

Sessions and Separation

Jonas Kastberg Hinrichsen
PhD Dissertation
IT University of Copenhagen

March 2021

Abstract

Session types and separation logic are two formalisms used to obtain strong guarantees about concurrent programs, each with their own benefits. Session types provide an intuitive protocol mechanism for specifying sequences of value exchanges, where each exchange is constrained to a specific type. In contrast, separation logic is often used for more involved specifications and proofs about shared-memory concurrency.

This thesis combines the two efforts by introducing *dependent separation protocols*—a session-type based reasoning mechanism in separation logic—and a *semantic session type system*—a session type system that combines the expressivity of separation logic with the simplicity of type checking.

Dependent separation protocols mimic the intuition behind session types, by describing sequences of value exchanges constrained by separation-logic propositions. The dependent separation protocols are formalised for a representative language with fine-grained concurrency, mutable state, and higher-order functions. This formalisation, which we call *Actris*, is built on top of the higher-order concurrent separation logic Iris. The use of Iris provides the ability to reason about recursion, substructural properties, and integration with existing concurrency models, such as shared-memory and lock-based concurrency. We demonstrate the expressivity of Actris by proving the correctness of several variants of a distributed merge sort algorithm, a distributed load-balancing mapper, and a variant of the map-reduce algorithm.

The semantic session type system is achieved using semantic typing and logical relations, where the types and judgements of the type system are defined in terms of Actris. The type system supports term- and session-type polymorphism, equi-recursion, and subtyping, as well as copyable types, shared and mutable reference types, and mutex types. The type system additionally supports manual typing proofs, where typing judgements that cannot be derived from the typing rules of the type system can instead be proven in the logic. This technique effectively lets us combine the full expressivity of the logic, with the simplicity of the syntax-directed type checking of the type system. We demonstrate this feature by manually proving the typing judgement of a program, where two threads operate on a single session-typed channel endpoint in parallel.

All of the results of the thesis have been mechanised in the interactive proof assistant Coq, on top of the existing mechanisation of Iris, along with tactic support for resolving proof obligations related to message passing in a style similar to working with session types.

Resumé

Sessionstyper og separationslogik er to metoder der benyttes til at opnå stærke garantier om samtidige programmer, med hver deres fordele. Sessionstyper giver en intuitiv protokolmekanisme til at specificere sekvenser af værdiudvekslinger, hvor hver udveksling er begrænset til en specifik type. Separationslogik er derimod ofte brugt i forbindelse med mere komplicerede specifikationer og beviser vedrørende delt hukommelse og samtidighed.

Denne afhandling kombinerer de to metoder ved at introducere *afhængige separationsprotokoller*—en sessionstype-baseret ræsonneringsmekanisme i separationslogik—og et *semantisk sessionstypesystem*—et sessionstypesystem der kombinerer ekspressiviteten af separationslogik med enkeltheden af type bestemmelse.

Afhængige separationsprotokoller efterligner intuitionen bag sessionstyper, da de beskriver sekvenser af værdiudvekslinger som er begrænset via separationslogikprædikater. De afhængige separationsprotokoller er formaliseret for et repræsentativt sprog med finkornet samtidighed, muterbar tilstand, og højere-ordens-funktioner. Denne formalisering, som vi kalder *Actris*, er bygget oven på den højere-ordens samtidige separationslogik Iris. Med brugen af Iris kan man ræsonnere om rekursion, substrukturelt ejerskab, og integration med eksisterende modeller om samtidighed, såsom delt hukommelse og låse-baseret samtidighed. Vi demonstrerer Actris ved at bevise korrektheden af flere varianter af en distribueret “merge sort”-algoritme, en distribueret belastnings-balanceret “mapper”, og en variant af “map-reduce” algoritmen.

Det semantiske sessionstypesystem opnås via semantisk typning og logiske relationer, hvor typerne og typebedømmelsesmekanismen af typesystemet er defineret via Actris. Typesystemet understøtter term- og sessionstype polymorfi, rekursion, og subtyper, såvel som kopierbare typer, delte og muterbare referencetyper og mutex-typer. Typesystemet understøtter derudover manuelle typningsbeviser, hvor typebedømmelser, der ikke kan udledes af type systemets typeregler, i stedet kan bevises i logikken. Denne metode lader os effektivt set kombinere logikkens fulde ekspressivitet med enkeltheden af typesystemets syntaksorienterede typebestemmelse. Vi demonstrerer denne funktionalitet ved manuelt at bevise typebedømmelsen af et program, hvor to tråde opererer på et enkelt sessionstypet kanal-slutpunkt i parallel.

Alle afhandlingens resultater er blevet mekaniseret i den interaktive bevisassistent Coq, oven på den eksisterende mekanisering af Iris, sammen med understøttelsen af taktikker til at løse bevisforpligtelser relateret til kommunikation, i en stil der minder om at arbejde med sessionstyper.

Acknowledgements

I would first and foremost like to thank Jesper Bengtson and Robbert Krebbers, for their tutelage as my supervisors, their collaboration as my co-authors, and their companionship as my friends. Of all the good memories that I have had during my time as a PhD student, I will treasure our intensive co-authoring sessions before a deadline the most, and our celebratory beer sessions only second to that.

I would then like to thank everyone at the IT University of Copenhagen, for creating a welcoming atmosphere. Especially so Frederik Madsen, my office mate for the longest time, with whom I shared the good times and endured the bad times. I would also like to thank the biking team, primarily spearheaded by Marco Carbone and Alessandro Bruni, who helped me achieve a better work life balance.

I extend my gratitude to the members of the committee, Derek Dreyer, Nobuko Yoshida, and Marco Carbone, for taking their time to review the thesis.

I would like to thank the hosts of my research stays abroad. First, Gregory Malecha of Bedrock Systems, who let me stay with his family until I found a temporary place to live in Boston, and to navigate the US infrastructure. A special thanks goes to my subsequent landlords Jordan and Nathan, and to my roommate Ben, who all made me feel at home. Secondly, Robbert Krebbers who hosted me twice at Delft University of Technology. Between chess matches, bouldering, and enjoying Belgian beer, it suffices to say that I had a great time staying in Delft, both during and outside of working hours. I can list many negative impacts that the corona pandemic had on my life, but cutting my time in Delft short is one of the worst offenders.

I would like to thank all of my friends. From the ones that I had before I started the degree, to the ones that I met during it, at the IT University of Copenhagen, the Oregon Programming Languages Summer School, the DeepSpec Summer School, and the Midlands Graduate School. Saying that I would not have been able to finish this thesis without them feels like an understatement.

A special thanks goes to Julie, who helped me through the most stressful time of my life, and who gave me the encouragement to seek the help that I needed.

Finally, I would like to thank my family. My sister and brother-in-law, who were the first to support my decision of pursuing a career in academia by “studying for three more years”. Lastly, my parents, who always urged me to “do my best”, while assuring that nothing more could be expected by anyone. These words always helped me push myself to improve my strengths, and occasionally to accept my weaknesses. I dedicate this thesis to them, my father Bjarne and my mother Susan.

Contents

1	Introduction	1
1.1	Message Passing	2
1.2	Safety and Session Types	3
1.3	Functional Correctness and Concurrent Separation Logic	4
1.4	Challenges	5
1.5	Problem Statement and Contributions	8
1.6	List of Publications and Manuscripts	10
2	Background	13
2.1	Operational Semantics	14
2.2	Safety, Functional Correctness, and Semantic Typing	21
2.3	The Iris Logic	26
3	Actris: Session-Type Based Reasoning in Separation Logic	39
3.1	Introduction	39
3.2	A Tour of Actris	48
3.3	Subprotocols	60
3.4	Manifest Sharing via Locks	69
3.5	Case Study: Map-Reduce	74
3.6	The Model of Actris	77
3.7	Coq Mechanisation	86
3.8	Related Work	95
3.9	Conclusion and Future Work	99
4	Semantic Session Typing	101
4.1	Introduction	101
4.2	A Tour of Semantic Session Typing	103
4.3	Extending the Type System	110
4.4	Manual Typing Proofs	118
4.5	Mechanisation in Coq	122
4.6	Related Work	125
4.7	Conclusion	127
4.A	The Complete Type System	127
	Bibliography	139

Introduction

The world is continuously moving towards concurrently executed programs, as we are reaching the physical limits of how fast a single chip can conduct computations. However, writing correct concurrent software is notoriously hard, and is therefore commonly approached using principled paradigms, such as message passing.

Even so, message-passing programs are prone to errors [Bagherzadeh et al. 2020], warranting increased correctness guarantees. A step in this direction is program testing, but as Dijkstra puts it “Program testing can be used to show the presence of bugs, but never to show their absence!”. This is especially true for concurrent programs, that are often non-deterministic, where only one of many possible executions may lead to an error. We therefore seek to verify such message-passing programs by proving their *safety*, *i.e.*, no execution of the program has bad behaviour (such as calling a function with incorrect arguments), and *functional correctness*, *i.e.*, any execution of the program outputs the correct result (based on some specification).

A prominent approach to showing safety of message-passing programs is that of *session types*, a type system originally invented for process calculi [Honda 1993] that is still an active research topic today, featuring extensions such as shared channels [Balzer et al. 2019], dependent protocols [Thiemann and Vasconcelos 2020], and protocol weakening [Bravetti et al. 2021]. Orthogonally, showing safety and functional correctness of concurrent programs has been achieved with *concurrent separation logic* [O’Hearn 2004; Brookes 2004], a school of logics for reasoning about safe distribution of resources in a concurrent setting, from which descendants have been used to verify *e.g.*, C programs [Appel 2011] and Rust programs [Jung et al. 2018a].

In this thesis we combine session types and concurrent separation logic to establish a new reasoning mechanism, *dependent separation protocols*, for showing safety and functional correctness of programs that use message passing along with other concurrency models, such as lock-based synchronisation. We additionally combine the simplicity of type checking with the expressivity of logic proofs, by defining a *semantic type system* [Ahmed 2004] for session types. Finally, we mechanise all our results on top of the Iris framework [Iris Development Team 2021] in the Coq interactive proof assistant [Coq Development Team 2021], which includes the verification of a map-reduce algorithm, and a message-passing-based producer-consumer.

In this chapter we provide an overview of the field of session types and separation logic, and present the problem that the thesis seek to address. In particular, we first provide a brief overview of message passing (Section 1.1), safety and session types (Section 1.2), functional correctness and concurrent separation logic (Section 1.3), and an outline of some of the existing challenges in the field (Section 1.4). We then present the problem statement, the key idea to addressing the problem, and the contributions of the thesis (Section 1.5). Finally, we provide an insight to the contents of the included papers, and my role in conducting the work (Section 1.6).

1.1 Message Passing

Message passing is a concurrency paradigm in which threads are treated as actors in a network that synchronise via messages, rather than having shared access to mutable data [Hewitt et al. 1973]. The paradigm has been integrated as a first-class primitive of state-of-the-art languages, such as Go [The Go Team 2021] and Erlang [The Erlang Team 2021], while other languages, *e.g.*, Java and C#, support message passing through libraries [Akka 2021; Akka.NET 2021]. In such high-level languages messages are often guaranteed to arrive and be delivered in order. For this thesis we therefore consider binary message passing (where channels are between two parties) with these properties.

Message passing is often realised using a means of creating channels, sending messages and receiving messages, for which we use the following syntax:

```
new_chan ()    send c v    recv c
```

Here, `new_chan ()` is used to allocate a new channel, creating two unbounded buffers (\vec{v}_1, \vec{v}_2) , while returning two *channel endpoints* c_1 and c_2 . The operation `send c_i v` is used to send a value v over the channel, by enqueueing it in the corresponding buffer \vec{v}_i . Conversely `recv c_i` receives values by dequeueing them from the other buffer, *i.e.*, \vec{v}_2 if $i = 1$ and \vec{v}_1 if $i = 2$. The receive operation blocks until a value is available.

While the message-passing paradigm seek to avoid shared access to mutable data, it has been shown that developers often choose to combine the two [Tasharofi et al. 2013]. We therefore consider a representative ML-like language that supports mutable state, fork-based concurrency, locks, and the three message-passing primitives described above. The operational semantics, *i.e.*, the mathematical model, of the language is described in Section 2.1.2.

To illustrate the combination of message passing with shared memory, we consider the interaction between a pizza shop and a customer. To order pizzas the customer has to hand over their credit card (a reference to an integer), from which the pizza shop deducts the cost of the pizzas, after which they send back a receipt (a list of integers). This interaction is modelled by the following program:

```
order_pizza cc c :=  
  send c [Margherita, Calzone];  
  send c cc;  
  let receipt := recv c in  
  (sum receipt, !cc)
```

More precisely, the customer first sends the order of a margherita and a calzone pizza over the communication channel (`send c [Margherita, Calzone]`), and then send their credit card (`send c cc`). They then await a response from the pizza shop (`recv c`), expected to be the receipt (`receipt`). Finally the customer inspects (returns) the sum of the receipt (`sum receipt`) and the dereferenced value of the credit card (`!cc`).¹

1.2 Safety and Session Types

We initially want to show safety of message-passing programs. However, safety is a general term informally described as “something [bad] will *not* happen” [Lampert 1977], and depends on the operational semantics of the language and what you want to prove about the programs. The safety that we consider is then whether the programs are *type-safe*, *e.g.*, that functions are applied to the right types of arguments, and that literal values are not used as functions. Safety of the `order_pizza` program presented in Section 1.1 then depends on a guarantee that the received message `xs` is a list of integers, as the operation `sum xs` can otherwise fail.

