

Composing Interfering Abstract Protocols

Filipe Militão^{1,2}, Jonathan Aldrich¹, and Luís Caires²

- 1 Carnegie Mellon University, Pittsburgh, USA
filipe.militao@cs.cmu.edu, aldrich@cs.cmu.edu
- 2 Universidade Nova de Lisboa, Lisboa, Portugal
lcaires@fct.unl.pt

Abstract

The undisciplined use of shared mutable state can be a source of program errors when aliases unsafely interfere with each other. While protocol-based techniques to reason about interference abound, they do not address two practical concerns: the decidability of protocol composition and its integration with protocol abstraction. We show that our composition procedure is decidable and that it ensures safe interference even when composing abstract protocols. To evaluate the expressiveness of our protocol framework for ensuring safe shared memory interference, we show how this same protocol framework can be used to model safe, typeful message-passing concurrency idioms.

1 Introduction

The interactions that can occur via shared mutable state can be a source of program errors. When different clients access the same mutable state, their actions can potentially *interfere*. For instance, the programmer may wrongly assume that a cell holds a particular type, when another part of the program has changed that cell to hold a different type. When this happens, the program may fault due to *unsafe* interference caused by unexpected actions through other aliases to that shared state. Thus, to reason about interference we must reason about how state is aliased and how the different aliases use the shared state.

Our technique builds on the use of linear capabilities [1] to track type-changing resource mutation within the framework of a linear type system. However, relying solely on linearity is often too restrictive. For instance, linearity enforces exclusive ownership of mutable state, which is incompatible with multithreading—i.e. linearity forbids sharing. To allow sharing, we extend the concept of *rely-guarantee protocols* [18]. By sequencing *steps* of “rely \Rightarrow guarantee” actions, each protocol characterizes an alias’s local, isolated perspective on interactions with a piece of shared state:

$$\underbrace{\text{“what I } \textit{assume} \text{ about the state”} \Rightarrow \text{“what I } \textit{guarantee} \text{ about the state”}}_{\textit{current step}}; \textit{next step}$$

Since the interactions performed by an alias may change over time, a rely-guarantee protocol is formed by a sequence of steps that specify each interfering action. Each step *relies* on the shared state having some type and then, after some private actions, *guarantees* that the shared state will now have some other type, which becomes visible to other aliases. By constraining the actions of each alias, we can make strong assumptions about the kind of interference that an alias may produce, in the spirit of rely-guarantee reasoning [13]. Naturally, not all protocols compose safely. While a protocol describes its own actions on a piece of shared state, protocol *composition* will ensure that those actions are safe w.r.t. the actions that can be done via other existing (and even future) protocols over that state. Composition is safe only if the set of protocols accounts for all possible run-time action interleavings.

Our main contribution is a decidable protocol composition procedure that also allows abstract protocols to be composed. We break down our contributions as follows:



© Filipe Militão, Jonathan Aldrich, Luís Caires;
licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICIS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$x \in \text{VARIABLES}$	$t \in \text{TAGS}$	$f \in \text{FIELDS}$	$\rho \in \text{LOCATION CONSTANTS}$
$e ::= v$	(value)		$\text{lock } \bar{v}$ (lock locations)
$v.f$	(field selection)		$\text{unlock } \bar{v}$ (unlock locations)
$v v$	(application)		$\text{fork } e$ (spawn thread)
$\text{let } x = e \text{ in } e \text{ end}$	(let)		
$\text{new } v$	(cell creation)	$v ::= \rho$	(address)
$\text{delete } v$	(cell deletion)	x	(variable)
$!v$	(dereference)	$\lambda x.e$	(function)
$v := v$	(assign)	$\{\bar{f} = v\}$	(record)
$\text{case } v \text{ of } t\#x \rightarrow e \text{ end}$	(case)	$t\#v$	(tagged value)

Notes: \bar{Z} is a potentially empty sequence of Z elements. ρ is not source-level.

■ **Figure 1** Values (v) and expressions (e).

- We adapt the existing constructs of rely-guarantee protocols [18] to work in a system with concurrent runtime semantics, and show that rely-guarantee protocols are useful to reason about safe interference in the concurrent setting.
- We give an axiomatic definition of protocol composition. We show that this procedure can be implemented in a sound and complete (w.r.t. the formal definition) algorithm that terminates on all legal inputs.¹ The protocol composition algorithm is implemented in a prototype.²
- We show that our use of type abstraction and bounded quantification at the protocol level enables us to model new, and more general, polymorphic forms of safe modular shared state interactions.
- We prove our system sound through progress and preservation theorems that show the absence of unsafe interference in correctly typed programs. Our design ensures *memory safety* and *data-race freedom*, where linear resources are shared via protocol composition (a partial commutative monoid [16, 6]).
- We evaluate the expressiveness of our system by discussing how our core shared memory protocol framework is capable of expressing safe, typeful message-passing idioms.

Next, we briefly introduce the language that “hosts” our protocols, with the remaining text focused on discussing new protocol-level features. Sections 2 and 3 introduce our novel definition of protocol composition and its extensions to support abstract protocols. Section 4 discusses technical results. The paper ends with discussions of expressiveness, related work, and conclusions.

1.1 Preliminaries: Language Overview

Our language supports fork/join concurrency combined with lock-based mutual exclusion, where all threads share a common heap. We use the variant of the polymorphic λ -calculus shown in Fig. 1. For convenience, the grammar is let-expanded [27] so that all constructs, except let, are defined over values. The language includes first-class functions (λ), records ($\{\bar{f} = v\}$) that label a value as \bar{f} , and tagged values ($t\#v$) to mark a value with a tag. Standard constructs are used for field selection, application, let blocks, memory allocation, deletion, assignment, dereference, and case analysis.

¹ Note that we have not proven the decidability of the entire type system, but only of the protocol composition algorithm which is at its core. The remainder of the type system is more conventional and we did not encounter difficulties with decidability when implementing similar rules in our prior work [18].

² See: <http://www.cs.cmu.edu/~foliveir/protocol-composition.html>

“lock \bar{v} ” atomically locks a non-empty set of locations (ensuring both mutual exclusion and forbidding re-entrant uses) and analogously with “unlock \bar{v} ”. “fork e ” executes the *expression* in a new thread while sharing access to the common global heap. The operational semantics are standard and, as such, are only shown in the companion *Technical Report* [20]. They produce the standard evaluation of the language’s constructs such as creating or deleting memory, spawning new threads, etc.

Mutable References

We type mutable references by following the design proposed in L^3 [1]. Therefore, a mutable cell is decomposed into two components: a pure *reference*, which can be freely copied; and a linear [10] *capability*, a resource that is used to track the contents of that cell. To link a reference to its respective capability, we use location-dependent types. For instance, a new cell has type $\exists l. (!\mathbf{ref} \ l) :: (\mathbf{rw} \ l \ A)$. This type abstracts the fresh location, l , that was created by the memory allocation. Furthermore, we are given a reference of type “! $\mathbf{ref} \ l$ ” to mean a pure/duplicable (!) **reference** to a location l , where the information about the contents of that location is stored in the linear capability for l . The permission to access (e.g. dereference) the contents of a cell requires both the reference and the capability to be available.

Our capabilities follow the format “ $\mathbf{rw} \ l \ A$ ”, meaning a **read-write** capability to a location l that currently has contents of type A (the type of the value, given in “new v ”, that initializes the new cell). We depart from [1] by making capabilities typing artifacts that only exist at the level of typing. Consequently, capabilities are managed implicitly by the type system rather than manually manipulated by the programmer via language constructs. However, we may still need to associate a capability with another type. For this reason, we use the notion of *stacking* [19]. In $\exists l. (!\mathbf{ref} \ l) :: (\mathbf{rw} \ l \ A)$ we see that the capability to l is stacked on top of “! $\mathbf{ref} \ l$ ” since the capability is to the right of the “ $::$ ”. This allows the capability to be bundled together with the **ref** type, but no action is required to unbundle them if they are needed separately. We refer to prior work [16, 19, 18, 1] for more details on the use of capabilities, locations, and stacking, as well as convenient abbreviations. Here, it suffices to assume that they are handled automatically by the type system, as our focus here is on safely sharing the linear resources.

In the scheme above, all the variables that reference the same location also share a single linear (i.e. “exclusively owned” or “unique”) capability that tracks the changes to that location’s contents (as shown in Fig. 2).

```

// assume 'y' in scope
y := "ok!";
let x = y in
  x := false;
  delete x;
  !y // Type Error: missing capability to location '!'.

```

y	$:\mathbf{ref} \ l, \mathbf{rw} \ l \ \mathbf{int}$
y	$:\mathbf{ref} \ l, \mathbf{rw} \ l \ \mathbf{string}$
x	$:\mathbf{ref} \ l, y:\mathbf{ref} \ l, \mathbf{rw} \ l \ \mathbf{string}$
x	$:\mathbf{ref} \ l, y:\mathbf{ref} \ l, \mathbf{rw} \ l \ \mathbf{boolean}$
x	$:\mathbf{ref} \ l, y:\mathbf{ref} \ l$

■ **Figure 2** Tracking linear capabilities.

However, this tracking relies on a compile-time approximation of how variables alias, which constrains how state can be used. Since the linear capability must be (linearly) threaded through the program, this scheme forbids aliasing idioms that require “simultaneous” access to aliased state, such as when multiple threads share access to a cell. To enable this form of sharing, we split a linear resource into multiple protocols. Each protocol controls how an alias interacts with the shared state, without depending on precise knowledge of which variables alias each other. We will continue with a brief presentation of the base language, before diving into the details of our sharing mechanism in Section 2.

Types

(Note that rely and guarantee types will only be discussed in the next section, when we present sharing). Our types (Fig. 3) follow the connectives of linear logic [10]. For this reason a function type uses \multimap (instead of \rightarrow) to denote a linear function. The linear restriction can be lifted when

$x \in \text{VARIABLES}$	$X \in \text{TYPE VARIABLES}$	$\rho \in \text{LOCATION CONSTANTS}$	$l \in \text{LOCATION VARIABLES}$
$\tau \in \text{TAGS}$	$\mathbf{f} \in \text{FIELDS}$	$p ::= \rho l$	$u ::= l X$
		$U ::= p A$	
$A ::= !A$	(pure/persistent)	$(\text{rec } X(\bar{u}).A)[\bar{U}]$	(recursive type)
$A \multimap A$	(linear function)	$A \oplus A$	(alternative)
$[\mathbf{f} : A]$	(record)	$A \& A$	(intersection)
$\sum_i \tau_i \# A_i$	(tagged sum)	$\text{rw } p A$	(read-write capability to p)
$\forall l.A$	(universal location)	none	(empty resource)
$\exists l.A$	(existential location)	top	(top)
$\forall X <: A.A$	(bounded universal type)	$A :: A$	(stacking)
$\exists X <: A.A$	(bounded existential type)	$A * A$	(separation)
ref p	(reference type)	$A \Rightarrow A$	(rely)
$X[\bar{U}]$	(type variable)	$A ; A$	(guarantee)

Notes: we simplify $X[]$ to X ; \oplus , $\&$, $*$, $+$ are commutative and associative.

■ **Figure 3** Types (A).

the type is preceded by a “bang”, such as in $!A$, which denotes a pure/duplicable type. Records are typed as $[\mathbf{f} : A]$ where each field \mathbf{f} types the value of the record with some type A . $\sum_i \tau_i \# A_i$ denotes a single tagged type or a sequence of tagged types separated by $+$ (such as “ $a \# A + b \# B + c \# C$ ”). We have separate existential and universal quantification over locations and types, since locations and types are of different kinds. Note that we leave \forall/\exists as typing artifacts and as such they do not have corresponding constructs in the language. Quantification over types can provide a type bound (on the right of $<$;) and where **top** is assumed by default when the bound is omitted.

Our recursive types (assumed to be non-bottom types) are equi-recursive, interpreted co-inductively, and satisfy the usual folding/unfolding principle:

$$(\text{rec } X(\bar{u}).A)[\bar{U}] = A\{(\text{rec } X(\bar{u}).A)/X\}\{\bar{U}/\bar{u}\} \quad (\text{EQ:REC})$$

Recursive types may include a list of type/location parameters (\bar{u}) that are substituted by some type/location (\bar{U}) on unfold, besides unfolding the recursive type variable (X).

