segments (linearisation points), starting at
 57 environment moves and ending with corresponding library moves.

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

```

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) )

```

```

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 ...
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

```

Figure 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_{mset2} (lines 8-19 as in LHS) with public methods $\text{count}, \text{reset} : \text{int} \rightarrow \text{int}$, $\text{update} : \text{int} \times (\text{int} \rightarrow \text{int}) \rightarrow \text{int}$; abstract method $\text{default} : \text{int} \rightarrow \text{int}$.

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

71 *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:

$$\text{count} : \text{int} \rightarrow \text{int}, \quad \text{update} : (\text{int} \times (\text{int} \rightarrow \text{int})) \rightarrow \text{int}.$$

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

81 **Example 1 (Multiset).** Consider the concurrent multiset library L_{mset} in Figure 2 on
82 the LHS (the RHS will be discussed only later). It uses a private reference for storing the
83 multiset’s characteristic function and reads *optimistically*, without locking (cf. [12, 13]).
84 The update method in particular reads the current multiplicity of the given element i (via
85 count) and computes its new multiplicity without acquiring a lock on the characteristic
86 function. It only acquires a lock when it is ready to write the new value (line 10) in the

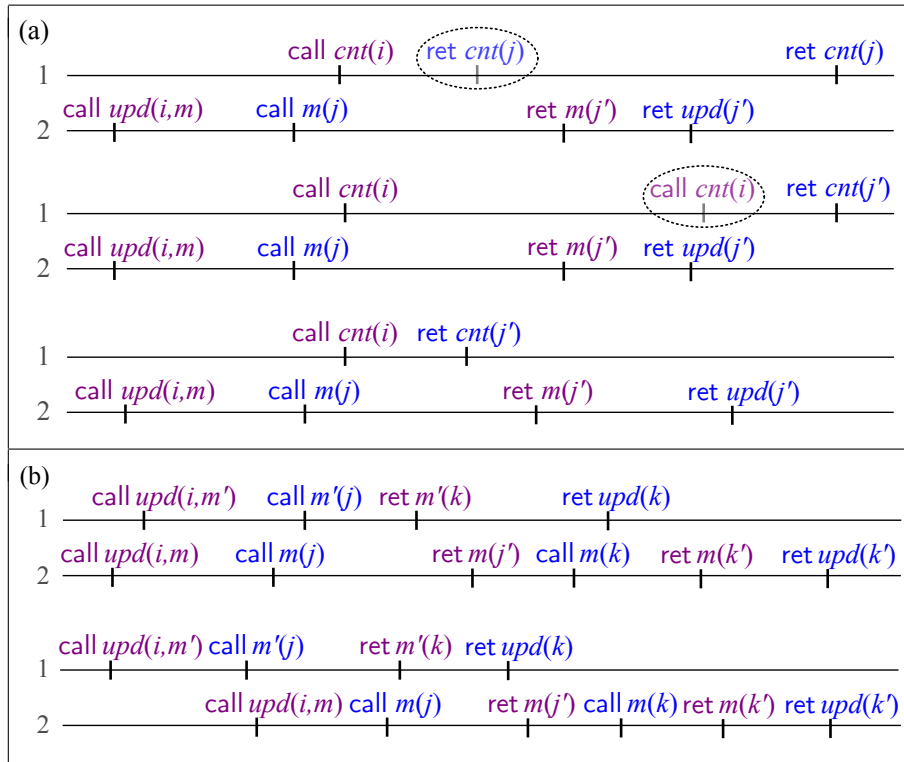


Figure 3: Example histories of L_{mset} .

87 hope that the value at i will still be the same and the update can proceed; if not, another
 88 attempt to update the value is made.

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

95 In the first execution, the client calls $\text{update}(i, m)$ in the second thread, and sub-
 96 sequently calls $\text{count}(i)$ in the first thread. The code for *update* stipulates that first
 97 $\text{count}(i)$ be called internally, returning some multiplicity j for i , and then $m(j)$ should
 98 be called. As soon m returns a value j' , *update* sets the multiplicity of i to j' and itself
 99 returns j' . The last event in this history is a return of *count* in the first thread with the
 100 old value j . According to our proposed definition, this history will be linearisable to
 101 another, intuitively correct one: the last return can be moved to the circled position. At
 102 this point the notion of linearisability is used informally, but it will be made precise
 103 in the following sections. In the second execution, the last return of *count* in the first
 104 thread returns the updated value. In this case, we will be able to move $\text{call cnt}(i)$ to the
 105 circled position to obtain a linearisation, which is obviously correct. Finally, in the third
 106 execution we have a history that will turn out non-linearisable to an intuitively correct
 107 history. Indeed, we should not be able to return the updated value in the first thread
 108 before m has returned it in the second one.

109 The two histories in block (b) in the same figure demonstrate the mechanism for

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

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

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

124 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 $\text{Meths} = \uplus_{\theta, \theta'} \text{Meths}_{\theta, \theta'}$,
 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.

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

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

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

135 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.

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

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

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

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

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

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

148 where, if $m \in \text{Meths}_{\theta, \theta'}$, the types of v, v' must match θ, θ' respectively. We let
 149 $\text{PreH} = \text{PreH}_O \cup \text{PreH}_P$.

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

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

- 157 1. The name of any method called in h must come from Ψ or Ψ' , or be introduced
 158 earlier in h as a higher-order argument or result (*no methods out of thin air*). In
 159 addition:
 - 160 • if the method is from Ψ' , the call must be tagged with O (i.e. issued by K);
 - 161 • if the method is from Ψ , the call must be tagged with P (i.e. issued by L
 162 towards L');
 - 163 • for a call of method $m \notin \Psi \cup \Psi'$ to be valid, m must be introduced in an earlier
 164 move of dual polarity (*calls dual to introductions*).
- 165 2. Any method name appearing inside a call or return argument in h must be *fresh*, i.e.
 166 not used earlier (*introductions always fresh*).
 - 167 • This reflects the assumption that methods can be called and returned from, but
 168 not compared for identity equality. It is therefore a requirement towards the
 169 *completeness* of histories as a semantics for concurrent libraries. For example,
 170 this ensures that rules like η -equality are preserved in the semantics.
 - 171 • The condition serves the additional purpose of making the setting described
 172 in Figure 1 robust, as it prevents method names in Ψ from being leaked to the
 173 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 :

$$t \times ((x_1)_{X_1} (x_2)_{X_2} \cdots (x_k)_{X_k}) = (t, x_1)_{X_1} (t, x_2)_{X_2} \cdots (t, x_k)_{X_k}.$$

174 We say that a move $(t, x)_X$ *introduces* a name $m \in \text{Meths}$ when $x \in \{\text{call } m'(v), \text{ret } m'(v)\}$
 175 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

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

where $(1 \times h_1) \mid \dots \mid (N \times h_N)$ 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

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

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 Figure 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:

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

while the next ones are not:

$$\begin{aligned}
& (1, \text{call } \text{upd}(5, m))_O (1, \text{call } m(5))_P (1, \text{call } \text{cnt}(5))_O (1, \text{ret } \text{cnt}(42))_P \cdots \\
& (1, \text{call } \text{upd}(5, m))_O (2, \text{call } \text{upd}(6, m'))_O \cdots \\
& (1, \text{call } \text{upd}(5, m))_O (1, \text{call } m(5))_P (2, \text{call } \text{upd}(5, m'))_O (2, \text{call } m'(5))_P \\
& \quad (1, \text{ret } m(42))_O (1, \text{ret } \text{upd}(42))_P (3, \text{call } \text{cnt}(5))_O (3, \text{ret } \text{cnt}(42))_P \\
& \quad (2, \text{ret } m'(24))_O (2, \text{ret } \text{upd}(24))_P (1, \text{call } \text{cnt}(5))_O (1, \text{ret } \text{cnt}(42))_P
\end{aligned}$$

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

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

218 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 Ξ_{pre} for the prefix relation, we define $A_{\text{mset}}^{\circ} = \{ \pi_1(s) \mid s \in A_{\text{mset}}^{\circ} \}$ where:

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

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

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

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

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

228 2.2. Three notions of linearisability

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

235 2.2.1. General linearisability

236 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 $t \neq t'$:

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

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

238 Intuitively, two histories h_1, h_2 are related by \triangleleft_{PO} if the latter can be obtained from
 239 the former by swapping two adjacent moves from different threads in such a way that,
 240 after the swap, a P -move will occur earlier or an O -move will occur later. Note that the
 241 relation always applies to adjacent moves of the same polarity. On the other hand, we
 242 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.

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

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

245 As another example, we can revisit the histories in Figure 2. There, O -moves are
 246 coloured purple and P -moves are blue. In part (a) we can see that:

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

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

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

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

258 Given libraries $L, L' : \Psi \rightarrow \Psi'$ and a set of sequential histories $A \subseteq \mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$, we write
 259 $L \triangleleft A$, and say that L can be linearised to A , if for any $h \in \llbracket L \rrbracket$ there exists $h' \in A$ such
 260 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
 261 sequential $h' \in \llbracket L' \rrbracket$ such that $h \triangleleft h'$).

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

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

271 For further comparison, recall that the classic definition allowed for call/call,
 272 ret/ret and call/ret swaps, but ret/call was forbidden. According to our definition,
 273 what is allowed depends on polarity, so a call/call swap may well be illegal if the
 274 first call is a P -move and the second call is an O -move. Similarly, a ret/call swap is
 275 allowed as long as both actions belong to the same player or the return is an O -action
 276 and the call is a P -action. For instance, Example 8 involves the following kinds of
 277 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} .

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

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

285 2.2.2. Encapsulated linearisability

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

- 290 • a move with side \mathcal{K} is played between the library L and the client K , while
- 291 • 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:

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

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

292 where, in the latter case, $(t, \text{call } m(v'))$ is the corresponding call of $(t, \text{ret } m(v))$. Thus,
 293 every move in h can be assigned a unique side polarity from $\{\mathcal{K}, \mathcal{L}\}$. For simplicity,

294 we shall be tagging moves with a second index $Y \in \{\mathcal{K}, \mathcal{L}\}$ corresponding to their side
 295 polarity.

296 In this more restrictive nature of interaction, in which K and L' are separated, in
 297 addition to sequentiality in every thread we shall insist that a move made by the library
 298 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

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

299 and moves from t are absent from s_2 then $Y = Y'$. Moreover, if $L : \Psi \rightarrow \Psi'$, we set
 300 $\mathcal{H}_{\Psi, \Psi'}^{\text{enc}} = \{h \in \mathcal{H}_{\Psi, \Psi'} \mid h \text{ encapsulated}\}$ 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 $t \neq t'$:

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

301 **Definition 13 (Encapsulated linearisability).** Given $h_1, h_2 \in \mathcal{H}_{\Psi, \Psi'}^{\text{enc}}$, we say that h_1 is
 302 *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.

303 A library $L : \Psi \rightarrow \Psi'$ can be *enc-linearised to* A , written $L \triangleleft_{\text{enc}} A$, if $A \subseteq \mathcal{H}_{\Psi, \Psi'}^{\text{seq}} \cap$
 304 $\mathcal{H}_{\Psi, \Psi'}^{\text{enc}}$, and for any $h \in \llbracket L \rrbracket_{\text{enc}}$ there exists $h' \in A$ such that $h \triangleleft_{\text{enc}} h'$. We write $L \triangleleft_{\text{enc}} L'$
 305 if $L \triangleleft_{\text{enc}} \llbracket L' \rrbracket_{\text{enc}} \cap \mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$.

306 **Remark 14.** Suppose $\Psi = \{m : \text{int} \rightarrow \text{int}\}$ and $\Psi' = \{m' : \text{int} \rightarrow \text{int}\}$. Histories from
 307 $\mathcal{H}_{\Psi, \Psi'}$ may contain the following actions only: call $m'(i)_{OK}$, ret $m(i)_{OL}$, call $m(i)_{PL}$,
 308 ret $m'(i)_{PK}$. Then $(\triangleleft_{PO} \cup \diamond)^*$ prevents call $m(i)_{PL}$ from being swapped with
 309 ret $m(i)_{OL}$ and, similarly, for ret $m'(i)_{PK}$ and call $m'(i)_{OK}$, i.e. it coincides with
 310 Definition 3 of [3].

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

317 **Example 16 (Parameterised multiset).** We revisit the multiset library of Example 1
 318 and extend it with a public method *reset*, which performs multiplicity resets to default
 319 values using an abstract method *default* as the default-value function (again, we use
 320 absolute values to avoid negative multiplicities). The extended library is shown in
 321 the RHS of Figure 2 and written $L_{\text{mset2}} : \Psi \rightarrow \Psi'$, with $\Psi = \{\text{default}\}$ and $\Psi' =$
 322 $\{\text{count}, \text{update}, \text{reset}\}$. In contrast to the *update* method of L_{mset} , *reset* is not optimistic:
 323 it retrieves the lock upon its call, and only releases it before return. In particular, the
 324 method calls *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 $A_{\text{mset2}} = \{\pi_1(s) \mid s \in A_{\text{mset2}}^\circ\}$ where:

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

```

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;

```

Figure 4: Flat combination library L_{fc} .