Showing safety is often achieved through type checking, *i.e.*, checking that the program is in compliance with the rules of a type system. Safety then follows under the condition that the type system satisfies *type safety*, *i.e.*, “[...] well-typed programs cannot ‘go wrong’” as coined by Milner [1978]. A state-of-the-art type system for message passing is session types [Honda 1993], where channel endpoints are assigned a protocol. These protocols specify obligations to (primarily) send (`!A. S`) or receive (`?A. S`) a value of type A (*e.g.*, `List Z`) and then continue as S . Session types guarantee *session fidelity*, *i.e.*, that any received message has been sent, and that it has the denoted type. This can in turn be used to show that the received messages are used safely. The following session type captures the interaction with the pizza shop:

$$pizza_styp \triangleq !(List\ Pizza).!(ref\ Z).?(List\ Z).pizza_styp$$

The session type specifically states that the customer must first send (`!`) the order, consisting of a list of pizzas (`List Pizza`), and their credit card (`ref Z`), after which they can receive (`?`) the receipt (`List Z`). The protocol then loops, awaiting a new order. The `order_pizza` program follows the protocol, as it first sends a list of pizzas (`[Margherita, Calzone] : List Pizza`), then the credit card reference (`cc : ref Z`), and then use the received value (`receipt : List Z`) in accordance with its type (`sum receipt`):

```
order_pizza cc c :=
  // c : !(List Pizza).!(ref Z).?(List Z).pizza_styp
  send c [Margherita, Calzone];
  // c : !(ref Z).?(List Z).pizza_styp
  send c cc;
  // c : ?(List Z).pizza_styp
  let receipt := recv c in
  // c : pizza_styp, receipt : List Z
  (sum receipt, !cc)
```

¹The hungry reader might notice that the ordered pizzas are never received. Sadly, pizzas cannot yet be sent digitally, and are thus not included in the protocol.

We can thus guarantee that the received receipt has the expected type, under the assumption that the pizza shop has similarly been type checked (thus complying with the protocol), and thereby guarantee that the program is type-safe.

However, the astute reader might have noticed that there is no guarantee that the credit card has not been charged more than the cost of the pizzas. To guarantee such a property we would instead need a proof of functional correctness.

1.3 Functional Correctness and Concurrent Separation Logic

Proving functional correctness can be achieved with Hoare logic [Hoare 1969], in which programs are ascribed with logical assertions, that specify propositions which hold before and after the execution of a program. These specifications are then proven by applying the rules of the logic, which capture how the assertions change throughout the evaluation of the program.

Extending the approach of Hoare logic to reason about stateful programs, *i.e.*, programs with mutable state, such as references, was later achieved with separation logic [Reynolds 2002; Ishtiaq and O’Hearn 2001]. Separation logic supports modular proofs of sub-programs, *i.e.*, proving the specifications of sub-programs in isolation. This is achieved by virtue of the separation property, which guarantees that the memory footprint of a sub-program is disjoint from the remaining program, thereby guaranteeing that they do not interfere with each other.

A later extension, concurrent separation logic [O’Hearn 2004; Brookes 2004], repurposed the separation property to implicitly capture the safe interleaving of concurrently executed threads, by guaranteeing that their executions are independent of each other. Concurrent separation logic additionally employ techniques for reasoning about the safe interleaving of threads with overlapping memory footprint. This direction has since spawned multiple descendants, for example VST [Appel 2011], FCSL [Nanevski et al. 2014], and Iris [Jung et al. 2015, 2018b]. Iris, in particular, is a higher-order concurrent separation logic, which supports the construction of high-level abstractions, such as locks and barriers, integrating them like first-class primitives, using higher-order ghost state [Jung et al. 2016].

Proving functional correctness of message-passing programs in the context of separation logic has been pursued in multiple efforts [Francalanza et al. 2011; Lozes and Villard 2012; Craciun et al. 2015; Oortwijn et al. 2016]. One particular effort is the work by Craciun et al. [2015], who combine separation logic with session types. They do so by replacing the message types with separation logic propositions, to specify additional details about the exchanged values, *e.g.*, to reason about exchanges of references and other channel endpoints.

The work by Craciun et al. [2015] is similar in spirit to prior work carried out by Bocchi et al. [2010]. They replace the message types with dependent first-order propositions, to specify additional details about the exchanged values, where the propositions can refer to previously exchanged values.

To prove that the program presented in Section 1.1 is correct, we use a technique similar to Bocchi et al. [2010] and Craciun et al. [2015], by extending session types with Iris propositions.

In particular, our protocols describe obligations to send ($!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$) or receive ($?\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$) a value v constrained by a proposition P in the context of the logical binders $\vec{x}:\vec{\tau}$, after which the protocol continues as $prot$. We can then use the following protocol, to prove that the credit card is charged according to the receipt:

$$\begin{aligned}
 & pizza_prot \triangleq \\
 & \quad ! (v_1 : \text{Val}) (\vec{x}_1 : \text{List Pizza}) \langle v_1 \rangle \{ \text{is_pizza_list } v_1 \vec{x}_1 * 0 < |\vec{x}_1| \}. \\
 & \quad ! (\ell : \text{Loc}) (x : \mathbb{Z}) \langle \ell \rangle \{ \ell \mapsto x \}. \\
 & \quad ? (v_2 : \text{Val}), (\vec{x}_2 : \text{List } \mathbb{Z}) \langle v_2 \rangle \left\{ \begin{array}{l} \text{is_int_list } v_2 \vec{x}_2 * |\vec{x}_1| = |\vec{x}_2| * \\ \ell \mapsto (x - (\text{sum } \vec{x}_2)) \end{array} \right\}. \\
 & \quad pizza_prot
 \end{aligned}$$

The protocol states that the customer must first send a value v_1 , which is an encoding of the list of pizzas \vec{x}_1 , as captured by the proposition `is_pizza_list` $v_1 \vec{x}_1$. The side-condition $0 < |\vec{x}_1|$ captures that the list must be non-empty (to avoid prank calls). The side-condition holds alongside the list proposition, denoted by the separation conjunction ($*$), which for now can be thought of as regular propositional conjunction (\wedge). The customer must then send a location ℓ , representing their credit card, while noting that the original amount is x , captured by the separation logic points-to connective $\ell \mapsto x$. The pizza shop then returns the receipt v_2 , which is a list of the cost of each pizza \vec{x}_2 , captured by `is_int_list` $v_2 \vec{x}_2$. The side-condition $|\vec{x}_1| = |\vec{x}_2|$ states that the number of deductions must correspond to the number of ordered pizzas. Finally, the received proposition $\ell \mapsto (x - (\text{sum } \vec{x}_2))$ captures that the new value stored in the credit card corresponds to the cost of the pizzas.

We can thus prove that the deducted amount is in accordance with the sum of the receipt, under the assumption that the pizza shop is in compliance with the protocol.

1.4 Challenges

This section presents some of the challenges related to reasoning about safety and functional correctness of message-passing programs. We first cover the importance of being able to reason about message passing in the context of other concurrency models, such as lock-based synchronisation (Section 1.4.1). We then illustrate the benefits of being able to combine the simplicity of safety checking via type systems, with the expressivity of interactive proofs in concurrent separation logic (Section 1.4.2). We finally discuss how mechanisation can be used to improve the reliability of type systems, logical systems, and the proofs carried out in them (Section 1.4.3).

1.4.1 Message Passing and other Concurrency Models

Message passing is often combined with mutable state and other concurrency models, such as lock-based synchronisation [Tasharofi et al. 2013]. To illustrate the benefits of combining message passing with locks consider two concurrent customers of the pizza shop, Jesper and Robbert. Jesper and Robbert are each having individual parties, and may want to order pizzas at some point during the evening, over the same channel endpoint. Sharing of channel endpoints is non-trivial as they are typically treated

as exclusive resources in the context of session types, meaning that they can only be accessed by one thread at a time. However, this exclusivity of channel endpoints is crucial to establish the guarantees of session types, as they rely on the fact that two threads do not operate on the same channel endpoint simultaneously.

To share channel endpoints between multiple threads we use locks, which guarantee safe shared access to a locked exclusive resource through mutual exclusion. In particular, a locked resource is guaranteed to remain unchanged outside of any critical section, *i.e.*, the sections where the shared exclusive resource is accessed. To access and momentarily modify the resources a threads must enter the critical section by acquiring the lock lk , using `acquire lk` , and afterwards release the lock, using `release lk` , to exit the critical section. The following program illustrates how the channel endpoint to the pizza shop is safely shared between both parties using a lock:

```
lock_example jesper_cc robbert_cc c :=
  let lk := new_lock () in
  (
    // Jesper's party
    acquire lk;
    send c [Margherita];
    send c jesper_cc;
    let jesper_receipt := recv c in
    release lk;
    // Jesper's party cont.
    (sum jesper_receipt, !jesper_cc)
  ) |||
  (
    // Robbert's party
    acquire lk;
    send c [Calzone];
    send c robbert_cc;
    let robbert_receipt := recv c in
    release lk;
    // Robbert's party cont.
    (sum robbert_receipt, !robbert_cc)
  )
```

The program executes the parties in parallel, which then race for the lock. Once the lock is acquired by either party, they can safely order pizzas, while the other party waits for the lock to be released.

Proving safety of programs which combine session types and lock-based synchronisation have been studied by Balzer and Pfenning [2017]; Balzer et al. [2019]. They add synchronisation primitives to their version of session types to guarantee mutual exclusion of channels during critical sections. However, this is different from having two separate mechanisms, one for locks and one for channels, that compose seamlessly. Additionally, their solution specifically address the combination of message passing and locks, and is not equipped to be extended with other concurrency models.

In regards to functional correctness, existing solutions for session-type based reasoning in separation logic [Craciun et al. 2015] do not integrate with other concurrency models, such as lock-based synchronisation. This is a result of building on top of more a traditional separation logic, as opposed to a logic which readily supports other concurrency models as first-class notions, *e.g.*, lock-based synchronisation, such as the higher-order concurrent separation logic Iris.

1.4.2 Safety via Functional Correctness

Proving safety of programs using type checking is often straightforward, and can in some cases even be done automatically using a decidable type checker. However, type checking is inherently not expressive enough to capture safety of some classes of racy yet safe programs, *i.e.*, programs which safely access exclusive resources concurrently.

Imagine that the pizza shop is having their annual pizza-palooza, in which a finite amount of people can get vouchers for free pizzas. Consider the following program and type, modelling two people concurrently requesting one of the last two pizzas:

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

The program receives twice in parallel on the channel endpoint c , and returns the resulting pizzas as a pair. The program is safe, as any order of execution of the receive instructions results in a pair of some pizzas, as expected by the type annotation. However, as the channel endpoint is considered exclusive in conventional session type systems, they are incapable of showing safety of this program. This remains the case in the context of locks, since they would require the locked protocol to remain the same outside of the critical sections, which is not the case here, as each access strips one, and only one, receive step from the protocol, *i.e.*²

$$?Pizza. ?Pizza. \text{end} \rightsquigarrow ?Pizza. \text{end} \quad \text{or} \quad ?Pizza. \text{end} \rightsquigarrow \text{end}$$

As it is imperative that the pizza shop can safely continue the pizza-palooza tradition, we want to prove safety of the program, which is possible with concurrent separation logic, since its functional correctness proofs support reasoning about safe distributions of exclusive resources. However, proofs in a logic are almost always more time-consuming than type checking. We therefore wish to be able to type check the majority of a given program, and only prove safety of sub-programs in the logic when the expressivity of type checking is insufficient.

Achieving such a composition of type checking and manual typing proofs has been done by the RustBelt project [Jung et al. 2018a], which combines manual proofs of typing judgements for unsafe Rust code, with the syntactically typeable part of the language. They do so by employing the technique of semantic typing [Milner 1978; Ahmed 2004; Ahmed et al. 2010], in which types are given a logical interpretation, and the typing judgement is defined in a way that captures safety. Specifically, Jung et al. [2018a] define their type system in terms of the higher-order concurrent separation logic Iris, to capture resourceful types, such as Rust lifetimes.

1.4.3 Mechanisation

Proving properties such as safety and functional correctness using a *meta-theory*, *e.g.*, a session type system or a concurrent separation logic, is a means of increasing the reliability of software. However, this begs the question of how reliable these meta-theories are themselves. A meta-theory can similarly be prone to errors, as 1) proofs in the meta-theory might not actually guarantee the expected properties, and 2) proofs of individual programs in the meta-theory might have been carried out incorrectly.

Addressing these concerns can be done using mechanisation, where manual proofs are automatically checked by a machine, *e.g.*, via an interactive proof assistant such as Coq [Coq Development Team 2021], Isabelle/HOL [Nipkow et al. 2002], or

²The charitable reader may have noticed that the program could in fact be typed given a recursive type, however handing out infinite free pizzas is arguably a dubious business model.

Lean [de Moura et al. 2015]. Mechanisation has become popular over the years, and has been used for verification of cornerstone systems such as compilers [Leroy 2006; Kumar et al. 2014], operating systems [Klein et al. 2009; Gu et al. 2011], security protocols [Beringer et al. 2015], type systems [Jung et al. 2018a], and is even seeing use in industry [Bedrock Systems A/S 2021; The CertiK Team 2021].

We thus seek to mechanise the work carried out in this thesis, to strengthen the claim that our results are correct. In particular, we want to mechanise the meta-theory behind the proposed protocol mechanism, *i.e.*, that a proof in the meta-theory in fact guarantees functional correctness.

