Machine-Checked Semantic Session Typing

Jonas Kastberg Hinrichsen
IT University of Copenhagen
Denmark

Daniël Louwink
University of Amsterdam
The Netherlands

Robbert Krebbers
Radboud University and Delft University of Technology
The Netherlands

Jesper Bengtson
IT University of Copenhagen
Denmark

Abstract

Session types—a family of type systems for message-passing concurrency—have been subject to many extensions, where each extension comes with a separate proof of type safety. These extensions cannot be readily combined, and their proofs of type safety are generally not machine checked, making their correctness less trustworthy. We overcome these shortcomings with a semantic approach to binary asynchronous affine session types, by developing a logical relations model in Coq using the Iris program logic. We demonstrate the power of our approach by combining various forms of polymorphism and recursion, asynchronous subtyping, references, and locks/mutexes. As an additional benefit of the semantic approach, we demonstrate how to manually prove typing judgements of racy, but safe, programs that cannot be type checked using only the rules of the type system.

CCS Concepts: • Theory of computation → Separation logic; Program verification; Programming logic.

Keywords: Message passing, concurrency, session types, separation logic, semantic typing, Iris, Coq

ACM Reference Format:

Jonas Kastberg Hinrichsen, Daniël Louwink, Robbert Krebbers, and Jesper Bengtson. 2021. Machine-Checked Semantic Session Typing. In *Proceedings of the 10th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP '21), January 18–19, 2021, Virtual, Denmark*. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3437992.3439914>

1 Introduction

Session types [26] guarantee that message-passing programs comply with a protocol (*session fidelity*), and do not crash (*type safety*). While session types are an active research area,

we believe the following challenges have not received the attention that they deserve:

1. There are many extensions of session types with *e.g.*, polymorphism [18], asynchronous subtyping [37], and sharing via locks [5]. While type safety has been proven for each extension in isolation, existing proofs cannot be readily composed with each other, nor with other substructural type systems like Affe, Alms, Linear Haskell, Plaid, Rust, Mezzo, Quill, or System F°.
2. Session types use substructural types to enforce a strict discipline of channel ownership. While conventional session-type systems can type check many functions, they inherently exclude some functions that do not obey the ownership discipline, even if they are safe.
3. Only few session-type systems and their safety proofs have been machine checked by a proof assistant, making their correctness less trustworthy.

We address these challenges by eschewing the traditional *syntactic approach* to type safety (using *progress and preservation*) and instead embrace the *semantic approach* to type safety [1–3], using *logical relations* defined in terms of a program logic [4, 14, 15].

The semantic approach addresses the challenges above as (1) typing judgements are definitions in the program logic, and typing rules are lemmas in the program logic (they are not inductively defined), which means that extending the system with new typing rules boils down to proving the corresponding typing lemmas correct; (2) safe functions that cannot be conventionally type checked can still be semantically type checked by manually proving a typing lemma (3) all of our results have been mechanised in Coq using the Iris framework for concurrent separation logic [29–34] giving us a high degree of trust that they are correct.

The syntactic approach to type safety requires global proofs of progress (well-typed programs are either values or can take a step) and preservation (steps taken by the program do not change types), culminating in type safety (well-typed programs do not get stuck). One key selling point of the semantic approach to type safety is that it does not require progress and preservation proofs, thereby allowing snippets of safe code to be type checked without requiring well-typed terms mid execution. Safety proofs are deferred to the program logic, whose adequacy/soundness theorem states that

This work is licensed under a Creative Commons Attribution International 4.0 License.

CPP '21, January 18–19, 2021, Virtual, Denmark
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8299-1/21/01.
<https://doi.org/10.1145/3437992.3439914>

proving a program correct for any postcondition implies that the code will never get stuck.

A concrete example of a racy program that can be semantically, but not conventionally, type checked is:

$$\lambda c. (\text{recv } c \parallel \text{recv } c) : \text{chan } (?Z. ?Z. \text{end}) \multimap (Z \times Z)$$

Two values are requested over channel c in parallel, and returned as a tuple (using the operator $\parallel$ for parallel composition, and the type $\text{chan } (?Z. ?Z. \text{end})$ for a channel that expects to receive two integers). This program cannot be type checked using conventional session-type systems as channels are substructural/ownership types and cannot be owned by multiple threads at the same time. Nevertheless, this program is safe¹—the order in which the values are received is irrelevant, as they have the same type.

The fact that this program cannot be type checked is not a shortcoming of conventional session-type systems. Since the correctness relies on a subtle argument (the `recv` is executed *exactly* twice in parallel), it is unreasonable to expect having syntactical typing rules that account for it. However, using the semantic approach, we can prove the corresponding typing lemma using the full power of the program logic.

An important prerequisite for proving typing lemmas such as the above is to use an expressive program logic. The Iris concurrent separation logic [29–31, 33] has proved to be sufficiently expressive to define semantic type systems for e.g., Rust [27, 28] and Scala [22], due to its state-of-the-art built-in support for e.g., resource ownership, recursion, polymorphism, and concurrency. In addition, we make use of the Actris framework for message passing in Iris [23, 24]. Actris includes the notion of *dependent separation protocols*, which are like session types in structure, but were developed to prove functional correctness of message-passing programs.

An additional advantage of Iris (and Actris) is that they come with an existing mechanisation in Coq. This mechanisation not only includes an adequacy/soundness theorem, but also tactical support for separation logic proofs [32, 34].

Contributions and Outline. This paper presents an extensive machine-checked and semantic account of binary (two-party) asynchronous (sends are non-blocking) affine (resources may be discarded) session types. It makes the following contributions:

- We define a semantic session-type system as a logical relation in Iris using Actris’s notion of dependent separation protocols (§2). As an additional conceptual contribution, this construction provides a concise connection between session types and separation logic.
- We demonstrate the extensibility of our approach by adding subtyping for term and session types, copyable types, equi-recursive term and session types, polymorphic term and session types, and mutexes (§3).

¹For simplicity, we assume `recv` to be atomic, or a lock is needed. Even with a lock, conventional session-type systems cannot handle this program.

- We demonstrate the benefit of our type system being semantic by integrating the manual verification of safe but not conventionally type-checkable programs (§4).
- We provide insight on the benefits of a semantic type system in regards to mechanisation efforts (§5). All of our results are mechanised in the Coq proof assistant and can be found in [25].

2 A Tour of Semantic Session Typing

We show how to build a semantic session-type system using logical relations on top of an untyped concurrent language with message passing (§2.1). We provide a brief overview of Iris (§2.2), and then present a lightweight affine type system (§2.3) as the core upon which we built our session-type system (§2.4). Our affine type system is inspired by RustBelt [27, 28], but drops Rust-specific features like borrowing and lifetimes to focus on session types.

2.1 Language

We use an untyped higher-order functional programming language with concurrency, mutable references, and binary asynchronous message passing, whose syntax is:

$$\begin{aligned} v \in \text{Val} &::= () \mid b \mid i \mid \ell \mid c \mid \text{rec } f \ x = e \mid \dots \\ e \in \text{Expr} &::= v \mid x \mid \text{rec } f \ x = e \mid e_1 \ e_2 \mid e_1 \parallel e_2 \mid \\ &\quad \text{ref } e \mid e_1 \leftarrow e_2 \mid !e \mid \\ &\quad \text{new_chan } () \mid \text{send } e_1 \ e_2 \mid \text{recv } e \mid \dots \end{aligned}$$

We let $b \in \mathbb{B}$, $i \in \mathbb{Z}$, $\ell \in \text{Loc}$, and $c \in \text{Chan}$, where Loc and Chan are countably infinite sets of identifiers. We omit the standard operations on pairs, sums, *etc.* We write $\lambda x. e$ for `rec  $\_$   $x = e$` , and `let  $x = e_1$  in  $e_2$` for $(\lambda x. e_2) \ e_1$, and $e_1; e_2$ for `let  $\_ = e_1$  in  $e_2$` . Message passing is given an asynchronous semantics: `new_chan` $()$ returns a pair (c_1, c_2) of channel endpoints that operate on buffers $(\vec{v}_1, \vec{v}_2)$ that are initially empty, `send` $c_i \ w$ enqueues message w in $\vec{v}_i$, while `recv` c_i blocks until a message w is available in $\vec{v}_i$ (if $i=2$ then 1 else 2), and then dequeues and returns w . Mutable references ℓ are allocated with `ref` e , updated using $e_1 \leftarrow e_2$, and dereferenced with `!e`. Parallel composition $e_1 \parallel e_2$ executes e_1 and e_2 in parallel and returns the results as a tuple, once they have terminated. The language also supports fork and compare-and-set.

2.2 Semantic Typing in Iris

The idea of semantic typing is to represent types as *logical relations*, which are predicates that describe the values that inhabit the type. To model type systems with features like references or session types, these predicates need to range over program states. To avoid threading through program states explicitly, we do not use ordinary set-theoretic predicates, but use predicates in a program logic, and use the connectives of the program logic to give concise definitions of types. The program logic that we use is Iris, whose propositions

Term types:	Judgements:
$\text{Type}_\star \triangleq \text{Val} \rightarrow \text{iProp}$ $\text{any} \triangleq \lambda w. \text{True}$ $1 \triangleq \lambda w. w \in \{()\}$ $\mathbb{B} \triangleq \lambda w. w \in \mathbb{B}$ $\mathbb{Z} \triangleq \lambda w. w \in \mathbb{Z}$ $\text{ref}_{\text{uniq}} A \triangleq \lambda w. \exists v. w \in \text{Loc} * (w \mapsto v) * \triangleright(A v)$ $A_1 \times A_2 \triangleq \lambda w. \exists v_1, v_2. w = (v_1, v_2) * \triangleright(A_1 v_1) * \triangleright(A_2 v_2)$ $A_1 + A_2 \triangleq \lambda w. \exists v. (w = \text{inl } v * \triangleright(A_1 v)) \vee (w = \text{inr } v * \triangleright(A_2 v))$ $A \multimap B \triangleq \lambda w. \forall v. \triangleright(A v) \multimap \text{wp } e \{w v\} \{B\}$ $\text{chan } S \triangleq \lambda w. w \multimap S$	$\Gamma \models \sigma \triangleq *_{(x,A) \in \Gamma}. A(\sigma(x))$ $\Gamma \models e : A \ni \Gamma' \triangleq \forall \sigma. (\Gamma \models \sigma) \multimap \text{wp } e[\sigma] \{v. A v * (\Gamma' \models \sigma)\}$
Session types:	
$\text{Type}_\diamond \triangleq \text{iProto}$ $\text{end} \triangleq \text{end}$ $!A. S \triangleq ! (v : \text{Val}) \langle v \rangle \{A v\}. S$ $?A. S \triangleq ? (v : \text{Val}) \langle v \rangle \{A v\}. S$ $\oplus \{\vec{S}\} \triangleq ! (l : \mathbb{Z}) \langle l \rangle \{l \in \text{dom}(\vec{S})\}. \vec{S}(l)$ $\&\{\vec{S}\} \triangleq ? (l : \mathbb{Z}) \langle l \rangle \{l \in \text{dom}(\vec{S})\}. \vec{S}(l)$	

Figure 1. Typing judgements and type formers of the semantic type system.

$P, Q \in \text{iProp}$ implicitly range over an extensible notion of resources, which includes the program state.

Iris is a higher-order separation logic, so it has the usual logical connectives such as conjunction ($P \wedge Q$), implication ($P \Rightarrow Q$), universal ($\forall x : \tau. P$) and existential ($\exists x : \tau. P$) quantification, as well as the connectives of separation logic:

- The *points-to connective* ($\ell \mapsto v$) asserts exclusive resource ownership of a heap location $\ell \in \text{Loc}$, stating that it holds the value $v \in \text{Val}$.
- The *separating conjunction* ($P * Q$) states that P and Q holds for disjoint sets of resources.
- The *separating implication* ($P \multimap Q$) states that by giving up ownership of the resources described by P , we obtain ownership of the resources described by Q . Separating implication is used similarly to implication since (P entails $Q \multimap R$) iff ($P * Q$ entails R).
- The *weakest precondition* ($\text{wp } e \{\Phi\}$) states that given a postcondition $\Phi : \text{Val} \rightarrow \text{iProp}$, the expression e is safe, and, if e reduces to a value v , then Φv holds. We write $\text{wp } e \{w. Q\}$ for $\text{wp } e \{\lambda w. Q\}$.

As we see in §2.3 these connectives match up with the type formers for unique references ($\ell \mapsto v$), products ($P * Q$), and affine functions ($P \multimap Q$ and $\text{wp } e \{\Phi\}$).

Iris's notion of resources is not limited to heap locations, but can be extended with custom resources. This feature is used by Actris to extend Iris with support for reasoning about functional correctness of message-passing programs (§2.4) by means of the connective ($c \multimap -$) that asserts exclusive resource ownership of the channel c . Moreover, Iris has an extensible mechanism of *ghost* resources, which we use in this paper to semantically type safe yet not conventionally type-checkable programs (§4).

To define recursive types semantically (§3.3), Iris provides the *later modality* ($\triangleright P$) and the *guarded fixpoint operator* ($\mu x : \tau. t$), which enable guarded recursive definitions of Iris propositions and terms. The guarded fixpoint operator requires all recursive occurrences of the variable x to occur *guarded* in t , where an occurrence is guarded if it appears below a $\triangleright$ modality. This ensures that t is *contractive* in the

variable x , which guarantees that a unique fixpoint exists. Guarded fixpoints can be folded and unfolded using the equality $\mu (x : \tau). t = t[(\mu (x : \tau). t)/x]$.

The proposition $\triangleright P$ is strictly weaker than P , since P entails $\triangleright P$, while the reverse does not hold. The $\triangleright$ modality can be eliminated by taking a program step, which is formalised by the Iris proof rule: ($\triangleright P$) $* \text{wp } e \{\Phi\} \multimap \text{wp } e \{w. P * \Phi w\}$ if $e \notin \text{Val}$ and $w \notin \text{FV}(P)$. This rule indicates that $\triangleright P$ can also be read as “ P holds after one more step of computation”, seeing as P is obtained without $\triangleright$ modality in the postcondition of the weakest precondition, denoting that at least one step has been taken.

In this paper we will not further detail the semantics of Iris, but refer the interested reader to Jung et al. [30] for an extensive account of the Iris model and proof rules.

2.3 Term Types

The definitions of our semantic type system are shown in Figure 1. Types Type_k are indexed by kinds; $\star$ for *term types*, and $\diamond$ for *session types*. Meta-variables $A, B \in \text{Type}_\star$ are used for term types, $S \in \text{Type}_\diamond$ for session types, and $K \in \text{Type}_k$ for types of any kind. Term types $\text{Type}_\star$ are defined as Iris predicates over values, and session types $\text{Type}_\diamond$ are defined as dependent separation protocols of Actris (§2.4).

Type Formers. The ground types (the unit type 1 , Boolean type $\mathbb{B}$, and integer type $\mathbb{Z}$) are defined through membership of the corresponding set ($\{()\}$, $\mathbb{B}$, and $\mathbb{Z}$, respectively). The type former $\text{ref}_{\text{uniq}} A$ for uniquely-owned references, $A_1 \times A_2$ for products, and $A \multimap B$ for affine functions nicely demonstrate the advantage of separation logic—since types are Iris predicates, they implicitly describe which resources are owned. The points-to connective ($w \mapsto v$) is used to describe that $\text{ref}_{\text{uniq}} A$ consists of the locations $w \in \text{Loc}$ that hold a value $v \in \text{Val}$ for which the resources $A v$ are owned. The separating conjunction ($*$) is used to describe that $A_1 \times A_2$ consists of tuples (w_1, w_2) , where the resources $A_1 w_1$ and $A_2 w_2$ are owned *separately*. The separating implication ($\multimap$) and weakest precondition are used to describe that the affine function type $A \multimap B$ consists of values w that

when applied to an argument v consume the resources $A v$, and in return, produce the resources B for the result of $w v$. Note that the weakest precondition $\text{wp } (w v) \{B\}$ is used so we can talk about the result of $w v$. We could not use $B (w v)$ since the term $w v$ is not a value.