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

$$\mathcal{S} \rightarrow (t, \text{call } \text{reset}(i), I)_{OK} (t, \text{call } \text{default}(i), I)_{PL} (t, \text{ret } \text{default}(j), I)_{OL} (t, \text{ret } \text{reset}(|j|), I')_{PK} \mathcal{S}$$

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

327 2.2.3. Relational linearisability

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

336 **Definition 17 (Relational linearisability).** Let $\mathcal{R} \subseteq \mathcal{H}_{\Psi, \Psi'} \times \mathcal{H}_{\Psi, \Psi'}$ be closed under
337 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
338 \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
339 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
340 $h \in \llbracket L \rrbracket$ there exists $h' \in A$ such that $h \triangleleft_{\mathcal{R}} h'$. We write $L \triangleleft_{\mathcal{R}} L'$ if $L \triangleleft_{\mathcal{R}} \llbracket L' \rrbracket \cap \mathcal{H}_{\Psi, \Psi'}^{\text{seq}}$.

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

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

355 $run = \lambda (f, x). (lock.acquire (); \text{let } result = f(x) \text{ in } lock.release (); result)$

356 where each function call in L_{spec} is protected by a lock. Observe that we cannot hope
 357 for $L_{fc} \triangleleft L_{spec}$, because clients may call library methods with functional arguments
 358 that recognise thread identity. Consequently, we can relate the two libraries only if
 359 context behaviour is guaranteed to be independent of thread identifiers. This can be
 360 expressed through $\triangleleft_{\mathcal{R}}$, where $\mathcal{R} \subseteq \mathcal{H}_{\emptyset, \Psi'} \times \mathcal{H}_{\emptyset, \Psi'}$ is a relation capturing thread-blind
 361 client behaviour (see Subsection 3.2 for details).

362 3. Library syntax and semantics

363 We now look at the concrete syntax of libraries and clients. Libraries comprise
 364 collections of typed methods whose argument and result types adhere to the grammar:
 365 $\theta ::= \text{unit} \mid \text{int} \mid \theta \rightarrow \theta \mid \theta \times \theta$.

366 We shall use three disjoint enumerable sets of names, referred to as Vars, Meths
 367 and Refs, to name respectively variables, methods and references. x, f (and their
 368 decorated variants) will be used to range over Vars; m will range over Meths; and r
 369 over Refs. Methods and references are implicitly typed, i.e. $\text{Meths} = \uplus_{\theta, \theta'} \text{Meths}_{\theta, \theta'}$
 370 and $\text{Refs} = \text{Refs}_{\text{int}} \uplus \uplus_{\theta, \theta'} \text{Refs}_{\theta, \theta'}$, where $\text{Meths}_{\theta, \theta'}$ contains names for methods of type
 371 $\theta \rightarrow \theta'$, Refs_{int} contains names of integer references and $\text{Refs}_{\theta, \theta'}$ contains names for
 372 references to methods of type $\theta \rightarrow \theta'$. We write \uplus for disjoint set union.

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

378 Terms M specify the shape of allowable method bodies. $()$ is the skip command,
 379 i ranges over integers, tid is the current thread identifier and \oplus represents standard
 380 arithmetic operations. Thanks to higher-order references, we can simulate divergence
 381 by $(!r)()$, where $r \in \text{Refs}_{\text{unit}, \text{unit}}$ is initialised with $\lambda x^{\text{unit}}. (!r)()$. Similarly, while $M N$
 382 can be simulated by $(!r)()$ after $r := \lambda x^{\text{unit}}. \text{let } y = M \text{ in } (\text{if } y \text{ then } (N; (!r)()) \text{ else } ())$.
 383 We also use the standard derived syntax for sequential composition, i.e. $M; N$ stands for
 384 $\text{let } x = M \text{ in } N$, where x does not occur in N . For each term M , we write $\text{Meths}(M)$
 385 for the set of method names occurring in M . We use the same notation for method
 386 names in blocks and libraries.

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

<p><i>Libraries</i> $L ::= B \mid \text{abstract } m; L \mid \text{public } m; L$</p> <p><i>Blocks</i> $B ::= \epsilon \mid m = \lambda x.M; B \mid r := \lambda x.M; B \mid r := i; B$</p> <p><i>Terms</i> $M ::= () \mid i \mid \text{tid} \mid x \mid m \mid M \oplus M \mid \langle M, M \rangle \mid \pi_1 M \mid \pi_2 M \mid \text{if } M \text{ then } M \text{ else } M$ $\mid \lambda x^\theta.M \mid xM \mid mM \mid \text{let } x = M \text{ in } M \mid r := M \mid !r$</p>	<p><i>Clients</i> $K ::= M \parallel \dots \parallel M$</p> <p><i>Values</i> $v ::= () \mid i \mid m \mid \langle v, v \rangle$</p>
$\frac{}{\Gamma \vdash () : \text{unit}} \quad \frac{}{\Gamma \vdash i : \text{int}} \quad \frac{}{\Gamma \vdash \text{tid} : \text{int}} \quad \frac{\Gamma(x) = \theta \quad m \in \text{Meths}_{\theta, \theta'}}{\Gamma \vdash x : \theta} \quad \frac{\Gamma \vdash M : \text{int} \quad \Gamma \vdash M_0, M_1 : \theta}{\Gamma \vdash \text{if } M \text{ then } M_1 \text{ else } M_0 : \theta}$ $\frac{\Gamma \vdash M : \theta_1 \times \theta_2}{\Gamma \vdash \pi_i M : \theta_i \quad (i = 1, 2)} \quad \frac{\Gamma \vdash M_i : \theta_i \quad (i = 1, 2)}{\Gamma \vdash \langle M_1, M_2 \rangle : \theta_1 \times \theta_2} \quad \frac{\Gamma \vdash M_1, M_2 : \text{int}}{\Gamma \vdash M_1 \oplus M_2 : \text{int}} \quad \frac{\Gamma, x : \theta \vdash M : \theta'}{\Gamma \vdash \lambda x^\theta.M : \theta \rightarrow \theta'}$ $\frac{\Gamma(x) = \theta \rightarrow \theta' \quad \Gamma \vdash M : \theta}{\Gamma \vdash xM : \theta'} \quad \frac{m \in \text{Meths}_{\theta, \theta'} \quad \Gamma \vdash M : \theta}{\Gamma \vdash mM : \theta'} \quad \frac{\Gamma \vdash M : \theta \quad \Gamma, x : \theta \vdash N : \theta'}{\Gamma \vdash \text{let } x = M \text{ in } N : \theta'}$ $\frac{r \in \text{Refs}_{\text{int}} \quad \Gamma \vdash M : \text{int}}{\Gamma \vdash r := M : \text{unit}} \quad \frac{r \in \text{Refs}_{\theta, \theta'} \quad \Gamma \vdash M : \theta \rightarrow \theta'}{\Gamma \vdash r := M : \text{unit}} \quad \frac{r \in \text{Refs}_{\text{int}}}{\Gamma \vdash !r : \text{int}} \quad \frac{r \in \text{Refs}_{\theta, \theta'}}{\Gamma \vdash !r : \theta \rightarrow \theta'}$	
$\frac{m \in \text{Meths}_{\theta, \theta'} \quad x : \theta \vdash M : \theta' \quad \vdash_B B : \Psi}{\vdash_B \epsilon : \emptyset} \quad \frac{\vdash_B m = \lambda x.M; B : \Psi \uplus \{m\}}{\vdash_B r := \lambda x.M; B : \Psi}$ $\frac{r \in \text{Refs}_{\text{int}} \quad \vdash_B B : \Psi}{\vdash_B r := i; B : \Psi} \quad \frac{\vdash_B B : \Psi}{\text{Meths}(B) \vdash_L B : \emptyset \rightarrow \Psi} \quad \frac{\Psi \uplus \{m\} \vdash_L L : \Psi' \rightarrow \Psi'' \quad m \in \Psi''}{\Psi \vdash_L \text{public } m; L : \Psi' \rightarrow \Psi''}$ $\frac{\Psi \uplus \{m\} \vdash_L L : \Psi' \rightarrow \Psi'' \quad m \notin \Psi''}{\Psi \vdash_L \text{abstract } m; L : \Psi' \uplus \{m\} \rightarrow \Psi''} \quad \frac{\vdash_{M_j} : \text{unit} \quad (j = 1, \dots, N)}{\Psi \vdash_K M_1 \parallel \dots \parallel M_N : \text{unit}} \quad \forall j. \text{Meths}(M_j) \subseteq \Psi$	

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

393 For simplicity, we do not include private methods, yet the same effect could be
394 achieved by storing them in higher-order references. As we explain in the next sec-
395 tion, references present in library definitions are de facto private to the library. Note
396 also that, according to our definition, sets of abstract and public methods are disjoint.
397 However, given $m, m' \in \text{Refs}_{\theta, \theta'}$, one can define a “public abstract” method with:
398 $\text{public } m; \text{abstract } m'; m = \lambda x^\theta.m'x$.

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

403 Finally, we type libraries using statements of the form $\Psi \vdash_L L : \Psi' \rightarrow \Psi''$, where
404 $\Psi, \Psi', \Psi'' \subseteq \text{Meths}$ and $\Psi' \cap \Psi'' = \emptyset$. The judgment $\emptyset \vdash_L L : \Psi' \rightarrow \Psi''$ guarantees that
405 any method occurring in L is present either in Ψ' or Ψ'' , that all methods in Ψ' are
406 declared as abstract and unimplemented, while all methods in Ψ'' are declared as public
407 and defined. Thus, $\emptyset \vdash_L L : \Psi \rightarrow \Psi'$ is a library in which Ψ, Ψ' are the abstract and
408 public methods respectively. In this case, we also write $L : \Psi \rightarrow \Psi'$.

409 3.1. Semantics

410 The semantics of our system is given in several stages. First, we define an operational
411 semantics for sequential and concurrent terms that may draw methods from a repository.
412 We then adapt it to capture interactions of concurrent clients with closed libraries (no
413 abstract methods). This notion is then used to define contextual approximation for

$$\begin{array}{c}
\begin{array}{c}
(L) \longrightarrow_{\text{lib}} (L, \emptyset, S_{\text{init}}) \quad (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) \quad (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) \quad (r := \lambda x.M; B, \mathcal{R}, S) \longrightarrow_{\text{lib}} (B, \mathcal{R}_{**}, S[r \mapsto m])
\end{array} \\
\hline
\begin{array}{c}
(E[\text{tid}], \mathcal{R}, S) \rightarrow_t (E[t], \mathcal{R}, S) \quad (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) \quad (E[\pi_j \langle v_1, v_2 \rangle], \mathcal{R}, S) \rightarrow_t (E[v_j], \mathcal{R}, S) \\
(E[!r], \mathcal{R}, S) \rightarrow_t (E[S(r)], \mathcal{R}, S) \quad (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]) \quad (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) \quad (E[mv], \mathcal{R}_*, S) \rightarrow_t (E[M\{v/x\}], \mathcal{R}_*, S)
\end{array} \\
\hline
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 \\
\hline
\begin{array}{c}
(M, \mathcal{R}, S) \rightarrow_t (M', \mathcal{R}', S') \\
\hline
(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') \quad (K_N)
\end{array}
\end{array}$$

Figure 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 $i_* = 0$.

414 arbitrary libraries. Finally, we introduce a trace semantics of arbitrary libraries, which
415 generates the histories on which our notions of linearisability are based.

416 3.1.1. Library-client evaluation

417 Libraries, terms and clients are evaluated in environments comprising:

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

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

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

430 First we evaluate the library to create an initial repository and store. This is achieved
431 by the first set of rules in Figure 6, where we assume that S_{init} is empty. Thus, library
432 evaluation produces a tuple $(\epsilon, \mathcal{R}_0, S_0)$ including a method repository and a store, which
433 can be used as the initial repository and store for evaluating $M_1 \parallel \dots \parallel M_N$ using the (K_N)
434 rule. We shall call the latter evaluation semantics for clients (denoted by \Longrightarrow) the
435 *multi-threaded operational semantics*. The latter relies on closed-term reduction (\rightarrow_t),
436 whose rules are given in the middle group, where t is the current thread index. Note
437 that the rules for $E[\lambda x.M]$ in the middle group, along with those for $m = \lambda x.M$ and
438 $r := \lambda x.M$ in the first group, involve the creation of a fresh method name m , which is
439 used to put the function in the repository \mathcal{R} . Name creation is non-deterministic: any
440 fresh m of the appropriate type can be chosen.

441 We define termination for clients linked with libraries that have no abstract methods.

442 Recall our convention (Remark 20) that L and M_1, \dots, M_N must access disjoint parts of
 443 the store. Terms M_1, \dots, M_N can share reference names, though.

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

447 We shall build a notion of contextual approximation of libraries on top of termination:
 448 one library will be said to approximate another if, whenever the former terminates when
 449 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:

$$\begin{aligned} L_1 \cup L_2 : (\Psi_1 \cup \Psi'_1) \setminus (\Psi_2 \cup \Psi'_2) \rightarrow \Psi_2 \cup \Psi'_2 &= (D_1; B_1) \cup (D_2; B_2) = D'_1; D'_2; B_1; B_2 \\ L_1 \setminus \Psi : \Psi_1 \rightarrow (\Psi_2 \setminus \Psi) &= (D_1; B_1) \setminus \Psi = D''_1; B'_1 \{!r_1/m_1\} \dots \{!r_n/m_n\} \\ L'; L_1 : \emptyset \rightarrow \Psi_2, \Psi' &= (L' \cup L_1) \setminus \Psi_1 \end{aligned}$$

450 where D'_1 is D_1 with any abstract m declaration removed for $m \in \Psi'_2$, dually for D'_2 ;
 451 and where D''_1 is D_1 without public m declarations for $m \in \Psi$ and each r_i is a fresh
 452 reference matching the type of m_i , and B'_1 is obtained from B_1 by replacing each
 453 $m_i = \lambda x.M$ by $r_i := \lambda x.M$. Thus, the union of libraries L_1 and L_2 corresponds to
 454 merging their code and removing any abstract declarations for methods defined in the
 455 union. On the other hand, the hiding of a public method simply renders it private via the
 456 use of references. Sequencing allows for the following notion.