Doing so requires mechanising the operational semantics of the language, the meta-theory, and the connection between them. This mechanisation effort grows increasingly non-trivial for each feature of the meta-theory, *e.g.*, exclusive ownership, concurrency, and recursion. However, the mechanisation effort can be alleviated by building on top of an existing framework. Such an existing framework is the higher-order concurrent separation logic Iris in Coq [Iris Development Team 2021; Coq Development Team 2021]. Iris is parametric on an operational semantics, and includes a concrete instantiation with the operational semantics for an untyped functional language with higher-order functions, garbage collected higher-order mutable references, and fork-based concurrency, called HeapLang.

Additionally, we seek to carry out concrete proofs using the meta-theory, as evidence of its expressivity and usability, *i.e.*, that it is feasible to carry out proofs of larger and more realistic programs using the meta-theory.

To do so it is imperative to have a scalable method for carrying out proofs at the level of the meta-theory and its abstractions. This is the case as interactive proof assistants, such as Coq, are typically designed towards reasoning in their own logic. As a result, proofs in an embedded logic, such as Iris in Coq, often require manual handling of details that would be considered implicit in the embedded logic, *e.g.*, the distribution of substructural propositions in separation logic. A scalable method is the MoSeL framework (formerly called the Iris Proof Mode) [Krebbers et al. 2017b, 2018], which allows for reasoning at the level of concurrent separation logic in Coq, as well as being extensible with custom language- and abstraction-specific tactics.

Finally, following the trend of it becoming more common practice for researchers to mechanise their work, some mechanisation efforts have been conducted within the field of session types [Tassarotti et al. 2017; Thiemann 2019; Rouvoet et al. 2020; Gay et al. 2020; Castro et al. 2020]. However, there is a lack of mechanisations for more expressive session type systems, that combine multiple session type features such as recursion, polymorphism, and subtyping, along with types of other concurrent models such as mutex types for lock-based synchronisation.

1.5 Problem Statement and Contributions

This thesis concretely seeks to address the following three problems:

1. Session types provide an intuitive foundation for proving functional correctness of message-passing programs, however, existing session-type-based mechanisms do not compose with existing solutions for other concurrency models.

2. It is not possible to combine the expressivity of concurrent separation logic, *e.g.*, for proving safety of racy yet safe programs, with the simplicity of type checking in a session type system
3. It is becoming more common practice to mechanise new work, however, there is a lack of mechanisation efforts in the field of session types

The key idea to solving problem 1 is to draw inspiration from existing session-type based protocols for functional verification [Bocchi et al. 2010; Craciun et al. 2015]. We define a new protocol mechanism, that is both *dependent*, *i.e.*, protocol steps can rely on prior exchanges, and *resourceful*, *i.e.*, protocols can describe exchanges of exclusive resources, such as references or other channel endpoints. The protocols additionally support recursion, higher-order quantification, and weakening.

The protocol mechanism is formalised in Iris, using its support for constructing high-level abstractions, to integrate the protocol mechanism with Iris’s expressive variant of concurrent separation logic, which include a first-class formalism of locks. To connect the protocol mechanism with an operational semantics, the message-passing primitives are then implemented on top of the HeapLang language, which has already been connected to the Iris meta-theory.

The key idea to solving problem 2 is to use the technique of *semantic typing*, drawing inspiration from the semantic type system for Rust in Iris by Jung et al. [2018a]. In particular, how we can integrate manual proofs of untypeable programs with the type system. Additionally, by building on top of Iris we inherit its affine properties for defining affine types, its recursion mechanisms for recursive types, and its lock library for mutex types. We use the proposed protocol mechanism as the semantic interpretation of session types with recursion, polymorphism, and subtyping.

Finally, the key idea to solving problem 3 is to mechanise all of our efforts on top of the Iris framework in Coq. In particular, we inherit the mechanisation of the operational semantics of HeapLang, and its connection to Iris’s meta-theory. This lets us focus on the mechanisation of our protocol mechanism, the semantic type system, and the examples carried out in both of the systems.

Contributions The contributions made by this thesis are as follows:

1. We construct a first-class session-type based mechanism, *dependent separation protocols*, for proving functional correctness of message-passing programs in the higher-order concurrent separation logic Iris, using its support for high-level abstractions. The first-class mechanism integrates with Iris’s existing concurrency mechanisms, such as shared mutable state and locks
2. We define a semantic type system for session types—based on the dependent separation protocols—along with shared- and mutable reference types, mutex types, term- and session subtyping, term- and session type polymorphism, recursive types, and support for manual proofs of “racy” yet safe programs
3. We fully mechanise both of the above results on top of the Iris mechanisation in Coq, with tooling support for proving typing judgements and functional specifications of concrete programs, along with various examples, including a functional proof of a map-reduce algorithm, and the manual typing proof of a message-passing-based producer-consumer

1.6 List of Publications and Manuscripts

Each of the contributions in Section 1.5 were achieved as a part of the work presented in the following collection of publications and manuscripts.

- Actris: Session-Type Based Reasoning in Separation Logic, ACM SIGPLAN Symposium on Principles of Programming Languages 2020 (POPL’20)
- Actris 2.0: Asynchronous Session-Type Based Reasoning in Separation Logic, Journal of Logical Methods in Computer Science (LMCS manuscript in review)
- Machine-Checked Semantic Session Typing, Certified Programs and Proofs 2021 (CPP’21) (recipient of Distinguished Paper Award)

More specifically, the first-class reasoning mechanism for message passing (contribution 1) is addressed by the POPL’20 paper and its LMCS journal version. The semantic session type system (contribution 2) is covered by the CPP’21 paper. Finally, the mechanisations (contribution 3) are included as artifacts of their respective papers. The LMCS manuscript and the CPP’21 paper have been reformatted and included in the thesis as Chapter 3 and Chapter 4, respectively. The remainder of this section is a brief overview of each paper (Sections 1.6.1 to 1.6.3), and concludes with references to the original papers and their mechanisation artifacts (Section 1.6.4).

1.6.1 Actris: Session-Type Based Reasoning in Separation Logic.

Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers, POPL’20

This paper presents the first-class separation logic construct of *dependent separation protocols*, which can be thought of as a logical interpretation of session types.

In contrast to session types, which describe the exchanged values in terms of types, the dependent separation protocols describe exchanges using separation logic propositions, that can capture more precise properties than the types. An example of a dependent separation protocol, along with its corresponding session type, is the following, describing the protocol of a service that computes the length of a list:

Session type:	Dependent separation protocol:
<code>?(List Z).!Z.end</code>	<code>?($v : \text{Val}$), ($\vec{x} : \text{List } \mathbb{Z}$) $\langle v \rangle \{ \text{is_int_list } v \vec{x} \}. ! \vec{x} . \text{end}$</code>

The dependent separation protocol states that the channel first receives a value v , which is an encoding of a list of integers \vec{x} , as determined by `is_int_list v \vec{x}` . It then sends back the length of the list $|\vec{x}|$, with no propositional restriction.

The protocol construct has been integrated with the Iris concurrent separation logic framework, allowing it to *e.g.*, carry exclusive ownership of references and other channel endpoints. Additionally, the integration allows combining the protocols with lock-based reasoning, which is demonstrated in the paper, along with a case study in which we verify a map-reduce algorithm.

Personal Contribution The publication is joint work, with close collaboration between all co-authors on all sections. As first author I spearheaded the writing of all

of the sections except the related work section (initial version by Jesper Bengtson), and the model and mechanisation sections (initial version by Robbert Krebbers).

The publication is accompanied by an artifact consisting of the mechanisation of the dependent separation protocols, the channel implementation, and the examples of the paper. The mechanisation was developed in collaboration with Robbert Krebbers. Krebbers particularly carried out the implementation of the recursive domain equation of the dependent separation protocols, and the proof of the load-balancing mapper and map-reduce examples. I did the rest under his supervision, *e.g.*, carried out the protocol definition, and connected it to the channel implementation.

1.6.2 Actris 2.0: Asynchronous Session-Type Based Reasoning in Separation Logic. Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers, LMCS manuscript in review

This paper extends the dependent separation protocol mechanism with a notion similar to asynchronous session subtyping, where protocols can be weakened to allow the composition of more programs that communicate. Subtyping achieved this by permitting (1) sending more, (2) receiving less, and (3) eagerly sending before receiving. In addition to conventional subtyping, our subprotocol mechanism is resourceful. It allows “framing” resources that are sent in excess, which can then be re-obtained during a subsequent receive. The mechanism also enjoys the full expressivity of Iris’s step-indexed recursion, allowing weakening under the recursive operator. An example of a subprotocol, along with its corresponding session type subtyping, is the following which captures how the list length service may disregard the type of the lists elements:

Subtype:	Subprotocol:
$?(\text{List } \mathbb{Z}). !\mathbb{Z}. \text{end}$ $<: ?(\text{List any}). !\mathbb{Z}. \text{end}$	$?(v : \text{Val}), (\vec{x} : \text{List } \mathbb{Z}) \langle v \rangle \{ \text{is_int_list } v \vec{x} \}. ! \langle \vec{x} \rangle. \text{end}$ $\sqsubseteq ?(v : \text{Val}), (\vec{w} : \text{List Val}) \langle v \rangle \{ \text{is_list } v \vec{w} \}. ! \langle \vec{w} \rangle. \text{end}$

This subprotocol relation captures that the channel endpoint can receive a list of arbitrary values \vec{w} , instead of a list of integers \vec{x} . It does so formally by weakening the received proposition `is_int_list v \vec{x}` , and its associated binder \vec{x} , to the strictly weaker proposition `is_list v \vec{w}` , along with its binder \vec{w} . The tail of the protocol is suitably updated, to send back the length of the list of arbitrary values $|\vec{w}|$.

The manuscript additionally extends the model and mechanisation sections of the POPL’20 publication, to provide further insight into how the Actris framework is defined on top of Iris, and how the mechanisation made use of the Iris Proof Mode.

Personal Contribution The manuscript is joint work where I wrote the initial draft of each of the three new/ revised sections, while my co-authors provided feedback and revisions. The extension of the mechanisation was achieved as a collaboration with Robbert Krebbers, where Krebbers solved technically challenging problems related updating the recursive domain equation, the step-indexing of Iris, and the tactic support for the subprotocol mechanism, while I worked on the rest under his supervision, such as conceptualising the subprotocol definition and integrating it with our existing protocol framework.

1.6.3 Machine-Checked Semantic Session Typing. Jonas Kastberg Hinrichsen, Daniël Louwink, Jesper Bengtson, and Robbert Krebbers, CPP’21, recipient of Distinguished Paper Award

This paper draws a formal connection between the dependent separation protocols and session types, by defining and mechanising a semantic session type system as a logical relation on top of Actris.

In doing so, we obtain a rich session type system, with term- and session type subtyping, polymorphism, and equi-recursion, in addition to copyable types, mutable and shared reference types, and mutex types. Furthermore, it integrates the expressiveness of the underlying logic with the simplicity of the session type system, by enabling manually proving typing judgements, of uncheckable programs, such as the one originally presented in Section 1.4.2:

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

As previously stated, the program is safe, as the order in which it receives does not matter, but even so, it cannot be type checked by conventional session type systems, as the channel type cannot be distributed to both parallel threads. However, we can manually prove the typing judgement, using the underlying logic, and then use it in the context of type checking for compound programs.

Personal Contribution The writing was joint work with close collaboration between the co-authors. I wrote the initial draft of all sections excluding the mechanisation section (initial version by Robbert Krebbers) and the related work section (initial version by Jesper Bengtson). The mechanisation was developed in collaboration with Daniël Louwink, who laid the ground work inspired by the RustBelt project, and Robbert Krebbers, who assisted with technical Coq challenges *e.g.*, related to applying polymorphic typing rules. I did the rest, such as extending the system with subtyping and carrying out the mechanisation of the examples in the paper.

1.6.4 Mechanisations and Artifacts

The Coq mechanisation for the three paper share the same code base and thus repository, for which the newest version is found at <https://gitlab.mpi-sws.org/iris/actris>. Additionally, each of the published papers have an associated self-contained persistently **archived** virtual machine, along with a **repo branch** of the repository reflecting the source code at their time of publication:

- POPL’20: <https://dl.acm.org/doi/10.1145/3371074>
 - archive: <https://dl.acm.org/doi/10.1145/3373096/full/>
 - repo branch: <https://gitlab.mpi-sws.org/iris/actris/-/tree/popl20>
- LMCS: <https://iris-project.org/pdfs/2020-actris2-submission.pdf> (pre-print)
 - repo branch: <https://gitlab.mpi-sws.org/iris/actris/-/tree/lmcs>
- CPP’21: <https://doi.org/10.1145/3437992.3439914>
 - archive: <https://zenodo.org/record/4322752>
 - repo branch: <https://gitlab.mpi-sws.org/iris/actris/-/tree/cpp21>

Background

The work of this thesis is carried out on top of the higher-order concurrent separation logic Iris in Coq [Jung et al. 2015, 2018b; Iris Development Team 2021]. Iris is language independent, *i.e.*, it is parametric on an operational semantics. For the purpose of this thesis we consider the instantiation with an untyped functional language with higher-order functions, garbage collected higher-order mutable references, and fork-based concurrency, called HeapLang.

In this chapter we provide a technical background for the remaining chapters. In particular, we first present the operational semantics of HeapLang, along with existing implementations of parallel composition and locks, as well as our own implementations of linked lists and binary channels for message passing (Section 2.1). We specifically consider the problem of capturing all of the possible interleavings of a concurrent program, and how this is modelled semantically.