We use Iris's later modality ($\triangleright$) to ensure that type formers are contractive, which is needed to model equi-recursive types using Iris's guarded fixpoint operator in §3.3.

Typing Judgement. As is common in substructural type systems with operations that perform strong updates, we use a typing judgement $\Gamma \vdash e : A \dashv \Gamma'$ (defined in Figure 1) with a *pre-* and *post-context* $\Gamma, \Gamma' \in \text{List}(\text{String} \times \text{Type}_\star)$. These contexts describe the types of variables before and after execution of the expression e .

As is standard in logical relations, we use *closing substitutions* to give a semantics to typing contexts. Closing substitutions $\sigma \in \text{String} \xrightarrow{\text{fin}} \text{Val}$ are finite partial functions that map the free variables of an expression to corresponding values. Closing substitutions come with a judgement $\Gamma \vdash \sigma$, which expresses that the closing substitution σ is well-typed in the context Γ . The definition of this judgement employs the iterated separation conjunction $\bigstar_{(x,A) \in \Gamma}$ to ensure that for each variable typing (x, A) in Γ , there is a corresponding value in the closing substitution $\sigma(x)$ for which the resources $A(\sigma(x))$ are owned separately.

The typing judgement $\Gamma \vdash e : A \dashv \Gamma'$ is defined in terms of Iris's weakest precondition. That is, given a closing substitution σ and resources $\Gamma \vdash \sigma$ for the pre-context Γ , the weakest precondition holds for e (under substitution with σ), with the postcondition stating that the resources $A v$ for the resulting value v are owned separately from the resources $\Gamma' \vdash \sigma$ for the post-context Γ' .

Typing Rules. Now that the type formers and the typing judgement are in place, we state the conventional typing rules as lemmas. We prove these lemmas by unfolding the definition of the semantic typing judgement $\Gamma \vdash e : A \dashv \Gamma'$, and proving the corresponding proposition in Iris using the rules for weakest preconditions. A selection of typing rules, along with Iris's weakest precondition rules used to prove them, is presented in Figure 2.

The typing rule for integer literals follows immediately from **wp-val**, which states that the weakest precondition of a value v holds if the postcondition Φv holds. The typing rule for variables also uses **wp-val**. Since the pre-context is $\Gamma, (x : A)$, we can assume ownership of $A v$ for some value v , and should prove a weakest precondition for v . After using **wp-val**, we prove the postcondition by giving up $A v$. Note that the post-context is $\Gamma, (x : \text{any})$ as ownership of A has been moved out. For substructural type systems this is crucial as in expressions such as **let** $x = y$ **in** e , it is generally not allowed to use y in e as ownership of the type of y has moved to x . This is formalised by giving the variable y type any in

e . The typing rules for load, let, and parallel composition are proved using the Iris rules **wp-load**, **wp-let**, and **wp-par**. The rule for parallel composition moreover relies on the property $(\Gamma_1 \cdot \Gamma_2 \vdash \sigma)$ iff $(\Gamma_1 \vdash \sigma) * (\Gamma_2 \vdash \sigma)$, which allows us to subdivide and recombine ownership of the pre- and post-contexts between both operands.

Type Safety. Type safety means: if $[\] \models e : A \dashv \Gamma$, then e is safe, i.e., e will not get stuck w.r.t. the language's operational semantics. For syntactic type systems, type safety is usually proven via the progress and preservation theorems. For our semantic type system, we get type safety from Iris's adequacy theorem, which states that a closed proof of a weakest precondition implies safety [30, 33]. Note that our type system is affine (resources are not explicitly deallocated), and thus the post-context Γ in the type safety statement need not be empty. We use an affine type system as that allows more practical safe programs to be typeable.

2.4 Session Types

We extend our core type system with the basic session-type formers for sending a message $!A.S$, receiving a message $?A.S$, the choice primitives for selection $\oplus\{\vec{S}\}$ and branching $\&\{\vec{S}\}$, and the terminator **end**. We let $\vec{S} : \mathbb{Z} \xrightarrow{\text{fin}} \text{Type}_\star$ be finite partial functions from labels to session types, and often write $\vec{S} = l_1 : S_1, \dots, l_n : S_n$. The term type **chan** S dictates that a term is a channel that follows the session type S .

Session types are defined in terms of Actris's *dependent separation protocols* [24], which are similar to session types in structure, but can express functional properties of the transferred data. Dependent separation protocols $\text{prot} \in \text{iProto}$ are streams of $! \vec{x} : \vec{\tau} \langle v \rangle \{P\}$. prot and $? \vec{x} : \vec{\tau} \langle v \rangle \{P\}$. prot constructors that are either infinite or finite. Here, v is the value that is being sent or received, P is an Iris proposition denoting the ownership of the resources being transferred as part of the message, and the logical variables $\vec{x} : \vec{\tau}$ bind into v, P , and prot to constrain the message v and the tail protocol prot . Finite protocols are ultimately terminated by an **end** constructor. As an example, the dependent separation protocol $!(\ell : \text{Loc}) (i : \mathbb{Z}) \langle \ell \rangle \{ \ell \mapsto i * 10 \leq i \}. ? \langle () \rangle \{ \ell \mapsto (i + 1) \}$. **end** expresses that an integer reference whose value is at least 10 is sent, after which the recipient increments it by one and sends back a unit token $()$ along with the reference ownership.

Actris's connective $c \multimap \text{prot}$ denotes ownership of a channel c with a dependent separation protocol prot . The Actris proof rules are shown in Figure 3. The rule for **new_chan** $()$ allows ascribing any protocol to a new channel, obtaining ownership of $c \multimap \text{prot}$ and $c' \multimap \overline{\text{prot}}$ for the respective endpoints. Here, $\overline{\text{prot}}$ is the *dual* of prot , in which any receive $(?)$ is turned into a send $(!)$, and *vice versa*. The rule for **send** $c w$ requires the head of the protocol to be a send $(!)$, and the value w that is sent to match up with the ascribed value. Concretely, to send a message w , one needs to give up ownership of $c \multimap ! \vec{x} : \vec{\tau} \langle v \rangle \{P\}$. prot , pick an appropriate

Selection of Iris's proof rules for weakest preconditions:

$$\begin{aligned}
& \Phi \vdash \text{wp } v \{ \Phi \} & (\text{WP-VAL}) \\
& \ell \mapsto v \vdash \text{wp } !\ell \{ w. (v = w) * (\ell \mapsto v) \} & (\text{WP-LOAD}) \\
& \text{wp } e_1 \{ v. \text{wp } e_2[v/x] \{ \Phi \} \} \vdash \text{wp } (\text{let } x = e_1 \text{ in } e_2) \{ \Phi \} & (\text{WP-LET}) \\
& \text{wp } e_1 \{ \Phi_1 \} * \text{wp } e_2 \{ \Phi_2 \} \vdash \text{wp } (e_1 \parallel e_2) \{ v. \exists v_1, v_2. (v = (v_1, v_2)) * \Phi_1 v_1 * \Phi_2 v_2 \} & (\text{WP-PAR})
\end{aligned}$$

Selection of semantic typing rules:

$$\begin{aligned}
& \Gamma \models i : \mathbf{Z} \models \Gamma & \Gamma, (x : A) \models x : A \models \Gamma, (x : \text{any}) & \Gamma, (x : \text{ref}_{\text{uniq}} A) \models !x : A \models \Gamma, (x : \text{ref}_{\text{uniq}} \text{any}) \\
& \frac{\Gamma_1 \models e_1 : A \models \Gamma_2 \quad \Gamma_2, (x : A) \models e_2 : B \models \Gamma_3}{\Gamma_1 \models (\text{let } x = e_1 \text{ in } e_2) : B \models \Gamma_3 \setminus x} & \frac{\Gamma_1 \models e_1 : A_1 \models \Gamma'_1 \quad \Gamma_2 \models e_2 : A_2 \models \Gamma'_2}{\Gamma_1 \cdot \Gamma_2 \models e_1 \parallel e_2 : A_1 \times A_2 \models \Gamma'_1 \cdot \Gamma'_2}
\end{aligned}$$

Figure 2. A selection of Iris's proof rules and semantic typing rules.**Actris's proof rules for dependent separation protocols:**

$$\begin{aligned}
& \text{wp } \text{new_chan} () \{ v. \exists c, c'. (v = (c, c')) * c \mapsto \text{prot} * c' \mapsto \overline{\text{prot}} \} & (\text{WP-NEWCHAN}) \\
& c \mapsto !\vec{x} : \vec{\tau} \langle v \rangle \{ P \}. \text{prot} * P[\vec{t}/\vec{x}] \vdash \text{wp } \text{send } c (v[\vec{t}/\vec{x}]) \{ c \mapsto \text{prot}[\vec{t}/\vec{x}] \} & (\text{WP-SEND}) \\
& c \mapsto ?\vec{x} : \vec{\tau} \langle v \rangle \{ P \}. \text{prot} \vdash \text{wp } \text{recv } c \{ w. \exists \vec{y}. (w = v[\vec{y}/\vec{x}]) * c \mapsto \text{prot}[\vec{y}/\vec{x}] * P[\vec{y}/\vec{x}] \} & (\text{WP-RECV})
\end{aligned}$$

Semantic typing rules for channels:

$$\begin{aligned}
& \Gamma \models \text{new_chan} () : \text{chan } S \times \text{chan } \bar{S} \models \Gamma \\
& \frac{\Gamma \models e : A \models \Gamma', (x : \text{chan } (!A. S))}{\Gamma \models \text{send } x e : \mathbf{1} \models \Gamma', (x : \text{chan } S)} & \Gamma, (x : \text{chan } (?A. S)) \models \text{recv } x : A \models \Gamma, (x : \text{chan } S) \\
& \frac{1 \leq i \leq n}{\Gamma, (x : \text{chan } (\oplus \{ l_1 : S_1, \dots, l_n : S_n \})) \models \text{select } x l_i : \mathbf{1} \models \Gamma, (x : \text{chan } S_i)} \\
& \frac{\Gamma, (x : \text{chan } S_1) \models e_1 : A \models \Gamma' \quad \dots \quad \Gamma, (x : \text{chan } S_n) \models e_n : A \models \Gamma'}{\Gamma, (x : \text{chan } (\& \{ l_1 : S_1, \dots, l_n : S_n \})) \models \text{branch } x \text{ with } l_1 \Rightarrow e_1 \mid \dots \mid l_n \Rightarrow e_n : A \models \Gamma'}
\end{aligned}$$

Figure 3. Actris's proof rules for dependent separation protocols and semantic typing rules for channels.

instantiation $\vec{t}$ for the variables $\vec{x} : \vec{\tau}$ so that $w = v[\vec{t}/\vec{x}]$, and give up ownership of the associated resources $P[\vec{t}/\vec{x}]$. Subsequently, one gets back ownership of the protocol tail $c \mapsto \text{prot}[\vec{t}/\vec{x}]$. The rule for `recv` c is essentially dual to the rule for `send` c w . One needs to give up ownership of $c \mapsto ?\vec{x} : \vec{\tau} \langle v \rangle \{ P \}. \text{prot}$, and in return acquires the resources $P[\vec{y}/\vec{x}]$, the return value w where $w = v[\vec{y}/\vec{x}]$, and finally the ownership of the protocol tail $\text{prot}[\vec{y}/\vec{x}]$, where $\vec{y}$ are instances of the variables of the protocol.

Semantics of Session Types. The definitions of session types are shown in Figure 1. Since session types ($\text{Type}_\bullet$) are defined as dependent separation protocols iProto , the channel type $\text{chan } S$ is defined in terms of Actris's connective for channel ownership $w \mapsto S$. The definition of the terminator (`end`), send (`!`), and receive (`?`) follow from their dependent separation protocol counterparts. For example $!A. S$ is defined as $!(v : \text{Val}) \langle v \rangle \{ A v \}. S$. It says that a value v is sent along with ownership of $A v$.

While the choice types $\oplus \{ \vec{S} \}$ and $\& \{ \vec{S} \}$ do not have a direct counterpart in Actris, they can be encoded. For example, $\oplus \{ \vec{S} \}$ is defined as $!(l : \mathbb{Z}) \langle l \rangle \{ l \in \text{dom}(\vec{S}) \}. \vec{S}(l)$. It expresses that a valid label $l \in \text{dom}(\vec{S})$ (modelled as an integer) is sent. This definition makes use of the fact that dependent separation protocols are *dependent*, as the protocol tail $\vec{S}(l)$ depends on the label l that is sent.

Duality. The duality $\bar{S}$ of session types S is inherited from Actris. We thus obtain the usual duality laws (on the left) from the Actris duality laws (on the right):

$$\begin{aligned}
& \overline{\text{end}} = \text{end} & \overline{\text{end}} = \text{end} \\
& \overline{!A. S} = ?A. \bar{S} & \overline{!\vec{x} : \vec{\tau} \langle v \rangle \{ P \}. \text{prot}} = ?\vec{x} : \vec{\tau} \langle v \rangle \{ P \}. \overline{\text{prot}} \\
& \overline{?A. S} = !A. \bar{S} & \overline{?\vec{x} : \vec{\tau} \langle v \rangle \{ P \}. \text{prot}} = !\vec{x} : \vec{\tau} \langle v \rangle \{ P \}. \overline{\text{prot}}
\end{aligned}$$

Similarly, our semantic definition of the branch ($\&$) and select ($\oplus$) operators in terms of Actris's send (`!`) and receive (`?`) protocols, enables us to use the Actris duality laws to prove

that the dual of a select is a branch, and *vice versa*:

$$\begin{aligned} \overline{\oplus\{l_1 : S_1, \dots, l_n : S_n\}} &= \&\{l_1 : \overline{S_1}, \dots, l_n : \overline{S_n}\} \\ \&\{l_1 : S_1, \dots, l_n : S_n\} &= \overline{\oplus\{l_1 : \overline{S_1}, \dots, l_n : \overline{S_n}\}} \end{aligned}$$

Session Typing Rules. The session typing rules are shown in Figure 3. Since the channel operations perform strong updates, the typing rules require channels to be variables so they can update the context. Given the close similarity between Actris and session typing, the typing rules follow from the Actris rules up to minor separation logic reasoning.

The rules for `select` and `branch` demonstrate the extensibility of our approach. Our language does not have these operations as primitives, but they can be defined as macros:

```
select x l  $\triangleq$  send x l
branch x with  $\triangleq$  let i = recv x in
  l1  $\Rightarrow$  e1 | ... | if i = l1 then e1 else ...
  ln  $\Rightarrow$  en if i = ln then en else () ()
```

The typing rule for `select` follows immediately from the proof rule for `send`. Similarly, the typing rule for `branch` follows from the proof rule for `recv`, but additionally requires some reasoning in Iris about the sequence of if-expressions. Note that the stuck expression `() ()` is used in case no matching branch for the label l_i has been found. While this stuck expression is obviously not safe, it is never executed because of the condition $l \in \text{dom}(\vec{S})$ in the semantic definition of the select ($\oplus$) and branch ($\&$) operators.

Type Safety. Since the extension with session types did not change the definition of the semantic typing judgement, but merely added new type formers and typing rules, the type safety result from §2.3 remains applicable without change.

3 Extending the Type System

We demonstrate the extensibility of our approach to session types by adding term- and session-level subtyping (§3.1 and §3.7), copyable types (§3.2), term- and session-level equi-recursive types (§3.3), term- and session-level polymorphism (§3.4 and §3.5), and locks/mutexes (§3.6). While we only present a small representative selection of rules associated with each extension, all rules can be found in Appendix A.

3.1 Term-Level Subtyping

Subtyping $A <: B$ indicates that any member of type A is also a member of type B . In a semantic type system, subtyping is defined in terms of the separating implication:

$$\begin{aligned} A <: B &\triangleq \forall v. A \ v \multimap B \ v \\ \Gamma <_{\text{ctx}} \Gamma' &\triangleq \forall \sigma. (\Gamma \models \sigma) \multimap (\Gamma' \models \sigma) \end{aligned}$$

The definition states that A is a subtype of B if for any value v , we can give up resources $A \ v$ to obtain resources $B \ v$. The *context subtyping relation* $\Gamma <_{\text{ctx}} \Gamma'$ is defined similarly. It is essentially the pointwise lifting of the subtyping relation,

applied to each type in the contexts Γ and Γ' . It expresses that when we hold resources $\Gamma \models \sigma$ for the context Γ , then we can give those up to obtain the resources $\Gamma' \models \sigma$ for Γ' .