457 **Definition 22.** Given $L_1, L_2 : \Psi \rightarrow \Psi'$, we say that L_1 *contextually approximates*
 458 L_2 , written $L_1 \sqsubseteq L_2$, if for all $L' : \emptyset \rightarrow \Psi, \Psi''$ and $\Psi', \Psi'' \vdash_{\mathcal{K}} M_1 \parallel \dots \parallel M_N : \text{unit}$, if
 459 $\text{link } L'; L_1$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$ then $\text{link } L'; L_2$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$. In this case, we also
 460 say that L_2 *contextually refines* L_1 .

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

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

471 3.1.2. Trace semantics

472 Building on the earlier semantics, we next introduce a trace semantics of libraries
 473 in the spirit of game semantics [14]. As mentioned in Section 2, the behaviour of a
 474 library will be represented as an exchange of moves between two players called P and
 475 O , representing the library and its corresponding context respectively. The context

476 consists of the client of the library as well as the parameter library, with an index on
 477 each move (\mathcal{K}/\mathcal{L}) specifying which of them is involved in the move.

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

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

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

489 where the component \mathcal{E} is an *evaluation stack*, that is, a stack of the form $[X_1, X_2, \dots, X_n]$
 490 with each X_i being either an evaluation context or a method name. On the other
 491 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
 492 $\mathcal{A} = (\mathcal{A}_{\mathcal{L}}, \mathcal{A}_{\mathcal{K}})$ is a pair of sets of *abstract* method names. \mathcal{P} will be used to record
 493 all the method names produced by P and passed to O : those passed to OK are stored
 494 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
 495 one: $\mathcal{A}_{\mathcal{K}}$ ($\mathcal{A}_{\mathcal{L}}$) stores the method names produced by OK (resp. OL) and passed to P .
 496 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.

497 Given a pair \mathcal{P} as above and a set $Z \subseteq \text{Meths}$, we write $\mathcal{P} \cup_{\mathcal{K}} Z$ for the pair
 498 $(\mathcal{P}_{\mathcal{L}}, \mathcal{P}_{\mathcal{K}} \cup Z)$. We define $\cup_{\mathcal{L}}$ in a similar manner, and extend it to pairs \mathcal{A} as well.
 499 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} :
 500 $\phi(\mathcal{P}, \mathcal{A}) = \text{Meths} \setminus (\mathcal{P}_{\mathcal{L}} \cup \mathcal{P}_{\mathcal{K}} \cup \mathcal{A}_{\mathcal{L}} \cup \mathcal{A}_{\mathcal{K}})$.

501 We give the rules generating the trace semantics in Figure 7. Note that the rules are
 502 parameterised by: P/O and Y , which together determine the polarity of the next move;
 503 C/R , which stands for the move being a call (C) or a return (R) respectively. The rules
 504 depict the intuition presented above. When in an *O*-configuration, the context may issue
 505 a call to a public method $m \in \mathcal{P}_Y$ and pass control to the library (rule (OCY)). Note that,
 506 when this occurs, the name m is added to the evaluation stack \mathcal{E} and a *P*-configuration
 507 is obtained. From there on, the library will compute internally using rule (INT), until: it
 508 either needs to evaluate an abstract method (i.e. some $m' \in \mathcal{A}_Y$), and hence issues a call
 509 via rule (PCY); or it completes its computation and returns the call (rule (PRY)). Calls
 510 to abstract methods, on the other hand, are met either by further calls to public methods
 511 (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-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

- (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 $\mathcal{A}' = \mathcal{A} \cup_Y \{m_1, \dots, m_k\}$.

Figure 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 .

concurrent traces are produced by the following two rules

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

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

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

514 We can now define the trace semantics of a library L . We call a configuration
515 component \mathcal{C}_i **final** if it is in one of the following forms, for O - and P -configurations
516 respectively: $\mathcal{C}_i = ([], -)$ or $\mathcal{C}_i = ([], ())$. We call $(\vec{\mathcal{C}}, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$ final just if $\vec{\mathcal{C}} =$
517 $\mathcal{C}_1 \parallel \dots \parallel \mathcal{C}_N$ 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:

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

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

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

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

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

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

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

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

529 3.2. Proofs of examples

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

532 Recall the multiset library L_{mset} and the specification A_{mset} of Example 1 and
 533 Figure 2. We show that $L_{\text{mset}} \triangleleft A_{\text{mset}}$. More precisely, taking an arbitrary history
 534 $h \in \llbracket L_{\text{mset}} \rrbracket$ we show that h can be rearranged using \triangleleft_{PO}^* to match an element of
 535 A_{mset} . We achieve this by identifying, for each O -move $(t, x)_O$ and its following
 536 P -move $(t, x')_P$ in h , a *linearisation point* between them, i.e. a place in h to which
 537 $(t, x)_O$ can be moved right and to which $(t, x')_P$ can be moved left so that they become
 538 consecutive and, moreover, the resulting history is still produced by L_{mset} . After all
 539 these rearrangements, we obtain a sequential history \hat{h} such that $h \triangleleft \hat{h}$ and \hat{h} is also
 540 produced by L_{mset} . It then suffices to show that $\hat{h} \in A_{\text{mset}}$.

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

542 *Proof.* Given some $h \in \llbracket L_{\text{mset}} \rrbracket$, let us assume that h has been generated by a sequence
 543 $\rho_1 \Rightarrow \rho_2 \Rightarrow \dots \Rightarrow \rho_k$ of atomic transitions and that the variable F of L_{mset} is instantiated
 544 with a reference r_F . We demonstrate the linearisation points for pairs of (O, P) moves in
 545 h , by case analysis on the moves (we drop \mathcal{K} indices from moves as they are ubiquitous).
 546 Line numbers below refer to the LHS of Figure 2.

- 547 1. $h = \dots (t, \text{call } \text{cnt}(i))_O s (t, \text{ret } \text{cnt}(i'))_P \dots$. Here the linearisation point (LP) is the
 548 configuration ρ_j that dereferences r_F as per line 5 in L_{mset} (the $!F$ expression).
- 549 2. $h = \dots (t, \text{call } \text{upd}(i, m))_O s (t, \text{call } m(j))_P \dots$. The LP is the dereferencing of r_F
 550 in line 5 (called from within *update*).
- 551 3. $h = \dots (t, \text{ret } m(j'))_O s (t, \text{ret } \text{upd}(|j'|))_P \dots$. The LP is the update of r_F in line 13.
- 552 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
 553 line 11.

554 Each of the linearisation points above specifies a PO -rearrangement of moves. For in-
 555 stance, for $h = s_0 (t, \text{call } \text{cnt}(i))_O s (t, \text{ret } \text{cnt}(i'))_P s'$, let $s = s_1 s_2$ where $s_0 (t, \text{call } \text{cnt}(i))_O s_1$
 556 is the prefix of h produced by $\rho_1 \Rightarrow \rho_2 \Rightarrow \dots \Rightarrow \rho_j$. The rearrangement of h is then
 557 $\hat{h} = s_0 s_1 (t, \text{call } \text{cnt}(i))_O (t, \text{ret } \text{cnt}(i'))_P s_2 s'$. We thus obtain $h \triangleleft_{PO}^* \hat{h}$.

558 The selection of linearisation points is such that it guarantees that $\hat{h} \in \llbracket L_{\text{mset}} \rrbracket$. E.g. in
 559 case 1, the transitions occurring in thread t between $(t, \text{call } \text{cnt}(i))_O$ and configuration
 560 ρ_j do not access r_F . Hence, we can postpone them and fire them in sequence just before
 561 ρ_j . After ρ_{j+1} and until $(t, \text{ret } \text{cnt}(i'))_P$ there is again no access of r_F in t and we
 562 can thus bring forward the corresponding transitions just after ρ_{j+1} . Similar reasoning

563 applies to case 2. In case 4, we reason similarly but also take into account that rendering
 564 the acquisition of the lock by t atomic is sound (i.e. the semantics can produce the
 565 rearranged history). Case 3 is similar, but we also use the fact that the access to r_F in
 566 lines 10-15 is inside the lock, and hence postponing dereferencing (line 11) to occur in
 567 sequence before update (line 13) is sound.

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

572 **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.

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

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

- 576 • Let us write s_3 as $s_3 = s_3^1 s_3^2$, where the split is at the linearisation point. Since the
 577 lock is constantly held by thread t in $s_2 s_3^1$, there can be no calls or returns to *default* in
 578 $s_2 s_3^1$. Hence, all moves in $s_2 s_3^1$ are in component \mathcal{K} and can be transposed with the \mathcal{L} -
 579 moves above, using \diamond^* , to obtain $h' = s(t, \text{call } \text{reset}(i))_{OK} s_1 s_2 s_3^1(t, \text{call } \text{default}(j))_{PL}$
 580 $(t, \text{ret } \text{default}(j'))_{OL} s_3^2(t, \text{ret } \text{reset}(|j'|))_{PK} \dots$
- 581 • Next, by *PO*-rearrangement we obtain $\hat{h} = s s_1 s_2 s_3^1(t, \text{call } \text{reset}(i))_{OK}(t, \text{call } \text{default}(j))_{PL}$
 582 $(t, \text{ret } \text{default}(j'))_{OL}(t, \text{ret } \text{reset}(|j'|))_{PK} s_3^2 \dots$. Thus, $h(\triangleleft_{PO} \cup \diamond)^* \hat{h}$.

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

587 For our last example, recall the flat combination library $L_{\text{fc}} : \emptyset \rightarrow \Psi'$ of Example 18,
 588 and Figure 4, along with its specification library $L_{\text{spec}} : \emptyset \rightarrow \Psi'$, where $\Psi' = \{\text{run} \in$
 589 $\text{Meths}_{(\theta \rightarrow \theta') \times \theta, \theta'}\}$.

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

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

- 598 • $s_1(t, \text{call})_{PS_2}(t, \text{ret})_{OS_3} < s_1(t', \text{call})_{PS_2}(t', \text{ret})_{OS_3}$
 - 599 • $s_1(t, \text{call})_{PS_2}(t, \text{call})_{OS_3}(t, \text{ret})_{PS_4}(t, \text{ret})_{OS_5} < s_1(t', \text{call})_{PS_2}(t', \text{call})_{OS_3}(t', \text{ret})_{PS_4}(t', \text{ret})_{PS_5}$
- 600 for any s_1, s_2, s_3, s_4, s_5 such that s_2, s_4 do not contain any t -moves.

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

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

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

- 607 • $(t, \text{call run}(f, x))_O (t, \text{call } f(x'))_P \cdots (t, \text{ret } f(v))_O (t, \text{ret run}(v'))_P$, or
- 608 • $(t', \text{call } w(v))_O (t', \text{call } w'(v'))_P$, where w is a name introduced in an earlier move
609 $(t'', x)_P$ and w' is a corresponding name introduced by the move preceding $(t'', x)_P$
610 in t'' , or
- 611 • $(t', \text{ret } w'(v''))_O (t', \text{ret } w(v'''))_P$ such that a segment $(t', \text{call } w(v))_O (t', \text{call } w'(v'))_P$
612 already occurred earlier.

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

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

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

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

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

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

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

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

632 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
633 preceded (in h_1) by a matching public call $(t', \text{call run}(f_j, x_j))_O$ followed by a cor-
634 responding return $(t', \text{ret run}(v_j))_P$, where t' need not be equal to t . We can obtain
635 the requisite h (for $\triangleleft_{\mathcal{R}}$) by changing t to t' in the whole of S_j for each S_j . Note that
636 run-moves are not affected and we get $h_1 \mathcal{R}^* h$.

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

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

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

648 4. Soundness

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

654 **Definition 31 (Encapsulated \approx).** Given libraries $L_1, L_2 : \Psi \rightarrow \Psi'$, we write $L_1 \approx_{\text{enc}} L_2$
655 when, for all $L' : \emptyset \rightarrow \Psi$ and $\Psi' \vdash_{\mathcal{K}} M_1 \parallel \dots \parallel M_N : \text{unit}$, if $\text{link } L' ; L_1$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$
656 then $\text{link } L' ; L_2$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$.

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

659 **Definition 32.** Let $\mathcal{R} \subseteq \mathcal{H}_{\Psi, \Psi'} \times \mathcal{H}_{\Psi, \Psi'}$ be a set closed under permutation of names in
660 $\text{Meths} \setminus (\Psi \cup \Psi')$. We say that a context formed by L' and M_1, \dots, M_N is \mathcal{R} -closed if, for
661 any $h \in \llbracket \text{link } L' ; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket$, $\bar{h} \mathcal{R} \bar{h}'$ implies $h' \in \llbracket \text{link } L' ; - \text{ in } (M_1 \parallel \dots \parallel M_N) \rrbracket$.
662 Given $L_1, L_2 : \Psi \rightarrow \Psi'$, we write $L_1 \approx_{\mathcal{R}} L_2$ if, for all \mathcal{R} -closed contexts formed
663 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
664 $\text{link } L' ; L_2$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$.