We then provide a precise definition of the safety and functional correctness properties, based on the operational semantics of HeapLang, that we consider (Section 2.2). We additionally show how these properties can be established using a type system or Hoare logic, and how the benefits of these techniques can be combined using the method of semantic typing. In doing so, we precisely capture the functional correctness property that we obtain with the Actris framework presented in Chapter 3, and the safety property that we obtain with the semantic type system built on top of Actris presented in Chapter 4.

We finally give an overview of the Iris framework, and demonstrate how it can be used to prove specifications of concurrent programs that share exclusive resources using locks (Section 2.3). That is to give an intuition for how we may later use Iris as the foundation for constructing our session-type based protocols, and for how we prove the specifications of our binary channels for message passing in the Actris framework in Chapter 3. We additionally present some of the features of Iris that we use to model various types in our semantic session type system presented in Chapter 4.

The remaining chapters of the thesis are self-contained in regards to session types, and thus we do not cover them in this chapter.

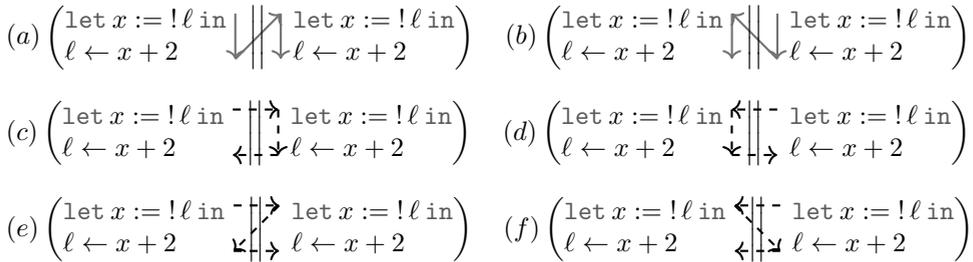


Figure 2.1: The possible interleavings of `interleaving_example`. The interleavings (a) and (b) yield the correct result, in contrast to (c), (d), (e), and (f) which yields an incorrect result. Arrows indicate the flow of time.

2.1 Operational Semantics

In this section we cover the operational semantics of HeapLang and the abstractions, that we use in the remaining chapters, built on top of it. In particular, we first give a brief account of the intricacies associated with concurrent operational semantics (Section 2.1.1). We then present the concrete operational semantics of the HeapLang language (Section 2.1.2). We then cover the existing implementations of locks and parallel composition (Section 2.1.3). Finally, we cover our own implementations of mutable linked lists (Section 2.1.4), and bidirectional channels for message passing (Section 2.1.5).¹

2.1.1 Interleavings in Concurrent Operational Semantics

Reasoning about concurrent software is notoriously hard, since the program instructions of individual threads are sporadically interleaved. As a consequence of this, even though some possible execution orders might produce the intended result, there can still exist interleavings that result in faulty behaviour. The following program illustrates this problem:

$$\text{interleaving_example } \ell := \left(\begin{array}{c} \text{let } x := !\ell \text{ in} \\ \ell \leftarrow x + 2 \end{array} \parallel \begin{array}{c} \text{let } x := !\ell \text{ in} \\ \ell \leftarrow x + 2 \end{array} \right)$$

The program concurrently updates a value x stored in the reference ℓ twice. It does so by executing two threads in parallel, that reads the value from the ℓ reference, adds 2 to the value, and stores the updated value back in the reference.

The problem is that four of its six possible interleavings result in an incorrect result, as shown in Figure 2.1.² The program correctly adds 4 to the referenced value

¹The content of (Section 2.1.5) is primarily replicated from the corresponding section on channel semantics in (Section 3.6.5), in order to give a complete overview of the language semantics.

²We assume that the language has a sequentially-consistent memory model.

in the first two cases (a) and (b), as the reference is both read and updated by one thread, before the other thread reads the updated value. However, in the remaining cases (c), (d), (e), and (f), both threads read the original value (x), update their local view of the value (to $x + 2$), and then store that in the reference. As a result only one update is effectively applied, resulting in the reference now containing $x + 2$, as opposed to $x + 4$. We fix this bug by protecting the critical sections with a lock in Section 2.1.3, and prove that the program is both safe and correct in Section 2.3.

To obtain proper guarantees about our concurrent programs, we want our proofs of safety and functional correctness to account for any possible interleaving. This is achieved by reflecting the interleaving behaviour in the operational semantics.

2.1.2 The Semantics of HeapLang

HeapLang is an untyped functional language with higher-order functions, garbage collected higher-order mutable references, and fork-based concurrency, for which the syntax is as follows:³

$$\begin{aligned}
 v \in \text{Val} &::= () \mid i \mid b \mid \ell \mid \text{rec } f \ x := e \mid && (i \in \mathbb{Z}, b \in \mathbb{B}, \ell \in \text{Loc}) \\
 &(v_1, v_2) \mid \text{inj}_1 v \mid \text{inj}_2 v \mid \dots \\
 e \in \text{Expr} &::= v \mid x \mid e_1 \ e_2 \mid (e_1, e_2) \mid \text{fst } e \mid \text{snd } e \mid \\
 &\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid \text{inj}_1 e \mid \text{inj}_2 e \mid \\
 &(\text{match } e_1 \text{ with } (\text{inj}_1 x) \Rightarrow e_2 \mid (\text{inj}_2 x) \Rightarrow e_3 \text{ end}) \mid \\
 &\text{ref } e \mid !e \mid e_1 \leftarrow e_2 \mid && (\text{Mutable state}) \\
 &\text{fork } \{e\} \mid \text{CAS } e_1 \ e_2 \ e_3 \mid \dots && (\text{Concurrency})
 \end{aligned}$$

We omit the the standard unary and binary operations such as equality, addition, and subtraction. The language includes the standard operations for heap manipulation. New references are created using $\text{ref } e$, dereferenced using $!e$, and assigned to using $e_1 \leftarrow e_2$. Concurrency is supported via $\text{fork } \{e\}$, which spawns a new thread, that executes the expression e in the background. The language also supports the atomic compare-and-set (CAS) operation, which is used to implement the lock in Section 2.1.3.

Derived Notions We derive various notions as syntactic sugar, *i.e.*, definitions in the meta-language by use of \triangleq , such as lambda expressions $(\lambda x. e)$, let-bindings $(\text{let } x := e_1 \text{ in } e_2)$, and options $(\text{None}, \text{Some } e)$, from the core syntax:

$$\begin{aligned}
 \lambda x. e &\triangleq \text{rec } _ \ x := e && \text{None} \triangleq \text{inj}_1 () \\
 \text{let } x := e_1 \text{ in } e_2 &\triangleq (\lambda x. e_2) e_1 && \text{Some } e \triangleq \text{inj}_2 e \\
 e_1; e_2 &\triangleq \text{let } _ := e_1 \text{ in } e_2 && \text{skip} \triangleq (\lambda _ . ()) ()
 \end{aligned}$$

We often write definitions as $f \ x_1 \cdots x_n := e$ rather than $f \triangleq \text{rec } f \ x_1 \cdots x_n := e$.

³HeapLang has more features, *e.g.*, arrays, but for brevity's sake we describe a more concise version of the operational semantics. However, the mechanisation effort of the thesis is built on top of the full version of HeapLang.

$$\begin{array}{l}
 ((\text{rec } f \ x := e)(v); \sigma) \rightarrow_h (e[v/x][(\text{rec } f \ x := e)/f]; \sigma; []) \\
 (\text{fst } (v_1, v_2); \sigma) \rightarrow_h (v_1; \sigma; []) \\
 (\text{snd } (v_1, v_2); \sigma) \rightarrow_h (v_2; \sigma; []) \\
 (\text{if true then } e_1 \text{ else } e_2; \sigma) \rightarrow_h (e_1; \sigma; []) \\
 (\text{if false then } e_1 \text{ else } e_2; \sigma) \rightarrow_h (e_2; \sigma; []) \\
 \left(\begin{array}{l} \text{match } (\text{inj}_i \ v) \text{ with} \\ (\text{inj}_1 \ x) \Rightarrow e_1 \\ | (\text{inj}_2 \ x) \Rightarrow e_2 \\ \text{end} \end{array} ; \sigma \right) \rightarrow_h (e_i[v/x]; \sigma; []) \quad \text{if } i \in \{1, 2\} \\
 (\text{ref } v; \sigma) \rightarrow_h (\ell; \sigma[\ell \leftarrow v]; []) \quad \text{if } \sigma(\ell) = \perp \\
 (!\ell; \sigma[\ell \leftarrow v]) \rightarrow_h (v; \sigma[\ell \leftarrow v]; []) \\
 (\ell \leftarrow w; \sigma[\ell \leftarrow v]) \rightarrow_h ((); \sigma[\ell \leftarrow w]; []) \\
 (\text{fork } \{e\}; \sigma) \rightarrow_h ((); \sigma; [e]) \\
 (\text{CAS } \ell \ w_1 \ w_2; \sigma[\ell \leftarrow v]) \rightarrow_h (\text{true}; \sigma[\ell \leftarrow w_2]; []) \quad \text{if } v = w_1 \\
 (\text{CAS } \ell \ w_1 \ w_2; \sigma[\ell \leftarrow v]) \rightarrow_h (\text{false}; \sigma[\ell \leftarrow v]; []) \quad \text{if } v \neq w_1 \\
 \\
 \frac{e_1; \sigma_1 \rightarrow_h e_2; \sigma_2; \vec{e}}{K[e_1]; \sigma_1 \rightarrow_t K[e_2]; \sigma_2; \vec{e}} \quad \frac{e_1; \sigma_1 \rightarrow_t e_2; \sigma_2; \vec{e}}{T \cdot [e_1] \cdot T'; \sigma_1 \rightarrow_{\text{tp}} T \cdot [e_2] \cdot T' \cdot \vec{e}; \sigma_2}
 \end{array}$$

Figure 2.2: The operational semantics of HeapLang.

Operational Semantics The small-step operational semantics of HeapLang can be found in Figure 2.2. It consists of three layers of reductions:

- *term reductions*, $(-; - \rightarrow_h -; -; -) \in (\text{Expr} \times \text{State}) \times ((\text{Expr} \times \text{State}) \times \text{List Expr})$ which capture how individual operations reduce
- *thread-local reduction*, $(-; - \rightarrow_t -; -; -) \in (\text{Expr} \times \text{State}) \times ((\text{Expr} \times \text{State}) \times \text{List Expr})$ which capture the evaluation order of sub-expressions
- *threadpool reduction*, $(-; - \rightarrow_{\text{tp}} -; -) \in (\text{List Expr} \times \text{State}) \times (\text{List Expr} \times \text{State})$ which capture the concurrent semantics of the language

Here $\sigma \in \text{State} \triangleq \text{Loc} \xrightarrow{\text{fin}} \text{Val}$, *i.e.*, a finite partial map from allocated locations to their stored values.

The *term reduction* \rightarrow_h captures how the language primitives reduce. The reduction is defined as a relation from a configuration $(\text{Expr} \times \text{State})$, to a configuration and a range of newly spawned threads $((\text{Expr} \times \text{State}) \times \text{List Expr})$. The individual primitive reductions are mostly standard, and are shown in Figure 2.2. The reduction of the fork-expression $(\text{fork } \{e\}; \sigma) \rightarrow_h ((); \sigma; [e])$ captures how a new thread e is spawned by adding it to the list of newly spawned threads $[e]$.

The *thread-local reduction* \rightarrow_t is used to capture the deterministic evaluation order of sub-expressions, using the following *evaluation context* syntax:

$$\begin{aligned}
K \in \text{Ctx} ::= & \bullet \mid e(K) \mid K(v) \mid (e_1, K) \mid (K, v_2) \mid \text{fst } (K) \mid \text{snd } (K) \mid \\
& \text{if } K \text{ then } e_1 \text{ else } e_2 \mid \text{inj}_1(K) \mid \text{inj}_2(K) \mid \\
& (\text{match } K \text{ with } (\text{inj}_1 x) \Rightarrow e_2 \mid (\text{inj}_2 x) \Rightarrow e_3 \text{ end}) \mid \\
& \text{ref } (K) \mid !K \mid e \leftarrow K \mid K \leftarrow v \mid \quad (\text{Mutable state}) \\
& \text{CAS } e_1 \ e_2 \ K \mid \text{CAS } e_1 \ K \ v_2 \mid \text{CAS } K \ v_1 \ v_2 \mid \dots \quad (\text{Concurrency})
\end{aligned}$$

In particular, expressions are represented in the form $K[e]$, capturing that e is the head-expression of a context K . HeapLang reduces from right-to-left, meaning that in expressions such as $e_1 \leftarrow e_2$ the expression e_2 reduces before e_1 . This is determined by its corresponding context syntax, from which we initially get $(e_1 \leftarrow \bullet)[e_2]$. If e_2 reduces to a value v_2 the context syntax dictates that the hole moves to e_1 yielding $(\bullet \leftarrow v_2)[e_1]$. If e_1 reduces to a value v_1 we finally end up with $v_1 \leftarrow v_2$, as there is no context syntax where both constituents are values.

Finally, the *threadpool reduction* \rightarrow_{tp} captures the concurrent semantics of HeapLang, being the top-level reduction of the language. The threadpool reduction is defined as a relation over all concurrently running threads and the heap state ($\text{List Expr} \times \text{State}$) to a new set of threads and an updated heap state ($\text{List Expr} \times \text{State}$). At each step a thread is picked arbitrarily and reduced once via the thread reduction \rightarrow_t , applying the updates to the heap state, while adding any forked-off threads \vec{e} to the threadpool $T \cdot [e_2] \cdot T' \cdot \vec{e}$.