We use \oplus to denote a union of alternative types, and $\&$ to denote a linear choice of different types. **none** is the empty resource. Finally we have “ $A_0 :: A_1$ ” for stacking resource A_1 on top of A_0 . Stacking is not commutative, so that it is not guaranteed that “ $A_0 :: A_1 :: A_2$ ” can be used in place of “ $A_0 :: A_2 :: A_1$ ”. To enable resource commutation, we use $*$ such that “ $A_0 :: (A_1 * A_2)$ ” and “ $A_0 :: (A_2 * A_1)$ ” are interchangeable via subtyping. For clarity, we will review these type annotations as we present examples further below. Note that we do not syntactically distinguish resources (such as capabilities or protocols) from value-inhabited types. However, the type system ensures that types such as **none** can never be used to type a value. Indeed, even though “wrong” types can be assumed (such as in a function’s argument) they can never actually be introduced as values.

Type System

To enable automatic threading of resources, we use a type-and-effect system with judgments of the form: $\Gamma \mid \Delta_0 \vdash e : A \vdash \Delta_1$ stating that with lexical environment Γ and linear resources Δ_0 we assign the expression e the type A , with effects resulting in the resources in Δ_1 .

$\Gamma \mid \Delta_0 \vdash e : A \dashv \Delta_1$ Typing rules, (τ :*)

$\frac{(\tau:\text{PURE}) \quad \Gamma \mid \cdot \vdash v : A \dashv \cdot}{\Gamma \mid \cdot \vdash v : !A \dashv \cdot}$	$\frac{(\tau:\text{PURE-ELIM}) \quad \Gamma, x : A_0 \mid \Delta_0 \vdash e : A_1 \dashv \Delta_1}{\Gamma \mid \Delta_0, x : !A_0 \vdash e : A_1 \dashv \Delta_1}$	$\frac{(\tau:\text{FRAME}) \quad \Gamma \mid \Delta_0 \vdash e : A \dashv \Delta_1}{\Gamma \mid \Delta_0, \Delta_2 \vdash e : A \dashv \Delta_1, \Delta_2}$
$\frac{(\tau:\text{FUNCTION}) \quad \Gamma \mid \Delta, x : A_0 \vdash e : A_1 \dashv \cdot}{\Gamma \mid \Delta \vdash \lambda x. e : A_0 \multimap A_1 \dashv \cdot}$	$\frac{(\tau:\text{APPLICATION}) \quad \Gamma \mid \Delta_0 \vdash v_0 : A_0 \multimap A_1 \dashv \Delta_1 \quad \Gamma \mid \Delta_1 \vdash v_1 : A_0 \dashv \Delta_2}{\Gamma \mid \Delta_0 \vdash v_0 v_1 : A_1 \dashv \Delta_2}$	
$\frac{(\tau:\text{NEW}) \quad \Gamma \mid \Delta_0 \vdash v : A \dashv \Delta_1}{\Gamma \mid \Delta_0 \vdash \text{new } v : \exists l. ((\text{ref } l) :: (\text{rw } l A)) \dashv \Delta_1}$	$\frac{(\tau:\text{DELETE}) \quad \Gamma \mid \Delta_0 \vdash v : \exists l. ((\text{ref } l) :: (\text{rw } l A)) \dashv \Delta_1}{\Gamma \mid \Delta_0 \vdash \text{delete } v : \exists l. A \dashv \Delta_1}$	
$\frac{(\tau:\text{ASSIGN}) \quad \Gamma \mid \Delta_0 \vdash v_1 : A_0 \dashv \Delta_1 \quad \Gamma \mid \Delta_1 \vdash v_0 : \text{ref } p \dashv \Delta_2, \text{rw } p A_1}{\Gamma \mid \Delta_0 \vdash v_0 := v_1 : A_1 \dashv \Delta_2, \text{rw } p A_0}$	$\frac{(\tau:\text{LET}) \quad \Gamma \mid \Delta_0 \vdash e_0 : A_0 \dashv \Delta_1 \quad \Gamma \mid \Delta_1, x : A_0 \vdash e_1 : A_1 \dashv \Delta_2}{\Gamma \mid \Delta_0 \vdash \text{let } x = e_0 \text{ in } e_1 \text{ end} : A_1 \dashv \Delta_2}$	
$\frac{(\tau:\text{DEREFERENCE-LINEAR}) \quad \Gamma \mid \Delta_0 \vdash v : \text{ref } p \dashv \Delta_1, \text{rw } p A}{\Gamma \mid \Delta_0 \vdash !v : A \dashv \Delta_1, \text{rw } p ![]}$	$\frac{(\tau:\text{LOCOPENBIND}) \quad \Gamma, l : \text{loc} \mid \Delta_0, x : A_1 \vdash e : A_2 \dashv \Delta_1}{\Gamma \mid \Delta_0, x : \exists l. A_1 \vdash e : A_2 \dashv \Delta_1}$	
$\frac{(\tau:\text{SUBSUMPTION}) \quad \Gamma \vdash \Delta_0 <: \Delta_1 \quad \Gamma \mid \Delta_1 \vdash e : A_0 \dashv \Delta_2 \quad \Gamma \vdash A_0 <: A_1 \quad \Gamma \vdash \Delta_2 <: \Delta_3}{\Gamma \mid \Delta_0 \vdash e : A_1 \dashv \Delta_3}$		

Notes: bounded variables of a construct and type/location variables of quantifiers must be fresh in the rule's conclusion.

■ **Figure 4** Typing rules (selected).

The typing environments are defined as follows:

$\Gamma ::= \cdot \quad (\text{empty})$ $\mid \Gamma, x : A \quad (\text{variable binding})$ $\mid \Gamma, p : \text{loc} \quad (\text{location assertion})$ $\mid \Gamma, X <: A \quad (\text{bound assertion})$ $\mid \Gamma, X : k \quad (\text{kind assertion})$	$\Delta ::= \cdot \quad (\text{empty})$ $\mid \Delta, x : A \quad (\text{linear binding})$ $\mid \Delta, A \quad (\text{linear resource})$ $k ::= \text{type} \mid \text{type} \rightarrow k \mid \text{loc} \rightarrow k \quad (\text{kinds})$
--	---

Recursive type variables are given an \rightarrow kind, where the left hand side tracks the type/location kind of a parameter of that recursive type.

Fig. 4 includes a few selected typing rules. Additional rules are shown below as they become relevant to the discussion on sharing, with the remainder left to the T.R.. (τ :PURE) types a value as pure if the value does not use any resources. If a variable is of a pure type, then (τ :PUREELIM) allows the binding to be moved to the linear context with its type explicitly “banged” with !. (τ :FRAME) enables framing [26] resources that are not used by an expression, just threaded through the expression. Since a function, (τ :FUNCTION), can depend on the resources inside of Δ (which the function captures), a functional value must be linear. However, the function can later be rendered pure (!) through the use of (τ :PURE) if the set of resources it captures is actually empty. (τ :APPLICATION) is the standard rule. As discussed above, (τ :NEW) and (τ :DELETE) manipulate types that abstract the underlying location that was created or that is to be deleted. (τ :ASSIGN) updates the contents of a location with the type of the newly assigned value. (τ :LET) threads the effects of e_0 to the initial linear resources of e_1 , sequencing

the evaluation of the expressions as usual. (τ :DEREFERENCE-LINEAR) removes the contents of a cell, leaving the residual “unit” type behind (the semantics leave the cell unchanged but unusable through typing). (τ :LOCOPENBIND) illustrates the non-syntax-directed opening of existential location packages.

The subtyping rules are deferred to the T.R., but it suffices to know the subtyping judgment, $\Gamma \vdash A_0 <: A_1$, which states that A_0 is a subtype of A_1 , meaning that A_0 can be used anywhere A_1 is expected. An analogous judgment governs subtyping between linear environments, $\Gamma \vdash \Delta_0 <: \Delta_1$. Thus, the (τ :SUBSUMPTION) rule simply states that we can type an expression while using weaker assumptions and ensuring a stronger result and effect, as these types cannot break the conclusion’s types expectations.

2 A Protocol for Modeling join

We begin by describing how non-abstracted protocols compose and how rely-guarantee protocols work in the concurrent setting. Our language supports the fork/join model of concurrency, in which a join is encoded via shared state interactions. There are two participants in this interaction: the Main thread and the Forked thread. The forked thread computes some *result*. When the main thread joins the forked thread it will *wait* until the result becomes available, if it is not yet ready. Our primitives to interact with shared state are reading/writing and locking/unlocking. Because of this, our protocols must explicitly model the “wait for result” cycle of a join.³ A thread scheduler could reduce or eliminate the spinning caused by this “busy-wait”, but this is beyond the scope of our discussion. We define the two protocols as follows:

$$\begin{aligned} F &\triangleq \text{Wait} \Rightarrow \text{Result} ; \mathbf{none} \\ M &\triangleq (\text{Wait} \Rightarrow \text{Wait} ; M) \oplus (\text{Result} \Rightarrow \text{Done} ; \text{Done}) \end{aligned}$$

Each protocol contains a sequence of steps that control the use of locks and specify the (type) assumptions on that locked state. Since locks hide all private actions, the protocols will only need to model the changes that become visible upon unlocking. These changes are bounded by a single lock-unlock block, which is mapped to a single rely \Rightarrow guarantee step in the protocol. When we lock a cell we will *assume* that the state is of some type and, when we eventually unlock that cell, we will *guarantee* that it changed to some other type. Multiple steps can be sequenced using the ; operator.

The forked thread will be given the F protocol. This protocol initially assumes that the shared state is of type *Wait* on locking. In order to legally unlock that cell, we must first fulfill the obligation to mutate the state to *Result*. Once that guarantee is obeyed the protocol continues as **none**. This empty resource type models termination since the forked thread will never be able to access that shared state again. Note that since subsequent steps may be influenced by the guarantee of the current step, a protocol step is to be interpreted as “*Wait* \Rightarrow (*Result* ; **none**)”.

The Main protocol includes two alternative (\oplus) steps that describe different uses of the shared state. If we find the shared cell containing the *Wait* type then the main thread must leave the state with the same type, before later retrying M. Otherwise, if we find the cell containing a *Result*, we know that F has already terminated and can no longer access the shared state. In that situation, we mutate the cell to *Done* and unlock it so that each lock always has a matching unlock. Afterwards, the protocol continues as *Done*, a type that is just a regular linear capability. Thus, M recovered ownership of the shared state and *Done* can continue to be used without locking since the cell is no longer shared. We can now give concrete definitions for *Wait*, *Result*, and *Done* as types describing a single capability

³ Each protocol must be aware of all valid states, as an omission would leave room for unsafe interference, such as when later re-splitting that protocol.

to location l as follows:

$\text{Wait} \triangleq \mathbf{rw} / \text{Wait}\#\![\]$ $\text{Result} \triangleq \mathbf{rw} / \text{Result}\#\!\mathbf{int}$ $\text{Done} \triangleq \mathbf{rw} / \![\]$

Wait is a capability to location l containing a tagged value, where Wait is the tag and “[]” (a pure empty record) is the type of the value. Result is a capability for l containing an integer value tagged with Result . The two tags will enable us to distinguish between the Wait and Result alternatives by using standard case analysis. With Done the content is an empty pure record (“unit”).

Each protocol describes an alias’s local, isolated view of the evolution of the shared state. Thus, we can discuss the uses of each protocol independently. Because a protocol is a linear resource, the forked thread will “consume” or “capture” F in its context, making it unavailable to the main thread. As with any linear resource, F is tracked by the linear typing environment (Δ) and is either used by an expression or threaded through to the next expression. However, the forked thread and main thread can share the enclosing lexical typing environment (Γ) because it only contains pure/duplicable assumptions. A possible use of the F protocol follows.