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

- 666 • $L_1 \triangleleft L_2$ implies $L_1 \approx L_2$,
- 667 • $L_1 \triangleleft_{\text{enc}} L_2$ implies $L_1 \approx_{\text{enc}} L_2$, and
- 668 • $L_1 \triangleleft_{\mathcal{R}} L_2$ implies $L_1 \approx_{\mathcal{R}} L_2$.

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

672 4.1. Correctness

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

679 **4.2. Compositionality**

680 The semantics we defined is compositional in the following ways:

- 681 • To compute the semantics of a library L inside a context C , it suffices to compose the
- 682 semantics of C with that of L , for a suitable notion of context-library composition
- 683 $(\llbracket C \rrbracket \otimes \llbracket L \rrbracket)$.
- 684 • To compute the semantics of a union library $L_1 \cup L_2$, we can compose the semantics
- 685 of L_1 and L_2 , for a suitable notion of library-library composition $(\llbracket L_1 \rrbracket \otimes \llbracket L_2 \rrbracket)$.

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

687 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
 688 following result, which we shall use for showing correctness.
 689
 690

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

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

691 **4.3. General linearisability**

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

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

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

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

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

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

- 704 • 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
 705 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
 706 ρ'_1, ρ'_2 . This is possible because the $((t', m')_P$ -labelled) transition does not access
 707 or modify the store, and none of the ϵ -transitions distinguished above can be in t' ,
 708 thanks to our earlier observations about the behaviour of the single-threaded system.
- 709 • Analogously, suppose $\rho \xrightarrow{(t', m')} \rho_1 \xrightarrow{\epsilon^*} \rho_2 \xrightarrow{(t, m)_O} \rho_3$. Then the $(t, m)_O$ -
 710 transition can be brought forward, i.e. $\rho \xrightarrow{(t, m)_O} \rho'_1 \xrightarrow{(t', m')} \rho'_2 \xrightarrow{\epsilon^*} \rho_3$, because it
 711 does not access or modify the store and the preceding ϵ -transitions cannot be from
 712 t . □

This, along with the fact that

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

713 lead us to the notion of linearisability defined in Def. 9.

714 We now prove the main theorem of this subsection.

```

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 ();

```

Figure 8: A library without a sequential history

715 **Theorem 35.** $L_1 \triangleleft L_2$ implies $L_1 \sqsupseteq L_2$.

716 *Proof.* Consider C such that $C[L_1] \Downarrow$. We need to show $C[L_2] \Downarrow$. Because $C[L_1] \Downarrow$,
717 Theorem 33 implies that there exists $h_1 \in \llbracket L_1 \rrbracket$ such that $\bar{h}_1 \in \llbracket C \rrbracket$. Because $L_1 \triangleleft L_2$,
718 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 Lem. 34, $\bar{h}_2 \in \llbracket C \rrbracket$.
719 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

720 **Remark 36.** A natural question to ask is whether the converse of Theorem 35 is
721 true. The answer is negative and can be traced back to the fact that \triangleleft is defined using
722 sequential histories: in order to establish $L_1 \triangleleft L_2$ (for $L_1, L_2 : \Psi \rightarrow \Psi'$) one needs to
723 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
724 that $L_1 \triangleleft A_2$.

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

730 Observe that if L does not generate any sequential histories then we vacuously have
731 $L \sqsupseteq L$, but cannot have $L \triangleleft L$. We conjecture that a completeness result would be
732 possible if we allowed for non-sequential specs in the definition of \triangleleft .

733 4.4. Encapsulated linearisability

734 In this case libraries are being tested by clients that do not communicate with the
735 parameter library explicitly. Recall from Definition 31 that, given libraries $L_1, L_2 : \Psi \rightarrow$
736 Ψ' , we write $L_1 \sqsupseteq_{\text{enc}} L_2$ when, for all $L' : \emptyset \rightarrow \Psi$ and $\Psi' \vdash_{\mathcal{K}} M_1 \parallel \dots \parallel M_N : \text{unit}$, if
737 $\text{link } L'; L_1$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$ then $\text{link } L'; L_2$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$.

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

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

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

743 **Remark 38.** Consider parameter library $L' : \emptyset \rightarrow \{m\}$ and client $\{m'\} \vdash_{\mathcal{K}} M : \text{unit}$
744 with $m, m' \in \text{Meths}_{\text{unit} \rightarrow (\text{unit} \rightarrow \text{unit})}$, and suppose we insert in their context a ‘‘copycat’’

745 library L which implements m' as $m' = \lambda x.mx$. Then the following scenario may
 746 seem to contradict encapsulation:

- 747 • M calls $m'()$;
- 748 • L calls $m()$;
- 749 • L' returns with $m(m'')$ to L ;
- 750 • and finally L copycats this return to M .

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

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

759 Moreover, we can show the following.

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

- 762 • *If $h \in X$ and $h (\triangleleft_{OP} \cup \diamond)^* h'$ then $h' \in X$.*
- 763 • *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*
 764 *t . Then $Y = Y'$, i.e. inside a thread only O can switch between \mathcal{K} and \mathcal{L} .*

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

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

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

779 **Theorem 40.** $L_1 \triangleleft_{\text{enc}} L_2$ implies $L_1 \overline{\varepsilon}_{\text{enc}} L_2$.

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

781 4.5. Relational linearisability

782 Finally, we examine relational linearisability (Definition 17).

783 **Theorem 41.** $L_1 \triangleleft_{\mathcal{R}} L_2$ implies $L_1 \overline{\varepsilon}_{\mathcal{R}} L_2$.

784 *Proof.* Consider \mathcal{R} -closed C such that $C[L_1] \Downarrow$. We need to show $C[L_2] \Downarrow$. Because
 785 $C[L_1] \Downarrow$, Theorem 33 implies that there exists $h_1 \in \llbracket L_1 \rrbracket$ such that $\overline{h_1} \in \llbracket C \rrbracket$. Because
 786 $L_1 \triangleleft_{\mathcal{R}} L_2$, 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} -
 787 closed by definition and closed under \triangleleft_{OP} by Lemma 34, we have $\overline{h_2} \in \llbracket C \rrbracket$. Because
 788 $h_2 \in \llbracket L_2 \rrbracket$ and $\overline{h_2} \in \llbracket C \rrbracket$, we can conclude $C[L_2] \Downarrow$. \square

789 5. Related and future work

790 Linearisability has been consistently used as a correctness criterion for concurrent
791 algorithms on a variety of data structures [15], and has inspired a variety of proof
792 methods [16]. An explicit connection between linearisability and refinement was made
793 in [17], where it was shown that, in base-type settings, linearisability and refinement
794 coincide. Similar results have been proved in [18, 19, 20, 3]. Our contributions are
795 notions of linearisability that serve as correctness criteria for libraries with methods of
796 arbitrary order and have a similar relationship to refinement. The next natural target is
797 to investigate proof methods for establishing linearisability of higher-order concurrent
798 libraries. The examples proved herein are only an initial step in that direction.

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

807 **Acknowledgements.** We thank the authors of [3] for bringing the higher-order
808 linearisability problem to our attention. We would also like to thank Kasper Svendsen
809 and Radha Jagadeesan for constructive comments. This work was partially funded by
810 the Engineering and Physical Sciences Research Council (EP/P004172/1).

811 References

- 812 [1] M. Herlihy, J. M. Wing, Linearizability: A correctness condition for concurrent
813 objects, *ACM Trans. Program. Lang. Syst.* 12 (3) (1990) 463–492.
- 814 [2] R. Jagadeesan, G. Petri, C. Pitcher, J. Riely, Quarantining weakness - composi-
815 tional reasoning under relaxed memory models (extended abstract), in: *ESOP*
816 2013, 2013, pp. 492–511.
- 817 [3] A. Cerone, A. Gotsman, H. Yang, Parameterised linearisability, in: *Proceedings of*
818 *ICALP’14*, Vol. 8573 of *Lecture Notes in Computer Science*, Springer, 2014, pp.
819 98–109.
- 820 [4] J. Laird, A game semantics of Idealized CSP, in: *Proceedings of MFPS’01*,
821 Elsevier, 2001, pp. 1–26, *ENTCS*, Vol. 45.
- 822 [5] D. R. Ghica, A. S. Murawski, Angelic semantics of fine-grained concurrency,
823 in: *Proceedings of FOSSACS*, Vol. 2987 of *Lecture Notes in Computer Science*,
824 Springer-Verlag, 2004, pp. 211–225.
- 825 [6] A. Jeffrey, J. Rathke, A fully abstract may testing semantics for concurrent objects,
826 *Theor. Comput. Sci.* 338 (1-3) (2005) 17–63.
- 827 [7] J. Laird, A fully abstract trace semantics for general references, in: *Proceedings*
828 *of ICALP*, Vol. 4596 of *Lecture Notes in Computer Science*, Springer, 2007, pp.
829 667–679.

- 830 [8] D. R. Ghica, N. Tzevelekos, A system-level game semantics, *Electr. Notes Theor.*
831 *Comput. Sci.* 286 (2012) 191–211.
- 832 [9] D. Hendler, I. Incze, N. Shavit, M. Tzafrir, Flat combining and the synchronization-
833 parallelism tradeoff, in: *Proceedings of SPAA 2010*, 2010, pp. 355–364.
- 834 [10] A. S. Murawski, N. Tzevelekos, Higher-order linearizability, in: *Proceedings of*
835 *CONCUR*, Vol. 85 of *Leibniz International Proceedings in Informatics (LIPIcs)*,
836 *Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik*, 2017, pp. 34:1–34:18.
- 837 [11] <http://c-cube.github.io/ocaml-containers/0.21/CCMultiSet.S.html>.
- 838 [12] S. Heller, M. Herlihy, V. Luchangco, M. Moir, W. N. Scherer III, N. Shavit, A lazy
839 concurrent list-based set algorithm, in: *OPODIS*, 2005, pp. 3–16.
- 840 [13] P. W. O’Hearn, N. Rinetzky, M. T. Vechev, E. Yahav, G. Yorsh, Verifying lineariz-
841 ability with hindsight, in: *PODC*, 2010, pp. 85–94.
- 842 [14] S. Abramsky, G. McCusker, Game semantics, in: H. Schwichtenberg, U. Berger
843 (Eds.), *Logic and Computation*, Springer-Verlag, 1998, proceedings of the 1997
844 Marktoberdorf Summer School.
- 845 [15] M. Moir, N. Shavit, Concurrent data structures, in: *Handbook of Data Structures*
846 *and Applications*, Chapman and Hall/CRC, 2004.
- 847 [16] B. Dongol, J. Derrick, Verifying linearisability: A comparative survey, *ACM*
848 *Comput. Surv.* 48 (2) (2015) 19.
- 849 [17] I. Filipovic, P. W. O’Hearn, N. Rinetzky, H. Yang, Abstraction for concurrent
850 objects, *Theor. Comput. Sci.* 411 (51-52) (2010) 4379–4398.
- 851 [18] J. Derrick, G. Schellhorn, H. Wehrheim, Mechanically verified proof obligations
852 for linearizability, *ACM Trans. Program. Lang. Syst.* 33 (1) (2011) 4.
- 853 [19] A. Gotsman, H. Yang, Liveness-preserving atomicity abstraction, in: *Automata,*
854 *Languages and Programming - 38th International Colloquium, ICALP 2011. Pro-*
855 *ceedings, Part II*, 2011, pp. 453–465.
- 856 [20] H. Liang, X. Feng, Modular verification of linearizability with non-fixed lineariza-
857 tion points, in: *ACM SIGPLAN Conference on Programming Language Design*
858 *and Implementation, PLDI ’13. Proceedings*, 2013, pp. 459–470.
- 859 [21] A. J. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, D. Dreyer, Logical relations
860 for fine-grained concurrency, in: *The 40th Annual ACM SIGPLAN-SIGACT*
861 *Symposium on Principles of Programming Languages, POPL ’13*, 2013, pp. 343–
862 356.
- 863 [22] A. Turon, D. Dreyer, L. Birkedal, Unifying refinement and Hoare-style reasoning in
864 a logic for higher-order concurrency, in: *ACM SIGPLAN International Conference*
865 *on Functional Programming, ICFP’13*, 2013, pp. 377–390.
- 866 [23] I. Sergey, A. Nanevski, A. Banerjee, Specifying and verifying concurrent algo-
867 rithms with histories and subjectivity, in: *Programming Languages and Systems -*
868 *24th European Symposium on Programming, ESOP 2015. Proceedings*, 2015, pp.
869 333–358.

- 870 [24] K. Svendsen, L. Birkedal, Impredicative concurrent abstract predicates, in: Pro-
871 gramming Languages and Systems - 23rd European Symposium on Programming,
872 ESOP 2014. Proceedings, 2014, pp. 149–168.
- 873 [25] K. Svendsen, L. Birkedal, M. J. Parkinson, Joins: A case study in modular
874 specification of a concurrent reentrant higher-order library, in: ECOOP 2013
875 - Object-Oriented Programming - 27th European Conference. Proceedings, 2013,
876 pp. 327–351.