As the reduced thread is picked arbitrarily, a program execution simulation can result in any possible interleaving, even ones where the same thread is infinitely picked repeatedly (assuming it loops), effectively accounting for an unfair scheduler. However, as will be detailed in Section 2.2 we are not concerned with termination, and thus an infinitely executing program is deemed well-behaved. As a result of the arbitrary reduction of threads, any proof over a program execution must consider every possible interleaving, even the unfair ones.

2.1.3 Implementation of Locks and Parallel Composition

Iris’s HeapLang library [Iris Development Team 2021, HeapLang] includes implementations of locks and the parallel composition, found in Figure 2.3.

The locks are encoded as a boolean reference, where `false` indicates that the lock is free, and `true` indicates that the lock is taken. New locks are constructed with `new_lock ()`, which does so by allocating a reference to `false`, *i.e.*, a free lock. Acquiring the lock is primarily done with `try_acquire lk`, which uses a compare-and-set (CAS) operation to atomically query whether the lock is free (`false`). If the lock is free, it is taken by atomically setting the reference to `true`. The function returns whether the CAS operation was successful. To make lock acquisition a *blocking operation*, *i.e.*, one that only continues once it succeeds, we implement the actual acquire function `acquire lk` by wrapping `try_acquire lk` in a loop, which only terminates once the CAS operation succeeds. Finally, the lock can be released with `release lk` which sets the reference back to `false`.

Implementation of locks:

```

new_lock () := ref false
try_acquire lk := CAS lk false true
acquire lk := if (try_acquire lk) then () else acquire lk
release lk := lk ← false

```

Implementation of parallel composition:

```

spawn f := let cb := ref None in      (e1 || e2)  $\triangleq$  let h := spawn ( $\lambda\_ . e_1$ ) in
  fork {cb ← Some (f ())};           let v2 := e2 in
  cb                                  let v1 := join h in
  (v1, v2)
join cb := match !cb with
  None   $\Rightarrow$  join cb
  | Some x  $\Rightarrow$  x
end

```

Figure 2.3: Implementation of locks and parallel composition on top of HeapLang.

The parallel composition $(e_1 \parallel e_2)$ computes the results of e_1 and e_2 in parallel, and returns the resulting values v_1 and v_2 as a tuple. The implementation depends on two auxiliary functions `spawn` and `join`. The `spawn` function takes a thunk, *i.e.*, a computation delayed by wrapping it in a closure, which it computes in a forked-off thread, while returning a callback `cb`. The callback is a reference to an option, where `None` indicates that the value has yet to be computed, and `Some x` means that the result x is ready. The `join` function takes a callback `cb`, and checks if the result is ready. If it is not, the function loops. If the result is ready, it is returned. The parallel composition $(e_1 \parallel e_2)$ is defined in the meta-language (using \triangleq), as opposed to using a lambda abstraction in HeapLang, to avoid eagerly computing e_1 and e_2 , when passing them to the function. The function spawns a new thread for computing the result of e_1 using `spawn ($\lambda_ . e_1$)`, after which it computes the result of e_2 locally. When the evaluation of e_2 is done, the function awaits the result of e_1 using `join`. Finally, the results v_1 and v_2 are returned as a tuple.

The Lock-Protected Interleaving Example By employing the locks we can define a correct variant of the `interleaving_example` example from Section 2.1.1:

```

locked_interleaving_example  $\ell$  :=
  let lk := new_lock () in
  (
    (
      acquire lk;
      let x := ! $\ell$  in
       $\ell$  ← x + 2;
      release lk
    ||
    (
      acquire lk;
      let x := ! $\ell$  in
       $\ell$  ← x + 2;
      release lk
    )
  );
  acquire lk

```

```

  lnil () := ref None
  lcons v l := l ← Some (v, ref (!l))
  lisnil l := match !l with
    None ⇒ true
    | Some p ⇒ false
    end
  lsnoc l v := match !l with
    None ⇒ l ← Some (v, ref None)
    | Some p ⇒ lsnoc (snd p) v
    end
  lpop l := match !l with
    None ⇒ () ()
    | Some p ⇒ l ← !(snd p); fst p
    end

```

Figure 2.4: A selection of the definitions for the implementation of linked lists on top of HeapLang. We omit the functions `llength`, `lappend`, `lreverse`, `lmerge`, and `lsplit`, used in the remaining chapters.

The lock guarantees that only one thread can access the location ℓ at a time. As a result, the possible interleavings of the program reduce to the correct cases (a) and (b), thus guaranteeing that the program correctly adds 4 to the stored value x .

We acquire the lock after the parallel composition, to simulate a deallocation of the lock. This is done to obtain a simpler proof of the functional correctness of the program in Section 2.3, inspired by [Jourdan and Krebbers 2018, Exercise 3].

2.1.4 Implementation of a Mutable Linked List

We implement mutable linked lists on top of HeapLang, for which a selection of definitions can be found in Figure 2.4. We take the common approach to linked lists, encoding them as references to options. A list may then either be empty (`None`), or a cell (`Some (v, l)`), where v is the value stored in the cell and l is a reference to the next cell. This is made apparent by the definition of `lnil ()`, which allocates a reference to a new empty list `ref None`. Consequentially, the definition `lcons v l` extends a linked list l by letting the reference point to a new cell with the new value v , and a newly allocated reference to the original tail `ref (!l)`.

The linked list implementation includes a small library of functions of which a selection is shown in Figure 2.4. In particular, we use the functions `lisnil`, `lsnoc`, and `lpop`, for the implementation of the binary channels for message passing in Section 2.1.5. The `lisnil l` function queries whether the list l is empty, by matching on its option constructor. The `lsnoc v l` function adds a value v to the end of the list l , by traversing the entire list. Finally, the `lpop l` function removes the first element of the list l and returns it. The function fails if the list is not empty, which is captured using the illegal function application `() ()`, *i.e.*, an application of the unit value `()` to itself, as if it were a function. As there are no valid reductions for a function application on values that are not functions, the program gets stuck, which will be discussed further in Section 2.2. Finally, note that we often use $|l|$ instead of `llength l`.

```

new_chan () := let (l, r, lk) := (lnil (), lnil (), new_lock ()) in
  ((l, r, lk), (r, l, lk))

send c v := let (l, r, lk) := c in
  acquire lk;
  lsnoc l v; skipN |r|;
  release lk

try_recv c := let (l, r, lk) := c in
  acquire lk;
  let ret := (if (lisnil r) then (None) else (Some (lpop l))) in
  release lk; ret

recv c := match (try_recv c) with
  None   => recv c
| Some v => v
end

```

Figure 2.5: Implementation of binary channels for message passing in HeapLang.

2.1.5 Implementation of Binary Channels for Message Passing

We implement our channels on top of HeapLang, using a pair of mutable linked lists protected by a lock, as shown in Figure 2.5. Each linked list represent a buffer of messages sent from one endpoint to the other.

New channels are created by the `new_chan` function, which allocates two empty mutable linked lists l and r using `lnil ()`, along with a lock lk using `new_lock ()`, and returns two channel endpoints, encoded as the tuples (l, r, lk) and (r, l, lk) , where the order of the linked lists l and r determines the side of the endpoints. We refer to the list in the left position as the endpoint’s own buffer, and the list in the right position as the other endpoint’s buffer.

Values are sent over a channel endpoint (l, r, lk) using the `send` function. This function operates in an atomic fashion by first acquiring the lock via `acquire lk`, thereby entering the critical section, after which the value is enqueued (*i.e.*, appended to the end) of the endpoint’s own buffer using `lsnoc l v`. The `skipN |r|` instruction is a no-op that is inserted to aid the proof. We explain the reason why this instruction is needed in Section 3.6.6.

Values are received over a channel endpoint (l, r, lk) using the `recv` function. The function performs a loop that repeatedly calls the helper function `try_recv`. It loops whenever `try_recv` fails (*i.e.*, returns `None`), and terminates when it succeeds (*i.e.*, returns `Some v`), returning the value v . The helper function `try_recv` attempts to receive a value atomically, and fails if there is no value in the other endpoint’s buffer. It does so by first acquiring the lock lk with `acquire lk`, after which it checks whether the other endpoint’s buffer is empty using `lisnil r`. If the buffer is empty nothing is returned (*i.e.*, `None`). If it is non-empty the value is dequeued and returned (*i.e.*, `Some (lpop l)`). Finally, the lock lk is released with `release lk`.

2.2 Safety, Functional Correctness, and Semantic Typing

This section elaborates on how safety and functional correctness of a program can be proven using type checking and Hoare logic. First, we formally define what we mean by safety, and how it can be shown via type checking (Section 2.2.1). In doing so, we illustrate a key limitation of conventional type systems, in that they cannot type check programs whose safety rely on runtime behaviour, *e.g.*, programs with unsafe branches that are never executed. We then formally define what we mean by functional correctness, how it can be shown through proofs in a program logic, and how it can be used to prove safety of programs beyond the scope of type checking (Section 2.2.2). Finally, we cover the technique of semantic typing, and how it can be used to combine the benefits of type checking and interactive proofs (Section 2.2.3).

2.2.1 Safety

Safety of a program has informally been described as “something [bad] will *not* happen” [Lampert 1977]. We interpret this as a program always being able to reduce, *i.e.*, being able to take a step according to the small-step operational semantics. This implies that the program is type-safe, as illegal (primitive) function applications cannot reduce. An example of an illegal function application is $()()$, as the literal unit value $()$ is used as a function. Additionally, the safety interpretation implies that programs are memory-safe, as the reduction rule for dereferencing requires that the location is in the heap to take a step. We capture the safety property as follows:

$$\begin{aligned} \text{safe } e \triangleq \forall \sigma, T, \sigma'. ([e]; \sigma \rightarrow_{\text{tp}}^* T; \sigma') \\ \text{implies } \forall e' \in T. (e' \in \text{Val}) \text{ or} \\ (\exists T', \sigma''. e'; \sigma' \rightarrow_{\text{tp}} T'; \sigma'') \end{aligned}$$

The property states that an expression e is safe, whenever any expression $e' \in T$ that it reduces to is either a value or can reduce further. As infinitely looping programs can always reduce, they are deemed safe.

Safety is often shown using a type system, which consists of types and typing rules, capturing the set of values that the program terms can have. The types and typing rules are most commonly defined as inductive definitions and relations, respectively. For demonstration purposes, we consider a simple type system that has ground types (A_g) for integers and booleans, as well as a term type (A) for function types:

$$A_g ::= \mathbf{Z} \mid \mathbf{B} \mid \dots \quad A ::= A_g \mid A_1 \rightarrow A_2 \mid \dots$$

The type system captures that an expression e has the type A with the typing judgement $\Gamma \vdash e : A$, where Γ is a typing context capturing the types of free variables. A selection of the typing rules for the type system is as follows:

$$\begin{array}{c} \overline{\Gamma \vdash i : \mathbf{Z}} \qquad \overline{\Gamma \vdash b : \mathbf{B}} \qquad \frac{\Gamma, x : A \vdash e : B}{\Gamma \vdash \lambda x. e : A \rightarrow B} \\ \\ \frac{\Gamma \vdash e_1 : A \rightarrow B \quad \Gamma \vdash e_2 : A}{\Gamma \vdash e_1 e_2 : B} \qquad \frac{\Gamma \vdash e_1 : \mathbf{B} \quad \Gamma \vdash e_2 : A \quad \Gamma \vdash e_3 : A}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : A} \end{array}$$

Safety of a closed expression e can then be established by checking that it is well-typed in the empty typing context ($\vdash e : A$), according to the typing rules.

Indeed, safety follows from type checking under the condition that the type system enjoys the property of *type safety*, stating that “well-typed programs cannot go wrong” [Milner 1978], which is formally captured as:

$$\vdash e : A \quad \text{implies} \quad \text{safe } e \tag{2.1}$$

This property states that if a program is well-typed ($\vdash e : A$) it is also safe ($\text{safe } e$). Proving that a type system satisfies the type safety property is most commonly done using using the technique of *progress* and *preservation* [Wright and Felleisen 1994].

- *progress*: well-typed programs are either values or they can take a step
- *preservation*: well-typed programs are still well-typed after taking a step

In particular, the properties are individually proven using induction, after which type safety follows directly.

Safety via Type Checking The following are examples of the type checking of two programs; The first being safe:

$$\frac{\frac{\overline{b : \mathbf{B} \vdash b : \mathbf{B}} \checkmark \quad \overline{b : \mathbf{B} \vdash 42 : \mathbf{Z}} \checkmark \quad \overline{b : \mathbf{B} \vdash 0 : \mathbf{Z}} \checkmark}{b : \mathbf{B} \vdash \text{if } b \text{ then } 42 \text{ else } 0 : \mathbf{Z}}}{\vdash \lambda b. \text{if } b \text{ then } 42 \text{ else } 0 : \mathbf{B} \rightarrow \mathbf{Z}}$$

And the second being unsafe:

$$\frac{\frac{\overline{b : \mathbf{B} \vdash b : \mathbf{B}} \checkmark \quad \overline{b : \mathbf{B} \vdash 42 : \mathbf{Z}} \checkmark \quad \overline{b : \mathbf{B} \vdash ()() : \mathbf{Z}} \times}{b : \mathbf{B} \vdash \text{if } b \text{ then } 42 \text{ else } ()() : \mathbf{Z}}}{\vdash \lambda b. \text{if } b \text{ then } 42 \text{ else } ()() : \mathbf{B} \rightarrow \mathbf{Z}}$$

In both programs we first abstract over the argument b of the boolean type \mathbf{B} , and then type check the if-else-expression by checking that the conditional argument b is of type \mathbf{B} , and that both branches have the same type (here the integer type \mathbf{Z}). The check on the conditional argument $b : \mathbf{B}$ passes in both examples, as per the typing context ($b : \mathbf{B}$). The check on the branches pass for the first example, as both branches are literal integers of type \mathbf{Z} . However, the second program is ill-typed, as the second branch $()()$ does not type check, since $()$ is not a function.