<pre> 3 fork 4 let r = work {} in 5 lock c; 6 c := Result#r; 7 unlock c 8 end </pre>	<table border="0"> <tr> <td style="padding-right: 10px;">$\Gamma = c : \mathbf{ref} \ l, l : \mathbf{loc}$</td> <td style="padding-right: 10px;"> </td> <td>$\Delta = \text{work} : \![\] \rightarrow \mathbf{int}, F$</td> </tr> <tr> <td>$\Gamma = r : \mathbf{int}, \dots$</td> <td> </td> <td>$\Delta = \text{Wait} \Rightarrow (\text{Result}; \mathbf{none})$</td> </tr> <tr> <td>$\Gamma = \dots$</td> <td> </td> <td>$\Delta = \text{Wait}, (\text{Result}; \mathbf{none})$</td> </tr> <tr> <td>$\Gamma = \dots$</td> <td> </td> <td>$\Delta = \text{Result}, (\text{Result}; \mathbf{none})$</td> </tr> <tr> <td>$\Gamma = \dots$</td> <td> </td> <td>$\Delta = \mathbf{none}$</td> </tr> <tr> <td>$\Gamma = \dots$</td> <td> </td> <td>$\Delta = \cdot$</td> </tr> </table>	$\Gamma = c : \mathbf{ref} \ l, l : \mathbf{loc}$		$\Delta = \text{work} : \![\] \rightarrow \mathbf{int}, F$	$\Gamma = r : \mathbf{int}, \dots$		$\Delta = \text{Wait} \Rightarrow (\text{Result}; \mathbf{none})$	$\Gamma = \dots$		$\Delta = \text{Wait}, (\text{Result}; \mathbf{none})$	$\Gamma = \dots$		$\Delta = \text{Result}, (\text{Result}; \mathbf{none})$	$\Gamma = \dots$		$\Delta = \mathbf{none}$	$\Gamma = \dots$		$\Delta = \cdot$
$\Gamma = c : \mathbf{ref} \ l, l : \mathbf{loc}$		$\Delta = \text{work} : \![\] \rightarrow \mathbf{int}, F$																	
$\Gamma = r : \mathbf{int}, \dots$		$\Delta = \text{Wait} \Rightarrow (\text{Result}; \mathbf{none})$																	
$\Gamma = \dots$		$\Delta = \text{Wait}, (\text{Result}; \mathbf{none})$																	
$\Gamma = \dots$		$\Delta = \text{Result}, (\text{Result}; \mathbf{none})$																	
$\Gamma = \dots$		$\Delta = \mathbf{none}$																	
$\Gamma = \dots$		$\Delta = \cdot$																	

Γ contains a reference (c) to the location (l) that is being shared by the protocol, and Δ contains a variable with the (linear) function that computes the `work` that the thread will do. (In this example both protocols refer to a well-known common location, but our technique also allows each protocol to \exists abstract its locations.) Line 4 consumes the function `work` by calling it, storing the result in variable `r`. At this point we want to update the shared state to signal that the result is ready. Since we are accessing shared state in a multi-threaded environment we first lock the shared location that is being referenced by `c`. To type a lock we must map the locations listed in the `lock` to those contained in the rely type of the protocol. Well-formedness conditions on the protocols ensure that, at each step, the rely and the guarantee types refer the same set of locations so that no lock on a location goes without a respective unlock.

$$\frac{\overline{\Gamma \mid \cdot \vdash v : \mathbf{ref} \ p \ \vdash \cdot} \quad \mathbf{locs}(A_0) = \bar{p}}{\Gamma \mid \Delta, A_0 \Rightarrow A_1 \vdash \mathbf{lock} \ \bar{v} : \![\] \ \vdash \Delta, A_0, A_1} \text{ (T:LOCK-RELY)}$$

When locking (line 5), the step of F is broken down into its two components: the rely type (Wait) and the guarantee type ($\text{Result}; \mathbf{none}$). While Wait describes the linear resources that are now available to use, the guarantee type is an obligation to mutate the state to fulfill the given type before unlocking. Indeed, line 7 is only valid because the shared state was modified to match the promised guarantee type (Result).

$$\frac{\overline{\Gamma \mid \cdot \vdash v : \mathbf{ref} \ p \ \vdash \cdot} \quad \mathbf{locs}(A_0) = \bar{p}}{\Gamma \mid \Delta, A_0, (A_0; A_1) \vdash \mathbf{unlock} \ \bar{v} : \![\] \ \vdash \Delta, A_1} \text{ (T:UNLOCK-GUARANTEE)}$$

(with parenthesis used for clarity). Once the guarantee is fulfilled, we can move on to the next step of the protocol (in the case of F , \mathbf{none} ; or A_1 , in the case of the rule above). The \mathbf{none} type is the empty resource that can be automatically discarded, leaving Δ empty (\cdot). Thus, the uses of protocols are mapped to the (T:LOCK-RELY) and (T:UNLOCK-GUARANTEE) rules that step a protocol. We now show the rest of the encoding:

<pre> 1 let newFork = λwork. 2 let c = new Wait#{ } in 3 fork ... // lines 3 to 8 shown above. </pre>	<table border="0"> <tr> <td style="padding-right: 10px;">$\Gamma = \cdot$</td> <td style="padding-right: 10px;"> </td> <td>$\Delta = \text{work} : \![\] \rightarrow \mathbf{int}$</td> </tr> <tr> <td>$\Gamma = c : \mathbf{ref} \ l, l : \mathbf{loc}$</td> <td> </td> <td>$\Delta = \mathbf{rw} / \text{Wait}\#\![\], \dots$</td> </tr> <tr> <td>$\Gamma = \dots$</td> <td> </td> <td>$\Delta = M, F, \dots$</td> </tr> </table>	$\Gamma = \cdot$		$\Delta = \text{work} : \![\] \rightarrow \mathbf{int}$	$\Gamma = c : \mathbf{ref} \ l, l : \mathbf{loc}$		$\Delta = \mathbf{rw} / \text{Wait}\#\![\], \dots$	$\Gamma = \dots$		$\Delta = M, F, \dots$
$\Gamma = \cdot$		$\Delta = \text{work} : \![\] \rightarrow \mathbf{int}$								
$\Gamma = c : \mathbf{ref} \ l, l : \mathbf{loc}$		$\Delta = \mathbf{rw} / \text{Wait}\#\![\], \dots$								
$\Gamma = \dots$		$\Delta = M, F, \dots$								

To simplify the presentation, our term language is stripped of type annotations. However, the `newFork` function has type $!(\![\] \multimap \mathbf{int}) \multimap (\![\] \multimap \mathbf{int})$ where the argument of this pure function is the `work` to be done by the thread, as was shown above. The resulting function is the `join` (shown below) that, once called, waits for the forked thread’s result. Line 2 creates the cell that will be shared by the main and forked threads. This new cell, although typed $\exists l.(\mathbf{ref} \ l) :: (\mathbf{rw} \ l \ \mathbf{Wait} \ \# \ ![\])$, is automatically opened by the type system via $(\mathbf{T}:\mathbf{LOCOPENBIND})$ to allow direct access to the `ref l` reference via variable `c`.

Line 3 shares the cell by splitting the capability to location `l` into the `M` and `F` protocols. This split is done in a non-syntax-directed way through $(\mathbf{T}:\mathbf{SUBSUMPTION})$ (of Fig. 4), combined with the following rule for subtyping on Δ ’s:

$$\frac{\Gamma \vdash \Delta_0 <: \Delta_1 \quad \Gamma \vdash A_0 \Rightarrow A_1 \parallel A_2}{\Gamma \vdash \Delta_0, A_0 <: \Delta_1, A_1, A_2} \text{ (SD:SHARE)}$$

Where the following resource split (\Rightarrow) is used:

$$\Gamma \vdash \mathbf{Wait} \Rightarrow \mathbf{M} \parallel \mathbf{F} \quad (\text{recall: } \mathbf{Wait} \triangleq \mathbf{rw} \ l \ \mathbf{Wait} \ \# \ ![\])$$

This split results in the capability to location `l` being replaced by the two protocols, `M` and `F`, in Δ . The composition check (described in the next subsection) relies on the knowledge that `M` and `F` share the same location. Once the protocols are known to compose safely, however, we no longer need to track this sharing—each protocol can abstract the location being accessed under a different name, and they can be used independently. The fork expression is typed by consuming the resources that the fork will use (such as `F` in the fork of line 3):

$$\frac{\Gamma \mid \Delta \vdash e : ![\] \dashv \cdot}{\Gamma \mid \Delta \vdash \mathbf{fork} \ e : ![\] \dashv \cdot} \text{ (T:FORK)}$$

This rule is somewhat similar to $(\mathbf{T}:\mathbf{FUNCTION})$, but the result type is unit because `fork` does not produce a result. Thus, a fork is executed for the effects it produces on the shared state. As such, to avoid leaking resources, the final residual resources of the forked expression must be empty and the resulting value pure (note that “ $!A <: ![\]$ ”).

Finally, we show the `join` function that will “busy-wait” for the forked thread to produce a result. Its use of both `recursion` and `case` analysis should be straightforward as they follow standard usage. The following text will focus on the less obvious details.

<pre> 9 λ_.rec R. 10 11 lock c; 12 case !c of 13 Wait#x → // must restore linear value 14 c := Wait#x; 15 unlock c; 16 R // retries 17 Result#x → 18 unlock c; 19 delete c; 20 x 21 end 22 end 23 end </pre>	<pre> Δ = (Wait ⇒ (Wait; M)) ⊕ (Result ⇒ (Done; Done)) [a]Δ = Wait ⇒ (Wait; M) [b]Δ = Result ⇒ (Done; Done) [a]Δ = Wait, (Wait; M) [b]Δ = Result, (Done; Done) [a]Δ = rw l ![], (Wait; M) [b]Δ = rw l ![], (Done; Done) [a]Δ = rw l ![], (Wait; M) [a]Δ = rw l ![], (Wait; M) [a]Δ = Wait, (Wait; M) [a]Δ = M [b]Γ = x : int, ... Δ = rw l ![], (Done; Done) [b]Γ = x : int, ... Δ = rw l ![] [b]Γ = x : int, ... Δ = · </pre>
--	--

We omit Γ to center the discussion on the contents of Δ . The alternative type (\oplus) lists a union of types that may be valid at that point in the program. To use such a type, an expression must consider

each alternative individually via (T:ALTERNATIVE-LEFT):

$$\frac{\Gamma \mid \Delta_0, A_0 \vdash e : A_2 \dashv \Delta_1 \quad \Gamma \mid \Delta_0, A_1 \vdash e : A_2 \dashv \Delta_1}{\Gamma \mid \Delta_0, A_0 \oplus A_1 \vdash e : A_2 \dashv \Delta_1} \text{ (T:ALTERNATIVE-LEFT)}$$

The breakdown of \oplus (line 10) is done automatically by the type system. Thus, the body of the recursion must be typed individually under each one of those alternatives, marked as **[a]** and **[b]**. The type of the resource on each alternative contains a sum type that matches different branches in the **case** of line 12. Note that it is safe for this sum type to only match a subset of the branches that the **case** lists. The remaining branches are simply ignored when typing the **case** with that sum type:

$$\frac{\Gamma \mid \Delta_0 \vdash v : \sum_i \tau_i \# A_i \dashv \Delta_1 \quad \overline{\Gamma \mid \Delta_1, x_i : A_i \vdash e_i : A \dashv \Delta_2} \quad i \leq j}{\Gamma \mid \Delta_0 \vdash \text{case } v \text{ of } \tau_j \# x_j \rightarrow e_j \text{ end} : A \dashv \Delta_2} \text{ (T:CASE)}$$

This enables the same **case** to produce different effects, such as obeying incompatible guarantees, based solely on the tagged contents of v . For instance, the **Result** branch will recover ownership and destroy the shared cell (line 19), while the **Wait** branch must restore the linear value of that cell (that was removed by the linear dereference of line 12, that left “**rw l ![]**” in Δ) before retrying. Although line 19 deletes the cell, we first unlock the cell to fulfill the “Done; Done” guarantee of the final protocol step.

A rely-guarantee protocol is a specification of each lock-unlock usage, modeled by a protocol type. Therefore, we will continue the discussion on interference by only looking at the protocols, while omitting the actual concrete programs that use them.

2.1 Checking Safe Protocol Composition

We now introduce our main contribution: a novel axiomatic definition of protocol composition, which is later extended to support abstraction. Composing protocols over some shared state requires considering all possible ways in which the use of these protocols may be interleaved. Thus, regardless of the non-deterministic way by which aliases are interleaved at run-time, a correct composition will ensure that all possible uses are safe.

Intuitively, a binary protocol split will generate an infinite binary tree representing all combinations of interleaved uses of the two new protocols. Each node of that tree has two children based on which protocol remains stationary while the other is stepped. Since this tree may be infinite, we must build a co-inductive proof of safe interference. We only consider binary splits when checking composition but since a protocol can be later re-split, there is no limit to how many protocols may share some state.

The two protocols, M and F , shown above contain a finite number of different positions. We call a *configuration* the combination of the positions of each protocol and the current type of the shared resources. Each configuration is of the form:

$$\langle \Gamma \vdash \text{Resources} \Rightarrow \text{Protocol} \parallel \text{Protocol} \rangle$$

Thus, when we split a **Wait** cell into protocols M^4 and F^5 , we get the following set of configurations that simulate the uses done via the protocols (seen as atomic public transitions of lock-unlock uses, corresponding to the respective rely and guarantee types):

$$\{ \textcircled{1} \langle \Gamma \vdash \text{Wait} \Rightarrow M \parallel F \rangle, \textcircled{2} \langle \Gamma \vdash \text{Result} \Rightarrow M \parallel \text{none} \rangle, \\ \textcircled{3} \langle \Gamma \vdash \text{Done} \Rightarrow \text{Done} \parallel \text{none} \rangle, \textcircled{4} \langle \Gamma \vdash \text{none} \Rightarrow \text{none} \parallel \text{none} \rangle \}$$

⁴ $M \triangleq (\text{Wait} \Rightarrow (\text{Wait}; M)) \oplus (\text{Result} \Rightarrow (\text{Done}; \text{Done}))$

⁵ $F \triangleq \text{Wait} \Rightarrow (\text{Result}; \text{none})$

$$\begin{aligned}
P, Q &::= (\mathbf{rec} X(\overline{u}).P)[\overline{U_P}] \mid X[\overline{U_P}] \mid P \oplus P \mid P \& P \mid \mathbf{none} \\
&\mid S \Rightarrow P \mid S ; P \mid \exists l.P \mid \forall l.P \mid \exists X <: A.P \mid \forall X <: A.P \\
S &::= (\mathbf{rec} X(\overline{u}).S)[\overline{U_S}] \mid X[\overline{U_S}] \mid S \oplus S \mid S \& S \mid \mathbf{none} \mid A * A \mid \mathbf{rw} \ p \ A \\
R &::= P \mid S
\end{aligned}$$

Note: that the structure of allowed protocols is further restricted via protocol composition, beyond the syntactical categories above. Namely, abstraction is only enabled by the rules of Section 3.3.