With these definitions at hand, we prove the usual subsumption rule as a lemma:

$$\frac{\Gamma_1 <_{\text{ctx}} \Gamma'_1 \quad \Gamma'_1 \models e : A \multimap \Gamma'_2 \quad A <: B \quad \Gamma'_2 <_{\text{ctx}} \Gamma_2}{\Gamma_1 \models e : B \multimap \Gamma_2}$$

The proof of the above lemma makes use of the Iris proof rule $(\forall v. \Phi_1 v \multimap \Phi_2 v) \multimap \text{wp } e \{ \Phi_1 \} \multimap \text{wp } e \{ \Phi_2 \}$, which states that separating implications can be applied in the postconditions of weakest preconditions.

In addition to the subsumption rule, we prove the conventional subtyping rules as lemmas. For example:

$$\begin{aligned} A <: A \quad \frac{A <: B \quad B <: C}{A <: C} \quad A <: \text{any} \\ \frac{C <: A \quad B <: D}{A \multimap B <: C \multimap D} \quad \frac{A <: C \quad B <: D}{A \times B <: C \times D} \end{aligned}$$

These lemmas are proved by unfolding the definition of the subtyping relation, and involve some trivial reasoning using separating implication in Iris. We will see more interesting subtyping rules in §3.2 and §3.7.

3.2 Copyable Types

Session-type systems are substructural, in the sense that some types are inhabited by values that can be used at most once. This becomes evident in the variable and load rules from §2.3, which move out ownership by turning the element type into the any type. While moving out ownership is necessary for soundness in general, this is too restrictive for types that do not assert ownership of any resources, such as $\mathbf{B}$, $\mathbf{Z}$, or $\mathbf{Z} * \mathbf{B}$. These types need not be moved out as their inhabitants can be used multiple times. We therefore extend the type system with a notion of *copyable types*. Concretely, we define a type former `copy` and a property `copyable`:

$$\begin{aligned} \text{copy } A &\triangleq \lambda w. \Box(A \ w) \\ \text{copyable } A &\triangleq A <: \text{copy } A \end{aligned}$$

The type `copy A` describes the values of type A that can be freely duplicated (used an arbitrary number of times). We thus have $A <: \text{copy } A$ for ground types $A \in \{\mathbf{1}, \mathbf{B}, \mathbf{Z}\}$, but not for types like $A \in \{\text{ref}_{\text{uniq}} B, \text{chan } S\}$ that assert ownership. Conversely, we have $\text{copy } A <: A$ for any type A , i.e., `copy A` is always a subtype of A . A type is *copyable* (written `copyable A`) if *all* of its values can be freely duplicated, i.e., when A is a subtype of `copy A`. Ground types ($\mathbf{1}$, $\mathbf{B}$, $\mathbf{Z}$) are copyable, and copyability is closed under products and sums.

An example of a type where some, but not all, values can be duplicated is the type $A \multimap B$ of affine functions: a function can only be duplicated if it has not captured ownership of exclusive resources from the context (through a free variable

that has a non-copyable type). Hence, we define $A \rightarrow B \triangleq \text{copy } (A \multimap B)$ as the type of *unrestricted* functions, that can be applied any number of times.

The type former `copy` is defined using the *persistence* modality ($\Box$) of Iris, where $\Box P$ means that the proposition P holds without ownership of (exclusively-owned) resources. Propositions that do not assert ownership of (exclusively-owned) resources are called *persistent*. In particular, $\Box P$ is always persistent, allowing the proposition P to be freely duplicated using the rule $\Box P \multimap (\Box P * P)$. This allows copyable types occurring in the context to be duplicated:

$$(x : A) <_{\text{ctx}} (x : A), (x : A) \quad \text{if copyable } A$$

Our approach of using Iris's notion of persistence to model copyability of types is similar to the approach used in RustBelt [27, 28] to model the substructural features of Rust. However, copyability in RustBelt is defined directly in Iris, and not reflected into the type system by means of a copy type former and a subtyping rule.

3.3 Equi-Recursive Term and Session Types

We extend our type system with equi-recursive types using Iris's fixpoint operator. Recall from §2.2 that Iris's fixpoint operator requires that recursive definitions are contractive, meaning that recursive occurrences appear below a later modality ($\triangleright$). A recursive occurrence is also considered guarded when it appears in:

- The postcondition Φ of an Iris weakest precondition $\text{wp } e \{ \Phi \}$ with $e \notin \text{Val}$.
- The tail *prot* of the dependent separation protocols $! \vec{x} : \vec{\tau} \langle v \rangle \{ P \}. \text{prot}$ and $? \vec{x} : \vec{\tau} \langle v \rangle \{ P \}. \text{prot}$.
- The protocol *prot* of the Actris connective $c \mapsto \text{prot}$ for channel ownership.

These occurrences are guarded because the corresponding constructs contain $\triangleright$ modalities internally.

We lift the guarded recursion operator of Iris into a *kinded* operator for equi-recursion in the type system:

$$\mu (X : k). K \triangleq \mu (X : \text{Type}_k). K \quad (K \text{ is contractive in } X)$$

From Iris's proof rule for fixpoints, we get that this is indeed a fixpoint, *i.e.*, we have $\mu (X : k). K = K[\mu (X : k). K/X]$.

We put later modalities in the definitions of type formers to ensure that they are contractive in all arguments. This allows construction of recursive term and session types, including examples from the session type literature [21], such as $\mu (X : \diamond). !(\text{chan } X). X$, where the recursion variable X occurs in the type of messages.

It is worth noting that most existing logical relation developments in Iris model iso-recursive types. Hence, instead of putting $\triangleright$ modalities in the definitions of type formers, they put a $\triangleright$ modality in the definition of the recursion operator. This avoids the contractive side-condition, but requires explicit fold and unfold operations in the language (to take an operational step to remove the $\triangleright$ modality).

3.4 Polymorphism in Term Types

We extend the type system with kinded parametric polymorphism, by introducing universal types $\forall (X : k). A$ and existential types $\exists (X : k). A$, which are polymorphic in a variable X of kind k . The kind k indicates whether the type is polymorphic over term types (kind $\star$) or session types (kind $\diamond$). Using polymorphism in term types, we can write types such as $\forall (X : \star). X \rightarrow X$ (for describing the polymorphic identity function). Using polymorphism in session types, we can write types such as $\forall (X : \diamond). \text{chan } (!\mathbb{Z}.X) \multimap \text{chan } X$ (for describing a function that reads an integer from a channel with an arbitrary tail X). Universal and existential types are defined as follows:

$$\begin{aligned} \forall (X : k). A &\triangleq \lambda w. \forall (X : \text{Type}_k). \text{wp } (w ()) \{ A \} \\ \exists (X : k). A &\triangleq \lambda w. \exists (X : \text{Type}_k). \triangleright (A w) \end{aligned}$$

As is custom for logical relations in Iris, these types are defined in the style of parametricity—they use Iris-level universal and existential quantifiers over semantic types $X : \text{Type}_k$. This is possible because Iris supports higher-order impredicative quantification (*i.e.*, quantification over Iris predicates and Actris protocols).

Note that universal types are inhabited by values w that produce a value of the instantiated type A when applied to the unit value $()$, as indicated by the weakest precondition in the definition. In other words, the inhabitants of universal types are *thunks*. This is since we consider a type system with explicit type abstraction and type instantiation constructs. Since the base language is untyped, we use term-level abstractions to indicate type abstraction and instantiation: type abstraction is written $\lambda _. e$ and type instantiation is written $w ()$ when w is a type abstraction. By using explicit thunks, we avoid having to impose an ML-like *value restriction* [50] to ensure type safety in the presence of imperative side-effects. The typing rules for term-level polymorphism are standard and can be found in Appendix A.

3.5 Polymorphism in Session Types

A more interesting extension is polymorphism in session types [18]. An example is the following type, which describes the interaction with a polymorphic computation service:

$$\begin{aligned} \text{compute_type} &\triangleq \mu (\text{rec} : \diamond). \\ &\quad \oplus \{ \text{cont} : ! (X : \star). (1 \multimap X). ?X. \text{rec}, \text{stop} : \text{end} \} \end{aligned}$$

The service can be used by sending computation requests $1 \multimap X$, and then awaiting their results X . Different types can be picked for the type variable X at each recursive iteration.

To extend our type system with polymorphism in session types, we redefine the send and receive session types to include binders $\vec{X}$ for type variables:

$$\begin{aligned} !_{\vec{X} : \vec{k}} A. S &\triangleq ! (\vec{X} : \vec{\text{Type}}_k) (v : \text{Val}) \langle v \rangle \{ A v \}. S \\ ?_{\vec{X} : \vec{k}} A. S &\triangleq ? (\vec{X} : \vec{\text{Type}}_k) (v : \text{Val}) \langle v \rangle \{ A v \}. S \end{aligned}$$

Iris's proof rules for locks:

$\text{isLock } lk \ R \multimap \Box(\text{isLock } lk \ R)$
 $R \multimap \text{wp } \text{newlock } () \{lk. \text{isLock } lk \ R\}$
 $\text{isLock } lk \ R \multimap \text{wp } \text{acquire } lk \ \{R\}$
 $\text{isLock } lk \ R \multimap R \multimap \text{wp } \text{release } lk \ \{\text{True}\}$

Semantic typing rules for mutexes:

$\text{copyable } (\text{mutex } A) \quad \Gamma \models \text{newmutex} : A \rightarrow \text{mutex } A \models \Gamma$
 $\Gamma, (x : \text{mutex } A) \models \text{acquiremutex } x : A \models \Gamma, (x : \overline{\text{mutex}} A)$

$$\frac{\Gamma \models e : A \models \Gamma', (x : \overline{\text{mutex}} A)}{\Gamma \models \text{releasemutex } x \ e : 1 \models \Gamma', (x : \text{mutex } A)}$$

Figure 4. Iris's proof rules for locks and semantic typing rules for mutexes.

The binders $\vec{X}$ are kinded so that we can quantify over both term types and session types.

This definition relies on the fact that binders $\vec{x} : \vec{\tau}$ in Actris's dependent separation protocols $! \vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}$ and $? \vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}$ are higher-order and impredicative (*i.e.*, they allow quantification over Iris predicates and Actris protocols). The typing rules are extended to allow instantiation of binders when sending a message, and elimination of type variables when receiving a message. Concretely, the rule for channel creation remains unchanged, while the rules for send and receive become:

$$\frac{\Gamma \models e : A[\vec{K}/\vec{X}] \models \Gamma', (x : \text{chan } (!_{\vec{X}, \vec{K}} A. S))}{\Gamma \models \text{send } x \ e : 1 \models \Gamma', (x : \text{chan } S[\vec{K}/\vec{X}])}$$

$$\frac{\Gamma, (x : \text{chan } S), (y : A) \models e : B \models \Gamma' \quad \vec{X} \notin FV(\Gamma, \Gamma', B)}{\Gamma, (x : \text{chan } (?_{\vec{X}, \vec{K}} A. S)) \models \text{let } y = \text{recv } x \text{ in } e : B \models \Gamma' \setminus \{y\}}$$

The second rule requires the result y of `recv` to be let-bound to ensure that the type variables $\vec{X}$ cannot escape into the context Γ' or the type B .

With this rule, we can type check the following function that follows the computation service type `compute_type`:

```

compute_service  $\triangleq$  rec go c =
  branch c with
    cont  $\Rightarrow$  let f = recv c in send (f ()); go c
  | stop  $\Rightarrow$  ()
end

```

We can prove the typing judgement $\Gamma \models \text{compute_service} : \text{chan } \text{compute_type} \rightarrow () \models \Gamma$ using only the typing rules of our semantic type system. In §4.2 we consider a client that uses this service, which cannot itself be type checked using our typing rules but rather requires a manual proof of its typing judgement.

3.6 Locks and Mutexes

The substructural nature of channels (of type `chan S`) ensure that they can be used by at most one thread at the same time. Balzer and Pfenning [5] proposed a more liberal extension of session types that allows channels to be shared between multiple threads via locks. We show that we can achieve a similar kind of sharing by extending our type system with a

type former `mutex A` of mutexes (*i.e.*, lock-protected values of type A) inspired by Rust's Mutex library. For example, mutexes make it possible to share the channel to the computation service `compute_type` from §3.5 between multiple clients—they can acquire the mutex (`mutex compute_type`), send any number of computation requests, retrieve the corresponding results, and then release the mutex.

The mutex type former is copyable, and comes with operations `newmutex` to allocate a mutex, `acquiremutex` to acquire a mutex by blocking until no other thread holds it, and `releasemutex` to release the mutex. The typing rules are shown in Figure 4 and include the type former `mutex`, which signifies that the mutex is acquired.

To extend our type system with mutexes we make use of the locks library that is available in Iris. This library consists of operations `newlock`, `acquire`, and `release`, which are similar to the mutex operations, but do not protect a value. The mutex operations are defined in terms of locks as follows:

```

newmutex  $\triangleq$   $\lambda y. (\text{newlock } (), \text{ref } y)$ 
acquiremutex  $\triangleq$   $\lambda x. \text{acquire } (\text{fst } x); !(\text{snd } x)$ 
releasemutex  $\triangleq$   $\lambda x \ y. (\text{snd } x) \leftarrow y; \text{release } (\text{fst } x)$ 

```

That is, `newmutex` creates a lock alongside a boxed value. The value can then be acquired with `acquiremutex`, which first acquires the lock. Finally, `releasemutex` moves the value back into the box, and releases the lock.

The Iris rules for locks are shown in Figure 4 and make use of the representation predicate `isLock lk R`, which expresses that a lock lk guards the resources R . When creating a new lock one has to give up ownership of R , and in turn, obtains the representation predicate `isLock lk R`. The representation is persistent, so it can be freely duplicated. When entering a critical section using `acquire lk`, a thread gets exclusive ownership of R , which has to be given up when releasing the lock using `release lk`. Using the lock representation predicate, we define type formers for mutexes:

```

mutex A  $\triangleq$   $\lambda w. \exists lk, \ell. (w = (lk, \ell)) \ast$ 
 $\text{isLock } lk \ (\exists v. (\ell \mapsto v) \ast \triangleright (A v))$ 
 $\overline{\text{mutex}} A \triangleq$   $\lambda w. \exists lk, \ell. (w = (lk, \ell)) \ast (\ell \mapsto -) \ast$ 
 $\text{isLock } lk \ (\exists v. (\ell \mapsto v) \ast \triangleright (A v))$ 

```

The mutex type former states that its values are pairs of locks and boxed values. The `mutex` type former additionally

Actris's proof rules for subprotocols:

$$\begin{array}{ll}
(\forall \vec{x}:\vec{\tau}. P * (prot_1 \sqsubseteq !\langle v \rangle\{\text{True}\}. prot_2)) * (prot_1 \sqsubseteq !\vec{x}:\vec{\tau} \langle v \rangle\{P\}. prot_2) & \text{if each } \tau \in \vec{\tau} \text{ is inhabited} \quad (\sqsubseteq\text{-SEND-OUT}) \\
(\forall \vec{x}:\vec{\tau}. P * (? \langle v \rangle\{\text{True}\}. prot_1 \sqsubseteq prot_2)) * (? \vec{x}:\vec{\tau} \langle v \rangle\{P\}. prot_1 \sqsubseteq prot_2) & \text{if each } \tau \in \vec{\tau} \text{ is inhabited} \quad (\sqsubseteq\text{-RECV-OUT}) \\
P[\vec{t}/\vec{x}] * (!\vec{x}:\vec{\tau} \langle v \rangle\{P\}. prot \sqsubseteq !\langle v[\vec{t}/\vec{x}] \rangle\{\text{True}\}. prot[\vec{t}/\vec{x}]) & (\sqsubseteq\text{-SEND-IN}) \\
P[\vec{t}/\vec{x}] * (? \langle v[\vec{t}/\vec{x}] \rangle\{\text{True}\}. prot[\vec{t}/\vec{x}] \sqsubseteq ? \vec{x}:\vec{\tau} \langle v \rangle\{P\}. prot) & (\sqsubseteq\text{-RECV-IN}) \\
? \langle v_1 \rangle\{P_1\}. ! \langle v_2 \rangle\{P_2\}. prot \sqsubseteq ! \langle v_2 \rangle\{P_2\}. ? \langle v_1 \rangle\{P_1\}. prot & (\sqsubseteq\text{-SWAP}) \\
\triangleright (prot_1 \sqsubseteq prot_2) * (! \langle v \rangle\{P\}. prot_1 \sqsubseteq ! \langle v \rangle\{P\}. prot_2) & (\sqsubseteq\text{-SEND-MONO}) \\
\triangleright (prot_1 \sqsubseteq prot_2) * (? \langle v \rangle\{P\}. prot_1 \sqsubseteq ? \langle v \rangle\{P\}. prot_2) & (\sqsubseteq\text{-RECV-MONO}) \\
(prot_1 \sqsubseteq prot_2) * (c \multimap prot_1) * (c \multimap prot_2) & (\sqsubseteq\text{-CHAN-MONO})
\end{array}$$

Semantic subtyping rules for session polymorphism:

$$\frac{S_1 <: !A. S_2}{S_1 <: !_{(\vec{X}:\vec{k})} A. S_2} \quad \frac{?A. S_1 <: S_2}{?_{(\vec{X}:\vec{k})} A. S_1 <: S_2} \quad !_{(\vec{X}:\vec{k})} A. S <: !A[\vec{K}/\vec{X}]. S[\vec{K}/\vec{X}] \quad ?A[\vec{K}/\vec{X}]. S[\vec{K}/\vec{X}] <: ?_{(\vec{X}:\vec{k})} A. S$$

Figure 5. A selection of the Actris's proof rules for subprotocols and semantic session subtyping rules.

asserts ownership of the reference, implying that the lock has been acquired. The typing rules for mutexes as shown in Figure 4 are proven as lemmas.