Limitations of Type Checking In both of the above examples the safety of the program corresponds directly to whether it is well-typed, but such a correspondence is not always the case, as stated in Section 1.4.2. Consider the following program which is safe, even though it is not well-typed:

$$\frac{\overline{\vdash \text{true} : \mathbf{B}} \checkmark \quad \overline{\vdash 42 : \mathbf{Z}} \checkmark \quad \overline{\vdash ()() : \mathbf{Z}} \times}{\vdash \text{if true then } 42 \text{ else } ()() : \mathbf{Z}}$$

As before, the second branch of the program $()()$ is unsafe. However, the program as a whole is safe as the unsafe branch is never executed, since the conditional argument `true` holds vacuously. While conventional type systems are agnostic to such runtime behaviour, they can be captured when proving functional correctness, which in turn can prove that the program is in fact safe.

2.2.2 Functional Correctness

Functional correctness of a program is a property that is dependent on a given specification, which denotes the expected output of the program. Such a specification can be thought of as a predicate $\varphi \in \text{Val} \rightarrow \text{Prop}$. We call this predicate the postcondition, by which the correctness property can be captured formally as follows:⁴

$$\text{correct } (e, \varphi) \triangleq (\text{safe } e) \text{ and} \\ \forall \sigma, (v \in \text{Val}), T, \sigma'. (e; \sigma \rightarrow_{\text{tp}}^* [v] \cdot T; \sigma') \\ \text{implies } (\varphi v)$$

The definition combines safety with postcondition validity, stating that the postcondition φ holds for any value v that the expression e reduces to. The correctness property subsumes our safety property, meaning that we have the following:

$$\text{correct } (e, \varphi) \text{ implies } \text{safe } e \tag{2.2}$$

A formal approach to showing functional correctness is using a program logic, *e.g.*, Hoare logic [Hoare 1969], where *Hoare triples* $\{P\} e \{\Phi\}$ are used to ascribe a program e with a logical precondition $P \in \text{iProp}$ and postcondition $\Phi \in \text{Val} \rightarrow \text{iProp}$. We often write $\{P\} e \{v.Q\}$ instead of $\{P\} e \{\lambda v. Q\}$, and $\{P\} e \{Q\}$, when $v = ()$. The propositions often range over a program logic iProp , that extends a meta-logic Prop , with connectives to describe the state of the program, *e.g.*, the locations of the heap. We often implicitly coerce predicates in the meta-logic φ into predicates of the program logic, *e.g.*, in $\{\text{True}\} e \{\varphi\}$. A common variant of a program logic is separation logic, which is covered in Section 2.3.

Proving a Hoare triple is achieved by applying the rules of the program logic. As an example, the following are the rules for the conditional expression:

$$\begin{array}{c} \text{HT-IF-TRUE} \\ \frac{\{P\} e_1 \{\Phi\}}{\{P\} \text{ if true then } e_1 \text{ else } e_2 \{\Phi\}} \end{array} \qquad \begin{array}{c} \text{HT-IF-FALSE} \\ \frac{\{P\} e_2 \{\Phi\}}{\{P\} \text{ if false then } e_1 \text{ else } e_2 \{\Phi\}} \end{array}$$

Specifically, they capture that we are only concerned with the branch of the program that corresponds to the truth value of the boolean condition.

Connecting a Hoare-style program logic to the functional correctness of a program relies on a proof of *adequacy*, capturing that Hoare triples with meta-logic postconditions φ imply correctness up to that postcondition:

$$\{\text{True}\} e \{\varphi\} \text{ implies } \text{correct } (e, \varphi) \tag{2.3}$$

⁴Note that we are not concerned with whether a program *terminates* (namely *total correctness*), but use the term correctness in place of the more apt term *partial correctness*, for brevity's sake.

The postcondition φ is restricted to the meta-logic fragment of the program logic as functional correctness $\text{correct}(e, \varphi)$ is a meta-logic property. This may seem restrictive at first, but adequacy is only relevant for the specifications of the top-level expression of a program, *e.g.*, the main-function. As such, we are not concerned with properties about *e.g.*, the state of the heap, as it is soon to be deallocated.

The adequacy theorem is often obtained by modelling the Hoare triple in terms of the operational semantics of the language under consideration, such that it captures safety directly.

Functional Correctness via Hoare Logic An example of a (simplified) proof of functional correctness is the following derivation, which captures that the safe yet ill-typed program from the prior section is safe and that it results in an integer:

$$\frac{\overline{\{\text{True}\} 42 \{v. v \in \mathbb{Z}\}} \checkmark}{\{\text{True}\} \text{if true then } 42 \text{ else } () \{v. v \in \mathbb{Z}\}} \text{HT-IF-TRUE}$$

The derivation follows from HT-IF-TRUE, by which we only need to focus on the first branch, as the conditional argument is vacuously `true`. After this step the program is a value, specifically the literal integer `42`, for which the postcondition holds trivially. We can thus conclude that the program is safe by Equation (2.3) and Equation (2.2).

While more expressive than type checking, verifying whether there exists a valid derivation for a Hoare-style logic specification is often undecidable. Showing functional correctness via this technique thus often requires interactive effort, even when we only seek to prove safety, and not full functional correctness.

2.2.3 Semantic Typing

The techniques of type checking and program logic proofs may seem orthogonal, but they can be related through the technique of semantic typing [Milner 1978; Ahmed 2004; Ahmed et al. 2010] (contrasting syntactic typing, discussed thus far).

In a semantic type system the types and the typing judgement are given a semantic interpretation, respectively capturing what it means for a value to belong to a type, and that the program is safe to execute. The typing rules of a semantic type system are consequentially proven as lemmas, instead of being defined inductively.

The type safety of a semantic type system is a direct consequence of the definition of the typing judgement, as opposed to being proven using a technique such as progress and preservation by Wright and Felleisen [1994]. As a result, a semantic type system is *open*, where rules are abstractions capturing safety of specific term structures, in contrast to capturing a notion of being well-typed, which may then imply safety.

This openness effectively lets us add “ad-hoc” rules, whenever we reach a typing judgement with no corresponding typing rule, by proving it manually in the underlying model. This technique was applied by Jung et al. [2018a], who created a semantic type system for the Rust language, in which they manually prove safety of unsafe Rust libraries, *i.e.*, programs that deploy a complex form of sharing, such as

reference-counted pointers, which are beyond the scope of the type system. Achieving a similar result with the technique of progress and preservation is infeasible, as it would require typing rules for all mid-execution reductions.

Finally, as demonstrated by Appel et al. [2007] and Dreyer et al. [2010], building on top of an existing system, *e.g.*, Iris, a semantic type system can inherit the features of the existing system, *e.g.*, mutable references or recursion, its adequacy theorem, and any existing mechanisation effort.

A semantic type system corresponding to the simple syntactic type system shown in Section 2.2.1 can be obtained by building on top of the Hoare logic presented in Section 2.2.2. In particular, we define the types as a logical interpretation of the syntactic types $[\cdot] \in \mathbf{Type} \rightarrow \mathbf{Val} \rightarrow \mathbf{iProp}$, where *e.g.*, the type of integers is defined as $[\mathbf{Z}] \triangleq \lambda v. v \in \mathbb{Z}$. Drawing inspiration from the approach taken by Krebbers et al. [2017b], who built a semantic type system for a language similar to HeapLang, the typing judgement can be defined in terms of the Hoare triple:

$$\vDash e : A \triangleq \{\mathbf{True}\} e \{[A]\} \quad (2.4)$$

The semantic typing rules are then proven as lemmas. As an example, the following is the typing rule for conditionals, which is expanded to the underlying Hoare triples:

$$\frac{\vDash e_1 : \mathbf{B} \quad \vDash e_2 : A \quad \vDash e_3 : A}{\vDash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : A} \rightsquigarrow$$

$$\frac{\{\mathbf{True}\} e_1 \{v. v \in \mathbb{B}\} \quad \{\mathbf{True}\} e_2 \{[A]\} \quad \{\mathbf{True}\} e_3 \{[A]\}}{\{\mathbf{True}\} \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \{[A]\}}$$

The proof of the underlying Hoare triple follows from case analysis on the boolean result of evaluating e_1 , after which the rules HT-IF-TRUE and HT-IF-FALSE apply.

The connection to syntactic typing is captured by the fundamental lemma, often being a direct consequence of the semantic type system definitions, which states that the semantic type system subsumes the syntactic one:

$$\vdash e : A \quad \text{implies} \quad \vDash e : A \quad (2.5)$$

From the fundamental lemma, and the implications between Hoare triples, functional correctness, and safety, presented in Section 2.2.2, we then have the following subsumption relationship for the ground types A_g :⁵

$$\begin{array}{lll} \vdash e : A_g & \text{implies} & (2.5) \\ \vDash e : A_g & \rightsquigarrow & (2.4) \\ \{\mathbf{True}\} e \{[A_g]\} & \text{implies} & (2.3) \\ \text{correct } (e, [A_g]) & \text{implies} & (2.2) \\ \text{safe } e & & \end{array}$$

With these definitions and implications, we can then prove safety of syntactically ill-typed programs via semantic typing.

⁵ $\{\mathbf{True}\} e \{[A_g]\}$ implies $\text{correct } (e, [A_g])$ for ground types A_g , as adequacy (2.3) only hold for meta-logic predicates, which the semantic interpretation of ground types commonly are.

Safety of Syntactically Ill-Typed Programs via Semantic Typing The semantic typing judgement for the safe yet ill-typed program presented in Section 2.2.1 correspond to the Hoare triple that we proved in Section 2.2.2:

$$\begin{array}{c} \vDash \text{if true then } 42 \text{ else } ()() : \mathbf{Z} \quad \rightsquigarrow \\ \{\text{True}\} \text{if true then } 42 \text{ else } ()() \{v. v \in \mathbb{Z}\} \end{array}$$

The judgement can thus be used as a part of an otherwise syntactically checked typing derivation, even though it is internally ill-typed. This demonstrates how we can add ad-hoc rules for term structures that might otherwise not have a derivation in the syntactic system, *e.g.*, for the racy yet safe program presented in Section 1.4.2.

The ability to add such ad-hoc rules is how semantic typing allows for combining the simplicity of type checking with the expressivity of program logic proofs for showing safety. In particular, one can deploy type checking for the fragment of the program that does not rely on runtime behaviour, and manually prove the uncheckable remainder of the typing derivation in the underlying program logic.

The Foundational Approach In our semantic type system we take the *foundational approach* [Ahmed 2004], as similarly done by Jung et al. [2018a] for their semantic type system for Rust built on top of Iris. Taking the foundational approach means that we sidestep the syntactic type system, by not defining the types as inductive definitions. Instead we define the types as standalone predicates, *i.e.*, $\text{Type} \in \text{Val} \rightarrow \text{iProp}$, rather than defining them as interpretations of syntactic types. Recovering proofs similar to the type checking of a syntactic type system can then be achieved by proving semantic typing rules (often called compatibility lemmas in the logical relations literature), corresponding to all of the rules that would otherwise be part of the syntactic type system.

By taking the foundational approach, we sidestep the effort of defining and mechanising a syntactic type system, and instead fully inherit the existing work of the underlying logic, which in our case is the features, adequacy, and mechanisation of the higher-order concurrent separation logic Iris.

2.3 The Iris Logic

In this section we cover the basics of the higher-order concurrent separation logic Iris [Jung et al. 2015, 2016] instantiated with HeapLang, and how it can be used to prove specifications of concurrent programs that safely share mutable data, such as the locked interleaving program presented in Section 2.1.3. An overview of grammar and inference rules of the logic can be found in Figure 2.6.

Program specifications in Iris can, alike standard Hoare and separation logics, be captured in terms of Hoare triples $\{P\} e \{Q\}$. Additionally, a proof of a Hoare triple in Iris similarly imply functional correctness for meta-logic postconditions $\varphi \in \text{Val} \rightarrow \text{Prop}$, as formally captured by Iris’s adequacy theorem:

$$\{\text{True}\} e \{\varphi\} \quad \text{implies} \quad \text{correct}(e, \varphi)$$

The meta-logic of Iris is that of Coq, which is a higher-order propositional logic.