■ **Figure 5** Grammar restrictions for checking safe protocol composition: *Protocols, States, and Resources*.

Configuration ❶ represents the initial split of Wait into M and F. Starting from some configuration, we will leave one of the protocols stationary while we simulate a use of the shared state (a *step*) with the remaining protocol. From ❶ if we step M we will stay in the same configuration. If instead F is stepped, we get to configuration ❷ that changed the state to Result and terminates the F protocol. By continuing to step M we have the two last configurations: ❸ where the last step of M is ready to recover ownership, and ❹ where the ownership of the shared resource was recovered and all protocols have terminated (i.e. all resources are empty, **none**).

Upon sharing, the ownership of the shared resources belongs to all intervening protocols; all protocols can access the shared resources through locking. Ownership recovery means that this ownership is given back to one single protocol and “revoked” from all remaining protocols. In our protocols, recovery is modeled via protocol termination, such that a step transitions to a state rather than to another protocol step. However, to be safe, we must be sure that this permanent ownership transfer only occurs on the *last* protocol to terminate, ensuring that no other protocol may accidentally assume that that shared state is still available. The ownership recovery in ❸ transfers Done from the “pool” of shared resources to the alias that uses the last step of the M protocol. We also see by ❹ that this stepping consumes both the shared resource (leaving it as **none**) and the final “step” of M (leaving the protocol position also as **none**).

All protocol configurations shown above can take a step. (Even **none** can take a vacuous step that remains in the same configuration since **none** cannot change the shared resources.) Therefore, each protocol will always find an expected state in the shared cell regardless of how protocols are interleaved—i.e. all interference is safe since no configuration is stuck. A stuck configuration occurs when at least one of the protocols cannot take a step with the current type of the shared resources. For instance, $\langle \Gamma \vdash \mathbf{Result} \Rightarrow M \parallel F \rangle$ cannot take a step with F since F does not rely on Result in any of its available steps. If such stuck configurations were allowed to occur, then a program could fault due to unexpected values stored in shared cells or due to attempts to access cells that were destroyed using wrong assumptions of ownership recovery.

Protocol composition ensures that a resource, R (capabilities or protocols), can be shared (split) as two protocols, P and Q , noted: $\Gamma \vdash R \Rightarrow P \parallel Q$. Fig. 5 lists the grammatical categories (for protocols, states and resources) that we consider when composing protocols. As exemplified above we use a set of *configurations*, C , to represent the positions of each protocol as we traverse all possible interleaved uses of the two new protocols. C is defined as:

$$\begin{aligned}
C &::= \langle \Gamma \vdash R \Rightarrow P \parallel Q \rangle \quad (\text{configuration}) \\
&\mid C \cdot C \quad (\text{configuration union})
\end{aligned}$$

Protocol composition, applied via (sd:SHARE), ensures that all configurations reachable through stepping are themselves able to take a step, as follows:

$$\frac{\langle \Gamma \vdash R \Rightarrow P \parallel Q \rangle \uparrow}{\Gamma \vdash R \Rightarrow P \parallel Q} \quad (\text{WF:SPLIT}) \qquad \frac{C_0 \mapsto C_1 \quad C_1 \uparrow}{C_0 \uparrow} \quad (\text{WF:CONFIGURATION})$$

$C \mapsto C$ **Composition, (c:*)**

$$\begin{array}{c}
\text{(c:STEP)} \\
\frac{\langle \Gamma \vdash R \Rightarrow \mathcal{R}_L[P] \rangle \mapsto C_0 \quad \mathcal{R}_L[\square] = \square \parallel Q \quad \langle \Gamma \vdash R \Rightarrow \mathcal{R}_R[Q] \rangle \mapsto C_1 \quad \mathcal{R}_R[\square] = P \parallel \square}{\langle \Gamma \vdash R \Rightarrow P \parallel Q \rangle \mapsto C_0 \cdot C_1} \\
\text{(c:ALLSTEP)} \\
\frac{C_0 \mapsto C_2 \quad C_1 \mapsto C_3}{C_0 \cdot C_1 \mapsto C_2 \cdot C_3}
\end{array}$$

Composition — Reduction Step, (c-rs:*)

$$\begin{array}{c}
\text{(c-rs:NONE)} \\
\frac{}{\langle \Gamma \vdash R \Rightarrow \mathcal{R}[\mathbf{none}] \rangle \mapsto \langle \Gamma \vdash R \Rightarrow \mathcal{R}[\mathbf{none}] \rangle} \\
\text{(c-rs:STATEINTERSECTION)} \quad \text{(c-rs:PROTOCOLALTERNATIVE)} \\
\frac{\langle \Gamma \vdash R_0 \Rightarrow \mathcal{R}[P] \rangle \mapsto C}{\langle \Gamma \vdash R_0 \& R_1 \Rightarrow \mathcal{R}[P] \rangle \mapsto C} \quad \frac{\langle \Gamma \vdash R \Rightarrow \mathcal{R}[P_0] \rangle \mapsto C}{\langle \Gamma \vdash R \Rightarrow \mathcal{R}[P_0 \oplus P_1] \rangle \mapsto C} \\
\text{(c-rs:PROTOCOLINTERSECTION)} \quad \text{(c-rs:STATEALTERNATIVE)} \\
\frac{\langle \Gamma \vdash R \Rightarrow \mathcal{R}[P_0] \rangle \mapsto C_0 \quad \langle \Gamma \vdash R \Rightarrow \mathcal{R}[P_1] \rangle \mapsto C_1}{\langle \Gamma \vdash R \Rightarrow \mathcal{R}[P_0 \& P_1] \rangle \mapsto C_0 \cdot C_1} \quad \frac{\langle \Gamma \vdash R_0 \Rightarrow \mathcal{R}[P] \rangle \mapsto C_0 \quad \langle \Gamma \vdash R_1 \Rightarrow \mathcal{R}[P] \rangle \mapsto C_1}{\langle \Gamma \vdash R_0 \oplus R_1 \Rightarrow \mathcal{R}[P] \rangle \mapsto C_0 \cdot C_1}
\end{array}$$

Composition — State Stepping, (c-ss:*)

$$\begin{array}{c}
\text{(c-ss:STEP)} \\
\frac{}{\langle \Gamma \vdash S_0 \Rightarrow \mathcal{R}[S_0 \Rightarrow S_1; P] \rangle \mapsto \langle \Gamma \vdash S_1 \Rightarrow \mathcal{R}[P] \rangle} \\
\text{(c-ss:RECOVERY)} \\
\frac{}{\langle \Gamma \vdash S \Rightarrow \mathcal{R}[S] \rangle \mapsto \langle \Gamma \vdash \mathbf{none} \Rightarrow \mathcal{R}[\mathbf{none}] \rangle}
\end{array}$$

Composition — Protocol Stepping, (c-ps:*)

$$\text{(c-ps:STEP)} \\
\frac{}{\langle \Gamma \vdash S_0 \Rightarrow S_1; Q \Rightarrow \mathcal{R}[S_0 \Rightarrow S_1; P] \rangle \mapsto \langle \Gamma \vdash Q \Rightarrow \mathcal{R}[P] \rangle}$$

■ **Figure 6** Basic protocol composition stepping rules.

Where $C \uparrow$ signals the *divergence* of stepping, consistent with the co-inductive nature of protocol composition. We use a double line, as in (wf:CONFIGURATION), to mean that a rule is to be interpreted co-inductively. This definition accounts for protocols that never terminate and also ensures that all protocols can take a step with a given resource.

We now discuss the basic protocol composition definition of Fig. 6. (c:ALLSTEP) synchronously steps all existing configurations, where each configuration is stepped through (c:STEP). We use \mathcal{R}_* (where $*$ is either L or R) to specify the configuration reduction context on one of the protocols of a configuration, while the remaining protocol remains stationary, i.e.:

$$\begin{array}{l}
\mathcal{R}_L[\square] = \square \parallel Q \quad (\text{for the Left protocol, } Q \text{ is stationary}) \\
\mathcal{R}_R[\square] = P \parallel \square \quad (\text{for the Right protocol, } P \text{ is stationary})
\end{array}$$

The subsequent stepping rules use \mathcal{R} to range over both \mathcal{R}_L and \mathcal{R}_R .

We use three distinct label prefixes to group the stepping rules based on whether a rule is stepping over a protocol (c-ps:*), stepping over some state (c-ss:*), or is applicable on both kinds of resource (c-rs:*). (c-rs:NONE) “spins” a configuration since a terminated protocol cannot use the shared resources

but must be stuck-free for consistency with our definition. The following (C-RS:*ALTERNATIVE) and (C-RS:*INTERSECTION) rules “dissect” a resource based on the alternative (\oplus) or choice ($\&$) presented. Each different alternative state must be individually considered by a protocol, while only one alternative step of a protocol needs to be valid. The situation is the reverse for choices: all choices of a protocol must have a valid step, but a step of a protocol can choose which resource to consider when stepping. State stepping, (C-SS:STEP), transitions the step of the protocol and changes the state of the shared resources to reflect the guaranteed state of the protocol. Ownership recovery, (C-SS:RECOVERY), “consumes” the shared state (leaving it as **none**) which models the transfer of ownership of that state back to the client context that uses the final step of the protocol. Protocol stepping, (C-PS:STEP), requires an exact simulation of the rely and guarantee types when stepping both the simulated protocol and the current stepping protocol.

Note that the rules above also enable the re-splitting of a protocol by extending an ownership recovery step. In this situation, we have that the simulation of the original protocol will seamlessly switch from the protocol stepping rules to the state stepping rules.

3 Polymorphic Protocol Composition

Up to this point, protocol composition does a strict stepping of protocols. Consequently, stepping requires each protocol to know the exact type representation of the shared resources. Ideally, to improve both locality and modularity, each protocol should only depend on the type information that is relevant to the actions done through that alias. For instance, the action done through the F protocol of page 6 does not need to know the precise type (*Wait*) that is initially stored in location *l*. Thus, we want to be able to abstract *Wait* as *X* such that the protocol only keeps the typing information that is relevant to that protocol’s *local perspective* on the shared resources:

$$\exists X. (\text{rw } l X \Rightarrow (\text{rw } l \text{Result}\#\text{int} ; \text{none}))$$

Similarly, the wait step of the M protocol only depends on the tag of the shared cell enabling everything else to be abstracted from its perspective:

$$\exists X. (\text{rw } l \text{Wait}\#X \Rightarrow (\text{rw } l \text{Wait}\#X ; M)) \oplus \dots$$

Since rely-guarantee protocols are first-class types, they can move outside the scope of a “module”. Without this form of abstraction, such a move would either expose potentially private information or limit how clients may later re-split the shared resources. While enabling protocols to abstract part of their uses based on their perspective of the shared resources improves modularity and increases flexibility, it also brings new challenges on defining safe protocol composition and ensuring its termination. We will focus the discussion on two new aliasing idioms that this kind of abstraction enables: a) *existential-universal interaction*, how a universally quantified guarantee can safely interact with an existentially quantified rely; and b) *step extensions over abstractions*, how abstractions enable existing protocol steps to be re-split (i.e. nested protocol re-splitting) yet without the risk of introducing unsafe interference on older protocols of that state.

Section 4 approaches the decidability problem. The remaining of this section starts by introducing the basic intuition of how protocol-level abstraction works, before extending our axiomatic definition of composition to account for abstraction.