3.7 Session-Level Subtyping

Session-level subtyping $S <: T$, originally presented by Gay and Hole [20], relates session subtypes S with session super-types T , that can be used in place of the subtype, captured by monotonicity with the subtyping of the channel type:

$$\frac{S <: T}{\text{chan } S <: \text{chan } T}$$

Subtyping in session types allows sending supertypes and receiving subtypes, as well as increasing and reducing the range of choices for branchings and selections, respectively:

$$\begin{array}{c}
\frac{A_2 <: A_1 \quad S_1 <: S_2}{!A_1. S_1 <: !A_2. S_2} \quad \frac{A_1 <: A_2 \quad S_1 <: S_2}{?A_1. S_1 <: ?A_2. S_2} \\
\\
\frac{\vec{S}_2 \subseteq \vec{S}_1}{\oplus\{\vec{S}_1\} <: \oplus\{\vec{S}_2\}} \quad \frac{\vec{S}_1 \subseteq \vec{S}_2}{\&\{\vec{S}_1\} <: \&\{\vec{S}_2\}}
\end{array}$$

This is essential for program reuse, e.g., any program that handles more choices than indicated by a branch type should be able to accept a channel with that branch type.

In asynchronous session types, one can further extend subtyping with a “swapping” rule $?A_1. !A_2. S <: !A_2. ?A_1. S$ that allows performing sends (!) ahead of receives (?), and similar rules that allow performing selects ($\oplus$) ahead of receives (?), sends (!) ahead of branches ($\&$), and selects ($\oplus$) ahead of branches ($\&$) [37]². For example, using swapping, a client of the computation service from §3.5, with type `compute_type`

can swap the selects and sends ahead of receives, to send multiple computation requests at once, and only then await the computed results.

To extend our semantic session types with session subtyping, we make use of Actris's notion of *subprotocols* [23], for which the rules are shown in Figure 5. The first four rules mimic the behaviour of session subtyping in how it is possible to send more and receive less, while accounting for the protocol-level binders of dependent separation protocols. In particular, we can (1) move out binders and propositions of right-hand side sending protocols ($\sqsubseteq\text{-SEND-OUT}$) and (2) left-hand side receiving protocols ($\sqsubseteq\text{-RECV-OUT}$), and (3) move in binders and propositions of right-hand side sending protocols ($\sqsubseteq\text{-SEND-IN}$) and (4) left-hand side receiving protocols ($\sqsubseteq\text{-RECV-IN}$). Rule $\sqsubseteq\text{-SWAP}$ accounts for the swapping of sends and receives that are independent of each other, as guaranteed by the omission of binders in the rule. If binders are present, the first four rules should be used first. Rules $\sqsubseteq\text{-SEND-MONO}$ and $\sqsubseteq\text{-RECV-MONO}$ account for the monotonicity of the subprotocol relation in the tails, and rule $\sqsubseteq\text{-CHAN-MONO}$ states that Actris's connective for channel ownership is closed under the subprotocol relation. The subprotocol relation is reflexive and transitive.

With Actris's subprotocol relation at hand, we define the semantic subtyping relation for session types as follows:

$$S <: T \triangleq S \sqsubseteq T$$

We then prove the conventional subtyping rules for asynchronous session types as lemmas using the rules in Figure 5 for Actris's subprotocol relation. These subtyping rules include, but are not limited to, contra- and covariance of the type A of the send $!A. S$ and receive $?A. S$ session types respectively, the various forms of swapping as described in the beginning of this section, and the rules for reducing and

²Discrepancies in the direction between the swapping rules of Mostrous et al. [37] and us will be discussed in §6.

increasing the range of choices for selecting and branching protocols as also shown in the beginning of this section.

As a new feature, which up to our knowledge is not present in existing session type systems, we prove the subtyping rules for polymorphic session types as shown in Figure 5. For sending session types, we can instantiate the polymorphic types of subtypes, and generalise over the polymorphic types for supertypes. Conversely, for receiving session types, we can instantiate the polymorphic types of supertypes, and generalise over the polymorphic types for subtypes.

Subtyping for polymorphic session types is useful to describe the interaction between generic services and concrete clients. For example, consider a mapping service to which one can send a function $A \multimap B$, a value A , and get back the mapped result B . The most generic session type for interacting with such a service would be the following:

$$!(X, Y, \star) (X \multimap Y). !X. ?Y. \text{end}$$

Now assume that we have type checked a concrete client with the following session type:

$$!(Z \multimap B). !Z. ?B. \text{end}$$

While this concrete session type is not compatible with the one expected by the service, we can use the subtyping relation to weaken the generic type into the concrete one:

$$!(X, Y, \star) (X \multimap Y). !X. ?Y. \text{end} \sqsubseteq !(Z \multimap B). !Z. ?B. \text{end}$$

The judgement follows from $!(\vec{X}, \vec{k}) A. S <: !A[\vec{K}/\vec{X}]. S[\vec{K}/\vec{X}]$. Conversely, one could allocate the types from the perspective of the concrete client, and then weaken the service type into the generic type, by generalising over the received types:

$$?(Z \multimap B). ?Z. !B. \text{end} \sqsubseteq ?(X, Y, \star) (X \multimap Y). ?X. !Y. \text{end}$$

This subtyping judgement follows from the rule for receive $?A[\vec{K}/\vec{X}]. S[\vec{K}/\vec{X}] <: ?(\vec{X}, \vec{k}) A. S$.

4 Manual Typing Proofs

We demonstrate how safe programs that are not typeable using the existing typing rules can be assigned a typing judgement via a manual proof in Iris/Actris. We call such proofs *manual typing proofs*. As advocated by Jung et al. [27, 28], such proofs are useful since typing judgements, regardless of whether they have been derived manually or by using our typing rules, are interchangeable. While Jung et al. use such proofs to verify low-level concurrent libraries, we use them to verify binary message-passing programs where the user of one endpoint is verified using existing typing rules, and the other via a manual typing proof.

We first provide an intuition for the manual typing proof approach by proving the typing judgement of the parallel receiving program from the introduction (§4.1), and then show a more realistic example by proving the typing judgement of a parallel client of the computation service from §3.5 that uses a producer/consumer pattern (§4.2).

4.1 Receiving in Parallel

Consider the example from the introduction (where the locks have been made explicit):

```
threadprog  $\triangleq$   $\lambda c$  lk. acquire lk;
                let  $x = \text{recv } c$  in
                release lk;
                 $x$ 

lockprog  $\triangleq$   $\lambda c$ . let lk = newlock () in
                (threadprog  $c$  lk || threadprog  $c$  lk)
```

We want to prove:

$$\Gamma \models \text{lockprog} : \text{chan } (?Z. ?Z. \text{end}) \multimap (Z \times Z) \models \Gamma$$

This typing judgement is not derivable from the typing rules we presented so far, even with mutexes instead of plain locks, as the channel type changes each time the lock/mutex is acquired and released. However, we can unfold the definition of the semantic typing judgement and types, which gives us the following proof obligation in Iris/Actris:

$$(c \multimap ?(v_1 : \text{Val}) \langle v_1 \rangle \{v_1 \in \mathbb{Z}\}. \\ ?(v_2 : \text{Val}) \langle v_2 \rangle \{v_2 \in \mathbb{Z}\}. \text{end}) \multimap \\ \text{wp lockprog } c \left\{ v. \exists v_1, v_2. (v = (v_1, v_2)) * \right. \\ \left. \triangleright (v_1 \in \mathbb{Z}) * \triangleright (v_2 \in \mathbb{Z}) \right\}$$

The proof of above obligation is carried out using Iris's support for fractional permissions $[\bar{q}]^\gamma$ where $q \in (0, 1]_{\mathbb{Q}}$ and γ is an identifier. The permission reflects how much of the channel protocol its owner is allowed to resolve, enforced by the following lock invariant:

$$\begin{aligned} \text{chaninv} &\triangleq \\ (c \multimap ?(v_1 : \text{Val}) \langle v_1 \rangle \{v_1 \in \mathbb{Z}\}. \\ &\quad ?(v_2 : \text{Val}) \langle v_2 \rangle \{v_2 \in \mathbb{Z}\}. \text{end}) \quad \vee \quad (i) \\ (c \multimap ?(v_2 : \text{Val}) \langle v_2 \rangle \{v_2 \in \mathbb{Z}\}. \text{end} * [\bar{1}/2]^\gamma) \quad \vee \quad (ii) \\ (c \multimap \text{end} * [\bar{1}]^\gamma) \quad \quad \quad (iii) \end{aligned}$$

The invariant describes that the channel is in one of three states: (i) no values have been received yet, (ii) one value has been received, or (iii) all values have been received. State (ii) and (iii) assert that the invariant (not the thread) has half and full ownership of the fractional permission respectively.

The proof is carried out by allocating a full fractional permission $[\bar{1}]^\gamma$ (with a fresh identifier γ), after which the lock predicate $\text{isLock } lk \text{ chaninv}$ is allocated by giving up ownership of the channel c , where chaninv is initially in state (i). The fractional permission is then split into two halves $[\bar{1}/2]^\gamma$, which are each delegated to a thread, along with the persistent lock predicate $\text{isLock } lk \text{ chaninv}$. Both threads have the same proof obligation:

$$(\text{isLock } lk \text{ chaninv} * [\bar{1}/2]^\gamma) \multimap \\ \text{wp threadprog } c \text{ lk } \{v. v \in \mathbb{Z}\}$$

First, the lock invariant is obtained by acquiring the lock. The channel can then either be in state (i) or (ii), as having

```

compute_client  $\triangleq$   $\lambda l c.$ 
  let  $n = \text{llength } l$  in
  let  $\text{cntr} = \text{ref } 0$  in
  let  $l' = \text{lnil } ()$  in
  let  $lk = \text{newlock } ()$  in
  (send_all  $l \text{ cntr } lk \ c \ ||$ 
   recv_all  $l' \ n \ \text{cntr } lk \ c$ );
  l'

send_all  $\triangleq$ 
  rec go  $l \ \text{cntr } lk \ c =$ 
    if  $\text{lisnil } l$  then
      acquire  $lk$ ;
      select  $c \ \text{stop}$ ;
      release  $lk$ 
    else
      acquire  $lk$ ;
      select  $c \ \text{cont}$ ;
      send  $c \ (\text{lpop } l)$ ;
       $\text{cntr} \leftarrow !\text{cntr} + 1$ ;
      release  $lk$ ;
      go  $l \ \text{cntr } lk \ c$ 

recv_all  $\triangleq$ 
  rec go  $l \ n \ \text{cntr } lk \ c =$ 
    if  $n = 0$  then  $()$  else
      acquire  $lk$ ;
      if  $!\text{cntr} = 0$  then
        release  $lk$ ;
        go  $l \ n \ \text{cntr } lk \ c$ 
      else
        let  $x = \text{recv } c$  in
         $\text{cntr} \leftarrow !\text{cntr} - 1$ ;
        release  $lk$ ;
        go  $l \ (n - 1) \ \text{cntr } lk \ c$ ;
        lcons  $x \ l$ 

```

Figure 6. A producer-consumer client for the computation service. (The operations on lists `llength`, `lnil`, `lisnil`, and `lpop`, are standard and their code have thus been elided).

half of the fractional permission excludes the possibility of the full fraction being in the lock (and thereby state (iii)).

If the invariant is in state (i), the thread takes a step of the protocol and surrenders its fractional permission $\frac{1}{2}$ leaving the invariant in state (ii); if the invariant is in state (ii) a similar step is taken leaving the invariant with the full fractional permission 1 in state (iii).

4.2 A Parallel Computation Client

In §3.5 we considered the session type `compute_type` for a client of a polymorphic recursive computation service. We now consider a client `compute_client`, shown in Figure 6, which interacts with the service by sending a list of computation requests and receiving their results in parallel, similar to the producer-consumer pattern.³ We want to prove:

$$\Gamma \models \text{compute_client} : \text{list } (1 \multimap A) \multimap \text{chan compute_type} \multimap \text{list } A \dashv \Gamma$$

where $\text{list } A \triangleq \mu \text{rec. ref}_{\text{uniq}} (1 + (A \times \text{rec}))$.

The client `compute_client` operates on a channel endpoint c , where the computation service has the other endpoint. The client creates a shared counter cntr to keep track of the number of requests that are being processed, a linked list l' for the results, and a lock lk . It runs the producer `send_all` and consumer `recv_all` in parallel, which both race for the lock lk to access the channel c and counter cntr . The producer processes the input list l one-by-one by sending each computation in l on the channel c , and increasing the shared counter cntr thereafter. The consumer `recv_all` adds the results one-by-one to the list l' by receiving them on the channel c , and decreasing the shared counter cntr thereafter. When both the producer and consumer terminate, the client returns the list l' that then contains the results.

³For simplicity, our producer and consumer just iterate through a list, whereas in reality they would perform some computations so there is a point in having the producer and consumer operate in parallel.

The type system cannot type check `compute_client`, as (1) its safety depends on the length of the list, which is not available from the type, and (2) the channel c is shared and the type changes between each concurrent access. To prove that `compute_client` is semantically typed, we unfold its typing judgement, and resolve each step of the program in sequence, by applying the related weakest precondition rules. We first use the weakest precondition rule for `llength`:

$$\text{list } B \vdash \text{wp } \text{llength } l \left\{ n. n = |\vec{v}| * l \xrightarrow{\text{list}} \vec{v} \right\}$$

This rule converts the type predicate `list B` of the linked list l into the separation-logic list representation predicate $l \xrightarrow{\text{list}} \vec{v}$, which additionally makes the contents $\vec{v}$ of the linked list l explicit. This predicate is defined as follows:

$$l \xrightarrow{\text{list}} \vec{v} \triangleq \begin{cases} \ell \mapsto \text{inl } () & \text{if } \vec{v} = [] \\ \exists \ell_2. \ell \mapsto \text{inr } (v_1, \ell_2) * A \ v_1 * \ell_2 \xrightarrow{\text{list}} \vec{v}_2 & \text{if } \vec{v} = [v_1] \cdot \vec{v}_2 \end{cases}$$

The remainder of the proof is similar to the proof of the parallel receive in §4.1—we establish a lock invariant `chaninv` to share the counter cntr and the channel c between the producer and consumer, and use a fractional permission $\frac{1}{2}$ to determine the state of the shared channel c :

$$\begin{aligned} \text{chaninv} &\triangleq \exists n. \text{cntr} \mapsto n * \\ &\quad (c \multimap ((?A)^n \cdot \text{compute_type}) \vee \quad (i) \\ &\quad (c \multimap ((?A)^n \cdot \text{end}) * \frac{1}{2})) \quad (ii) \end{aligned}$$

The lock invariant states that the session type of the channel starts with a sequence of receive actions $(?A)^n$, where n is the value of the shared counter cntr . Here, the notation S^n denotes S appended to itself n times (the append operation $\cdot$ is inherited from Actris). The invariant expresses that either (i) the channel is still open, which permits unfolding the recursive definition to send additional requests, or (ii) the channel terminates with `end`, after the n receive steps have

been resolved. State (ii) requires the full fractional permission $\frac{1}{2}^y$, which must be released before closing the channel.