877 Appendix

878 Appendix A. Big-step vs small-step reorderings

879 [3] defines linearisation in the general case using a “big-step” relation that applies a
 880 single permutation to the whole sequence. This contrasts with our definition as \triangleleft_{PO}^* ,
 881 in which we combine multiple adjacent swaps. We show that the two definitions are
 882 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 $i < j$:

$$\begin{aligned} & ((\exists t. h_1(i) = (t, -) \wedge h_1(j) = (t, -)) \\ & \vee (\exists t_1, t_2. h_1(i) = (t_1, -)_P \wedge h_1(j) = (t_2, -)_O)) \implies h_2(i) < h_2(j) \end{aligned}$$

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

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

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

888 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
 889 that X_{h_1, h_2} is not empty, because $h_1 \in X_{h_1, h_2}$.

890 For two histories h', h'' , define $\delta(h', h'')$ to be the length of the longest common
 891 prefix of h' and h'' . Let $N = \max_h \{\delta(h, h_2) \mid h \in X_{h_1, h_2}\}$. Note that $N \leq |h_1| = |h_2|$.

892 • If $N = |h_2|$ then we are done, because $N = |h_2|$ implies $h_2 \in X_{h_1, h_2}$ and, thus,
 893 $h_1 \triangleleft_{PO}^* h_2$.

894 • Suppose $N < |h_2|$ and consider h such that $N = \delta(h, h_2)$. We are going to arrive at
 895 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

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

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

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

904 □

905 Appendix B. Auxiliary lemmas about histories

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

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

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

910 **Lemma 45.** • For all $h \in \mathcal{H}_{\Psi, \Psi'}$ we have $h \upharpoonright \mathcal{L} \in \mathcal{H}_{\emptyset, \Psi}^{co}$ and $h \upharpoonright \mathcal{K} \in \mathcal{H}_{\emptyset, \Psi'}$.

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

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

914 *Proof.* The relevant sequences of moves are clearly alternating and well-bracketed,
915 when projected on single threads, because the LTS is bipartite (O - and P -configurations)
916 and separate evaluation stacks control the evolution in each thread. Other conditions
917 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
918 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
919 $\llbracket C \rrbracket$. \square

920 Appendix C. Trace compositionality

921 In this section we demonstrate how the semantics of a library inside a context can
922 be drawn by composing the semantics of the library and that of the context. The result
923 played a crucial role in our arguments about linearisability and contextual refinement in
924 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.

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

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

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

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

- 936 • $\mathcal{C}_i = (\mathcal{E}_C, M)$ and $\mathcal{C}'_i = (\mathcal{E}_L, -)$, or $\mathcal{C}_i = (\mathcal{E}'_C, -)$ and $\mathcal{C}'_i = (\mathcal{E}'_L, M)$.
- 937 • 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
938 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
939 $\mathcal{A}_{\mathcal{K}} = \mathcal{P}'_{\mathcal{K}}$.
- 940 • 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
941 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 $\rho \asymp \rho'$ 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 \otimes 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.

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

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).

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

963 Above we write M^i for the term M with all its methods and lambdas tagged (or re-
964 tagged) with i . Moreover, we use the convention e.g. $\vec{\mathcal{R}} \uplus_1 (m \mapsto \lambda x.M) = (\mathcal{R}_1 \uplus (m \mapsto$
965 $\lambda x.M), \mathcal{R}_2)$. Note that the repositories need not contain tags as, whenever a method is
966 looked up, we subsequently tag its body explicitly.

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

973 Defining parallel reduction \Longrightarrow' in an analogous way to \Longrightarrow , we can show the
974 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
975 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;
976 these notions are defined for $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S)$ in the same manner.

977 **Lemma 47.** *For any legal $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S)$, we have that $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S) \Downarrow$*
978 *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 $\rho_1 \asymp \rho_2$, 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:

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

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

980 We proceed to fleshing out the correspondence. We observe that an L -configuration
981 ρ can be the final configuration of a trace just if all its components are O -configurations
982 with empty evaluation stacks. On the other hand, for C -configurations, we need to
983 reach P -configurations with terms $()$. Thus, we call an N -configuration ρ *final* if
984 $\rho = (\vec{C}, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$ and either $\mathcal{C}_i = ([\], -)$ for all i , or $\mathcal{C}_i = ([\], ())$ for all i .

985 Let us write $(\mathcal{S}_1, \hookrightarrow_1, \mathcal{F}_1)$ for the transition system induced from external composi-
986 tion, and $(\mathcal{S}_2, \hookrightarrow_2, \mathcal{F}_2)$ be the transition system derived from internal composition:

- 987 • $\mathcal{S}_1 = \{\rho \otimes \rho' \mid \rho \succ \rho'\}$, $\mathcal{F}_1 = \{\rho \otimes \rho' \in \mathcal{S}_1 \mid \rho, \rho' \text{ final}\}$, and \hookrightarrow_1 the transition relation
988 \longrightarrow defined previously.
- 989 • $\mathcal{S}_2 = \{(M_1 \parallel \dots \parallel M_N, \vec{\mathcal{R}}, S) \mid (M_1 \parallel \dots \parallel M_N, \mathcal{R}_1 \uplus \mathcal{R}_2, S) \text{ valid}\}$, $\mathcal{F}_2 = \{x \in \mathcal{S}_2 \mid$
990 $x \text{ final}\}$, and \hookrightarrow_2 the transition relation \Longrightarrow' defined above.

991 A relation $R \subseteq \mathcal{S}_1 \times \mathcal{S}_2$ is called a *bisimulation* if, for all $(x_1, x_2) \in R$:

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

995 Given $(x_1, x_2) \in \mathcal{S}_1 \times \mathcal{S}_2$, we say that x_1 and x_2 are *bisimilar*, written $x_1 \sim x_2$, if
996 $(x_1, x_2) \in R$ for some bisimulation R .

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

998 Recall we write \bar{h} for the O/P complement of the history h . We can now prove
999 Theorem 33, which states that the behaviour of a library L inside a context C can be
1000 deduced by composing the semantics of L and C .

1001 **Theorem 33** *Let $L : \Psi \rightarrow \Psi'$, $L' : 1 \rightarrow \Psi, \Psi_1$ and $\Psi', \Psi_1 \vdash M_1, \dots, M_N : \text{unit}$,
1002 with L, L' and $M_1; \dots; M_N$ accessing pairwise disjoint parts of the store. Then,
1003 $\text{link } L'; L \text{ in } (M_1 \parallel \dots \parallel M_N) \Downarrow$ 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 $(L) \xrightarrow{*}_{\text{lib}} (\epsilon, \mathcal{R}_0, S_0)$
and $(L') \xrightarrow{*}_{\text{lib}} (\epsilon, \mathcal{R}'_0, S'_0)$ with $\text{dom}(\mathcal{R}_0) \cap \text{dom}(\mathcal{R}'_0) = \text{dom}(S_0) \cap \text{dom}(S'_0) = \emptyset$. We set:

$$\begin{aligned} \rho_0 &= (([\], -) \parallel \dots \parallel ([\], -), \mathcal{R}_0, (\emptyset, \Psi'), (\Psi, \emptyset), S_0) \\ \rho'_0 &= (([\], M_1) \parallel \dots \parallel ([\], M_N), \mathcal{R}'_0, (\Psi, \emptyset), (\emptyset, \Psi'), S'_0) \end{aligned}$$

1004 We pick these as the initial N -configurations for $\llbracket L \rrbracket_N$ and $\llbracket C \rrbracket$ respectively. Moreover,
1005 we have that $(L'; L) \xrightarrow{*}_{\text{lib}} (\epsilon, \mathcal{R}'_0, S'_0)$ where $\mathcal{R}'_0 = \{(m, (\mathcal{R}_0 \uplus \mathcal{R}'_0)(m) \{!r/\bar{m}\}) \mid$

1006 $m \in \text{dom}(\mathcal{R}_0 \uplus \mathcal{R}'_0)$ and $S''_0 = (S_0 \uplus S'_0) \{!\vec{r}/\vec{m}\} \uplus_s \{(r_i, m_i) \mid i = 1, \dots, n\}$, assuming
 1007 $\Psi = \{m_1, \dots, m_n\}$ and r_1, \dots, r_n are fresh references of corresponding types. Hence,
 1008 the initial triple for $\llbracket C[L] \rrbracket$ is taken to be $\phi_0 = (([], M_1) \parallel \dots \parallel ([], M_N), \mathcal{R}''_0, S''_0)$. On
 1009 the other hand, $\rho'_0 \bowtie \rho_0 = (([], M_1) \parallel \dots \parallel ([], M_N), \mathcal{R}'_0, \mathcal{R}_0, S_0 \uplus S'_0)$ and, using also
 1010 Lemma 47, we have that $\phi_0 \Downarrow$ iff $\rho'_0 \bowtie \rho_0 \Downarrow$.

1011 Then, for the forward direction of the claim, from $\phi_0 \Downarrow$ we obtain that $\rho'_0 \bowtie \rho_0 \Downarrow$. From
 1012 the previous lemma, we have that so does $\rho'_0 \otimes \rho_0$. From the latter reduction we obtain
 1013 the required common history. Conversely, suppose $h \in \llbracket L \rrbracket_N$ and $\bar{h} \in \llbracket C \rrbracket$. WLOG,
 1014 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
 1015 alpha-covert \mathcal{R}_0 and \mathcal{R}'_0 for this). Then, ρ_0 and ρ'_0 both produce h , with opposite
 1016 polarities. By definition of the external composite reduction, we then have that $\rho'_0 \otimes \rho_0$
 1017 reduces to some final state. By the previous lemma, we have that $\rho'_0 \bowtie \rho_0$ reduces to some
 1018 final quadruple, which in turn implies that $\phi_0 \Downarrow$, i.e. $\text{link } L'; L$ in $(M_1 \parallel \dots \parallel M_N) \Downarrow$. \square

1019 We conclude this section with the proofs of the last two lemmata used.

1020 Appendix C.1. Proof of Lemma 47

1021 We purpose to show that, for any legal $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S)$, $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, S) \Downarrow$
 1022 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:

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

1023 where $\mathcal{R}' \{\mathcal{R}\} = \{(m, \mathcal{R}'(m) \{\mathcal{R}\}) \mid m \in \text{dom}(\mathcal{R}' \setminus \mathcal{R}) \cup \mathcal{L}(\mathcal{R})\}$, $S \{\mathcal{R}\} = (S \upharpoonright$
 1024 $\text{Refs}_{\text{int}}) \cup \{(r, S(r) \{\mathcal{R}\}) \mid r \in \text{dom}(S) \setminus \text{Refs}_{\text{int}}\}$, and $E[M]$ is the term obtained from
 1025 $E[M]$ by removing all tagging.

1026 We next define a notion of indexed bisimulation between the transition systems
 1027 produced from quadruples and triples respectively. Given an expansion class \mathcal{R} , a
 1028 relation $R_{\mathcal{R}}$ between quadruples and triples is called an \mathcal{R} -bisimulation if, whenever
 1029 $\phi_1 R_{\mathcal{R}} \phi_2$:

- 1030 • ϕ_1 final implies ϕ_2 final
- 1031 • ϕ_2 final implies $\phi_2 \Downarrow$
- 1032 • $\phi_1 \Longrightarrow' \phi'_1$ implies $\phi_2 \Longrightarrow^= \phi'_2$ and $\phi'_1 R_{\mathcal{R}'} \phi'_2$ for some expansion class $\mathcal{R}' \supseteq \mathcal{R}$
- 1033 • $\phi_2 \Longrightarrow \phi'_2$ implies $\phi_1 \Longrightarrow'^* \phi'_1$ and $\phi'_1 R_{\mathcal{R}'} \phi'_2$ for some expansion class $\mathcal{R}' \supseteq \mathcal{R}$.

1034 Thus, Lemma 47 directly follows from the next result.

Lemma 49. For all expansion classes \mathcal{R} , the relation $R_{\mathcal{R}} =$

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

1035 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:

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

1036 so $\phi \Downarrow$ by acyclicity of $\mathcal{G}(\mathcal{R})$.

1037 Suppose now $\phi \Longrightarrow \phi'$, say due to $(E_1[M_1], \mathcal{R}_1, \mathcal{R}_2, S) \rightarrow'_1 (E_1[M'_1], \mathcal{R}'_1, \mathcal{R}'_2, S')$.

1038 In case the reduction is not a function call or return, then it can be clearly simulated by
1039 $\phi^{\#\mathcal{R}}$. 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 $(\underline{E}_1[mv], \mathcal{R}_{12}, S)\{\mathcal{R}\} \rightarrow_1 (\underline{E}_1[M\{v/x\}], \mathcal{R}_{12}, S)\{\mathcal{R}\}$. If, on the other hand, $m \in \text{dom}(\mathcal{R})$, suppose $\vec{\mathcal{R}}_i(m) = \lambda x.m'x$, then $M = m'x$ and $m\{\mathcal{R}\} = m'\{\mathcal{R}\}$ so we have:

$$\underline{E}_1[M\{v/x\}^i]\{\mathcal{R}\} = \underline{E}_1[(m'v)^i]\{\mathcal{R}\} = \underline{E}_1[(mv)^i]\{\mathcal{R}\}$$

1040 and $E_1[(mv)^i] = E_1[m^i v]$ by the way the semantics was defined, so $\phi^{\#\mathcal{R}} = \phi^{\#\mathcal{R}}$.