3.1 Existential-Universal Interaction

Enabling existential abstraction over the contents of the shared state will naturally allow a greater decoupling from the actions done by other aliases to that shared state. However, since a protocol

encodes *sequences* of steps, ensuring safety must also account for the validity of the scope of the opaque type. For instance, consider the composition:

$$\Gamma \vdash \mathbf{rw} \ p \ \mathbf{int} \Rightarrow \exists X. (\mathbf{rw} \ p \ X \Rightarrow \mathbf{rw} \ p \ X ; \mathbf{rw} \ p \ X \Rightarrow \dots) \parallel (\mathbf{rw} \ p \ \mathbf{int} \Rightarrow \mathbf{rw} \ p \ \mathbf{boolean} ; \dots)$$

On the left protocol, the assumption of X extends beyond a single step. Because the right protocol can change the underlying representation of X , this composition cannot be ruled safe. Indeed, if X were of a pure type, the left protocol could potentially reintroduce a type that would unsafely interfere with the right protocol's assumptions on the shared state. Thus, while the left protocol depends on an opaque type, it still requires that the “lifetime” of X extends to the next step.

We now discuss the core ideas that enable the safe composition of protocols that interact over abstractions. First the interaction will only occur via the “lifetime” of the stored type (as it changes on each step), and then we will use bounded quantification to enable types that are less opaque. Consider the following protocols that are sharing a location p :

$$\begin{aligned} \mathbf{Nothing} &\triangleq \exists X. (\mathbf{rw} \ p \ X \Rightarrow \mathbf{rw} \ p \ X ; \mathbf{Nothing}) \\ \mathbf{Full}[Y] &\triangleq \mathbf{rw} \ p \ Y \Rightarrow \forall Z. (\mathbf{rw} \ p \ Z ; \mathbf{Full}[Z]) \end{aligned}$$

The **Nothing** protocol is defined using X to abstract the contents of the shared cell on a single step, while also guaranteeing that X is restored before repeating the protocol. Thus, **Nothing** cannot publicly modify the shared state, although p can undergo private changes. Conversely, **Full** is able to arbitrarily modify the shared state by allowing its clients to pick any type to apply to the \forall of the guarantee. **Full** itself is parametric on the type that is currently stored in the shared cell, Y . Each step of **Full** can exploit the precise local information on how the state was modified, by remembering its own changes to cell p . However, the “lifetime” of X in **Nothing** is restricted to a single step. Naturally, to be able to check this composition in a finite number of steps, we must check the changes done by **Full** abstractly. To illustrate how composition works in this case, consider the following split where p initially holds a value of type **int**:

$$p : \mathbf{loc} \vdash \mathbf{rw} \ p \ \mathbf{int} \Rightarrow \mathbf{Full}[\mathbf{int}] \parallel \mathbf{Nothing}$$

Protocol composition results in the following set of configurations:

$$\begin{aligned} \textcircled{1} \langle p : \mathbf{loc} \quad \vdash \mathbf{rw} \ p \ \mathbf{int} \Rightarrow \mathbf{Full}[\mathbf{int}] \parallel \mathbf{Nothing} \rangle, \\ \textcircled{2} \langle p : \mathbf{loc}, Z : \mathbf{type} \vdash \mathbf{rw} \ p \ Z \Rightarrow \mathbf{Full}[Z] \parallel \mathbf{Nothing} \rangle \end{aligned}$$

The use of abstraction will mean that each configuration may have different assumptions of type (and location) variables. Configuration **1** is the initial configuration given by the split above, which includes the assumption that p is a known **location**. To step **Nothing** from **1**, we must first find a representation type to open the existential. This type is found by unifying the current state of the shared state ($\mathbf{rw} \ p \ \mathbf{int}$) with the rely type of **Nothing** ($\mathbf{rw} \ p \ X$). Thus, we see that X is abstracting **int**. After we open the existential, by exposing the **int** type, we see that the step will preserve **int** resulting in **Nothing** yielding the same **1** configuration. To step **1** with **Full**[**int**], we must consider that its resulting guarantee is abstract. The new configuration, **2**, must consider a fresh type variable to represent that new type that a client can pick. In this case, we used Z to represent that new type. It is straightforward to see that if we were to step **Nothing** from **2** we would remain in configuration **2** following similar reasoning to that done for **1**. Perhaps the surprising aspect is that further steps with **Full** will also yield configurations that are *equivalent* to **2**.

The typing environment plays a crucial role in enabling us to close the proof of safe composition. Although each step of **Full** must consider a fresh type due to the \forall , stepping results in configurations that are equivalent up to renaming of variables and weakening of Γ . Weakening allows us to ignore

variables that no longer occur free in a configuration. This means that further steps with `Full` result in configurations that are equivalent to already seen configurations. Thus, although the set of different types that can be applied to `Full`'s guarantee is infinite, the number of *distinct interactions* that can legally occur through that shared state is finite if we model those interactions abstractly. Lifetime conflicts cannot occur with this technique as even if we open an existential, we must still step the new configuration. Consequently, the problematic composition above would be detected via stepping.

We can use bounded quantification to provide more expressive abstractions that go beyond the fully opaque types used above (which are equivalent to a “<: **top**” bound), and convert this example into one of more practical use. By using appropriate bounds, we can give concrete roles to the `Nothing` and `Full` protocols. Consider that we want to share access to some data structure among several different threads. However, depending on how these threads dynamically use that data structure, it may become important to switch its representation (such as change from a linked list to a binary tree, etc.). Furthermore, we want one specialized thread (the *Controller*) to retain precise control over the data structure and to be allowed to monitor and change its representation. Concurrently, an arbitrary number of other threads (the *Workers*) also have access to the data structure but are limited to only access its Basic operations.

$$\begin{aligned} \bar{W} &\triangleq \exists X <: B. (\mathbf{rw} \ p \ X \Rightarrow \mathbf{rw} \ p \ X ; \bar{W}) \\ C[Y] &\triangleq \mathbf{rw} \ p \ Y \Rightarrow \forall Z <: B. (\mathbf{rw} \ p \ Z ; C[Z]) \end{aligned}$$

As before \bar{W} is committed to preserve the representation type of X although it now has sufficient room to use that type as B . C is now more constrained than before since it is forced to guarantee a type that is compatible with B . However, C retains the possibility of both changing the representation type contained in the shared state, and also of “remembering” the precise (representation) type that was the result of its own local action. Finally, note that we can safely re-split \bar{W} arbitrarily (i.e. $\bar{W} \Rightarrow \bar{W} \parallel \bar{W}$). Protocol composition yields similar set of configurations, but with the bound assumption on Z .

This form of asymmetric interaction over shared state relates to the `full` – `pure` interaction of *access permissions* [3]. A `full` permission allows exclusive write permission to an object, but also enables read-only permissions (`pure`) to co-exists. Consequently, each `pure` permission must assume that other permissions can modify the shared object up to a certain type, the *state guarantee*. While their work focuses on the read-write distinction, and our work is centered on modeling type-changing mutations (so all aliases can write), the example shows that we are able to naturally model similar asymmetric interaction within our protocol framework.

3.2 Inner Step Extension with Specialization

Re-splitting an existing protocol while specializing its interference is possible, provided that its effects remain consistent with those of the original protocol. Namely we can append new steps to an otherwise ownership recovery step, or produce effects that are more precise than those of the original protocol. The first case allows us to connect two protocols together by that recovery step. The latter case is more interesting: when combined with abstraction it allows specialization *within* an existing step (i.e. nested re-splits), enabling new forms of shared state interaction through that abstraction.

To illustrate the expressiveness gains, we revisit the join protocol of Section 2. However, instead of spawning a single thread to compute the work, we re-split the join protocol in two symmetric workers that share the workload. The last of the workers to complete merges the two results together and “signals” the waiting main thread. First, we rewrite the two protocols to enable abstraction on the \bar{M} protocol, and add a choice ($\&$) to the \bar{F} protocol that enables \bar{F} to use the state more than once until it provides a result.

$$\begin{aligned} \bar{F}[X] &\triangleq (\mathbf{rw} \ p \ \bar{W}\#X \Rightarrow \forall Y. (\mathbf{rw} \ p \ \bar{W}\#Y ; \bar{F}[Y])) \ \& \ (\mathbf{rw} \ p \ \bar{W}\#X \Rightarrow \mathbf{rw} \ p \ \mathbf{R}\#\mathbf{int} ; \mathbf{none}) \\ \bar{M} &\triangleq \exists Z. (\mathbf{rw} \ p \ \bar{W}\#Z \Rightarrow \mathbf{rw} \ p \ \bar{W}\#Z ; \bar{M}) \ \oplus \ (\mathbf{rw} \ p \ \mathbf{R}\#\mathbf{int} \Rightarrow \mathbf{rw} \ p \ \mathbf{int} ; \mathbf{rw} \ p \ \mathbf{int}) \end{aligned}$$

As before, M will Wait until there is a Result in p . At that point, M will recover ownership of that cell. Unlike before, M no longer depends on the value tagged as W since it is abstracted as Z . The F protocol now holds two choices ($\&$): the old step that transitions from Wait to Result, and a new step that changes the representation of the value tagged as W and used during the wait phase. The F protocol of Section 2 is a specialization of this protocol since it includes only one of the choices. In here, we specialize F into two symmetric worker protocols. To simplify the presentation, we assume that the worker thread will receive the work parameters through some other mean (such as a pure value shared among threads). Once a worker finishes its job, it will push the resulting `int` to the shared state. If it notices it is the last worker to finish, it will merge the two results together and flag the state as ready, so that Main can proceed.

$$K \triangleq (\text{rw } p \text{ W\#}(E\#[]) \Rightarrow \text{rw } p \text{ W\#}(R\#\text{int}) ; \text{none}) \oplus (\text{rw } p \text{ W\#}(R\#\text{int}) \Rightarrow \text{rw } p \text{ R\#\text{int}} ; \text{none})$$

It is important to note that the new tags/values are nested *inside* the old W tag. This ensures that the new usages remain hidden from M and “look” just like the previous F usage. (There are also no lifetime conflicts since M does not preserve its type assumption on the abstraction beyond a single step.) However, these inner tags are used by the two workers for coordination: the $W\#\text{Empty}$ tag means that neither thread has finished, and $W\#\text{Result}$ means that one of the threads has already finished. We can then re-split F as follows (note the required initial type in F , $E\#[]$, for this split to be valid):

$$\Gamma \vdash F[E\#[]] \Rightarrow K \parallel K$$

Protocol composition follows analogous principles to above, except that we are now simulating the steps of the original F protocol with the steps of the two new K protocols:

$$\{ \langle \Gamma \vdash F[E\#[]] \Rightarrow K \parallel K \rangle, \langle \Gamma \vdash F[R\#\text{int}] \Rightarrow K \parallel \text{none} \rangle, \\ \langle \Gamma \vdash F[R\#\text{int}] \Rightarrow \text{none} \parallel K \rangle, \langle \Gamma \vdash \text{none} \Rightarrow \text{none} \parallel \text{none} \rangle \}$$

Each simulation will match the rely and guarantee types of a step in F with a step in K , even if specializing a \forall of F to a specific type in K . As before, K can choose which step to simulate when given a choice ($\&$) of F steps. Similarly, at least one alternative (\oplus) of K must match a step in F . Therefore, the new K protocols work within the interference of the original F protocol, but specialize its uses of the shared state.

3.3 Composing Abstract Protocols

The composition rules of Fig. 7 complement those of Fig. 6 to enable composing abstract protocols. Weakening on a configuration (up to renaming), ($c\text{-rs:WEAKENING}$), is the crucial mechanism that enables us to close the co-inductive proof when using quantifiers. Thus, when we reach a configuration that is equivalent up to renaming of variables and weakening of Γ , we can close the proof. The ($c\text{-ss:FORALL}^*$) rules do similar stepping to ($c\text{-ss:STEP}$) but considering an abstracted guarantee, which results in a typing environment with the opened abstraction. ($c\text{-ss:OPEN}^*$) exposes the representation type/location (if it exists) before doing a regular step. ($c\text{-ps:FORALL}^*$) and ($c\text{-ps:EXISTS}^*$) open their respective abstraction before doing a regular simulation step. More interestingly, ($c\text{-ps:APP}^*$) enables a simulated step to pick a particular type/location to apply before that regular simulation stepping, enabling step specialization during simulation. In the T.R, we also consider a straightforward extension to protocol composition that enables subtyping over stepping.