The proof is carried out by allocating the fractional permission $\frac{1}{2}^y$ (with a fresh identifier y), after which the weakest precondition rules for parallel composition (see Figure 2), the producer `send_all`, and consumer `recv_all` are used:

$$\begin{aligned} & \text{isLock } lk \text{ chaninv} * \frac{1}{2}^y * l \mapsto_{(1 \multimap A)} \vec{v} \multimap \\ & \quad \text{wp send_all } l \text{ cntr } lk \text{ } c \left\{ l \mapsto_{(1 \multimap A)} [] \right\} \\ & \text{isLock } lk \text{ chaninv} * l \mapsto_A [] \multimap \\ & \quad \text{wp recv_all } l \text{ n cntr } lk \text{ } c \left\{ \exists \vec{w}. |\vec{w}| = n * l \mapsto_A \vec{w} \right\} \end{aligned}$$

The proof of `send_all` proceeds as follows. Owning $\frac{1}{2}^y$ means the lock invariant is in state (i). Therefore, after unfolding the recursive tail and instantiating the polymorphic binder in the type of c , we have:

$$c \mapsto (?A)^n \cdot \oplus \left\{ \begin{array}{l} \text{cont} : !(1 \multimap A). ?A. \text{compute_type} \\ \text{stop} : \text{end} \end{array} \right\}$$

The select and send actions can then be swapped ahead of the receives, which results in:

$$c \mapsto \oplus \left\{ \begin{array}{l} \text{cont} : !(1 \multimap A). (?A)^{n+1} \cdot \text{compute_type} \\ \text{stop} : (?A)^n \cdot \text{end} \end{array} \right\}$$

If the list is non-empty, we resolve a computation step (by selecting the `cont` branch and sending a computation with type $1 \multimap A$) resulting in $c \mapsto (?A)^{n+1} \cdot \text{compute_type}$. After incrementing the shared counter `cntr`, we reestablish the lock invariant in state (i). If the list is empty, we close the channel (by selecting the `stop` branch), resulting in $c \mapsto (?A)^n \cdot \text{end}$. We reestablish the lock invariant in state (ii) by giving up the fractional permission $\frac{1}{2}^y$.

The proof of `recv_all` proceeds as follows. We only perform a receive operation when the shared counter `cntr` is positive, which means we have $c \mapsto ?A. (?A)^{n-1} \cdot S$. Here, S is `compute_type` or `end`, depending on whether the lock invariant is in state (i) or (ii), respectively. After the receive operation we have $c \mapsto (?A)^{n-1} \cdot S$, so after decrementing the shared counter `cntr` we can reestablish the lock invariant.

To finalise the proof of the client `compute_client`, we weaken $l \mapsto_A \vec{w}$ returned by `recv_all` to list $A \ l$ by forgetting about the contents $\vec{w}$ of the linked list l .

5 Mechanisation in Coq

In this paper, we have used what is often called the “foundational approach” to semantic type safety [1–3]. That means that contrary to conventional logical relation developments, types are not defined syntactically, and then given a semantic interpretation. Instead, types are defined as combinators in terms of their semantic interpretation. This approach gives rise to an “open” system that can easily be extended with new type formers, and is thus particularly suitable for mechanisation in a proof assistant like Coq. Furthermore, as we

will show in this section, the foundational approach makes it possible to reuse Coq’s variables to model type-level binding, avoiding boilerplate that would be necessary with a first-order representation of variable binding.

Our mechanisation is built on top of the mechanisation of Iris and Actris in Coq, which provides a number of noteworthy advantages. First, we can reuse their libraries for various programming constructs, such as locks (from Iris) and channels (from Actris). Second, we avoid reasoning about explicit resources in Coq by making use of the MoSeL framework (formerly, Iris Proof Mode), which provides tactics tailored for reasoning about the connectives of separation logic, and thereby hides unnecessary details related to the embedding of separation logic in Coq [32, 34].

Typing Judgments. Term and session types are represented as a dependent type indexed by a kind:⁴

Inductive kind := tty_kind | sty_kind. (* ★ or ♦ *)

Inductive lty Σ : kind $\rightarrow$ Type :=
| Ltty : (val $\rightarrow$ iProp Σ) $\rightarrow$ lty Σ tty_kind
| Lsty : iProto $\Sigma \rightarrow$ lty Σ sty_kind.

Notation ltty Σ := (lty Σ tty_kind).

Notation lsty Σ := (lty Σ sty_kind).

Typing contexts are represented as association lists:

Inductive ctx_item Σ := CtxItem {
ctx_item_name : string;
ctx_item_type : lty Σ }.
Notation ctx Σ := (list (ctx_item Σ)).

The semantic term typing judgement is defined as:

(* ltty_car: lty $\Sigma \rightarrow$ (val $\rightarrow$ iProp Σ) is the
inverse of Ltty *)
Definition ltyped ($\Gamma_1 \ \Gamma_2$: ctx Σ)
(e : expr) (A : lty Σ) : iProp Σ :=
■ $\forall$ vs, ctx_ltyped vs Γ_1 $\multimap$
WP subst_map vs e { { v,
ltty_car A v * ctx_ltyped vs Γ_2 } }.
Notation " $\Gamma_1 \models e : A \models \Gamma_2$ " :=
(ltyped $\Gamma_1 \ \Gamma_2$ e A) : bi_scope.
Notation " $\Gamma_1 \models e : A \models \Gamma_2$ " :=
($\vdash$ ltyped $\Gamma_1 \ \Gamma_2$ e A) : type_scope.

The typing judgement is defined for the deeply-embedded expressions `expr` of the (untyped) language `HeapLang`, which is the default language shipped with Iris, and is extended by the Actris framework with connectives for message passing. `HeapLang` use strings for variables, and hence our typing contexts `ctx` do that too. Compared to e.g., De Bruijn indices or locally nameless, the use of strings makes it possible to write programs in a human-readable way.⁵

⁴As is common in Iris, all definitions are parameterised by a Σ , which describes the resources that are available. For the purpose of this paper, this technicality can be ignored.

⁵Since `HeapLang`’s operational semantics is defined on closed terms, the use of strings does not cause issues with variable capture. See also [41, Section STLC] for a discussion on the use of strings for variables.

The typing judgement is identical to the definition in §2.3, but is defined as an internal notion in Iris, *i.e.*, it is an Iris proposition **iProp** instead of a Coq proposition **Prop**. This provides some additional flexibility in manual typing proofs. For example, it makes it possible to prove typing judgements using Löb induction, without having to unfold their definition. To ensure that the typing judgement can be used as an ordinary proposition of higher-logic in Iris, it contains the *plainly modality* ($\blacksquare$), which ensures that it does not capture any separation logic resources.⁶ We define two notations so that the typing judgement can be used internally and externally. The second notation uses the validity predicate of Iris ($\vdash$), which turns an **iProp** into a **Prop**.

Typing Lemmas. As an example of how a semantic typing rule looks like in Coq, consider the lemma corresponding to the typing rule for let-expressions:

```
Lemma ltyped_let  $\Gamma_1 \Gamma_2 \Gamma_3 x e_1 e_2 A_1 A_2 :$ 
  ( $\Gamma_1 \models e_1 : A_1 \dashv \vdash \Gamma_2$ )  $\rightarrow$ 
  ( $\text{ctx\_cons } x A_1 \Gamma_2 \models e_2 : A_2 \dashv \vdash \Gamma_3$ )  $\rightarrow$ 
  ( $\Gamma_1 \models (\text{let } x := e_1 \text{ in } e_2) : A_2 \dashv \vdash$ 
     $\text{ctx\_filter\_eq } x \Gamma_2 ++ \text{ctx\_filter\_ne } x \Gamma_3$ ).
```

The typing rule of let shows the handling of shadowing of variables: $\text{ctx_cons } x A_1 \Gamma_2$ removes all bindings of x from Γ_2 before adding the new binding, and $\text{ctx_filter_eq } x \Gamma_2$ makes sure that potentially overshadowed variables are preserved. Dealing with shadowing in the proof is trivial due to some general-purpose lemmas for $\Gamma \models \sigma$ ($\text{ctx_ltyped } \Gamma$ vs in Coq). The proof of the typing rule is 9 lines of Coq code.

The term type for kinded universal types is defined as:

```
Definition lty_forall {k}
  ( $C : \text{lty } \Sigma k \rightarrow \text{lty } \Sigma$ ) :  $\text{lty } \Sigma :=$ 
   $\text{Lty } (\lambda w, \forall X, \text{WP } w \#() \{ \{ \text{lty\_car } (C X) \} \} )$ .
Notation " $\forall X, C$ " := ( $\text{lty\_forall } (\lambda X, C)$ ):  $\text{lty\_scope}$ .

Lemma ltyped_tlam  $\Gamma_1 \Gamma_2 \Gamma' e k (C : \text{lty } \Sigma k \rightarrow \text{lty } \Sigma) :$ 
  ( $\forall K, \Gamma_1 \models e : C K \dashv \vdash []$ )  $\rightarrow$ 
  ( $\Gamma_1 ++ \Gamma_2 \models (\lambda : \langle \rangle, e) : (\forall X, C X) \dashv \vdash \Gamma_2$ ).
Lemma ltyped_tapp  $\Gamma \Gamma_2 e k (C : \text{lty } \Sigma k \rightarrow \text{lty } \Sigma) K :$ 
  ( $\Gamma \models e : (\forall X, C X) \dashv \vdash \Gamma_2$ )  $\rightarrow$ 
  ( $\Gamma \models e \#() : C K \dashv \vdash \Gamma_2$ ).
```

The universal type shows how the semantic approach allows binders to be modelled using Coq's binders. The argument c of lty_forall is a Coq function, and thus the binding in the notation $\forall X, c$ is simply achieved using a Coq lambda abstraction $\lambda X, c$. This approach gives the same feeling of working with higher-order abstract syntax [40], albeit being semantical instead of syntactical. The typing rule for type abstraction similarly uses Coq's binders, where the $\forall K$ in the premise implicitly ensures that K is fresh. The proof of the two typing rules are 4 and 3 lines of code, respectively.

The session type for selection ($\oplus$) and branching ($\&$) is:

```
Inductive action := Send | Recv.
Definition lty_choice (a : action)
  (Ss :  $\text{gmap } Z (\text{lty } \Sigma)$ ) :  $\text{lty } \Sigma :=$ 
   $\text{Lty } (\langle a@(i : Z) \rangle \text{MSG } \#x \{ \{ \ulcorner \text{is\_Some } (Ss !! i) \urcorner \} \} ;$ 
     $\text{lty\_car } (Ss !! i) \}$ ).
Notation lty_select := ( $\text{lty\_choice } \text{Send}$ ).
```

```
Lemma ltyped_select  $\Gamma x i S Ss :$ 
  ( $\Gamma !! x = \text{Some } (\text{chan } (\text{lty\_select } Ss)) \rightarrow$ 
     $Ss !! i = \text{Some } S \rightarrow$ 
     $\Gamma \models \text{select } x \#i : () \dashv \vdash \text{env\_cons } x (\text{chan } S) \Gamma$ ).
```

Since $\oplus$ and $\&$ are dual, this definition (as well as in many other dual definitions, lemmas, and proofs) are factorised using the inductive type `action`. The syntax $\langle a@(\vec{x}:\vec{\tau}) \rangle \text{MSG } v \{ \{ P \} \}$; *prot* expands to Actris's $!\vec{x} : \vec{\tau} \langle v \rangle \{ P \}$. *prot* or $?\vec{x} : \vec{\tau} \langle v \rangle \{ P \}$. *prot* depending on the action a . The definition uses the finite map library `gmap` of `std++` [45], to represent the choices Ss . The notation $Ss !! i$ is the lookup function on maps, whose result is only well-defined if i is in the map Ss , as required by $\text{is_Some } (Ss !! i)$. The notation $\ulcorner _ \urcorner$ embeds a Coq **Prop** into Iris. The typing rule for select requires the label i to be in the map Ss , and updates the channel s based on the label. The proof makes use of Actris's proof rules, and is 6 lines of code.

Type Safety. The type safety lemma is stated as follows:

```
Lemma ltyped_safety  $e \sigma \text{es } \sigma' e' :$ 
  ( $\exists A, [] \models e : A \dashv \vdash []$ )  $\rightarrow$ 
   $\text{rtc } \text{erased\_step } ([e], \sigma) (\text{es}, \sigma') \rightarrow e' \in \text{es} \rightarrow$ 
   $\text{is\_Some } (\text{to\_val } e') \vee \text{reducible } e' \sigma'$ .
```

This lemma states that if we have a typing judgment for a closed expression e , and we start execution of the single thread e to obtain a list of resulting threads es after any number of execution steps (modeled using the reflexive-transitive closure, `rtc`, of `HeapLang`'s small-step reduction relation), then any thread e' in es is either a value or can take a step.

6 Related Work

Session Types. Seminal work on subtyping for binary recursive session types for a synchronous pi-calculus was done by Gay and Hole [20]. Mostrous et al. [37] expand on this work by adding support for multi-party asynchronous recursive session types, and later for higher-order process calculi [36]. These two works present the session subtyping relation with inverted orientations, inverting the sub- and supertypes, which has been discussed by Gay [19]. Our semantic session subtyping relation uses the same orientation as Gay and Hole. Mostrous et al. [37] also present an output-input swapping rule, which inspired our swapping rule in §3.7, even though their type system is multi-party, as the idea is compatible with both session type variants. They additionally claim that their subtyping is decidable, it was later proven to not be the case by Bravetti et al. [7], precisely because of the swapping rule.

⁶The plainly modality ($\blacksquare$) is like the persistent modality ($\Box$), but additionally makes sure no persistent resources are captured.

Gay [18] introduced bounded polymorphic session types where branches contain type variables for term types with upper and lower bounds. This work neither supports recursive types, session subtyping, nor delegation, but Gay hypothesised that recursion could be done. Dardha et al. [13] expanded on this work by adding subtyping and delegation, while still only conjecturing that recursion was a possible extension. Caires et al. [8] devised a polymorphic session type system for the synchronous pi-calculus with existential and universal quantifiers at the type-level, but not at the session-level. However, like Gay's work, their system supports neither recursive types nor subtyping.

Thiemann and Vasconcelos [47] introduced label dependent session types, where tails can depend on the communicated message, which allows for encoding choice using send and receive. This is similar to the encoding of our semantic choice types in terms of Actris's dependent send and receive. While their work does not have asynchronous subtyping or polymorphism, it supports recursive types over natural numbers, with a recursor for type checking of such types.

Balzer and Pfenning [5] and Balzer et al. [6] proposed a session-type system that allows sharing of channels via locks. Their system contains unrestricted types that can be shared, linear types that cannot, and modalities to move between the two through the use of locks. Our mutex type works similarly with copyable types, but our system is more general, as the copyable types tie into Iris's general-purpose mechanisms for sharing. We can also impose mutexes on only one endpoint of a channel, while they require mutual locking on both ends, and integrate manual typing proofs of racy programs. They provide proofs for subject reduction and type preservation, not just to obtain type safety, but also to obtain deadlock freedom, which we do not consider.