Grammar:

$$\begin{aligned}
\tau, \sigma &::= x \mid 0 \mid 1 \mid \mathbb{B} \mid \mathbb{N} \mid \mathbb{Z} \mid \text{Type} \mid \forall x : \tau. \sigma \mid \text{Loc} \mid \text{Val} \mid \text{Expr} \mid \text{iProp} \mid \dots \\
t, u, P, Q &::= x \mid \lambda x : \tau. t \mid t(u) \mid t(\tau) \mid \quad (\text{Polymorphic lambda-calculus}) \\
&\text{True} \mid \text{False} \mid P \wedge Q \mid P \vee Q \mid P \Rightarrow Q \mid \quad (\text{Propositional logic}) \\
&\forall x : \tau. P \mid \exists x : \tau. P \mid t = u \mid \quad (\text{Higher-order logic with equality}) \\
&\mu x : \tau. t \mid \triangleright P \mid \quad (\text{Guarded recursion}) \\
&P * Q \mid P \multimap Q \mid \ell \mapsto v \mid \{P\} e \{v. Q\} \mid \quad (\text{Separation logic}) \\
&\Box P \mid P \Rightarrow Q \mid \dots \quad (\text{Iris connectives})
\end{aligned}$$

Ordinary affine separation logic:

$$\begin{array}{c}
\text{AFFINE} \\
P * Q \Rightarrow P \\
\text{ADJOINT} \\
\frac{P * Q \vdash R}{P \vdash Q \multimap R} \\
\text{HT-FRAME} \\
\frac{\{P\} e \{w. Q\}}{\{P * R\} e \{w. Q * R\}} \\
\text{HT-VAL} \\
\{\text{True}\} v \{w. w = v\} \\
\text{HT-CSQ} \\
\frac{\Box(P \multimap P') \quad \{P'\} e \{w. Q'\} \quad \Box(\forall w. Q' \multimap Q)}{\{P\} e \{w. Q\}} \\
\text{HT-FORK} \\
\frac{\{P\} e \{\text{True}\}}{\{P\} \text{fork } \{e\} \{\text{True}\}} \\
\text{HT-BIND} \\
\frac{\{P\} e \{v. Q\} \quad \forall v. \{Q\} K[v] \{w. R\}}{\{P\} K[e] \{w. R\}}
\end{array}$$

Heap manipulation:

$$\begin{array}{c}
\text{HT-ALLOC} \\
\{\text{True}\} \text{ref } v \{\ell. \ell \mapsto v\} \\
\text{HT-LOAD} \\
\{\ell \mapsto v\} !\ell \{w. (w = v) * \ell \mapsto v\} \\
\text{HT-STORE} \\
\{\ell \mapsto v\} \ell \leftarrow w \{\ell \mapsto w\}
\end{array}$$

Recursion:

$$\begin{array}{c}
\triangleright\text{-INTRO} \\
P \Rightarrow \triangleright P \\
\text{LÖB} \\
(\triangleright P \Rightarrow P) \Rightarrow P \\
\mu\text{-UNFOLD} \\
(\mu x. t) = t[\mu x. t/x] \\
\text{HT-REC} \\
\frac{\{P\} e[v/x][(\text{rec } f \ x := e)/f] \{w. Q\}}{\{\triangleright P\} (\text{rec } f \ x := e) v \{w. Q\}} \\
\text{HT-PERS} \\
\frac{Q \Rightarrow \{P\} e \{w. R\}}{\{P * Q\} e \{w. R\}} \text{persistent}(Q)
\end{array}$$

Figure 2.6: The grammar and a selection of rules of Iris for HeapLang.

The derivation of the ill-typed yet safe program presented in Section 2.2.2 is valid in Iris, as it follows from the rules HT-IF-TRUE and HT-VAL (up to some simplification):

$$\frac{\frac{\{\text{True}\} \text{42 } \{v. v \in \mathbb{Z}\}}{\text{HT-VAL}}}{\{\text{True}\} \text{if true then 42 else } ()() \{v. v \in \mathbb{Z}\}} \text{HT-IF-TRUE}$$

The expressive power of Iris comes from the logical propositions $P, Q \in \text{iProp}$ used in the Hoare triples. In this section we cover the features that are used in the remaining chapters of the thesis.

We first cover the separation logic properties of Iris (Section 2.3.1), which inherently guarantee that programs do not interfere with each other, as a result of Iris treating propositions affinely. This treatment is fundamental to how we reason about our programs, and is the foundation of how we model affine types in our semantic type system in Chapter 4.

We subsequently present the notion of persistent resources (Section 2.3.2), which captures how some propositions may be duplicated, and can thereby be distributed to multiple sub-proofs. This allows us to share *e.g.*, specifications of higher-order functions multiple times, and is what lets us model copyable types in our semantic type system in Section 4.3.2.

We then detail how Iris is a step-indexed separation logic (Section 2.3.3), which allows us to reason about recursive programs, and construct recursive predicates. This is essential to how we achieve the recursion of our session-type based protocols in Section 3.6.1, and consequentially how we obtain recursive session types of our semantic type system in Section 4.3.3.

We additionally demonstrate how Iris allows us to prove specifications of concurrent program, by using locks (Section 2.3.4) and ghost state (Section 2.3.5). These are both crucial parts of how we prove the high-level specifications of our channel implementation, which will be presented in Section 3.6.

We finally provide a more detailed overview of how we will use Iris to achieve the contributions of the remaining chapters (Section 2.3.6).

2.3.1 Separation Logic in Iris

Iris is an affine separation logic, in which propositions can be thought of as resources, that can be used at most once. The propositions for one range over the heap of the program. In particular, the points-to predicate $\ell \mapsto v$ asserts exclusive ownership of a location ℓ , while stating that it contains the value v . Meanwhile, the separating conjunction $P * Q$ captures that the propositions P and Q hold for disjoint parts of the heap. The exclusive ownership of a location ℓ can thereby only ever validly exist on one side of such a separating conjunction, and otherwise result in a contradiction:

$$(\ell \mapsto v) * (\ell \mapsto w) \Rightarrow \text{False}$$

Finally, the separating implication $P -* Q$ is similar to regular implication $P \Rightarrow Q$, for the separating conjunction, as captured by the ADJOINT rule in Figure 2.6.

Separation logic allows for modular verification of subcomponents of a program, as specifications inherently guarantee that programs do not interfere with each other,

since their heap footprints are disjoint. We call this the separation property. The property is formally captured by the `HT-FRAME` rule, which states that a frame of resources described by the proposition R , that is disjoint from the resources described by the proposition P , is preserved throughout execution of e , and thus also hold disjoint from the resources described by the proposition Q . As a result, we obtain small-footprint specifications, where we only describe the exact resources needed, as we can thread the disjoint frame of resources through our proofs with `HT-FRAME`.

Examples of such small-footprint specifications are the rules `HT-ALLOC`, `HT-LOAD`, and `HT-STORE`. The `HT-ALLOC` rule yields exclusive ownership of the newly allocated reference $l \mapsto v$. The rule `HT-LOAD` states that we can obtain the value stored in an exclusively owned reference. Finally, the rule `HT-STORE` allows us to update the stored value of an exclusively owned reference. The exclusivity is crucial, as the view of the stored value could otherwise be inconsistent with other threads.

The separation property similarly ensures that concurrently running threads will not race for exclusive resources. This is made precise by the `HT-FORK` rule, which states that the subproof is modularly proven based on some disjoint part of the precondition. Finally, Iris is affine, meaning that resources can be dropped, as is standard for program logics over garbage collected languages. This property is made precise by the `AFFINE` rule.

Linked List Abstraction As an example of how separation logic can be used to implicitly capture disjointness of resources, consider the following definition, capturing what it means to be a linked list up to the interpretation predicate I :

$$\text{list_interp } (I : T \rightarrow \text{Val} \rightarrow \text{iProp}) (l : \text{Loc}) (\vec{x} : \text{List } T) \triangleq \begin{cases} l \mapsto \text{None} & \text{if } \vec{x} = [] \\ \exists v_1 l_2. I \ x_1 \ v_1 * l \mapsto \text{Some } (v_1, l_2) * \text{list_interp } I \ l_2 \ \vec{x}_2 & \text{if } \vec{x} = [x_1] \cdot \vec{x}_2 \end{cases}$$

Here, T is a polymorphic type that is inferred from the I argument. We often use the notation $l \mapsto_I \vec{x}$ for $\text{list_interp } I \ l \ \vec{x}$.

The definition is constructed as a fixpoint over a meta-level list \vec{x} . The definition captures the following:

1. The number of cells correspond to the length of the list \vec{x}
2. Separate exclusive ownership is asserted for each cell
3. The interpretation predicate I holds for each meta-value $x \in \vec{x}$, and its corresponding cell value v in the linked list

In particular, the predicate $l \mapsto_I \vec{x}$ asserts ownership of the linked list head-reference l , using the points-to connective. In the case that the list is empty ($\vec{x} = []$), it asserts that the reference points to nothing $l \mapsto \text{None}$. If the list is non-empty ($\vec{x} = [x_1] \cdot \vec{x}_2$ for some values x_1 and \vec{x}_2), it asserts that the reference points to a cell $l \mapsto \text{Some } (v_1, l_2)$, and that the value v_1 can be interpreted as x_1 via $I \ x_1 \ v_1$. Finally, the definition recursively captures separate ownership of the rest of the linked list $\text{list_interp } I \ l_2 \ \vec{x}_2$. The list predicate implicitly guarantees that the list does not form a cycle, as each of the value cells are guaranteed to be disjoint.

$$\begin{array}{l}
 \text{HT-LNIL} \\
 \{\text{True}\} \text{lnil } \{l. l \mapsto_I []\} \\
 \\
 \text{HT-LCONS} \\
 \{l \mapsto_I \vec{x} * I x v\} \text{lcons } v l \{l \mapsto_I ([x] \cdot \vec{x})\} \\
 \\
 \text{HT-LISNIL} \\
 \{l \mapsto_I \vec{x}\} \text{lisnil } l \{b. b = (\vec{x} = []) * l \mapsto_I \vec{x}\} \\
 \\
 \text{HT-LSNOC} \\
 \{l \mapsto_I \vec{x} * I x v\} \text{lsnoc } l v \{l \mapsto_I (\vec{x} \cdot [x])\} \\
 \\
 \text{HT-LPOP} \\
 \{l \mapsto_I ([x] \cdot \vec{x})\} \text{lpop } l \{w. I x w * l \mapsto_I \vec{x}\}
 \end{array}$$

Figure 2.7: Specifications of a selection of the functions for the mutable linked list implementation presented in Section 2.1.4.

The specifications of a selection of the the linked list functions is found in Figure 2.7. To demonstrate how we can prove such specifications of programs with multiple sub-expression, consider the specification for `lcons`, which is defined as follows:

$$\text{lcons } v l := l \leftarrow \text{Some } (v, \text{ref } (!l))$$

We first need to do case analysis on the meta-list \vec{x} , to obtain the reference ownership of l , which is present in both the empty and non-empty case. As the proof of both cases are similar, we only consider the empty case for simplicity sake:

$$\{l \mapsto \text{None} * I x v\} \text{lcons } v l \{l \mapsto_I [x]\}$$

As the program consists of multiple sub-expressions, *e.g.*, `!l`, we use the HT-BIND rule to string the sub-proofs of each sub-expression together. We thus start at the first sub-expression `!l`, being at the head-position of the program, captured by the surrounding evaluation context $K \triangleq (l \leftarrow \text{Some } (v, \text{ref } \bullet))$:

$$\{l \mapsto \text{None} * I x v\} (l \leftarrow \text{Some } (v, \text{ref } \bullet))[!l] \{l \mapsto_I [x]\}$$

We use the rule HT-LOAD to resolve the reference deallocation `!l`, which returns `None`. We then continue to the next sub-expression, being the reference allocation `ref None`:

$$\{l \mapsto \text{None} * I x v\} (l \leftarrow \text{Some } (v, \bullet))[\text{ref None}] \{l \mapsto_I [x]\}$$

Here we use the rule HT-ALLOC to resolve the reference allocation `ref None`, which returns a new location l' , for which we obtain $l' \mapsto \text{None}$. We then continue to the final sub-expression $l \leftarrow \text{Some } (v, l')$:

$$\{l \mapsto \text{None} * l' \mapsto \text{None} * I x v\} \bullet [l \leftarrow \text{Some } (v, l')] \{l \mapsto_I [x]\}$$

We resolve the step with the HT-STORE rule, which updates the reference predicate of l to $l \mapsto \text{Some } (v, l')$. We finally have:

$$\{l \mapsto \text{Some } (v, l') * l' \mapsto \text{None} * I x v\} () \{l \mapsto_I [x]\}$$

This final specification then follows directly from the definition of $l \mapsto_I \vec{x}$.

2.3.2 Persistent Resources in Iris

As Iris propositions are treated as affine resources they can by default be used at most once. However, some propositions, *e.g.*, equality, do not assert any properties about resource ownership, and are thus safe to reuse multiple times. It therefore makes sense to differentiate such propositions from affine propositions, as we might want to distribute them to multiple subproofs. In Iris we call these propositions persistent, which essentially mean that they can be duplicated. This is made precise by the \Box modality, which enjoys the following rules:

$$\text{persistent}(P) \triangleq P \Rightarrow \Box P \quad \begin{array}{l} \Box\text{-DUP} \\ \Box P \Rightarrow (\Box P) * (\Box P) \end{array} \quad \begin{array}{l} \Box\text{-ELIM} \\ \Box P \Rightarrow P \end{array}$$

A proposition P is considered to be persistent ($\text{persistent}(P)$), if it can be freely extended with the \Box modality, as captured by its definition. The rule $\Box\text{-DUP}$ captures that propositions under the \Box modality can be duplicated freely. Finally, The rule $\Box\text{-ELIM}$ captures that the \Box modality can always be eliminated, letting us obtain the persistent proposition. A persistent proposition can then be freely extended with the always modality \Box , after which it can be duplicated with $\Box\text{-DUP}$ and $\Box\text{-ELIM}$:

$$\begin{array}{l} P \Rightarrow (\text{persistent}(P)) \\ \Box P \Rightarrow (\Box\text{-DUP}) \\ \Box P * \Box P \Rightarrow (\Box\text{-ELIM}) \\ P * P \end{array}$$

The persistent propositions presented thus far are $t = u$, **True**, **False**, and $\Box P$. Moreover, the persistent modality commutes with conjunctions ($P \wedge Q$, $P * Q$), disjunction ($P \vee Q$), universal and existential quantification ($\forall x. P$, $\exists x. P$). The set of persistent propositions is therefore closed under these connectives.