3.4 Discussion & Brief Examples

Above, we showed how our local, isolated protocol types can model core interference concepts over a relatively small and simple calculus. We refrained from adding support for more precise states

$$\begin{array}{c}
\frac{(c\text{-rs:WEAKENING})}{\langle \Gamma_0 \vdash R \Rightarrow \mathcal{R}[P] \rangle \mapsto C} \quad \frac{(c\text{-ss:FORALLLOC})}{\langle \Gamma, l : \mathbf{loc} \vdash S \Rightarrow \mathcal{R}[S \Rightarrow P] \rangle \mapsto C} \\
\frac{}{\langle \Gamma_0, \Gamma_1 \vdash R \Rightarrow \mathcal{R}[P] \rangle \mapsto C} \quad \frac{}{\langle \Gamma \vdash S \Rightarrow \mathcal{R}[S \Rightarrow \forall l.P] \rangle \mapsto C} \\
\frac{(c\text{-ss:OPENLOC})}{\langle \Gamma \vdash S \Rightarrow \mathcal{R}[P\{p/l\}] \rangle \mapsto C} \quad \frac{(c\text{-ss:FORALLTYPE})}{\langle \Gamma, X : \mathbf{type}, X <: A \vdash S \Rightarrow \mathcal{R}[S \Rightarrow P] \rangle \mapsto C} \\
\frac{}{\langle \Gamma \vdash S \Rightarrow \mathcal{R}[\exists l.P] \rangle \mapsto C} \quad \frac{}{\langle \Gamma \vdash S \Rightarrow \mathcal{R}[S \Rightarrow \forall X <: A.P] \rangle \mapsto C} \\
\frac{(c\text{-ss:OPENTYPE})}{\Gamma \vdash A_1 <: A_0 \quad \langle \Gamma \vdash S \Rightarrow \mathcal{R}[P\{A_1/X\}] \rangle \mapsto C} \\
\frac{}{\langle \Gamma \vdash S \Rightarrow \mathcal{R}[\exists X <: A_0.P] \rangle \mapsto C} \\
\frac{(c\text{-ps:EXISTS TYPE})}{\langle \Gamma, X : \mathbf{type}, X <: A \vdash P \Rightarrow \mathcal{R}[Q] \rangle \mapsto C} \quad \frac{(c\text{-ps:EXISTSLOC})}{\langle \Gamma, l : \mathbf{loc} \vdash P \Rightarrow \mathcal{R}[Q] \rangle \mapsto C} \\
\frac{}{\langle \Gamma \vdash \exists X <: A.P \Rightarrow \mathcal{R}[\exists X <: A.Q] \rangle \mapsto C} \quad \frac{}{\langle \Gamma \vdash \exists l.P \Rightarrow \mathcal{R}[\exists l.Q] \rangle \mapsto C} \\
\frac{(c\text{-ps:FORALLTYPE})}{\langle \Gamma, X : \mathbf{type}, X <: A \vdash S \Rightarrow P \Rightarrow \mathcal{R}[S \Rightarrow Q] \rangle \mapsto C} \\
\frac{}{\langle \Gamma \vdash S \Rightarrow \forall X <: A.P \Rightarrow \mathcal{R}[S \Rightarrow \forall X <: A.Q] \rangle \mapsto C} \\
\frac{(c\text{-ps:FORALLLOC})}{\langle \Gamma, l : \mathbf{loc} \vdash S \Rightarrow P \Rightarrow \mathcal{R}[S \Rightarrow Q] \rangle \mapsto C} \quad \frac{(c\text{-ps:LOCAPP})}{\langle \Gamma \vdash S \Rightarrow P\{p/l\} \Rightarrow \mathcal{R}[S \Rightarrow Q] \rangle \mapsto C} \\
\frac{}{\langle \Gamma \vdash S \Rightarrow \forall l.P \Rightarrow \mathcal{R}[S \Rightarrow \forall l.Q] \rangle \mapsto C} \quad \frac{}{\langle \Gamma \vdash S \Rightarrow \forall l.P \Rightarrow \mathcal{R}[S \Rightarrow Q] \rangle \mapsto C} \\
\frac{(c\text{-ps:TYPEAPP})}{\Gamma \vdash A_1 <: A_0 \quad \langle \Gamma \vdash S \Rightarrow P\{A_1/X\} \Rightarrow \mathcal{R}[S \Rightarrow Q] \rangle \mapsto C} \\
\frac{}{\langle \Gamma \vdash S \Rightarrow \forall X <: A_0.P \Rightarrow \mathcal{R}[S \Rightarrow Q] \rangle \mapsto C}
\end{array}$$

$P\{A/X\} \triangleq$ “substitution, in P , of X for A ”

Note: bound type/location variables of a type must be fresh in that rule’s conclusion.

■ **Figure 7** Protocol composition abstraction extension.

and refined data abstractions of others (such as [15]), and focus instead on tpestates [19, 29, 30]. However, this is not an intrinsic limitation of our model. If we consider more precise states, we can (for instance) model monotonic counters from prior work [24, 11] where each counter shares state symmetrically. Our local protocols model these uses solely from the perspective of a single alias as:

$$\text{MC} \triangleq \exists \underbrace{\{j : \mathbf{int}\}}_J . (\mathbf{rw } p \underbrace{j}_J \Rightarrow \forall \underbrace{\{i : \mathbf{int} \mid i \geq j\}}_I . (\mathbf{rw } p \underbrace{i}_I ; \text{MC}))$$

The protocol models a monotonically increasing counter on location p . The step relies on location p initially containing some integer, j , and modifying the cell to store some other value, i , that is greater or equal than j . This interaction can be reduced to the core existential-universal protocol interaction discussed above (but, in our calculus, using less precise types: J and I) and where the protocol can be re-split indefinitely.

While our states are less precise, we can enforce more precise uses of that shared state. The semantics of prior work [24, 11] differed on whether the counter was forcefully used by clients, or whether the action was simply available to be used. We can model the two cases explicitly:

$$\exists p . ((!\mathbf{ref } p) :: \text{MC} \multimap [] :: \text{MC})$$

```

1 let newMVar = λ _.
2   let m = new Empty#{ } in
3     // "shares" the new cell using single MVar protocol
4     Γ ⊢ (rw / Empty#{ }) ⇒ MVar[l] || none
5     {
6       putMVar = λ val.
7         rec R.
8           lock m;
9           case !m of
10            Empty#x → m := Full#val;
11                  unlock m
12            | Full#value → m := Full#value;
13                       unlock m;
14                       R // retries
15          end
16        end,
17      splitMVar = λ _.
18        Γ ⊢ MVar[l] ⇒ MVar[l] || MVar[l]
19        {
20          takeMVar = λ _.
21            rec R.
22              lock m;
23              case !m of
24                Empty#x → m := Empty#x;
25                          unlock m;
26                          R // retries
27                | Full#value → m := Empty#{ };
28                              unlock m;
29                              value
30              end
31            end
32          }
33        end
34      end

```

■ Figure 8 MVar example.

Enables clients to use the counter an arbitrary number of times or simply thread it through, unused.

$$\forall X. \exists p. (!\text{ref } p) :: \exists J. (\text{rw } p \ J \Rightarrow \forall I. (\text{rw } p \ I ; X)) \multimap [] :: X)$$

By unfolding the protocol, the function guarantees that a single step of the protocol will be used. Since we (intentionally) abstract subsequent steps, the function cannot use the counter beyond that single use. Analogous reasoning can be used to enforce specific, finite, usages.

Adding support for dependent refinement types, and ensuring its decidability (even without interference), is beyond the scope of our work as we focus on the core composition problem. However, we believe that the underlying decidability insights made here will carry to a system with decidable dependent refinement types; even if perhaps requiring more fine-grained conditions to close the co-inductive proof of safe interference—that are only relevant once more precise typing is considered.

While we use a relatively simple calculus to keep the theory focused on the core of interference-control, we can for instance model MVars [23]. Fig. 8 shows an MVar, a structure that contains a single shared cell which is either empty or contains a value of some type. Notable operations include: `putMVar`, that waits until the cell is empty before inserting the given value; and `takeMVar` which waits until the cell is full to remove the cell’s value, leaving the cell empty. MVars can be shared by many aliases, each assigned a protocol such as:

$$\begin{aligned} \text{MVar}[m] \triangleq & \exists Y. ((\text{rw } m \ \text{Empty}\#Y \Rightarrow (\text{rw } m \ \text{Empty}\#Y) ; \text{MVar}[m]) \& \\ & ((\text{rw } m \ \text{Empty}\#Y \Rightarrow (\text{rw } m \ \text{Full}\#\text{int}) ; \text{MVar}[m])) \\ \oplus & (((\text{rw } m \ \text{Full}\#\text{int} \Rightarrow (\text{rw } m \ \text{Empty}\#[]) ; \text{MVar}[m]) \& \\ & \exists Y. ((\text{rw } m \ \text{Full}\#Y \Rightarrow (\text{rw } m \ \text{Full}\#Y) ; \text{MVar}[m]))) \end{aligned}$$

The T.R. [20] includes additional examples, including modeling examples of prior work with our more local protocol types. We can also model a shared pair where each alias keeps its own, local, precise knowledge on one of the two components of the pair stored in that shared state. The two aliases, L and R, share a common cell but keep part of that state private to itself. While both can do private actions over the shared cell, they are guaranteed to not interfere with the precise assumptions of the remaining alias.

$$\begin{aligned} P[A][B] & \triangleq \text{rw } p \ [A, B] \\ L[A] & \triangleq \exists X. (P[A][X] \Rightarrow \forall Y. (P[Y][X] ; L[Y])) & \Gamma \vdash P[X][Y] \Rightarrow L[X] \parallel R[Y] \\ R[A] & \triangleq \exists X. (P[X][A] \Rightarrow \forall Y. (P[X][Y] ; R[Y])) \end{aligned}$$

Thus, we can use the different perspectives of each protocol to model local knowledge that is hidden from other aliases, within our core protocol framework without needing additional mechanisms.

Since our types express sharing, we can use standard techniques to abstract the components of a protocol type after safe composition is checked. This enables an abstraction to expose a type interface that indirectly manipulates the shared state, such as indirectly locking/unlocking state (Fig. 9). We can type the record in such a way to hide the type in x but still expose some information on sharing that is useful for later enabling other tpestate functions [19]. For instance:

$$\exists A. \exists B. \exists C. [\dots, \text{lockMe} : [] :: (A \Rightarrow B; C) \multimap [] :: (A * (B; C)), \text{add} : (\text{int} :: A \multimap [] :: A), \dots]$$

Clients can only call `add` once the type A is available. This could model, for instance, a global lock on a collection to enable more coarse-grained control over the interference to that collection—but without exposing the lock to clients. Thus, when `lockMe` returns, the client receives a type that expresses that A is available and that a guarantee $(B; C)$ is expected to be fulfilled. However, this fulfillment can only occur indirectly via the wrapper record as clients do not have a direct way of accessing or mutating the internals of that shared state.

While we do not guarantee dead-lock freedom, it is possible to type more fine-grained locking schemes such as *hand-over-hand* locking (Fig. 10). Consider the protocol of a list’s node:

$$L[q] \triangleq \exists l. ((\text{rw } q \text{ !ref } l) * L[l] \Rightarrow (\text{rw } q \text{ !ref } l); \dots)$$

L is defined over a location q that contains the (abstracted) reference to the next element of the sequence of locations to be locked. Locking will enable access to that `ref l` which can then be locked to gain access to $L[l]$, the next element in the sequence of locations to lock. For brevity, we make each step simply consume the L protocol of the element in that sequence, instead of (for instance) re-splitting.

```
let x = new 0 in
// share 'x' as via some
// protocols
{
  lockMe = λ_.lock x,
  // ...
}
```

■ **Figure 9** Indirect locking.

```
lock @a; Γ = a : ref @a
let b = !a;
lock @b; // locks loc. of 'b'
unlock @a;
let c = !b;
lock @c;
unlock @b;
...
```

■ **Figure 10** Hand-over-hand locking example.

4 Composition Decidability & Other Technical Results

We now show decidability of protocol composition and discuss the remaining technical results of our language. The decidability statement comes as a direct consequence of ensuring a regular type structure via syntactic well-formedness constraints on recursive types. Although applied in the context of protocol composition, we follow ideas from prior work on ensuring decidable subtyping over bounded quantification [28, 4]. The main novelty is in extending this kind of reasoning to account for recursive types with parameters, in order to ensure a regular type structure over our more flexible recursive types. To achieve this, we apply well-formedness conditions which ensure that there is only a finite number of reachable (abstract) protocol states. We focus the discussion on decidability of protocol composition, and point interested readers to T.R. where these conditions are properly motivated and discussed. Crucially, these well-formedness conditions enable us to state the following:

► **Lemma 1** (Finite Uses). *Given a well-formed recursive type $(\text{rec } X(\bar{u}).A)[\bar{U}]$ the number of possible uses of X in A such that $\Gamma \vdash X[\bar{U}']$ type is bounded.*

► **Lemma 2** (Finite Unfolds). *Unfolding a well-formed recursive type $(\text{rec } X(\bar{u}).A)[\bar{U}]$ produces a finite set of variants of that original recursive type that (at most) contains: permutations of \bar{U} , or a set of mixtures of \bar{U} with some type/location variables representing a class of equivalent (\equiv) types.*

► **Lemma 3** (Finite Sub-Terms). *Given a well-formed type A , such that $\Gamma \vdash A$ **type**, the set of sub-terms of A is finite up to renaming of variables and weakening of Γ .*

4.1 Composition Properties, Algorithm, and Decidability

Informally, correctness of protocol composition is based on the two properties: 1) a split results in protocols that can always take a step with the current state of the shared resources, thus are never stuck; and, 2) protocol composition is a partial commutative monoid (associative, commutative, and with **none** as the identity element). Because of property 2), iterative splittings of existing protocols remain struck-free, unable to cause unsafe interference. We now state these properties formally but leave the proofs to the T.R.. The following two lemmas show stuck freedom by properties that resemble progress and preservation but over protocols:

► **Lemma 4.** *If $\Gamma \vdash R \Rightarrow P \parallel Q$ then $\langle \Gamma \vdash R \Rightarrow P \parallel Q \rangle \mapsto C$.*

Meaning that if two protocols, P and Q , compose safely then their configuration can take a step to another set of configurations, C .