Logical Relations. Logical relations have been studied extensively in the context of Iris, for type safety of type systems [22, 27, 34], program refinement [16, 34, 35, 44, 48], robust safety [43], and non-interference [17]. The most immediately related work in this area is the RustBelt project [27], which uses logical relations to prove type safety and datarace-freedom of a large subset of Rust and its standard libraries, focusing on Rust's lifetime and borrowing mechanism. RustBelt employs the foundational approach to logical relations in its Coq development, from which we have drawn much inspiration. Giarrusso et al. [22] used logical relations in Iris to prove type safety of a version of Scala's core calculus DOT, which has a rich notion of subtyping, but is different in nature from session subtyping.

The connection between logic and session types has been studied through the Curry-Howard correspondence by e.g., Caires and Pfenning [9], Wadler [49], Carbone et al. [10], and Dardha and Gay [12]. As part of this line of work, Perez et al. used logical relations to prove termination [38] and confluence [39] of session-based concurrent systems.

Mechanisation of Session Types. Mechanisations pertaining to session types are all fairly recent. There are two other mechanisations of session types in Iris. Tassarotti et al. [44] proved termination preserving refinements for a compiler from a session-typed language to a functional language where message buffers are modelled on the heap. Hinrichsen et al. [23, 24] developed the Actris mechanisation that this work is built on top of. Both lines of work focus on different properties than type safety.

Gay et al. [21] explored various notions of duality, mechanising their results in Agda, and demonstrate that allowing duality to distribute over the recursive μ -operator yields an unsound system when type variables appear in messages, as the message type could change in tandem with the dualisation of the recursion, making endpoints disagree on the type of exchanged values. In our setup duality does not distribute over μ . Instead recursive definitions must be unfolded to expose the session type before duality can be applied, rendering the recursion and message type unchanged. Even so, we can drop down to Actris and use Löb induction to prove (subtyping) properties of recursive types and their duals.

Castro et al. [11] focused on the metatheory of binary session types for synchronous communication, and prove in Coq, using the locally nameless approach to variable binding, subject reduction and that typing judgements are preserved by structural congruence.

Thiemann [46] mechanised an intrinsically-typed definitional interpreter for a session-typed language with recursive types and subtyping in Agda. The mechanisation did, however, require a substantial amount of manual bookkeeping, in particular for properties about resource separation. Rouvoet et al. [42] streamlined the intrinsically-typed approach by developing separation logic-like abstractions in Agda. They applied these abstractions to a small session-typed language without recursive types, subtyping, or polymorphism.

7 Conclusion

In this paper we demonstrate how the foundational semantic approach to type safety can be applied to session typing and how to construct and mechanise an extensible session-type system with support for manual typing proofs. The crux of the semantic approach is to use a program logic that is expressive enough to model all intended features (e.g., channels, subtyping, polymorphism, recursion, locks/mutexes) while satisfying the required properties (e.g., type safety). By building on top of the Iris and Actris frameworks we are able to inherit their constructs to mechanise such an extensible session-type system with little proof effort.

Acknowledgments

We thank the anonymous reviewers for their helpful feedback. Robbert Krebbers was supported by the Dutch Research Council (NWO), project 016.Veni.192.259.

A Type System

This appendix includes an extensive overview of the mechanised semantic session-type system. Like the paper, all of the definitions and rules have been mechanised in the Coq proof assistant, and can be found in [25].

In particular, the appendix shows the type and judgement definitions in Figure 7, the typing rules in Figures 8 and 9, and the subtyping rules in Figures 11 and 12.

As some of the details of the type system were omitted in the main text, we preface the overview with a cursory clarification of these. In particular, we introduce a streamlined approach for handling copyable versus uncopyable types, which allows unifying various typing rules (Appendix A.1). We furthermore describe kinded subtyping and type equivalence (Appendix A.2), shared reference types (Appendix A.3), and discuss the internal versions of all judgements of the type system (Appendix A.4).

We omit the typing rule for polymorphic sends from §3.5 because it can be derived from the original rule for send (TY-CHANSEND) along with the subsumption (TY-SUB) and the subtyping for instantiating the binders of the send session type (SUBTY-SEND-IN).

A.1 Uncopy

To handle copyable types, one typically has two rules for each construct that might move out ownership (one for non-copyable types and one for copyable types). For example:

$$\begin{array}{c} \text{TY-REFUNIQLoad-Move} \\ \Gamma, (x : \text{ref}_{\text{uniq}} A) \models !x : A \dashv \Gamma, (x : \text{ref}_{\text{uniq}} \text{any}) \\ \\ \text{TY-REFUNIQLoad-Copy} \\ \text{copyable } A \\ \hline \Gamma, (x : \text{ref}_{\text{uniq}} A) \models !x : A \dashv \Gamma, (x : \text{ref}_{\text{uniq}} A) \end{array}$$

The full version of our type system unifies these rules as the single rule TY-REFUNIQLoad using the uncopy type former. The uncopy type former acts as an inverse of the copy type former. When uncopy is applied to copy A , copy and uncopy cancel out, leaving the type A , as expressed by the subtyping rule SUBTY-UNCOPY-ELIM. In combination with the rule SUBTY-COPY-INTRO, this means that the uncopy type former has no effect on copyable types A , *i.e.*, if copyable A , then $\text{uncopy } A <: A$. However, when applied to a non-copyable type A , the uncopy type former has an effect, and thus cannot be stripped. This prevents the value from being used again, similarly to replacing the type by any.

The uncopy type former is defined in terms of the coreP modality of Iris (coreP itself is defined in terms of other logical primitives), which acts as a similar “inverse” to the persistence modality ($\Box$). The definition and proof rules of the coreP modality can be found at <https://gitlab.mpi-sws.org/iris/iris/-/blob/master/theories/bi/lib/core.v>.

A.2 Kinded Subtyping and Type Equivalence

The subtyping relation $<:$ is kinded, *i.e.*, it takes arguments of type Type_k and its definition depends on the kind k . By making the subtyping relation kinded, we can unify subtyping rules that are identical for both type kinds, such as the rule SUBTY-REFL for reflexivity.

Additionally, to unify subtyping rules that go in both directions, such as the rule SUBTY-REC-UNFOLD for unfolding recursive types, we define a relation for *type equivalence* $K <:> L$ as the symmetric closure of the subtyping relation:

$$K <:> L \triangleq K <: L \wedge L <: K$$

Similar to the subtyping relation, the relation for type equivalence is kinded so it applies to both term and session types.

A.3 Shared References

We also have an additional type former ref_{shr} , which is not mentioned in the main text of the paper. This is the type of *shared references*, or references that can be freely duplicated and shared between threads, but whose type is not allowed to change by writing new values. Moreover, shared references can only hold values of a copyable type, to prevent values from being copied by reading and writing to a reference.

The definition of the type former ref_{shr} for shared references is standard in logical relation developments in Iris. It is defined in terms of Iris *invariants*, written $\boxed{P}$, which contain a proposition P . Invariants are always persistent (even if the proposition P itself is not), meaning they can be freely duplicated. Moreover, it is possible to *open* an invariant to gain access to the proposition P inside, as long as that is restricted to an atomic program step, and the invariant is re-established by reproving P at the end of the atomic step. In practice, this means that it is only possible to apply atomic read and write operations to shared references, and the fact that invariants must be re-established ensures that we cannot *change* the type of the value contained in the reference, in contrast to the store rule for unique references ref_{uniq} .

A.4 Internal Judgements

In §5 we remarked that in the Coq mechanisation we defined the typing judgement as an internal definition in Iris, instead of as an external definition in the meta logic. In the full version of the type system, we use the same treatment for the typing judgements. To make sure that the judgements behave like ordinary propositions of higher-order logic (instead of propositions that hold ownership), their definitions include the *plainly* modality ($\blacksquare$). This modality carves out the step-indexed subset of the Iris logic. The rules of the plainly modality can be found in <https://gitlab.mpi-sws.org/iris/iris/-/blob/master/theories/bi/plainly.v>.

As a result of defining all judgements as internal notions, all typing rules are in fact implications in the Iris logic.

Term Types:

$\text{Type}_\star \triangleq \text{Val} \rightarrow \text{iProp}$
 $\text{any} \triangleq \lambda w. \text{True}$
 $1 \triangleq \lambda w. w \in \{()\}$
 $\mathbb{B} \triangleq \lambda w. w \in \mathbb{B}$
 $\mathbb{Z} \triangleq \lambda w. w \in \mathbb{Z}$
 $\text{ref}_{\text{uniq}} A \triangleq \lambda w. \exists v. w \in \text{Loc} * (w \mapsto v) * \triangleright(Av)$
 $\text{ref}_{\text{shr}} A \triangleq \lambda w. (w \in \text{Loc}) * \boxed{\exists v. (w \mapsto v) * \Box(Av)}$
 $A_1 \times A_2 \triangleq \lambda w. \exists w_1, w_2. w = (w_1, w_2) * \triangleright(A_1 w_1) * \triangleright(A_2 w_2)$
 $A_1 + A_2 \triangleq \lambda w. \exists v. (w = \text{inl } v * \triangleright(A_1 v)) \vee (w = \text{inr } v * \triangleright(A_2 v))$
 $A \multimap B \triangleq \lambda w. \forall v. \triangleright(Av) \multimap \text{wp } (w v) \{B\}$
 $\text{chan } S \triangleq \lambda w. w \multimap S$
 $\text{copy } A \triangleq \lambda w. \Box(Aw)$
 $A \rightarrow B \triangleq \text{copy } (A \multimap B)$
 $\text{uncopy } A \triangleq \lambda w. \text{coreP } (Aw)$
 $\mu(X : k). K \triangleq \mu(X : \text{Type}_k). K \quad (K \text{ is contractive in } X)$
 $\forall(X : k). A \triangleq \lambda w. \forall(X : \text{Type}_k). \text{wp } (w ()) \{A\}$
 $\exists(X : k). A \triangleq \lambda w. \exists(X : \text{Type}_k). \triangleright(Aw)$
 $\text{mutex } A \triangleq \lambda w. \exists lk, \ell. (w = (lk, \ell)) * \text{isLock } lk (\exists v. (\ell \mapsto v) * \triangleright(Av))$
 $\overline{\text{mutex}} A \triangleq \lambda w. \exists lk, \ell. (w = (lk, \ell)) * (\ell \mapsto -) * \text{isLock } lk (\exists v. (\ell \mapsto v) * \triangleright(Av))$

Typing Judgement:

$\Gamma \models \sigma \triangleq *_{(x,A) \in \Gamma}. A(\sigma(x))$
 $\Gamma \models e : A \dashv \Gamma' \triangleq \blacksquare(\forall \sigma. (\Gamma \models \sigma) \multimap \text{wp } e[\sigma] \{v.Av * (\Gamma' \models \sigma)\})$

Session Types:

$\text{Type}_\diamond \triangleq \text{iProto}$
 $\text{end} \triangleq \text{end}$
 $!A.S \triangleq ! (v : \text{Val}) \langle v \rangle \{Av\}. S$
 $?A.S \triangleq ? (v : \text{Val}) \langle v \rangle \{Av\}. S$
 $!\vec{X}:\vec{k}. A.S \triangleq ! (\vec{X} : \vec{\text{Type}}_k) (v : \text{Val}) \langle v \rangle \{Av\}. S$
 $? \vec{X}:\vec{k}. A.S \triangleq ? (\vec{X} : \vec{\text{Type}}_k) (v : \text{Val}) \langle v \rangle \{Av\}. S$
 $\oplus\{\vec{S}\} \triangleq ! (l : \mathbb{Z}) \langle l \rangle \{l \in \text{dom}(\vec{S})\}. \vec{S}(l)$
 $\&\{\vec{S}\} \triangleq ? (l : \mathbb{Z}) \langle l \rangle \{l \in \text{dom}(\vec{S})\}. \vec{S}(l)$

Subtyping:

$A <: B \triangleq \blacksquare(\forall v. Av \multimap Bv)$
 $S <: T \triangleq \blacksquare(S \sqsubseteq T)$
 $K <:> L \triangleq K <: L \wedge L <: K$
 $\Gamma <:\text{ctx } \Gamma' \triangleq \blacksquare(\forall \sigma. (\Gamma \models \sigma) \multimap (\Gamma' \models \sigma))$

Other:

$\text{copyable } A \triangleq A <: \text{copy } A$

Figure 7. Typing judgements and type formers.**Basics:**

$\text{TY-UNIT} \quad \Gamma \models () : 1 \dashv \Gamma$
 $\text{TY-BOOL} \quad \Gamma \models b : \mathbb{B} \dashv \Gamma$
 $\text{TY-INT} \quad \Gamma \models i : \mathbb{Z} \dashv \Gamma$
 $\text{TY-NEG} \quad \frac{\Gamma \models e : B \dashv \Gamma'}{\Gamma \models \neg e : B \dashv \Gamma'}$
 $\text{TY-ARITH} \quad \frac{\Gamma \models e_2 : Z \dashv \Gamma' \quad \Gamma' \models e_1 : Z \dashv \Gamma'' \quad \text{op} \in \{+, -\}}{\Gamma \models e_1 \text{ op } e_2 : Z \dashv \Gamma''}$
 $\text{TY-COND} \quad \frac{\Gamma \models e_2 : Z \dashv \Gamma' \quad \Gamma' \models e_1 : Z \dashv \Gamma'' \quad \text{op} \in \{=, \leq\}}{\Gamma \models e_1 \text{ op } e_2 : B \dashv \Gamma''}$
 $\text{TY-IF} \quad \frac{\Gamma \models e_1 : B \dashv \Gamma' \quad \Gamma' \models e_2 : A \dashv \Gamma'' \quad \Gamma' \models e_3 : A \dashv \Gamma''}{\Gamma \models \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : A \dashv \Gamma''}$
 $\text{TY-REC} \quad \frac{\Gamma = (x_1 : A_1), \dots, (x_n : A_n) \quad \Gamma_{\text{copy}} = (x_1 : \text{uncopy } A_1), \dots, (x_n : \text{uncopy } A_n) \quad \Gamma_{\text{copy}}, (f : A \rightarrow B), (x : A) \models e : B \dashv \Gamma''}{\Gamma \cdot \Gamma' \models \text{rec } f \ x = e : A \rightarrow B \dashv \Gamma'}$
 $\text{TY-LAM} \quad \frac{\Gamma, (x : A) \models e : B \dashv \Gamma''}{\Gamma \cdot \Gamma' \models \lambda x. e : A \multimap B \dashv \Gamma'}$
 $\text{TY-VAR} \quad \Gamma, (x : A) \models x : A \dashv \Gamma, (x : \text{uncopy } A)$
 $\text{TY-APP} \quad \frac{\Gamma \models e_2 : A \dashv \Gamma' \quad \Gamma' \models e_1 : A \multimap B \dashv \Gamma''}{\Gamma \models e_1 e_2 : B \dashv \Gamma''}$
 $\text{TY-LET} \quad \frac{\Gamma_1 \models e_1 : A \dashv \Gamma_2 \quad \Gamma_2, (x : A) \models e_2 : B \dashv \Gamma_3}{\Gamma_1 \models \text{let } x = e_1 \text{ in } e_2 : B \dashv \Gamma_3 \setminus x}$
 $\text{TY-PAR} \quad \frac{\Gamma_1 \models e_1 : A_1 \dashv \Gamma'_1 \quad \Gamma_2 \models e_2 : A_2 \dashv \Gamma'_2}{\Gamma_1 \cdot \Gamma_2 \models e_1 \parallel e_2 : A_1 \times A_2 \dashv \Gamma'_1 \cdot \Gamma'_2}$
 $\text{TY-SUB} \quad \frac{\Gamma_1 <:\text{ctx } \Gamma'_1 \quad \Gamma'_1 \models e : A \dashv \Gamma'_2 \quad A <: B \quad \Gamma'_2 <:\text{ctx } \Gamma_2}{\Gamma_1 \models e : B \dashv \Gamma_2}$

Figure 8. Term typing rules.