- $(E_1[m^i v], \vec{\mathcal{R}}, S) \rightarrow'_1 (E_1[\langle m^i \rangle M\{v'/x\}^{3-i}], \vec{\mathcal{R}}', S)$, with $\mathcal{R}_{3-i}(m) = \lambda x.M$, $\text{Meths}(v) = \{m_1, \dots, m_k\}$, $v' = \{\tilde{m}'/\tilde{m}\}$ 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 $(\underline{E}_1[mv], \mathcal{R}_{12}, S)\{\mathcal{R}\} \rightarrow_1 (\underline{E}_1[M\{v'/x\}], \mathcal{R}_{12}, S)\{\mathcal{R}\}$, and we have:

$$\begin{aligned} \underline{E}_1[\langle m^i \rangle M\{v'/x\}^{3-i}]\{\mathcal{R}'\} &= \underline{E}_1[M\{v'/x\}]\{\mathcal{R}'\} \\ &= \underline{E}_1[M\{v/x\}]\{\mathcal{R}\} \end{aligned}$$

Moreover, $\mathcal{R}_{12}\{\mathcal{R}\} = (\mathcal{R}'_1 \cup \mathcal{R}'_2)\{\mathcal{R}'\}$ and $S\{\mathcal{R}\} = S\{\mathcal{R}'\}$, so $\phi^{\#\mathcal{R}} = (\underline{E}_1[M\{v/x\}], \mathcal{R}_{12}, S)\{\mathcal{R}\}$.
On the other hand, if $\mathcal{R}(m) = \lambda x.m''x$ then:

$$\begin{aligned} \underline{E}[\langle m^i \rangle M\{v'/x\}^{3-i}]\{\mathcal{R}'\} &= \underline{E}[m''v']\{\mathcal{R}'\} \\ &= \underline{E}[m''v]\{\mathcal{R}\} = \underline{E}[mv]\{\mathcal{R}\} \end{aligned}$$

1041 so $\phi^{\#\mathcal{R}} = \phi^{\#\mathcal{R}'}$.

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

1042 Suppose now $\phi^{\#\mathcal{R}} \Longrightarrow \phi'$, 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 $\phi \Longrightarrow \phi''$ 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:

$$\begin{aligned} \text{size}_{\mathcal{R}}(E_1[M_1], \dots, E_N[M_N]) &= \text{size}_{\mathcal{R}}(E_1[M_1]) + \dots + \text{size}_{\mathcal{R}}(E_N[M_N]) \\ \text{size}_{\mathcal{R}}(E[M]) &= \sum_{m \in X_1} 2|m|_{\mathcal{R}} + \sum_{m \in X_2} 1 \end{aligned}$$

1043 where X_1 is the multiset containing all occurrences of methods $m \in \text{dom}(\mathcal{R})$ inside
1044 $E[M]$ in call position (e.g. mM'), and X_2 contains all occurrences of methods $m \in$
1045 $\text{dom}(\mathcal{R})$ inside $E[M]$ in return position (i.e. $\langle m^i \rangle \dots$). We write $|m|_{\mathcal{R}}$ for the length of
1046 the unique directed path from m to a leaf in $\mathcal{G}(\mathcal{R})$. The fact that X_1, X_2 are multisets
1047 reflects that we count all occurrences of m in call/return positions. Suppose WLOG
1048 that the reduction to ϕ' is due to some $(E_1[M_1], \mathcal{R}_{12}, S)\{\mathcal{R}\} \rightarrow_1 (E'[M'], \mathcal{R}', S')$.
1049 If the reduction happens inside $\underline{M}_1\{\mathcal{R}\}$ (this case also encompasses the base case of the
1050 induction) then the only case we need to examine is that of the reduction being a method
1051 call. In such a case, suppose we have $E_1[M_1]\{\mathcal{R}\} = E[mv]$, $E' = E$, $M' = M\{v/x\}$
1052 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
1053 $\tilde{m}\{\mathcal{R}\} = m$, $\tilde{v}\{\mathcal{R}\} = v$ and $\tilde{E}\{\mathcal{R}\} = E$. If $m \neq \tilde{m}$ then, supposing $\mathcal{R}(\tilde{m}) = \lambda x.\tilde{m}'x$ we
1054 have the following cases:

- 1055 • $(\tilde{E}[\tilde{m}^i \tilde{v}], \tilde{\mathcal{R}}, S) \rightarrow'_1 (\tilde{E}[\tilde{m}'^i \tilde{v}], \tilde{\mathcal{R}}, S) = \phi''_1$
- 1056 • $(\tilde{E}[\tilde{m}^i \tilde{v}], \tilde{\mathcal{R}}, S) \rightarrow'_1 (\tilde{E}[\langle \tilde{m}^i \rangle (\tilde{m}'v')^{3-i}], \tilde{\mathcal{R}}', S) = \phi''_1$, with $\tilde{\mathcal{R}}' = \tilde{\mathcal{R}} \uplus_{3-i} \{m'_j \mapsto$
1057 $\lambda x.m_jx \mid 1 \leq j \leq k\}$, etc.

1058 Let ϕ'' be the extension of ϕ''_1 to an N -quadruple by using the remaining $E_i[M_i]$'s of
1059 ϕ , so that $\phi \Longrightarrow' \phi''$. In the first case above we have that $\phi''^{\#\mathcal{R}} = \phi$, and in the latter
1060 that $\phi''^{\#\mathcal{R}'} = \phi$ (with $\mathcal{R}' = \mathcal{R} \uplus \{m'_j \mapsto \lambda x.m_jx \mid 1 \leq j \leq k\}$), and we appeal to the IH.
1061 Suppose now that $\tilde{m} = m$ and $\mathcal{R}_{12}(m) = \lambda x.\tilde{M}$. Then, one of the following is the case:

- 1062 • $(\tilde{E}[\tilde{m}^i \tilde{v}], \tilde{\mathcal{R}}, S), \tilde{\mathcal{R}}, S) \rightarrow'_1 (\tilde{E}[\tilde{M}\{\tilde{v}/x\}^i], \tilde{\mathcal{R}}, S) = \phi''_1$
- 1063 • $(\tilde{E}[\tilde{m}^i \tilde{v}], \tilde{\mathcal{R}}, S) \rightarrow'_1 (\tilde{E}[\langle \tilde{m}^i \rangle \tilde{M}\{v'/x\}^{3-i}], \tilde{\mathcal{R}}', S) = \phi''_1$, with $\tilde{\mathcal{R}}' = \tilde{\mathcal{R}} \uplus_{3-i} \{m'_j \mapsto$
1064 $\lambda x.m_jx \mid 1 \leq j \leq k\}$, 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, $(E_2[E'[v]], \tilde{\mathcal{R}}, S) \rightarrow'_1$
 $(E_2[E''[v^i]], \tilde{\mathcal{R}}', S)$ whereby $E' = E''[\langle m^i \rangle \bullet]$ and $\tilde{\mathcal{R}}' = \tilde{\mathcal{R}} \uplus_{3-i} \{m'_j \mapsto \lambda x.m_jx \mid$
 $1 \leq j \leq k\}$, etc. We have that

$$\phi''_1 = (E_2[E''[v^i]], \tilde{\mathcal{R}}', S)\{\mathcal{R}'\} = (E_2[E'[v]], \tilde{\mathcal{R}}, S)\{\mathcal{R}\}$$

1065 and hence, extending ϕ''_1 to ϕ'' , we have $\phi''^{\#\mathcal{R}'} = \phi^{\#\mathcal{R}}$. We can now appeal to the
1066 IH. \square

1067 Appendix C.2. . Proof of Lemma 48

1068 Let $\rho \asymp \rho'$ be compatible N -configurations. Then, $(\rho \circ \rho') \sim (\rho \bowtie \rho')$.

1069 We prove that the relation $R = \{(\rho_1 \circ \rho_2, \rho_1 \bowtie \rho_2) \mid \rho_1 \asymp \rho_2\}$ is a bisimulation. Let us
1070 suppose that $(\rho_1 \circ \rho_2, \rho_1 \bowtie \rho_2) \in R$.

- Suppose $\rho_1 \circ \rho_2 \hookrightarrow_1 \rho'_1 \circ \rho'_2$. 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

$\rho_1 = (C_1^1 \parallel \dots, \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1)$ and $\rho_2 = (C_1^2 \parallel \dots, \mathcal{R}_2, \mathcal{P}_2, \mathcal{A}_2, S_2)$, we have that either of the following scenarios holds, for some $x \in \{\mathcal{K}, \mathcal{L}\}$: $C_1^1 = (\mathcal{E}_1, E[mm'])$, $C_1^2 = (\mathcal{E}_2, -)$ and

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

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

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

1071 The case for (RETN) is treated similarly.

- 1072 • Suppose $\rho_1 \bowtie \rho_2 = (E[M_1] \parallel M_2 \parallel \dots \parallel M_N, \bar{\mathcal{R}}, S) \hookrightarrow_2 (E[M'_1] \parallel M_2 \parallel \dots \parallel M_N, \bar{\mathcal{R}}', S')$
1073 and let $\rho_1 = ((\mathcal{E}_1, M''_1) \parallel \dots, \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1)$ and $\rho_2 = ((\mathcal{E}_2, -) \parallel \dots, \mathcal{R}_2, \mathcal{P}_2, \mathcal{A}_2, S_2)$,
1074 where $(\mathcal{E}_1 \bowtie \mathcal{E}_2)[M''_1] = E[M_1]$. If the redex M_1 is not of the forms $M_1 = m^1 v$
1075 or $M_1 = \langle m^1 \rangle v$, with $m \in \text{dom}(\mathcal{R}_2)$, then the reduction can clearly be simulated
1076 by $\rho_1 \circ \rho_2$ (internally, by ρ_1). Otherwise, similarly as above, the reduction can be
1077 simulated by a mutual call/return of m .

1078 Finally, it is clear that $\rho_1 \circ \rho_2$ is final iff $\rho_1 \bowtie \rho_2$ is final. \square

1079 Appendix D. Library Compositionality

1080 This compositionality result will allow us to compose histories of component li-
1081 braries in order to obtain those of their composite library. Let $L_1 : \Psi_1 \rightarrow \Psi_2$ and
1082 $L_2 : \Psi'_1 \rightarrow \Psi'_2$. The semantic composition will be guided by two sets of names Π, P .
1083 Π contains method names that are shared between by the respective libraries and their
1084 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
1085 for private communication between L_1 and L_2 . Consequently, $\Pi \cap P$ consists of names
1086 that can be used both for internal communication between L_1 and L_2 , and for contextual
1087 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:

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

1088 along with the dual rules for the last four cases (i.e. where we schedule 2 in each case).
 1089 Note that the definition uses sequences of moves that are suffixes of histories (such as
 1090 s_i). The above equations are subject to the following side conditions:

- 1091 • $\text{Meths}(v) \cap (\Pi \cup P) = \emptyset$, $\Pi' = \Pi \uplus \text{Meths}(v)$ and $P' = P \uplus \text{Meths}(v)$;
- 1092 • $m \in P$ in the 0-scheduling cases;
- 1093 • $m \in \Pi$ in the 1-scheduling cases and, also, $m \in \Pi \setminus P$ in the third case (the P -call);
- 1094 • in the 1-scheduling cases, we also require that the leftmost move with thread index t
 1095 in s_2 is not a P -move.

1096 History composition is a partial function: if the conditions above are not met, or h_1, h_2, σ
 1097 are not of the appropriate form, then the composition is undefined. The above conditions
 1098 ensure that the composed histories are indeed compatible and can be produced by
 1099 composing actual libraries. For instance, the last condition corresponds to determinacy
 1100 of threads: there can only be at most one component starting with a P -move in each
 1101 thread t . We then have the following correspondence.

Theorem 50. *If $L_1 : \Psi_1 \rightarrow \Psi_2$ and $L_2 : \Psi'_1 \rightarrow \Psi'_2$ access disjoint parts of the store then*

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

1102 with $\Pi_0 = \Psi_1 \cup \Psi_2 \cup \Psi'_1 \cup \Psi'_2$ and $P_0 = (\Psi_1 \cup \Psi'_1) \cap (\Psi_2 \cup \Psi'_2)$.

1103 The rest of this section is devoted in proving the Theorem.

1104 Recall that we examine library composition in the sense of union of libraries. This
 1105 scenario is more general than the one of Appendix C as, during composition via union,
 1106 the calls and returns of each of the component libraries may be caught by the other
 1107 library or passed as a call/return to the outer context. Thus, the setting of this section
 1108 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
 1109 relating their semantics to that of their union $L_1 \cup L_2 : (\Psi_1 \cup \Psi'_1) \setminus (\Psi_2 \cup \Psi'_2) \rightarrow \Psi_2 \cup \Psi'_2$.