► **Lemma 5.** *If $\langle \Gamma \vdash R \Rightarrow P \parallel Q \rangle \mapsto \langle \Gamma' \vdash R' \Rightarrow P' \parallel Q' \rangle \cdot C$ and $\Gamma \vdash R \Rightarrow P \parallel Q$ then $\Gamma' \vdash R' \Rightarrow P' \parallel Q'$.*

The lemma ensures that if two protocols compose safely, then any of the next configurations that result from stepping will also be safe.

Note that protocol composition does not enforce that the shared resources are not lost. Instead our concern is on safe interference. Indeed, resources that are never used will never be able to unsafely interfere. To avoid losing resources, we must forbid the use of (c-rs:None) on non-terminated protocols and that both P and Q cannot have both simultaneously terminated if there are non-**none** resources left. Once that restriction is considered, our splitting induces a monoid in the sense that for any P and Q for which $\Gamma \vdash R \Rightarrow P \parallel Q$ is defined there is a single such R (defined up to subtyping and equivalent protocol/state interference specification). Since for any two protocols there may not always exist an R that can be split into P and Q , this is a partial monoid.

► **Lemma 6.** *Protocol composition obeys the following properties:*

1. (identity) $\Gamma \vdash R \Rightarrow R \parallel \mathbf{none}$.
2. (commutativity) *If $\Gamma \vdash R \Rightarrow P_0 \parallel P_1$ then $\Gamma \vdash R \Rightarrow P_1 \parallel P_0$.*
3. (associativity) *If we have $\Gamma \vdash R \Rightarrow P_0 \parallel P$ and $\Gamma \vdash P \Rightarrow P_1 \parallel P_2$ then exists Q such that $\Gamma \vdash R \Rightarrow Q \parallel P_2$ and $\Gamma \vdash Q \Rightarrow P_0 \parallel P_1$.*
(i.e. If $\Gamma \vdash R \Rightarrow P_0 \parallel (P_1 \parallel P_2)$ then $\Gamma \vdash R \Rightarrow (P_0 \parallel P_1) \parallel P_2$)

Protocol composition is defined as a “split”, left-to-right (\Rightarrow). Simply reading the rules as right-to-left (\Leftarrow) to compute a “merge” is not safe. For instance, it would enable merging to arbitrary choices with (c-rs:StateIntersection). Intuitively, merging needs to intertwine the uses of both protocols. However, since we do not track copies (as we target sharing when that tracking is not possible), merging cannot “collapse” a protocol into a non-protocol type. In this case “merging” is equivalent to simply having the two non-merged protocols available in Δ or bundled using the $*$ type.

The composition algorithm is shown in the T.R. and is a straightforward implementation of the axiomatic definitions shown above. The algorithm uses a set of visited configurations to remember past configurations and ensure that once all different protocol configurations are exhausted (up to renaming and weakening of Γ), the algorithm can terminate. We now state our technical lemmas on the composition algorithm but leave the proofs to the T.R..

► **Lemma 7.** *Given well-formed types and environment, we have that:*

1. (soundness) if $c(\Gamma, R, P, Q)$ then $\Gamma \vdash R \Rightarrow P \parallel Q$.
2. (completeness) if $\Gamma \vdash R \Rightarrow P \parallel Q$ then $c(\Gamma, R, P, Q)$.
3. (decidability) $c(\Gamma, R, P, Q)$ terminates.

4.2 Correctness Properties

The main safety theorems, progress and preservation, that are defined over valid program configurations such that:

$$\frac{\Gamma \mid \Delta_i \vdash e_i : ![] \vdash \cdot \quad i \in \{0, \dots, n\} \quad n \geq 0}{\Gamma \mid \Delta_0, \dots, \Delta_n \vdash e_0 \cdot \dots \cdot e_n} \text{ (WF:PROGRAM)}$$

Stating that a thread pool $(e_0 \cdot \dots \cdot e_n)$ is well formed if each thread can be assigned a “piece” of the linear typing environment (containing resources), and if each individual expression has type $![]$ without leaving any residual resources (\cdot) . Note that the conditions on each thread (e_i) are identical to those imposed by $(\tau:\text{FORK})$. For clarity, both safety theorems are supported by auxiliary theorems over a single expression, besides the main theorem over the complete thread pool.

We now state progress over programs:

► **Theorem 8.** *If $\Gamma \mid \Delta \vdash T_0$ and $\text{live}(T_0)$ and if exists H_0 such that $\Gamma \mid \Delta \vdash H_0$ then $H_0 ; T_0 \mapsto H_1 ; T_1$.*

$\text{live}(T)$ means that the thread pool T contains at least one “live” thread such that the thread is neither a value nor is waiting for a lock to be released (which includes deadlocks). $\Gamma \mid \Delta \vdash H$ ensures that the *Heap* is well-defined according to Γ and Δ .

We define $\text{Wait}(H, e)$ over a thread e and heap H such that the *Evaluation* context is reduced to evaluating the configuration: $H ; \mathcal{E}[\text{lock } \rho, \bar{\rho}] \cdot T$ where $\rho \hookrightarrow v \notin H$ which contains at least one location (ρ) that is currently locked or was deleted and, therefore, the thread must block waiting (potentially indefinitely) for that lock to be available before continuing. “Early” deletion of shared resources results in a pending guarantee. Since well-formed threads cannot leave residual resources, this situation is ruled out for correct programs, but may occur on the theorem below.

Progress over expressions is defined as follows:

► **Theorem 9.** *If $\Gamma \mid \Delta_0 \vdash e_0 : A \vdash \Delta_1$ then we have that either:*

- e_0 is a value, or;
- if exists H_0 and Δ such that $\Gamma \mid \Delta, \Delta_0 \vdash H_0$ then either:
 - (steps) $H_0 ; e_0 \mapsto H_1 ; e_1 \cdot T$, or;
 - (waits) $\text{Wait}(H_0, e_0)$.

Preservation ensures that a reduction step will preserve both the type and the effects of the expression that is being reduced (so that each thread’s type, $![]$, and effect, \cdot , remains unchanged). As above, we use a preservation theorem over programs that makes use of an auxiliary theorem on preservation over expressions:

► **Theorem 10.** *If we have $\Gamma_0 \mid \Delta_0 \vdash H_0$ and $\Gamma_0 \mid \Delta_0 \vdash T_0$ and $H_0 ; T_0 \mapsto H_1 ; T_1$ then, for some Δ_1 and Γ_1 , we have: $\Gamma_0, \Gamma_1 \mid \Delta_1 \vdash H_1$ and $\Gamma_0, \Gamma_1 \mid \Delta_1 \vdash T_1$.*

So that a well-formed pool of threads (T_0) remains well-formed after stepping one of these threads (resulting in T_1). Preservation over a single expression must still account for the resources (Δ_T) that may be consumed by a newly spawned thread (T) :

<pre> receive(c) ≜ rec R. lock c; case !c of // 1. waiting states (A..Z) A#n → ... // analogous to case below Z#n → // restore linear content c := Z#n; unlock c; R // retry // 2. desired (receive) state ReadyToReceive#v → c := Idle#{ }; // "received" unlock c; v // value received from "channel" end end </pre>	<pre> send(c,v) ≜ rec R. lock c; case !c of // 1. waiting states (A..Z) A#n → ... // analogous to case below Z#n → // restore linear content c := Z#n; unlock c; R // retry // 2. desired (idle) state Idle#_ → c := ReadyToReceive#v; // "sent" unlock c; {} // result of send is empty end end </pre>
--	---

■ **Figure 11** Simple encoding of send and receive functions via a shared cell.

► **Theorem 11.** *If we have $H_0 ; e_0 \mapsto H_1 ; e_1 \cdot T$ and $\Gamma_0 \mid \Delta_0, \Delta_T, \Delta_2 \vdash H_0$ and $\Gamma_0 \mid \Delta_0, \Delta_T \vdash e_0 : A \dashv \Delta$ then, for some Δ_1 and Γ_1 , we have: $\Gamma_0, \Gamma_1 \mid \Delta_1, \Delta_T, \Delta_2 \vdash H_1$ and $\Gamma_0, \Gamma_1 \mid \Delta_1 \vdash e_1 : A \dashv \Delta$ and $\Gamma_0, \Gamma_1 \mid \Delta_T \vdash T$.*

We complement our main results with the following “Error Freedom” corollary to show that our system cannot type programs that allow data races and the dereference of destroyed memory cells, i.e. that our system ensures memory safety and race freedom.

► **Corollary 12.** *The following program states cannot be typed:*

1. (Data Race) *Simultaneous read/modify by two thread over the same location (we also ensure read-exclusive accesses):*

$$H; \mathcal{E}_0[\rho := v] \cdot \mathcal{E}_1[!\rho] \cdot T \quad H; \mathcal{E}_0[\rho := v] \cdot \mathcal{E}_1[\rho := v'] \cdot T$$

2. (Memory Fault) *Accessing a non-existing/deleted location:*

$$H; \mathcal{E}[\rho := v] \cdot T \quad H; \mathcal{E}[!\rho] \cdot T \quad (\text{where } \rho \notin H)$$

3. (Ownership Fault) *Attempt to delete a non-existing location:*

$$H; \mathcal{E}[\text{delete } \rho] \cdot T \quad (\text{where } \rho \notin H)$$

The proof is straightforward due to our use of locks to ensure mutual exclusion and the fact that our protocols discipline the use of shared state. Thus, these errors are ruled out by either protocol composition or by the resource tracking of the core linear system.

5 Protocol Expressiveness

We show the expressiveness of our protocols by modeling typeful message-passing concurrency, using a straightforward encoding of message-passing via shared memory interference (Fig. 11). The encoding itself should be unsurprising as it follows well-known ideas from the literature, so we defer less important details to the T.R. to focus instead on the most interesting aspect of this example: how our protocol framework is able to type such uses and ensure their safety.

We encode a more primitive, “low-level” view of typeful message-passing concurrency via the causality of shared memory interference. We focus on the non-distributed setting where a channel can be precisely encoded as a low-level shared cell. Channel communication and its changing session properties are emulated indirectly via inspection of or interference over the contents of that shared cell. Thus, our functions to send/receive a value simply hides the underlying Waiting states that may be needed when the cell is not available. A receiving step can be modeled by a protocol of the form:

$$\text{Wait}[A..Z] \oplus (\text{rw } c \text{ ReadyToReceive}\#V \Rightarrow \text{rw } c \text{ Idle}\#[\] ; \text{NextStep})$$

where Wait is a sequence of retry steps that leave the state unmodified, until a value of type V is “received”. Sending uses a similar protocol but where we must wait for an Idle cell before “sending”.

The T.R. includes the complete “Buyer-Seller-Shipper” example (the canonical and simple example used in session-based concurrency works) while in here we only take a look at the main aspects of the Buyer’s interaction with the channel (Fig. 12).

We model a channel using a capability to location c . For brevity, we omit “ $\text{rw } c$ ” from “ $\text{rw } c A$ ” since all changes occur over that same location. The Buyer’s type uses standard π -calculus [21] notations where $!$ sends and $?$ receives a value. In our protocols, these actions are mapped to the rely type (receive) and the guarantee type (send).

```
let c = connectSeller() in
  c : buy!(prod) ; price?(p) ; details?(d)
  send(c, GET_USER_PRODUCT() );
  let price = receive(c) in
    let details = receive(c) in
      close(c)
    end
  end
end
```

Figure 12 Buyer code.

$$\underbrace{\text{buy}!(\text{prod})}_{\text{idle}\#0\#[\]} \Rightarrow \text{buy}\#\text{prod} \quad ; \quad \underbrace{\text{price}?(p)}_{\text{price}\#p} \Rightarrow \text{idle}\#2\#[\] \quad ; \quad \underbrace{\text{details}?(d)}_{\text{details}\#d} \Rightarrow [\]$$

Buyer starts by sending a request to buy some *product*, then waits for the *price*, and finally receives the *details* of that product. Under that interaction protocol, we simply map sends to a guarantee type of a step, and receives the a rely type of a step.