Product and Sums:

$$\begin{array}{c}
\text{TY-PAIR} \\
\frac{\Gamma \models e_2 : A_2 \ni \Gamma' \quad \Gamma' \models e_1 : A_1 \ni \Gamma''}{\Gamma \models (e_1, e_2) : A_1 \times A_2 \ni \Gamma''}
\end{array}
\quad
\begin{array}{c}
\text{TY-INL} \\
\frac{\Gamma \models e : A \ni \Gamma'}{\Gamma \models \text{inl } e : A + B \ni \Gamma'}
\end{array}
\quad
\begin{array}{c}
\text{TY-INR} \\
\frac{\Gamma \models e : B \ni \Gamma'}{\Gamma \models \text{inr } e : A + B \ni \Gamma'}
\end{array}$$

$$\begin{array}{c}
\text{TY-FST} \\
\Gamma, (x : A_1 \times A_2) \models \text{fst } x : A_1 \ni \Gamma, (x : \text{uncopy } A_1 \times A_2)
\end{array}
\quad
\begin{array}{c}
\text{TY-SND} \\
\Gamma, (x : A_1 \times A_2) \models \text{snd } x : A_2 \ni \Gamma, (x : A_1 \times \text{uncopy } A_2)
\end{array}$$

$$\begin{array}{c}
\text{TY-CASE} \\
\frac{\Gamma \models e_1 : A + B \ni \Gamma' \quad \Gamma' \models e_2 : A \multimap C \ni \Gamma'' \quad \Gamma' \models e_3 : B \multimap C \ni \Gamma''}{\Gamma \models \text{case } e_1 \text{ } e_2 \text{ } e_3 : C \ni \Gamma''}
\end{array}$$

Polymorphism:

$$\begin{array}{c}
\text{TY-TLAM} \\
\frac{\Gamma \models e : A \ni \Gamma' \quad X \notin FV(\Gamma, \Gamma')}{\Gamma \cdot \Gamma' \models \lambda _ . e : \forall X. A \ni \Gamma'}
\end{array}
\quad
\begin{array}{c}
\text{TY-TAPP} \\
\frac{\Gamma \models e : \forall X. A \ni \Gamma'}{\Gamma \models e () : A[K/X] \ni \Gamma'}
\end{array}
\quad
\begin{array}{c}
\text{TY-PACK} \\
\frac{\Gamma \models e : A[K/X] \ni \Gamma'}{\Gamma \models e : \exists X. A \ni \Gamma'}
\end{array}$$

$$\begin{array}{c}
\text{TY-UNPACK} \\
\frac{\Gamma \models e_1 : \exists X. A \ni \Gamma' \quad \Gamma', (x : A) \models e_2 : B \ni \Gamma'' \quad X \notin FV(\Gamma, \Gamma'', B)}{\Gamma \models \text{let } x = e_1 \text{ in } e_2 : B \ni \Gamma'' \setminus x}
\end{array}$$

References:

$$\begin{array}{c}
\text{TY-TOREFSHR} \\
\frac{\Gamma \models e : \text{ref}_{\text{uniq}}(\text{copy } A) \ni \Gamma'}{\Gamma \models e : \text{ref}_{\text{shr}} A \ni \Gamma'}
\end{array}
\quad
\begin{array}{c}
\text{TY-REFSHRLOAD} \\
\frac{\Gamma \models e : \text{ref}_{\text{shr}} A \ni \Gamma'}{\Gamma \models !e : A \ni \Gamma'}
\end{array}
\quad
\begin{array}{c}
\text{TY-REFSHRSTORE} \\
\frac{\Gamma \models e_2 : \text{copy } A \ni \Gamma' \quad \Gamma' \models e_1 : \text{ref}_{\text{shr}} A \ni \Gamma''}{\Gamma \models e_1 \leftarrow e_2 : 1 \ni \Gamma''}
\end{array}$$

$$\begin{array}{c}
\text{TY-REFUNIQUALLOC} \\
\frac{\Gamma \models e : A \ni \Gamma'}{\Gamma \models \text{ref } e : \text{ref}_{\text{uniq}} A \ni \Gamma'}
\end{array}
\quad
\begin{array}{c}
\text{TY-REFUNIQUFREE} \\
\frac{\Gamma \models e : \text{ref}_{\text{uniq}} A \ni \Gamma'}{\Gamma \models \text{free } e : 1 \ni \Gamma'}
\end{array}
\quad
\begin{array}{c}
\text{TY-REFUNIQUSTORE} \\
\frac{\Gamma \models e : B \ni \Gamma', (x : \text{ref}_{\text{uniq}} A)}{\Gamma \models x \leftarrow e : 1 \ni \Gamma', (x : \text{ref}_{\text{uniq}} B)}
\end{array}$$

$$\begin{array}{c}
\text{TY-REFUNIQLoad} \\
\Gamma, (x : \text{ref}_{\text{uniq}} A) \models !x : A \ni \Gamma, (x : \text{ref}_{\text{uniq}}(\text{uncopy } A))
\end{array}$$

Channels:

$$\begin{array}{c}
\text{TY-CHANALLOC} \\
\Gamma \models \text{new_chan} : 1 \rightarrow \text{chan } S \times \text{chan } \bar{S} \ni \Gamma
\end{array}
\quad
\begin{array}{c}
\text{TY-CHANSEND} \\
\frac{\Gamma \models e : A \ni \Gamma', (x : \text{chan } (!A. S))}{\Gamma \models \text{send } x \text{ } e : 1 \ni \Gamma', (x : \text{chan } S)}
\end{array}$$

$$\begin{array}{c}
\text{TY-CHANRECV} \\
\Gamma, (x : \text{chan } (?A. S)) \models \text{recv } x : A \ni \Gamma, (x : \text{chan } S)
\end{array}
\quad
\begin{array}{c}
\text{TY-CHANRECVPOLY} \\
\frac{\Gamma, (x : \text{chan } S), (y : A) \models e : B \ni \Gamma' \quad \vec{X} \notin FV(\Gamma, \Gamma', B)}{\Gamma, (x : \text{chan } (?_{\vec{X}. \vec{k}} A. S)) \models \text{let } y = \text{recv } x \text{ in } e : B \ni \Gamma' \setminus \{y\}}
\end{array}$$

$$\begin{array}{c}
\text{TY-SELECT} \\
\frac{1 \leq i \leq n}{\Gamma, (x : \text{chan } (\oplus \{l_1 : S_1, \dots, l_n : S_n\})) \models \text{select } x \text{ } l_i : 1 \ni \Gamma, (x : \text{chan } S_i)}
\end{array}$$

$$\begin{array}{c}
\text{TY-BRANCH} \\
\frac{\Gamma, (x : \text{chan } S_1) \models e_1 : A \ni \Gamma' \quad \dots \quad \Gamma, (x : \text{chan } S_n) \models e_n : A \ni \Gamma'}{\Gamma, (x : \text{chan } (\& \{l_1 : S_1, \dots, l_n : S_n\})) \models \text{branch } x \text{ with } l_1 \Rightarrow e_1 \mid \dots \mid l_n \Rightarrow e_n : A \ni \Gamma'}
\end{array}$$

Figure 9. Term typing rules (cont.)

Locks:

Ty-MutexAlloc $\Gamma \models \text{newmutex} : A \rightarrow \text{mutex } A \models \Gamma$	Ty-MutexAcquire $\Gamma, (x : \text{mutex } A) \models \text{acquiremutex } x : A \models \Gamma, (x : \overline{\text{mutex } A})$
Ty-MutexRelease $\frac{\Gamma \models e : A \models \Gamma', (x : \overline{\text{mutex } A})}{\Gamma \models \text{releasemutex } x \ e : 1 \models \Gamma', (x : \text{mutex } A)}$	

Figure 10. Term typing rules (cont.)**Subtyping Properties:**

SubTy-Refl $\frac{}{K <: K}$	SubTy-Trans $\frac{K <: L \quad L <: M}{K <: M}$	SubTy-Bi $\frac{K <: L \quad L <: K}{K <:> L}$	SubTy-Bi-Refl $\frac{}{K <:> K}$	SubTy-Bi-Trans $\frac{K <:> L \quad L <:> M}{K <:> M}$
$\text{SubTy-Bi-Trans-Left}$ $\frac{K <:> L \quad L <: M}{K <: M}$	$\text{SubTy-Bi-Trans-Right}$ $\frac{K <: L \quad L <:> M}{K <: M}$	SubTy-Bi-Sym $\frac{L <:> K}{K <:> L}$	SubTy-Rec-Unfold $\mu X. K <:> K(\mu X. K)$	

Term Subtyping:

SubTy-Any $\frac{}{A <: \text{any}}$	SubTy-Lolli $\frac{C <: A \quad B <: D}{A \multimap B <: C \multimap D}$	SubTy-Arr $\frac{C <: A \quad B <: D}{A \rightarrow B <: C \rightarrow D}$	SubTy-Product $\frac{A <: C \quad B <: D}{A \times B <: C \times D}$	SubTy-Sum $\frac{A <: C \quad B <: D}{A + B <: C + D}$
SubTy-Forall $\frac{\forall X. (A <: B)}{\forall X. A <: \forall X. B}$	SubTy-Exist $\frac{\forall X. (A <: B)}{\exists X. A <: \exists X. B}$	SubTy-Exist-Elim $\frac{}{A[K/X] <: \exists X. A}$	SubTy-Ref-Uniq $\frac{A <: B}{\text{ref}_{\text{uniq}} A <: \text{ref}_{\text{uniq}} B}$	SubTy-Ref-Shr $\frac{A <: B}{\text{ref}_{\text{shr}} A <: \text{ref}_{\text{shr}} B}$
SubTy-Chan $\frac{S <: T}{\text{chan } S <: \text{chan } T}$	SubTy-Mutex $\frac{A <:> B}{\text{mutex } A <: \text{mutex } B}$	SubTy-MutexGuard $\frac{A <:> B}{\text{mutex } A <: \text{mutex } B}$		

Copyable Types:

SubTy-Copy $\frac{A <: B}{\text{copy } A <: \text{copy } B}$	SubTy-Copy-Intro $\frac{}{\text{copyable } A}$	SubTy-Copy-Elim $\frac{}{\text{copy } A <: A}$	SubTy-Uncopy $\frac{A <: B}{\text{uncopy } A <: \text{uncopy } B}$	$\text{SubTy-Uncopy-Intro}$ $\frac{}{A <: \text{uncopy } A}$
SubTy-Uncopy-Elim $\text{uncopy } (\text{copy } A) <: A$	$\text{SubTy-Copyable-Copy}$ $\text{copyable } (\text{copy } A)$	$\text{SubTy-Copyable-Uncopy}$ $\text{copyable } (\text{uncopy } A)$	$\text{SubTy-Copyable-Any}$ copyable any	
$\text{SubTy-Copyable-Unit}$ $\text{copyable } 1$	$\text{SubTy-Copyable-Bool}$ $\text{copyable } B$	$\text{SubTy-Copyable-Int}$ $\text{copyable } Z$	$\text{SubTy-Copyable-Product}$ $\frac{\text{copyable } A \quad \text{copyable } B}{\text{copyable } (A \times B)}$	
$\text{SubTy-Copyable-Sum}$ $\frac{\text{copyable } A \quad \text{copyable } B}{\text{copyable } (A + B)}$	$\text{SubTy-Copyable-Exists}$ $\frac{\forall X. \text{copyable } A}{\text{copyable } (\exists X. A)}$	$\text{SubTy-Copyable-RefShr}$ $\text{copyable } (\text{ref}_{\text{shr}} X)$	$\text{SubTy-Copyable-Mutex}$ $\text{copyable } (\text{mutex } X)$	

Figure 11. Subtyping rules.

Context Subtyping:

$\frac{\text{CTX-PERMUTE} \quad \Gamma' \text{ is a permutation of } \Gamma}{\Gamma <:\text{ctx} \Gamma'}$	$\frac{\text{CTX-REFL}}{\Gamma <:\text{ctx} \Gamma}$	$\frac{\text{CTX-TRANS} \quad \Gamma_1 <:\text{ctx} \Gamma_2 \quad \Gamma_2 <:\text{ctx} \Gamma_3}{\Gamma_1 <:\text{ctx} \Gamma_3}$	$\frac{\text{CTX-NIL}}{\Gamma <:\text{ctx} []}$	$\frac{\text{CTX-CONS} \quad A <: B \quad \Gamma <:\text{ctx} \Gamma'}{(x:A), \Gamma <:\text{ctx} (x:B), \Gamma'}$
$\frac{\text{CTX-APP} \quad \Gamma_1 <:\text{ctx} \Gamma_2 \quad \Gamma'_1 <:\text{ctx} \Gamma'_2}{\Gamma_1 \cdot \Gamma'_1 <:\text{ctx} \Gamma_2 \cdot \Gamma'_2}$	$\frac{\text{CTX-COPY} \quad (x:A) <:\text{ctx} (x:A), (x:\text{uncopy } A)}{(x:A) <:\text{ctx} (x:A), (x:A)}$	$\frac{\text{CTX-COPYABLE} \quad \text{copyable } A}{(x:A) <:\text{ctx} (x:A), (x:A)}$		

Session Subtyping:

$\frac{\text{SUBTY-SEND} \quad B <: A \quad S <: T}{!A.S <: !B.T}$	$\frac{\text{SUBTY-RECV} \quad A <: B \quad S <: T}{?A.S <: ?B.T}$	$\frac{\text{SUBTY-SEND-IN}}{!_{(\vec{X}:\vec{k})} A.S <: !A[\vec{K}/\vec{X}].S[\vec{K}/\vec{X}]}$	$\frac{\text{SUBTY-RECV-IN}}{?A[\vec{K}/\vec{X}].S[\vec{K}/\vec{X}] <: ?_{(\vec{X}:\vec{k})} A.S}$
$\frac{\text{SUBTY-SEND-OUT} \quad S <: !A.T}{S <: !_{(\vec{X}:\vec{k})} A.T}$	$\frac{\text{SUBTY-RECV-OUT} \quad ?A.S <: T}{?_{(\vec{X}:\vec{k})} A.S <: T}$	$\frac{\text{SUBTY-SELECT} \quad \forall i. \vec{S}_i <: \vec{T}_i}{\oplus\{\vec{l}_i : \vec{S}_i\}_{i \in \vec{i}} <: \oplus\{\vec{l}_i : \vec{T}_i\}_{i \in \vec{i}}}$	$\frac{\text{SUBTY-SELECT-SUBSETEQ} \quad \vec{j} \subseteq \vec{i}}{\oplus\{\vec{l}_i : \vec{S}_i\}_{i \in \vec{i}} <: \oplus\{\vec{l}_j : \vec{S}_j\}_{j \in \vec{j}}}$
$\frac{\text{SUBTY-BRANCH} \quad \forall i. \vec{S}_i <: \vec{T}_i}{\&\{\vec{l}_i : \vec{S}_i\}_{i \in \vec{i}} <: \&\{\vec{l}_i : \vec{T}_i\}_{i \in \vec{i}}}$	$\frac{\text{SUBTY-BRANCH-SUBSETEQ} \quad \vec{i} \subseteq \vec{j}}{\&\{\vec{l}_i : \vec{S}_i\}_{i \in \vec{i}} <: \&\{\vec{l}_j : \vec{S}_j\}_{j \in \vec{j}}}$	$\frac{\text{SUBTY-SWAP-RECV-SEND}}{?A. !B.S <: !B. ?A.S}$	
$\frac{\text{SUBTY-SWAP-BRANCH-SEND} \quad \&\{l_1 : !A.S_1, \dots, l_n : !A.S_n\} <: !A. \&\{l_1 : S_1, \dots, l_n : S_n\}}{\&\{l_1 : !A.S_1, \dots, l_n : !A.S_n\} <: !A. \&\{l_1 : S_1, \dots, l_n : S_n\}}$	$\frac{\text{SUBTY-SWAP-RECV-SELECT} \quad ?A. \oplus\{l_1 : S_1, \dots, l_n : S_n\} <: \oplus\{l_1 : ?A.S_1, \dots, l_n : ?A.S_n\}}{?A. \oplus\{l_1 : S_1, \dots, l_n : S_n\} <: \oplus\{l_1 : ?A.S_1, \dots, l_n : ?A.S_n\}}$		
$\frac{\text{SUBTY-SWAP-BRANCH-SELECT} \quad \&\{l_1 : \oplus\{l'_1 : S_{(1,1)}, \dots, l'_m : S_{(1,m)}\}, \dots, l_n : \oplus\{l'_1 : S_{(n,1)}, \dots, l'_m : S_{(n,m)}\}\} <: \&\{l'_1 : \&\{l_1 : S_{(1,1)}, \dots, l_n : S_{(n,1)}\}, \dots, l'_m : \&\{l_1 : S_{(n,1)}, \dots, l_n : S_{(n,m)}\}\}}{\&\{l_1 : \oplus\{l'_1 : S_{(1,1)}, \dots, l'_m : S_{(1,m)}\}, \dots, l_n : \oplus\{l'_1 : S_{(n,1)}, \dots, l'_m : S_{(n,m)}\}\} <: \&\{l'_1 : \&\{l_1 : S_{(1,1)}, \dots, l_n : S_{(n,1)}\}, \dots, l'_m : \&\{l_1 : S_{(n,1)}, \dots, l_n : S_{(n,m)}\}\}}$			