1110 Given configurations for L_1 and L_2 , in order to be able to reduce them together we
 1111 need to determine which of their methods can be used for communication between them,
 1112 and which for interacting with the external context, which represents player O in the
 1113 game. We will therefore employ a set of method names, denoted by Π and variants, to
 1114 register those methods used for interaction with the external context. Another piece of
 1115 information we need to know is in which component in the composition was the last
 1116 call played, or whether it was an internal call instead. This is important so that, when O
 1117 (or P) has the choice to return to both components, in the same thread, we know which
 1118 one was last to call and therefore has precedence. We use for this purpose sequences
 1119 $w = (w_1, \dots, w_N)$ where, for each i , $w_i \in \{0, 1, 2\}^*$. Thus, if e.g. $w_1 = 2w'_1$, this would
 1120 mean that, in thread 1, the last call to O , was done from the second component; if, on
 1121 the other hand, $w_1 = 0w'_1$ then the last call in thread 1 was an internal one between the
 1122 two components. Given such a w and some $j \in \{0, 1, 2\}$, for each index t , we write
 1123 $j +_t w$ for $w[t \mapsto (jw_t)]$.

1124 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,
 1125 and in particular L -configurations, and suppose that $\rho_1 = (\bar{C}, \mathcal{R}, \mathcal{P}, \mathcal{A}, S)$ and $\rho_2 =$
 1126 $(\bar{C}', \mathcal{R}', \mathcal{P}', \mathcal{A}', S')$. Moreover, let $\Psi_1 \cup \Psi_2 \cup \Psi'_1 \cup \Psi'_2 \subseteq \Pi$. We say that ρ_1 and ρ_2 are
 1127 (w, Π) -**compatible**, written $\rho_1 \approx_w^\Pi \rho_2$, if S, S' have disjoint domains and, for each i ;

- 1128 • $C_i = (\mathcal{E}'_L, M)$ and $C'_i = (\mathcal{E}_L, -)$, or $C_i = (\mathcal{E}_L, -)$ and $C'_i = (\mathcal{E}'_L, M)$, or $C_i = (\mathcal{E}_{L1}, -)$
 1129 and $C'_i = (\mathcal{E}_{L2}, -)$.

- We have $\Psi_1 \subseteq \mathcal{A}_l, \Psi_2 \subseteq \mathcal{P}_K, \Psi'_1 \subseteq \mathcal{A}'_L, \Psi'_2 \subseteq \mathcal{P}'_K$ and, setting

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

1130 we also have:

$$1131 \quad - (\mathcal{P}_L \uplus \mathcal{P}_K \uplus \mathcal{A}_l \uplus \mathcal{A}_K) \cap (\mathcal{P}'_L \uplus \mathcal{P}'_K \uplus \mathcal{A}'_l \uplus \mathcal{A}'_K) = P \uplus (\Psi_1 \cap \Psi'_1),$$

$$1132 \quad - \Pi \cap P = (\Psi_2 \cup \Psi'_2) \cap (\Psi_1 \cup \Psi'_1),$$

$$1133 \quad - \Pi \cup P = \mathcal{P}_L \cup \mathcal{P}_K \cup \mathcal{P}'_L \cup \mathcal{P}'_K \cup \mathcal{A}_L \cup \mathcal{A}_K \cup \mathcal{A}'_L \cup \mathcal{A}'_K.$$

- 1134 • The private names of \mathcal{R} do not appear in ρ_2 , and dually for the private names of \mathcal{R}' .
- 1135 • 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
1136 $\mathcal{E} = \mathcal{E}' = []$ or:

$$1137 \quad - \mathcal{E} = m :: \mathcal{E}_1 \text{ and } \mathcal{E}' \in \mathcal{E}_L, \text{ with } m \in \Pi, w_i = 1u \text{ and } \mathcal{E}_1, \mathcal{E}' \text{ are } u\text{-compatible,}$$

$$1138 \quad - \text{or } \mathcal{E} = m :: \mathcal{E}_1 \text{ and } \mathcal{E}' = m :: E :: \mathcal{E}_2, \text{ with } m \in P, w_i = 0u \text{ and } \mathcal{E}_1, \mathcal{E}_2 \text{ are}$$

$$1139 \quad u\text{-compatible,}$$

$$1140 \quad - \text{or } \mathcal{E} = m :: E :: \mathcal{E}_1 \text{ and } \mathcal{E}' \in \mathcal{E}_L, \text{ with } m \in \Pi \setminus P, w_i = 1u \text{ and } \mathcal{E}_1, \mathcal{E}' \text{ are}$$

$$1141 \quad u\text{-compatible,}$$

1142 or the dual of one of the three conditions above holds.

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

$$\begin{array}{c} \frac{\rho_1 \Longrightarrow \rho'_1}{\rho_1 \otimes_{\Pi}^w \rho_2 \longrightarrow \rho'_1 \otimes_{\Pi}^w \rho_2} \text{INT}_1 \\ \frac{\rho_1 \xrightarrow{(t, \text{call } m(v))} \rho'_1 \quad \rho_2 \xrightarrow{(t, \text{call } m(v))} \rho'_2}{\rho_1 \otimes_{\Pi}^w \rho_2 \longrightarrow \rho'_1 \otimes_{\Pi}^{0+t} \rho'_2} \text{CALL } (m \in P) \\ \frac{\rho_1 \xrightarrow{(t, \text{ret } m(v))} \rho'_1 \quad \rho_2 \xrightarrow{(t, \text{ret } m(v))} \rho'_2}{\rho_1 \otimes_{\Pi}^{0+t} \rho_2 \longrightarrow \rho'_1 \otimes_{\Pi}^w \rho'_2} \text{RETN } (m \in P) \\ \frac{\rho_1 \xrightarrow{(t, \text{call } m(v))_{PY}} \rho'_1}{\rho_1 \otimes_{\Pi}^w \rho_2 \xrightarrow{(t, \text{call } m(v))_{PY}} \rho'_1 \otimes_{\Pi'}^{1+t} \rho_2} \text{PCALL}_1 (m \in \Pi \setminus P) \\ \frac{\rho_1 \xrightarrow{(t, \text{ret } m(v))_{PY}} \rho'_1}{\rho_1 \otimes_{\Pi}^{1+t} \rho_2 \xrightarrow{(t, \text{ret } m(v))_{PY}} \rho'_1 \otimes_{\Pi'}^w \rho_2} \text{PRETN}_1 (m \in \Pi) \\ \frac{\rho_1 \xrightarrow{(t, \text{call } m(v))_{OY}} \rho'_1}{\rho_1 \otimes_{\Pi}^w \rho_2 \xrightarrow{(t, \text{call } m(v))_{OY}} \rho'_1 \otimes_{\Pi'}^{1+t} \rho_2} \text{OCALL}_1 (m \in \Pi) \\ \frac{\rho_1 \xrightarrow{(t, \text{ret } m(v))_{OY}} \rho'_1}{\rho_1 \otimes_{\Pi}^{1+t} \rho_2 \xrightarrow{(t, \text{ret } m(v))_{OY}} \rho'_1 \otimes_{\Pi'}^w \rho_2} \text{ORETN}_1 (m \in \Pi \setminus P) \end{array}$$

1143 along with their dual counterparts ($\text{INT}_2, \text{XCALL}_2, \text{XRETN}_2$). The internal rules above
1144 have the same side-conditions on name privacy as before. Moreover, in (XRETN_i) and

1145 (XCALL i), for $X=O,P$, we let $\Pi' = \Pi \uplus_t \text{Meths}(v)$ and impose that the t -th component
 1146 of ρ_{3-i} be an O -configuration and $\text{Meths}(v) \cap \text{Meths}(\rho_{3-i}) = \emptyset$.

1147 We can now show the following.

1148 **Lemma 51.** *Let $\rho_1 \simeq_{\Pi}^w \rho_2$ and suppose $\rho_1 \otimes_{\Pi}^w \rho_2 \xrightarrow{s}^* \rho'_1 \otimes_{\Pi'}^w \rho'_2$ for some sequence s
 1149 of moves. Then, $\rho'_1 \simeq_{\Pi'}^w \rho'_2$.*

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:

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

1150 where E is again from the extended evaluation contexts.

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

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

1151 and $[] \mathcal{K}^e [] = (\bullet, [])$.

For each pair $\rho_1 \simeq_{\Pi}^w \rho_2$, 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 $\mathcal{E}_i \mathcal{K}^u \mathcal{E}'_i = (E_i, \mathcal{E}''_i)$, we set:

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

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

$$\rho_1 \mathcal{K}_{\Pi}^w \rho_2 = (\mathcal{C}_1 \mathcal{K}^{w_1} \mathcal{C}'_1 \parallel \dots \parallel \mathcal{C}_N \mathcal{K}^{w_N} \mathcal{C}'_N, \mathcal{R}_1, \mathcal{R}_2, \mathcal{P}', \mathcal{A}', S_1 \uplus S_2)$$

1152 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 $\mathcal{A}' = ((\mathcal{A}_{1\mathcal{L}} \cup \mathcal{A}_{2\mathcal{L}}) \cap$
 1153 $(\Pi \setminus \mathcal{P}), (\mathcal{A}_{1\mathcal{K}} \uplus \mathcal{A}_{2\mathcal{K}}) \cap \Pi)$.

Now, as expected, the definition of \Longrightarrow' builds upon \rightarrow'_t . The definition of the

latter is given by the following rules.

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

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

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

1154 just if $(\mathcal{C}_t, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{x_{XY}}_t (\mathcal{C}', \vec{\mathcal{R}}', \mathcal{P}', \mathcal{A}', S')$.

1155 We next relate the transition systems induced by external (via \otimes) and internal
1156 composition (via \mathbb{M}). Let us write $(\mathcal{S}_1, \hookrightarrow_1, \mathcal{F}_1)$ for the transition system induced by
1157 external composition of compatible N -configurations (so \hookrightarrow_1 is \longrightarrow), and $(\mathcal{S}_2, \hookrightarrow_2$
1158 $, \mathcal{F}_2)$ be the one for internal composition (so \hookrightarrow_2 is \Longrightarrow). Finality of extended N -
1159 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
1160 relation $R \subseteq \mathcal{S}_1 \times \mathcal{S}_2$ is called a *bisimulation* if, for all $(x_1, x_2) \in R$:

- 1161 • $x_1 \in \mathcal{F}_1$ iff $x_2 \in \mathcal{F}_2$,
- 1162 • if $x_1 \hookrightarrow_1 x'_1$ then $x_2 \hookrightarrow_2 x'_2$ and $(x'_1, x'_2) \in R$,
- 1163 • if $x_1 \xrightarrow{(t,x)_{XY}}_1 x'_1$ then $x_2 \xrightarrow{(t,x)_{XY}}_2 x'_2$ and $(x'_1, x'_2) \in R$,
- 1164 • if $x_2 \hookrightarrow_2 x'_2$ then $x_1 \hookrightarrow_1 x'_1$ and $(x'_1, x'_2) \in R$,
- 1165 • if $x_2 \xrightarrow{(t,x)_{XY}}_2 x'_2$ then $x_1 \xrightarrow{(t,x)_{XY}}_1 x'_1$ and $(x'_1, x'_2) \in R$.

1166 Again, we say that x_1 and x_2 are *bisimilar*, and write $x_1 \sim x_2$, if there exists a bisimula-
1167 tion R such that $(x_1, x_2) \in R$.

1168 **Lemma 52.** *Let $\rho \succ_{\Pi}^w \rho'$ be compatible N -configurations. Then, $(\rho \otimes_{\Pi}^w \rho') \sim (\rho \mathbb{M}_{\Pi}^w \rho')$.*

1169 *Proof.* We prove that the relation $R = \{(\rho_1 \otimes_{\Pi}^w \rho_2, \rho_1 \mathbb{M}_{\Pi}^w \rho_2) \mid \rho_1 \succ_{\Pi}^w \rho_2\}$ is a bisimula-
1170 tion. 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 $\rho_1 \sim$

$(\mathcal{C}_1^1 \parallel \dots, \mathcal{R}_1, \mathcal{P}_1, \mathcal{A}_1, S_1)$, we have that one of the following holds, for some $x \in \{\mathcal{K}, \mathcal{L}\}$:

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

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

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

1171 with \mathcal{P}, \mathcal{A} as in the definition of composition and $\mathcal{P}' = \mathcal{P} \cup_x \text{Meths}(v)$, and the latter
1172 N -configuration equals $\rho'_1 \varkappa_{\Pi'}^{w'} \rho_2$. The other case is treated in the same manner,
1173 and we work similarly for (RET N_1).

1174 • On the other hand, if the transition is due to (CALL) or (RET N) then we work as in
1175 the proof of Lemma 48.

- Suppose $\rho_1 \varkappa_{\Pi}^w \rho_2 = (\mathcal{C}_1 \parallel \dots, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S) \xrightarrow{(1, \text{call } m(v))'} (\mathcal{C}'_1 \parallel \dots, \vec{\mathcal{R}}', \mathcal{P}', \mathcal{A}', S)$. 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 $i \in \{1, 2\}$:

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

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 $(\mathcal{C}_1^1, \mathcal{C}_1^2) = ((\mathcal{E}_1, E'[mv']), (\mathcal{E}_2, -))$, for some $E, \mathcal{E}_1, \mathcal{E}_2$ such that $\mathcal{E}_1 \varkappa^{w_1} \mathcal{E}_2 = (E'', \mathcal{E})$ 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 $\text{Meths}(v) = \{m'_1, \dots, m'_k\}$),

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

1176 and $\rho'_1 \varkappa_{\Pi'}^{w'} \rho_2 = (\mathcal{C}'_1 \parallel \dots, \vec{\mathcal{R}}', \mathcal{P}', \mathcal{A}', S)$ as required. The case for return transitions
1177 is similar.

1178 • On the other hand, if the transition out of $\rho_1 \varkappa_{\Pi}^w \rho_2$ does not have a label then we
1179 work as in the proof of Lemma 48.