Our protocol interactions are both non-deterministic and may contain an arbitrary number of simultaneous participants. To ensure that the desired participant (Buyer) is the only one allowed to received (take) the price, we must mark the contents with a specific tag so that only Buyer has permission to change that state. To handle the non-deterministic interleaving of protocols, we must introduce explicit “wait states” that allow a participant to check if the communication has reached the desired point to that participant or if it should continue waiting. We abstract these steps as Wait as they simply recur on that same step of the protocol (i.e. “busy-wait”).

$$\text{idle}\#0\#[\] \Rightarrow \text{buy}\#\text{prod} ; \text{Wait} \oplus (\text{price}\#p \Rightarrow \text{idle}\#2\#[\]) ; \text{Wait} \oplus (\text{details}\#d \Rightarrow [\])$$

The richness of our shared state interactions means that we can immediately support fairly complex session-based mechanisms (such as delegation, asynchronous communication, “messages to self”, multiparty interactions, internal/external choices, etc.) within our small protocol framework. However, this flexibility comes at the cost of requiring a more complex composition mechanism. Protocol composition accounts for both non-deterministic protocol interleaving and “multi-way” communication, features which are usually absent from strictly choreographed session-based concurrency (favoring instead strong liveness properties over more deterministic, linear compositions).

Naturally, more complex examples are possible. In here our focus is on showing the core insights that enable us to relate the two techniques: 1) mapping receive/send to our rely/guarantee types; 2) adding explicit waiting states to account for non-deterministic protocol interleaving; and 3) tag the content of a cell in order to ensure that only the right participant will be able to mutate the state at that point in the interaction. (Recall that we do not guarantee deadlock freedom, nor termination.)

6 Related Work

Our work relates to prior work on *rely-guarantee protocols* [18]. We show that these protocols are useful to reason about concurrency and significantly improve the flexibility of protocol composition. Namely, we allow the composition of abstract protocols (enabling more local typing as in Section 3), show that our composition is decidable, and provide a novel axiomatic definition of composition that is straightforward to implement. Since thread-based interference is rooted in alias-related interference, the technique itself is mostly indifferent to whether sharing occurs in the sequential or concurrent setting. Still, we address all technicalities that make concurrency possible (such as adding support to arbitrarily many threads, locking locations, well-formedness changes to the rely/guarantee types to ensure matching lock/unlock of locations, etc). Furthermore, by considering the concurrent setting we are able to express and relate rely-guarantee protocols to typeful message-passing concurrency.

Our work is also related to recent work on more precise tracking of interference. *Chalice* [17] uses a simplified form of rely-guarantee to reason about shared state interference by constraining a thread’s changes to a two-state invariant, relating the previous and current states. *Monotonic* [8, 24] uses of shared state (where all changes converge to more precise states) are less dependent on aliasing information, which simplifies checking at the expense of expressiveness. Dynamic ownership recovery mechanisms [33, 25] choose some run-time overhead and dynamic safety guarantees to enable more flexible ownership recovery than purely static approaches. *Rely-guarantee references* [11] adapt the use of rely-guarantee to individual reference cells with support for dependent refinement types in a sequential language. Although the use of refinements adds expressiveness to the description of sharing, they do not support ownership recovery, nor address decidability, and typechecking can require manual assistance in Coq. *Access permissions* [33, 3, 2] control alias interference by categorizing read-write uses into different permission kinds. Our design omits the read-write distinction to focus exclusively on structuring alias interference using more fundamental protocol primitives. Interestingly, although we only model write-exclusive uses, our types can enforce effectively read-exclusive semantics by ensuring that any private change in a cell will be reverted to its original public value. However, this simpler form of read-only cannot capture their multiple, simultaneous readers semantics. Still, by modeling interference in a more fundamental way, we gain additional expressiveness beyond their most permissive share permission as we can model uses beyond invariant-based sharing. In [5] Crafa and Pavodani introduce a high-level (actor-like) model for sharing (type)state via join patterns. We target a more low-level programming paradigm (which builds typestates through type abstraction rather than as a first-class language feature), enabling us to introduce abstraction at the level of protocols and support protocol re-splitting in ways that are not expressible in their work.

Several recent works use *partial commutative monoids* [7, 16, 6] to model sharing by leveraging the concept of fictional separation [7, 12]. Commutative monoids offer the underlying general principle for splitting resources, enabling seemingly unrelated components to interact via aliasing under a layer of (fictional) separation. We compare more closely to [16] due to our common use of L^3 [1] and type-based approach. In [16], Krishnaswami *et al.* define a generic sharing rule based on programmer-supplied commutative monoids for safe sharing of state in a single-threaded environment. Their work does not approach the issue of decidability of resource splitting, and requires wrapping access to shared state in an module abstraction that serves as an intermediary to access shared state. Our work focuses on a custom commutative monoid that enables first-class sharing without (necessarily) needing a wrapping module abstraction. Although our protocol splitting is a specialized monoid, we showed that this mechanism is relatively flexible, decidable, and give an algorithmic implementation. Other technical differences between our works abound such as their use of affine refinement types (enabling more fine-grained types), our use of multi-threaded semantics and allowing inconsistent states (i.e. locked cells) to be moved around as first-class, etc.

Protocol-based mechanisms for safe interference are also used by other approaches, such as in program logic-based systems (e.g. [14, 31, 32, 22, 9]). By generally targeting manual proofs (and somewhat more involved specifications) these works generally fit into a different design space than ours, although share some interesting similarities. While we make concessions on expressiveness to achieve decidable protocol composition and re-splitting, these works focus instead on the expressiveness of their concurrency specification. LRG [9] supports lock-free structures but requires a special frame-rule to support framing over rely-guarantee conditions. We simply integrate protocols into the language (as linear resources) meaning that the standard frame-rule suffices. Supporting lock-free concurrency in our system would require reinterpreting a \Rightarrow step as a single-cell, atomic, conditional operation; with the shared resource (stored in the cell) being immediately extracted/inserted from/into the cell, rather than just accessible after locking. CaReSL [32] and Iris [14] support “islands”/regions of memory that are shared together and whose imprecise state must be considered when using. Our composition rules enforce that a protocol carries all information on imprecise states, which is then deconstructed via (T:ALTERNATIVE-LEFT) and case analysis. Our protocols can group shared state using the $*$ operator to define shallow “regions”, while their works allow for far more rich specifications of atomic regions of any depth. Iris [14] further supports a form of re-splitting via a “view shifting” mechanism, to repartition (or create) shared regions. FCSL [22] encodes protocols via auxiliary/ghost state. Although done in a compositional way, it can require checking for safe interference (“stability”) after a split since a safe split does not necessarily imply safe interference in all situations. Our protocol composition mechanism is essentially a form of checking for safe interference early, at the moment a protocol is split, by checking that all possible future uses are safe resembling a form of “pre-computed” stability check.

Protocol composition itself can also be seen as a form of model checking (to check that each state has a successor) that uses abstract states to ensure a finite state space, but in a system that is more intimately integrated with the language. Our protocols are first-class resources that can be specialized by clients, even abstracting (leaving out) later steps. Thus, our protocols guide the programmer on how to reason locally about (safe) interference by mapping its uses of locks to a local protocol type that models the alias perspective on the shared state. While our work focuses on modeling the core interference phenomenon within a small calculus, rather than precisely typing existing programs, we still showed that extensions may be used to model at least some existing programs within our model.

7 Conclusions

We defined a flexible and decidable procedure that ensures the safe composition of interfering abstract protocols that share access to mutable state. While employing a relatively small protocol framework, we are able to model the core interference principles of complex shared state interactions within our core calculus. Finally, we showed the expressiveness of our protocol framework by discussing how it can capture in a unified way both shared memory interference and typeful message-passing concurrency.

Acknowledgments. This material is supported in part by AFRL and DARPA under agreement #FA8750-16-2-0042, and by NSA label contract #H98230-14-C-0140; by NOVA LINCS UID/CEC/04516/2013; and by Fundação para a Ciência e Tecnologia (Portuguese Foundation for Science and Technology) through the Carnegie Mellon Portugal Program under grant SFRH / BD / 33765 / 2009 and the Information and Communication Technology Institute at CMU. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of any of the sponsors. We thank the anonymous reviewers for their helpful comments.

References

- 1 A. Ahmed, M. Fluet, and G. Morrisett. L3: A linear language with locations. *Fundam. Inf.*, 77(4):397–449, Dec. 2007.
- 2 N. E. Beckman, K. Bierhoff, and J. Aldrich. Verifying correct usage of atomic blocks and typestate. In *OOPSLA 2008*.
- 3 K. Bierhoff and J. Aldrich. Modular typestate checking of aliased objects. In *OOPSLA 2007*.
- 4 G. Castagna and B. C. Pierce. Decidable bounded quantification. In *POPL 1994*.
- 5 S. Crafa and L. Padovani. The chemical approach to typestate-oriented programming. In *OOPSLA 2015*.
- 6 T. Dinsdale-Young, L. Birkedal, P. Gardner, M. Parkinson, and H. Yang. Views: compositional reasoning for concurrent programs. In *POPL 2013*.
- 7 T. Dinsdale-Young, M. Dodds, P. Gardner, M. J. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In *ECOOP 2010*.
- 8 M. Fähndrich and K. R. M. Leino. Heap monotonic typestate. In *IWACO 2003*.
- 9 X. Feng. Local rely-guarantee reasoning. In *POPL '09*.
- 10 J.-Y. Girard. Linear logic. *Theor. Comput. Sci.*, 50:1–102, 1987.
- 11 C. S. Gordon, M. D. Ernst, and D. Grossman. Rely-guarantee references for refinement types over aliased mutable data. In *PLDI 2013*.
- 12 J. B. Jensen and L. Birkedal. Fictional separation logic. In *ESOP 2012*.
- 13 C. B. Jones. Tentative steps toward a development method for interfering programs. *ACM Trans. Program. Lang. Syst.* 1983.
- 14 R. Jung, D. Swasey, F. Sieczkowski, K. Svendsen, A. Turon, L. Birkedal, and D. Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *POPL 2015*.
- 15 N. R. Krishnaswami, P. Pradic, and N. Benton. Integrating linear and dependent types. In *POPL 2015*.
- 16 N. R. Krishnaswami, A. Turon, D. Dreyer, and D. Garg. Superficially substructural types. In *ICFP 2012*.
- 17 K. R. Leino and P. Müller. A basis for verifying multi-threaded programs. In *ESOP 2009*.
- 18 F. Militão, J. Aldrich, and L. Caires. Rely-guarantee protocols. In *ECOOP 2014*.
- 19 F. Militão, J. Aldrich, and L. Caires. Substructural typestates. In *PLPV 2014*.
- 20 F. Militão, J. Aldrich, and L. Caires. Composing interfering abstract protocols (technical report). CMU-CS-16-103, 2016.
- 21 R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, i. *Inf. Comput.*, Sept. 1992.
- 22 A. Nanevski, R. Ley-Wild, I. Sergey, and G. A. Delbianco. Communicating state transition systems for fine-grained concurrent resources. In *ESOP 2014*.
- 23 S. Peyton Jones, A. Gordon, and S. Finne. Concurrent haskell. In *POPL 1996*.
- 24 A. Pilkiewicz and F. Pottier. The essence of monotonic state. In *TLDI 2011*.
- 25 F. Pottier and J. Protzenko. Programming with permissions in mezzo. In *ICFP 2013*.
- 26 J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. Logic in Computer Science*, pages 55–74, 2002.
- 27 A. Sabry and M. Felleisen. Reasoning about programs in continuation-passing style. In *Proc. LISP and Functional Programming*, 1992.
- 28 J. Seco and L. Caires. Subtyping first-class polymorphic components. In *ESOP 2005*.
- 29 R. E. Strom. Mechanisms for compile-time enforcement of security. In *POPL 1983*.
- 30 R. E. Strom and S. Yemini. Typestate: A programming language concept for enhancing software reliability. *IEEE Trans. Software Eng.*, 12(1):157–171, 1986.
- 31 K. Svendsen and L. Birkedal. Impredicative concurrent abstract predicates. In *ESOP 2014*.
- 32 A. Turon, D. Dreyer, and L. Birkedal. Unifying refinement and hoare-style reasoning in a logic for higher-order concurrency. In *ICFP '13*.
- 33 R. Wolff, R. Garcia, E. Tanter, and J. Aldrich. Gradual typestate. In *ECOOP 2011*.