Append Subtyping:

$\frac{\text{SUBTY-APP} \quad S <: U \quad T <: V}{S \cdot T <: U \cdot V}$	$\frac{\text{SUBTY-APP-ASSOC} \quad S \cdot (T \cdot U) <:> (S \cdot T) \cdot U}{S \cdot (T \cdot U) <:> (S \cdot T) \cdot U}$	$\frac{\text{SUBTY-APP-SEND} \quad (!_{\vec{X}} A.S) \cdot T <:> !_{\vec{X}} A.(S \cdot T)}{(!_{\vec{X}} A.S) \cdot T <:> !_{\vec{X}} A.(S \cdot T)}$	$\frac{\text{SUBTY-APP-RECV} \quad (?_{\vec{X}} A.S) \cdot T <:> ?_{\vec{X}} A.(S \cdot T)}{(?_{\vec{X}} A.S) \cdot T <:> ?_{\vec{X}} A.(S \cdot T)}$
$\frac{\text{SUBTY-APP-SELECT} \quad (\oplus\{l_1 : S_1, \dots, l_n : S_n\}) \cdot T <:> \oplus\{l_1 : S_1 \cdot T, \dots, l_n : S_n \cdot T\}}{(\oplus\{l_1 : S_1, \dots, l_n : S_n\}) \cdot T <:> \oplus\{l_1 : S_1 \cdot T, \dots, l_n : S_n \cdot T\}}$	$\frac{\text{SUBTY-APP-BRANCH} \quad (\&\{l_1 : S_1, \dots, l_n : S_n\}) \cdot T <:> \&\{l_1 : S_1 \cdot T, \dots, l_n : S_n \cdot T\}}{(\&\{l_1 : S_1, \dots, l_n : S_n\}) \cdot T <:> \&\{l_1 : S_1 \cdot T, \dots, l_n : S_n \cdot T\}}$		
$\frac{\text{SUBTY-APP-END-RIGHT} \quad S \cdot \text{end} <:> S}{S \cdot \text{end} <:> S}$	$\frac{\text{SUBTY-APP-END-LEFT} \quad \text{end} \cdot S <:> S}{\text{end} \cdot S <:> S}$		

Duality Subtyping:

$\frac{\text{SUBTY-DUAL} \quad T <: S}{\overline{S} <: \overline{T}}$	$\frac{\text{SUBTY-DUAL-LEFT} \quad \overline{T} <: S}{\overline{S} <: T}$	$\frac{\text{SUBTY-DUAL-RIGHT} \quad T <: \overline{S}}{S <: \overline{T}}$	$\frac{\text{SUBTY-DUAL-SEND} \quad !_{\vec{X}} A.S <:> ?_{\vec{X}} A.\overline{S}}{!_{\vec{X}} A.S <:> ?_{\vec{X}} A.\overline{S}}$	$\frac{\text{SUBTY-DUAL-RECV} \quad ?_{\vec{X}} A.S <:> !_{\vec{X}} A.\overline{S}}{?_{\vec{X}} A.S <:> !_{\vec{X}} A.\overline{S}}$
$\frac{\text{SUBTY-DUAL-SELECT} \quad \oplus\{l_1 : S_1, \dots, l_n : S_n\} <:> \&\{l_1 : \overline{S}_1, \dots, l_n : \overline{S}_n\}}{\oplus\{l_1 : S_1, \dots, l_n : S_n\} <:> \&\{l_1 : \overline{S}_1, \dots, l_n : \overline{S}_n\}}$	$\frac{\text{SUBTY-DUAL-BRANCH} \quad \&\{l_1 : S_1, \dots, l_n : S_n\} <:> \oplus\{l_1 : \overline{S}_1, \dots, l_n : \overline{S}_n\}}{\&\{l_1 : S_1, \dots, l_n : S_n\} <:> \oplus\{l_1 : \overline{S}_1, \dots, l_n : \overline{S}_n\}}$	$\frac{\text{SUBTY-DUAL-END} \quad \text{end} <:> \text{end}}{\text{end} <:> \text{end}}$		

Figure 12. Subtyping rules (cont.)

References

- [1] Amal Ahmed. 2004. *Semantics of types for mutable state*. Ph.D. Dissertation. Princeton University.
- [2] Amal Ahmed, Andrew W. Appel, Christopher D. Richards, Kedar N. Swadi, Gang Tan, and Daniel C. Wang. 2010. Semantic foundations for typed assembly languages. *TOPLAS* 32, 3 (2010), 7:1–7:67. <https://doi.org/10.1145/1709093.1709094>
- [3] Andrew W. Appel and David A. McAllester. 2001. An indexed model of recursive types for foundational proof-carrying code. *TOPLAS* 23, 5 (2001), 657–683. <https://doi.org/10.1145/504709.504712>
- [4] Andrew W. Appel, Paul-André Melliès, Christopher D. Richards, and Jérôme Vouillon. 2007. A very modal model of a modern, major, general type system. In *POPL*. 109–122. <https://doi.org/10.1145/1190216.1190235>
- [5] Stephanie Balzer and Frank Pfenning. 2017. Manifest Sharing with Session Types. *PACMPL* 1, ICFP (2017), 37:1–37:29. <https://doi.org/10.1145/3110281>
- [6] Stephanie Balzer, Bernardo Toninho, and Frank Pfenning. 2019. Manifest Deadlock-Freedom for Shared Session Types. In *ESOP (LNCS, Vol. 11423)*. 611–639. https://doi.org/10.1007/978-3-030-17184-1_22
- [7] Mario Bravetti, Marco Carbone, and Gianluigi Zavattaro. 2017. Undecidability of asynchronous session subtyping. *Information and Computation* 256 (2017), 300–320. <https://doi.org/10.1016/j.ic.2017.07.010>
- [8] Luís Caires, Jorge A. Pérez, Frank Pfenning, and Bernardo Toninho. 2013. Behavioral Polymorphism and Parametricity in Session-Based Communication. In *ESOP (LNCS, Vol. 7792)*. 330–349. https://doi.org/10.1007/978-3-642-37036-6_19
- [9] Luís Caires and Frank Pfenning. 2010. Session Types as Intuitionistic Linear Propositions. In *CONCUR (LNCS, Vol. 6269)*. 222–236. https://doi.org/10.1007/978-3-642-15375-4_16
- [10] Marco Carbone, Fabrizio Montesi, Carsten Schürmann, and Nobuko Yoshida. 2017. Multiparty session types as coherence proofs. *Acta Informatica* 54, 3 (2017), 243–269.
- [11] David Castro, Francisco Ferreira, and Nobuko Yoshida. 2020. EMTST: Engineering the Meta-theory of Session Types. In *TACAS (LNCS, Vol. 12079)*. 278–285. https://doi.org/10.1007/978-3-030-45237-7_17
- [12] Ornela Dardha and Simon J. Gay. 2018. A New Linear Logic for Deadlock-Free Session-Typed Processes. In *FOSSACS (LNCS, Vol. 10803)*. 91–109. https://doi.org/10.1007/978-3-319-89366-2_5
- [13] Ornela Dardha, Elena Giachino, and Davide Sangiorgi. 2012. Session Types Revisited. In *PPDP*. 139–150. https://doi.org/10.1007/978-3-030-17184-1_22
- [14] Derek Dreyer, Amal Ahmed, and Lars Birkedal. 2009. Logical Step-Indexed Logical Relations. In *LICS*. 71–80. <https://doi.org/10.1109/LICS.2009.34>
- [15] Derek Dreyer, Amin Timany, Robbert Krebbers, Lars Birkedal, and Ralf Jung. 2019. What Type Soundness Theorem Do You Really Want to Prove? SIGPLAN blog post, available at <https://blog.sigplan.org/2019/10/17/what-type-soundness-theorem-do-you-really-want-to-prove/>.
- [16] Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2018. ReLoC: A Mechanised Relational Logic for Fine-Grained Concurrency. In *LICS*. 442–451. <https://doi.org/10.1145/3209108.3209174>
- [17] Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2020. Compositional Non-Interference for Fine-Grained Concurrent Programs. To appear in S&P'21.
- [18] Simon J. Gay. 2008. Bounded polymorphism in session types. *MSCS* 18, 5 (2008), 895–930. <https://doi.org/10.1017/S0960129508006944>
- [19] Simon J. Gay. 2016. Subtyping Supports Safe Session Substitution. In *A List of Successes That Can Change the World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*. 95–108. https://doi.org/10.1007/978-3-319-30936-1_5
- [20] Simon J. Gay and Malcolm Hole. 2005. Subtyping for session types in the pi calculus. *Acta Informatica* 42, 2-3 (2005), 191–225. <https://doi.org/10.1007/s00236-005-0177-z>
- [21] Simon J. Gay, Peter Thiemann, and Vasco T. Vasconcelos. 2020. Duality of Session Types: The Final Cut. In *PLACES (EPTCS, Vol. 314)*. 23–33. <https://doi.org/10.4204/EPTCS.314.3>
- [22] Paolo G. Giarrusso, Léo Stefanescu, Amin Timany, Lars Birkedal, and Robbert Krebbers. 2020. Scala step-by-step: soundness for DOT with step-indexed logical relations in Iris. *PACMPL* 4, ICFP (2020), 114:1–114:29. <https://doi.org/10.1145/3408996>
- [23] Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. 2020. Actris 2.0: Asynchronous session-type based reasoning in separation logic. (2020). Manuscript in preparation.
- [24] Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. 2020. Actris: Session-type based reasoning in separation logic. *PACMPL* 4, POPL (2020), 6:1–6:30. <https://doi.org/10.1145/3371074>
- [25] Jonas Kastberg Hinrichsen, Daniël Louwink, Robbert Krebbers, and Jesper Bengtson. 2021. Coq Mechanization of “Machine-Checked Semantic Session Typing”. Archived version at <https://zenodo.org/record/4322752>, latest version at <https://gitlab.mpi-sws.org/iris/actris>.
- [26] Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Discipline for Structured Communication-Based Programming. In *ESOP (LNCS, Vol. 1381)*. 122–138. <https://doi.org/10.1007/BF0053567>
- [27] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018. RustBelt: Securing the Foundations of the Rust Programming Language. *PACMPL* 2, POPL (2018), 66:1–66:34. <https://doi.org/10.1145/3158154>
- [28] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2020. Safe systems programming in Rust: The promise and the challenge. To appear in CACM.
- [29] Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-Order Ghost State. In *ICFP*. 256–269. <https://doi.org/10.1145/2951913>
- [30] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris From the Ground Up: A Modular Foundation for Higher-Order Concurrent Separation Logic. *JFP* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>
- [31] Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL*. 637–650. <https://doi.org/10.1145/2676726.2676980>
- [32] Robbert Krebbers, Jacques-Henri Jourdan, Ralf Jung, Joseph Tassarotti, Jan-Oliver Kaiser, Amin Timany, Arthur Charguéraud, and Derek Dreyer. 2018. MoSeL: A General, Extensible Modal Framework for Interactive Proofs in Separation Logic. *PACMPL* 2, ICFP (2018), 77:1–77:30. <https://doi.org/10.1145/3236772>
- [33] Robbert Krebbers, Ralf Jung, Ales Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017. The Essence of Higher-Order Concurrent Separation Logic. In *ESOP (LNCS, Vol. 10201)*. 696–723. https://doi.org/10.1007/978-3-662-54434-1_26
- [34] Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017. Interactive Proofs in Higher-Order Concurrent Separation Logic. In *POPL*. 205–217. <https://doi.org/10.1145/3093333.3009855>
- [35] Morten Krogh-Jespersen, Kasper Svendsen, and Lars Birkedal. 2017. A relational model of types-and-effects in higher-order concurrent separation logic. In *POPL*. 218–231. <https://doi.org/10.1145/3093333.3009877>
- [36] Dimitris Mostrous and Nobuko Yoshida. 2015. Session typing and asynchronous subtyping for the higher-order π -calculus. *Information and Computation* 241 (2015), 227–263. <https://doi.org/10.1016/j.ic.2015.02.002>
- [37] Dimitris Mostrous, Nobuko Yoshida, and Kohei Honda. 2009. Global Principal Typing in Partially Commutative Asynchronous Sessions. In *ESOP (LNCS, Vol. 5502)*. 316–332. https://doi.org/10.1007/978-3-642-00590-9_23

- [38] Jorge A. Pérez, Luís Caires, Frank Pfenning, and Bernardo Toninho. 2012. Linear Logical Relations for Session-Based Concurrency. In *ESOP (LNCS, Vol. 7211)*. 539–558. https://doi.org/10.1007/978-3-642-28869-2_27
- [39] Jorge A. Pérez, Luís Caires, Frank Pfenning, and Bernardo Toninho. 2014. Linear logical relations and observational equivalences for session-based concurrency. *Information and Computation* 239 (2014), 254–302. <https://doi.org/10.1016/j.ic.2014.08.001>
- [40] Frank Pfenning and Conal Elliott. 1988. Higher-Order Abstract Syntax. In *PLDI*. 199–208. <https://doi.org/10.1145/53990.54010>
- [41] Benjamin C. Pierce et al. 2020. Programming Language Foundations. <https://softwarefoundations.cis.upenn.edu/plf-current/index.html>
- [42] Arjen Rouvoet, Casper Bach Poulsen, Robbert Krebbers, and Eelco Visser. 2020. Intrinsically-typed definitional interpreters for linear, session-typed languages. In *CPP*. ACM, 284–298. <https://doi.org/10.1145/3372885.3373818>
- [43] David Swasey, Deepak Garg, and Derek Dreyer. 2017. Robust and compositional verification of object capability patterns. *PACMPL* 1, OOPSLA (2017), 89:1–89:26. <https://doi.org/10.1145/3133913>
- [44] Joseph Tassarotti, Ralf Jung, and Robert Harper. 2017. A Higher-Order Logic for Concurrent Termination-Preserving Refinement. In *ESOP (LNCS, Vol. 10201)*. 909–936. https://doi.org/10.1007/978-3-662-54434-1_34
- [45] The Coq-std++ Team. 2020. An extended “standard library” for Coq. Available online at <https://gitlab.mpi-sws.org/iris/stdpp>.
- [46] Peter Thiemann. 2019. Intrinsically-Typed Mechanized Semantics for Session Types. In *PPDP*. 19:1–19:15. <https://doi.org/10.1145/3354166.3354184>
- [47] Peter Thiemann and Vasco T. Vasconcelos. 2020. Label-dependent session types. *PACMPL* 4, POPL (2020), 67:1–67:29. <https://doi.org/10.1145/3371135>
- [48] Amin Timany, Léo Stefanescu, Morten Krogh-Jespersen, and Lars Birkedal. 2018. A logical relation for monadic encapsulation of state: Proving contextual equivalences in the presence of runST. *PACMPL* 2, POPL (2018), 64:1–64:28. <https://doi.org/10.1145/3158152>
- [49] Philip Wadler. 2012. Propositions as sessions. In *ICFP*. 273–286. <https://doi.org/10.1145/2364527.2364568>
- [50] Andrew K. Wright. 1995. Simple Imperative Polymorphism. *Lisp and Symbolic Computation* 8, 4 (1995), 343–355.