1180 Moreover, by definition of syntactic composition, $\rho_1 \otimes_{\Pi}^w \rho_2$ is final iff $\rho_1 \varkappa_{\Pi}^w \rho_2$ is. \square

1181 Given an N -configuration ρ and a history h , let us write $\rho \Downarrow h$ if $\rho \xrightarrow{h} \rho'$ for
 1182 some final configuration ρ' . Similarly if ρ is of the form $(\vec{C}, \vec{\mathcal{R}}, \mathcal{P}, \mathcal{A}, S)$. We have the
 1183 following connections in history productions. The next lemma is proven in a similar
 1184 fashion as Lemma 47.

1185 **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*
 1186 $(M_1 \parallel \dots \parallel M_N, \mathcal{R}_1, \mathcal{R}_2, \mathcal{P}, \mathcal{A}, S) \Downarrow h$ 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 $\rho_1 \approx_{\Pi}^w \rho_2$ and history h , $(\rho_1 \otimes_{\Pi}^w \rho_2) \Downarrow h$ iff:*

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

1187 where P is computed from ρ_1, ρ_2 and Π as before.

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

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

1188 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 $m \in P$:

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

1189 By $\rho_i \Downarrow s_i$ and $\rho_1 \approx_{\Pi}^w \rho_2$ we have that $\rho_1 \otimes_{\Pi}^w \rho_2 \longrightarrow \rho'_1 \otimes_{\Pi}^{w'} \rho'_2$, where $w' = 0 +_t w$ and
 1190 $\rho'_1 \approx_{\Pi}^{w'} \rho'_2$. Also, $\rho'_i \Downarrow s'_i$ and $s = s'_1 \mathfrak{M}_{\Pi, P}^{\sigma'} s'_2$ so, by IH, $(\rho'_1 \otimes_{\Pi}^{w'} \rho'_2) \Downarrow s$. \square

1191 We can now prove the correspondence between the traces of component libraries
 1192 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,*

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

1193 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 $(L_i) \xrightarrow{*}_{\text{lib}} (\epsilon, \mathcal{R}_i, S_i)$, 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:

$$\begin{aligned} \rho_1 &= (([], -) \parallel \dots \parallel ([[], -), \mathcal{R}_1, (\emptyset, \Psi_2), (\Psi_1, \emptyset), S_1) \\ \rho_2 &= (([], -) \parallel \dots \parallel ([[], -), \mathcal{R}_2, (\emptyset, \Psi'_2), (\Psi'_1, \emptyset), S_2) \end{aligned}$$

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

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

1194 as the initial N -configuration for $\llbracket L_1 \cup L_2 \rrbracket_N$. On the other hand, we have $\rho_1 \mathfrak{M}_{\Pi_0}^{\epsilon} \rho_2 =$
 1195 $(([], -) \parallel \dots \parallel ([[], -), \mathcal{R}_1, \mathcal{R}_2, (\emptyset, \Psi_2 \cup \Psi'_2), ((\Psi_1 \cup \Psi'_1) \setminus P_0, S_0)$. From Lemma 53, we
 1196 have that $\rho_0 \Downarrow h$ iff $\rho_1 \mathfrak{M}_{\Pi_0}^{\epsilon} \rho_2 \Downarrow h$, for all h .

1197 Pick a history h . For the forward direction of the claim, $\rho_0 \Downarrow h$ implies $\rho_1 \varkappa_{\Pi_0}^\epsilon \rho_2 \Downarrow h$
1198 which, from Lemma 52, implies $\rho_1 \otimes_{\Pi_0}^\epsilon \rho_2 \Downarrow h$. We now use Lemma 54 to obtain
1199 h_1, h_2, σ such that $\rho_i \Downarrow h_i$ and $h = h_1 \varkappa_{\Pi_0, P_0}^\sigma h_2$. Conversely, suppose that $h_i \in \llbracket L_i \rrbracket_N$
1200 and $h = h_1 \varkappa_{\Pi_0, P_0}^\sigma h_2$. WLOG assume that $(\text{Meths}(h_1) \cup \text{Meths}(h_2)) \cap (\text{dom}(\mathcal{R}_1) \cup$
1201 $\text{dom}(\mathcal{R}_2)) \subseteq \Pi_0$ (or we appropriately alpha-covert \mathcal{R}_1 and \mathcal{R}_2). Then, $\rho_i \Downarrow h_i$, for
1202 $i = 1, 2$, and therefore $\rho_1 \otimes_{\Pi_0}^\epsilon \rho_2 \Downarrow h$ by Lemma 54. By Lemma 52 we have that
1203 $\rho_1 \varkappa_{\Pi_0}^\epsilon \rho_2 \Downarrow h$, which in turn implies that $\rho_0 \Downarrow h$, i.e. $h \in \llbracket L_1 \cup L_2 \rrbracket_N$. \square

1204 Appendix E. Composition congruence

1205 **Theorem 55.** *If $L_1 \triangleleft L_2$ then, for suitably typed L accessing disjoint part of the store*
1206 *than L_1 and L_2 , we have $L \cup L_1 \triangleleft L \cup L_2$.*

1207 *Proof.* Assume $L_1 \triangleleft L_2$ and suppose $h_1 \in \llbracket L \cup L_1 \rrbracket$. By Theorem 50, $h_1 = h' \varkappa_{\Pi, P}^\sigma h'_1$,
1208 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
1209 $h'_1 \triangleleft h'_2$, i.e. $h'_1 \triangleleft_{PO}^* h'_2$. Note that some of the rearrangements necessary to transform
1210 h'_1 into h'_2 may concern actions shared by h'_1 and h' ; their polarity will then be different
1211 in h' . Let h'' be obtained by applying such rearrangements to h' . We claim that
1212 $h' \triangleleft_{OP}^* h''$. Indeed, suppose that $(t', x')(t, x)_P$ are consecutive in h'_1 , but swapped in
1213 order to obtain h'_2 , and $(t, x)_P$ appears in h' as $(t, x)_O$. Now, the move (t', x') either
1214 appears in h_1 , or it appears in h' and gets hidden in h_1 . In every case, let s contain
1215 the moves of h' that are after (t', x') in the composition to h_1 , and before $(t, x)_O$. We
1216 have that $s(t, x)_O$ is a subsequence of h' and $h' \triangleleft_{OP}^* h''$ holds just if s contains no
1217 moves from t . But, if s contained moves from t then the rightmost one such would be
1218 some $(t, y)_P$. Moreover, in the composition towards h_1 , the move would be scheduled
1219 with 1. The latter would break the conditions for trace composition as, at that point, the
1220 corresponding subsequence of h'_1 has as leftmost move in t the P-move $(t, x)_P$. We
1221 can show similarly that $h' \triangleleft_{OP}^* h''$ holds in the case that the permutation in h'_1 is on
1222 consecutive moves $(t, x)_O(t', x')$. Finally, the rearrangements in h'_1 that do not affect
1223 moves shared with h' can be treated in a simpler way: e.g. in the case of $(t', x')(t, x)_P$
1224 consecutive in h'_1 and swapped in h'_2 , if $(t, x)_P$ does not appear in h' then we can check
1225 that h' cannot contain any t -moves between (t', x') and (t, x) as the conditions for trace
1226 composition impose that only O is expected to play in that part of h' (and any t -move
1227 would swap this polarity).

1228 Now, since $h' \in \llbracket L \rrbracket$, Lemma 34 implies $h'' \in \llbracket L \rrbracket$. Take h_2 to be $h'' \varkappa_{\Pi, P}^{\sigma'} h'_2$, where σ'
1229 is obtained from σ following these move rearrangements. We then have $h_2 \in \llbracket L \cup L_2 \rrbracket$.
1230 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,
1231 $L \cup L_1 \triangleleft L \cup L_2$. \square

1232 We next examine the behaviour of $\triangleleft_{\text{enc}}$ with respect to library composition. In
1233 contrast to general linearisability, we need to restrict composition for it to be compatible
1234 with encapsulation.

1235 *Remark 56.* The general case of union does not conform with encapsulation in the sense
1236 that encapsulated testing of $L \cup L_i$ ($i = 1, 2$) according to Def. 31 may subject L_i to
1237 unencapsulated testing. For example, because method names of L and L_i are allowed
1238 to overlap, methods in L may call public methods from L_i as well as implementing
1239 abstract methods from L_i . This amounts to L playing the role of both \mathcal{K} and \mathcal{L} , which
1240 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:

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

which decomposes as $h_i = h' \mathbb{A}_{\Pi, \emptyset}^\sigma h'_i$, 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:

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

1241 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} .

1242 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-
1243 order methods. Given $L : \Psi_1 \rightarrow \Psi_2$ and $L' : \Psi'_1 \rightarrow \Psi'_2$, we define their *disjoint*
1244 *union* $L \uplus L' = L \cup L' : (\Psi_1 \cup \Psi'_1) \rightarrow (\Psi_2 \cup \Psi'_2)$ under the assumption that
1245 $(\Psi_1 \cup \Psi_2) \cap (\Psi'_1 \cup \Psi'_2) = \emptyset$.
1246

1247 **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:*

- 1248 • *assuming $\Psi'_2 = \Psi_1$, we have $L ; L_1 \triangleleft_{\text{enc}} L ; L_2$ and $L_1 ; L \triangleleft_{\text{enc}} L_2 ; L$;*
- 1249 • *if $\Psi_1, \Psi_2, \Psi'_1, \Psi'_2$ are first-order then $L \uplus L_1 \triangleleft_{\text{enc}} L \uplus L_2$.*

1250 *Proof.* Let us consider the first sequencing case (the second one is dual), and assume
1251 that $L_1, L_2 : \Psi \rightarrow \Psi'$ and $L : \Psi'' \rightarrow \Psi$. Assume $L_1 \triangleleft_{\text{enc}} L_2$ and suppose $h_1 \in \llbracket L ; L_1 \rrbracket_{\text{enc}}$.
1252 By Theorem 50, $h_1 = h' \mathbb{A}_{\Pi, P}^\sigma h'_1$, where $h' \in \llbracket L \rrbracket$, $h'_1 \in \llbracket L_1 \rrbracket$ and method calls
1253 from Ψ are always scheduled with 0. The fact that O cannot switch between \mathcal{L}/\mathcal{K}
1254 components in (threads of) h_1 implies that the same holds for h', h'_1 , hence $h' \in \llbracket L \rrbracket_{\text{enc}}$
1255 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 $h'_1 \triangleleft h'_2$,
1256 i.e. $h'_1 (\triangleleft_{PO} \cup \diamond)^* h'_2$. As before, some of the rearrangements necessary to transform
1257 h'_1 into h'_2 may concern actions shared by h'_1 and h' ; we need to check that these can
1258 lead to compatible $h'' \in \llbracket L \rrbracket_{\text{enc}}$. Let h'' be obtained by applying such rearrangements
1259 to h' . We claim that $h' \triangleleft_{OP}^* h''$. The transpositions covered by \triangleleft_{PO} are treated as in
1260 Lemma 55. Suppose now that $(t', x')_{P\mathcal{K}}(t, x)_{O\mathcal{L}}$ are consecutive in h'_1 but swapped
1261 in order to obtain h'_2 , and $(t, x)_{O\mathcal{L}}$ appears in h' as $(t, x)_{P\mathcal{K}}$. Now, the move (t', x')
1262 cannot appear in h' as it is in L_1 's \mathcal{K} -component (L is the \mathcal{L} -component of L_1). Let
1263 s contain the moves of h' that are after (t', x') in the composition to h_1 , and before
1264 $(t, x)_{P\mathcal{K}}$. We claim that s contains no moves from t , so h' can be directly composed
1265 with h'_2 as far as this transposition is concerned. Indeed, if s contained moves from t
1266 then, taking into account the encapsulation conditions, the leftmost one such would be
1267 some $(t, y)_{O\mathcal{K}}$. But the \mathcal{K} -component of L is L_1 , which contradicts the fact that the
1268 moves we consider are consecutive in h'_1 . Hence, taking h_2 to be $h'' \mathbb{A}_{\Pi, P}^{\sigma'} h'_2$, where σ'
1269 is obtained from σ following the \triangleleft_{PO} move rearrangements, we have $h_2 \in \llbracket L ; L_2 \rrbracket_{\text{enc}}$
1270 and $h_1 \triangleleft_{\text{enc}} h_2$. Thus, $L ; L_1 \triangleleft_{\text{enc}} L ; L_2$.

1271 The case of $L \uplus L_1 \triangleleft_{\text{enc}} L \uplus L_2$ is treated in a similar fashion. In this case, because
1272 of disjointness, the moves transposed in h'_1 do not have any counterparts in h' . Again,
1273 we consider consecutive moves $(t', x')_{P\mathcal{K}}(t, x)_{O\mathcal{L}}$ in h'_1 that are swapped in order to
1274 obtain h'_2 . Let s contain the moves of h' that are after (t', x') in the composition to
1275 h_1 , and before (t, x) . As Ψ_1, Ψ'_1 is first-order, $(t, x)_{O\mathcal{L}}$ must be a return move and the
1276 t -move preceding it in h_1 must be the corresponding call. The latter is a move in h'_1 ,

¹²⁷⁷ which therefore implies that there can be no moves from t in s . Similarly for the other
¹²⁷⁸ transposition case. \square