Weakest Precondition We have thus far presented the program logic of Iris in terms of Hoare triples. However, in Coq the program logic of Iris is actually defined in terms of the weakest precondition $\text{wp } e \{ \Phi \}$. The weakest precondition captures that the expression e does not get stuck, and that if it reduces to a value v , then Φv holds. The actual adequacy statement of Iris is thus:

$$\text{wp } e \{ \varphi \} \text{ implies } \text{correct } (e, \varphi)$$

Like the Hoare triples, the weakest preconditions are first-class members of the logic, and can thus be nested.

The Hoare triples of Iris are used as an intuitive way of defining program specifications and rules, and is defined in terms of the weakest precondition:

$$\{ P \} e \{ \Phi \} \triangleq \Box (P \multimap \text{wp } e \{ \Phi \})$$

In particular, they state that under the assumption that the precondition P holds, as captured by the separating implication \multimap , the weakest precondition for the expression e and postcondition Φ holds. The definition is wrapped in the \Box modality to capture

that Hoare triples do not depend on external resources, beyond what is provided in their precondition P . This in turn makes them persistent. This is crucial as Hoare triples denote specifications of programs, which can be executed repeatedly, and should therefore be duplicable.

We omit the definition and rules for the weakest precondition. The rules are similar in spirit to the corresponding Hoare triple rules. For more details on the definition of the weakest precondition of Iris we refer the interested reader to Jung et al. [2018b, Section 7.3].

For presentation purposes, we primarily use the Hoare triples in Chapter 3, while we instead use weakest preconditions in Chapter 4.

2.3.3 Step-Indexing in Iris

Iris is a step-indexed separation logic [Appel et al. 2007; Hobor et al. 2008; Svendsen et al. 2010; Dreyer et al. 2010], where propositions are implicitly indexed by a natural number called a step-index. The step-index allows us to create structurally recursive propositions, in which the number gets smaller. Finally, the step-index is associated with the number of steps taken by a program, to capture that propositions only hold after a certain number of steps. As a result, we can differentiate between propositions that hold now, and those that hold later. The step-indexing of Iris gives us:

1. A way of constructing recursive predicates, with the guarded fixpoint operator $\mu x. t$, which *e.g.*, lets us define infinite channel protocols as demonstrated in Section 3.6.2
2. Löb induction, which lets us reason about recursive functions and predicates, demonstrated momentarily
3. Impredicative invariants, which lets us prove high-level specifications of concurrent programs that safely share exclusive resources, in terms of custom abstract predicates, which are seemingly first-class citizens of the logic.

The step-indexing of Iris manifests with the later modality \triangleright , which is used to guard proposition *e.g.*, $\triangleright P$, capturing that they hold after one step of computation. This property is demonstrated in the HT-REC rule:

$$\frac{\text{HT-REC} \quad \{P\} e[v/x][(\text{rec } f \ x := e)/f] \{w. Q\}}{\{\triangleright P\} (\text{rec } f \ x := e) v \{w. Q\}}$$

The rule states that the guarded precondition $\triangleright P$ is stripped of its later, when the function application has been resolved, *i.e.*, when a step has been taken.

Löb Induction The step-indexing of Iris allows for reasoning about recursive programs via Löb induction, an induction principle over the steps of the program. The induction principle manifests with the LÖB rule:

$$(\triangleright P \Rightarrow P) \Rightarrow P$$

The rule states that when proving a proposition P , we can assume that the proposition holds one program step later $\triangleright P$.

To demonstrate Löb induction, consider the following infinitely looping program:

$$\text{loop_prog } () := \text{loop_prog } ()$$

The program repeatedly calls itself recursively, taking a single step (the application of the function), at each iteration. We want to show that the program is safe, even though it never terminates. This can be captured with a Hoare triple where the postcondition is `False`. The proof of the program is as follows:

$$\frac{\frac{\frac{\overline{\{\text{True}\} \text{loop_prog } () \{\text{False}\} \Rightarrow \{\text{True}\} \text{loop_prog } () \{\text{False}\}} \text{HT-PERS}}{\{(\{\text{True}\} \text{loop_prog } () \{\text{False}\})\} \text{loop_prog } () \{\text{False}\}} \text{HT-REC}}{\{\triangleright(\{\text{True}\} \text{loop_prog } () \{\text{False}\})\} \text{loop_prog } () \{\text{False}\}} \text{LÖB, HT-PERS}}{\{\text{True}\} \text{loop_prog } () \{\text{False}\}}$$

We first use the LÖB rule, to obtain the Löb induction hypothesis, along with HT-PERS, to move the guarded specification induction hypothesis into the Hoare triple precondition. We then strip the later of the hypothesis by resolving the function application with HT-REC. Finally, the proof follows immediately by moving the specification out of the precondition with HT-PERS.

2.3.4 Reasoning about Concurrent Programs in Iris

The separation property of separation logic guarantees that resources are disjoint. This lets us reason about concurrently running threads that work on disjoint parts of the heap, as they are guaranteed to not interfere with each other.

However, we often wish to reason about concurrent threads that safely access overlapping memory footprints. Consider the `locked_interleaving_example` program from Section 2.1.3, which safely share a location ℓ by protecting it using a lock around the critical sections of the parallel threads in which ℓ is accessed:

$$\begin{aligned} \text{locked_interleaving_example } \ell := & \\ & \text{let } lk := \text{new_lock } () \text{ in} \\ & \left(\begin{array}{c|c} \text{acquire } lk; & \text{acquire } lk; \\ \text{let } x := !\ell \text{ in} & \text{let } x := !\ell \text{ in} \\ \ell \leftarrow x + 2; & \ell \leftarrow x + 2; \\ \text{release } lk & \text{release } lk \end{array} \right); \\ & \text{acquire } lk \end{aligned}$$

For starters we simply want to prove safety, specified as follows:

$$\{\ell \mapsto x\} \text{locked_interleaving_example } \ell \{\text{True}\}$$

To prove such a specification, we can use the high-level specifications for the lock and parallel composition abstractions shown in Figure 2.8. These specifications are inspired by the lock implementation by Jourdan and Krebbers [2018, Exercise 3] and the implementation of the parallel composition by Iris Development Team [2021, HeapLang], who also prove their respective specifications on top of Iris.

Specifications of locks:

<p>HT-NEWLOCK $\{R\} \text{new_lock } () \{lk.\text{is_lock } lk R\}$</p>	<p>HT-ACQUIRE $\{\text{is_lock } lk R\} \text{acquire } lk \{R\}$</p>
<p>HT-RELEASE $\{\text{is_lock } lk R * R\} \text{release } lk \{\text{True}\}$</p>	<p>LOCK-PERS $\text{persistent}(\text{is_lock } lk R)$</p>

Specifications of parallel composition:

HT-PAR	$\frac{\{P_1\} e_1 \{\Phi_1\} \quad \{P_2\} e_2 \{\Phi_2\}}{\{P_1 * P_2\} e_1 \parallel e_2 \{w. \exists w_1, w_2. w = (w_1, w_2) * \Phi_1 w_1 * \Phi_2 w_2\}}$
--------	---

Figure 2.8: Specifications of the parallel composition and lock implementations.

The rule HT-PAR for parallel composition ($e_1 \parallel e_2$) states that we can prove the threads e_1 and e_2 separately, by splitting the precondition into two disjoint sets of resources P_1 and P_2 . As a result we obtain the postconditions Φ_1 and Φ_2 for the respective threads and their resulting values w_1 and w_2 .

The specifications of the lock operate on the *abstract predicate* $\text{is_lock } lk R$, which conceptually captures that the value lk is a lock which governs the proposition R . Additionally, the predicate $\text{is_lock } lk R$ is persistent, allowing us to duplicate it and share it between multiple threads. Finally, the predicate $\text{is_lock } lk R$ is built on top of Iris’s impredicative invariants, and can therefore be thought of as a first-class citizen of the logic, which lets us put locks inside of locks.

The rule HT-NEWLOCK states that we can allocate a new lock by giving up R , in turn giving us the abstract predicate $\text{is_lock } lk R$ for the return value lk . We can then access the locked proposition R with HT-ACQUIRE, which waits for the lock to be free, to then acquire it, after which we obtain R . Finally, we can release the lock with HT-RELEASE, under the condition that we give back the locked proposition R .

The proof of the `locked_interleaving_example` safety specification is carried out using HT-BIND to string the sub-proofs of the individual sub-expressions together. The individual steps are carried out as follows:

- Allocate the lock lk , protecting the exclusive reference, to get the abstract predicate $P_{lk} \triangleq \text{is_lock } lk (\exists x. \ell \mapsto x)$, using HT-NEWLOCK
- Duplicate P_{lk} using LOCK-PERS, \square -DUP, and \square -ELIM
- Distribute P_{lk} to the subproofs of the parallel threads with HT-PAR
- For each of the parallel threads:
 - Acquire the lock to obtain the exclusive reference $\ell \mapsto x$ with HT-ACQUIRE
 - Deallocate and update the reference with HT-LOAD and HT-STORE
 - Release the lock by putting the updated exclusive reference $\ell \mapsto (x + 2)$ back into the lock (with x instantiated as $x + 2$) with HT-RELEASE
- Acquire the lock, and throw away P_{lk} with HT-ACQUIRE and AFFINE

In P_{lk} we existentially quantify the stored value, as it must remain unchanged outside of the critical sections. This limitation restricts us from proving a stronger functional correctness specification, namely that the program adds exactly 4 to the stored value, using only the lock and parallel composition specifications. However, we can achieve such a proof with Iris’s notion of *ghost state*, which lets us keep track of the changes made to resources owned by invariants.

2.3.5 Ghost State in Iris

Ghost state is similar in spirit to heap state, as it carries ownership, but instead of describing the state of the heap it describes the state of proof-related variables, which are not a part of the actual program [Cohen et al. 2009; Krishnaswami et al. 2012].

Ghost state manifests as ghost resources, of the form $\llbracket x \rrbracket^\gamma$, where γ is the resource identifier, and x is its value. The values governed by ghost resources can take many forms, each with an individual set of rules for constructing and updating the resources. We call these sets of rules ghost theories. Iris comes equipped with a library of so-called resource algebras, for constructing various kinds of ghost theories. For the purpose of this thesis we do not cover how we construct our ghost theories, and instead simply present their rules. We refer the interested reader to Jung et al. [2018b, Section 3.1].

To prove functional correctness of the `locked_interleaving_example` program we consider a ghost theory consisting of two resources $\llbracket x \rrbracket^\gamma$ and $\llbracket \circ_\pi x \rrbracket^\gamma$. The resource $\llbracket x \rrbracket^\gamma$, which we call the authoritative piece, conceptually captures the global view of $x \in \mathbb{Z}$. The resource $\llbracket \circ_\pi y \rrbracket^\gamma$, which we call the permission piece, conceptually captures the permission to update the value recorded by the authoritative piece. Here, $y \in \mathbb{Z}$ is a local view of the updates that has been made to the authoritative piece in the presence of this permission piece. The fraction $\pi \in (0, 1]$ denotes how big a part of the full view that the local view represents. These properties are made precise by the following rules:

$$\begin{array}{c}
 \text{VS-CSQ} \\
 \frac{P \Rightarrow_{\mathcal{E}} P' \quad \{P'\} e \{w. Q'\} \quad (\forall w. Q' \Rightarrow_{\mathcal{E}} Q)}{\{P\} e \{w. Q\}}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{VS-FRACAUTH-ALLOC} \\
 \text{True} \Rightarrow_{\mathcal{E}} \exists \gamma. \llbracket x \rrbracket^\gamma * \llbracket \circ_1 x \rrbracket^\gamma
 \end{array}$$

$$\begin{array}{c}
 \text{VS-FRACAUTH-ADD} \\
 \llbracket x \rrbracket^\gamma * \llbracket \circ_\pi y \rrbracket^\gamma \Rightarrow_{\mathcal{E}} \llbracket (x+z) \rrbracket^\gamma * \llbracket \circ_\pi (y+z) \rrbracket^\gamma
 \end{array}$$

$$\begin{array}{c}
 \text{FRACAUTH-OP} \\
 \llbracket \circ_{\pi_1 + \pi_2} (x+y) \rrbracket^\gamma ** \llbracket \circ_{\pi_1} x \rrbracket^\gamma * \llbracket \circ_{\pi_2} y \rrbracket^\gamma
 \end{array}
 \qquad
 \begin{array}{c}
 \text{FRACAUTH-EQ} \\
 \llbracket x \rrbracket^\gamma * \llbracket \circ_1 y \rrbracket^\gamma \text{ -* } x = y
 \end{array}$$

Here the rule VS-CSQ states that we can always update the ghost state before or after the execution of any program. That is by employing the viewshift $\Rightarrow_{\mathcal{E}}$, which for the purpose of this thesis can be thought of as an implication for the ghost state. We can thus always allocate an authoritative piece $\llbracket x \rrbracket^\gamma$ and a full-fraction permission piece $\llbracket \circ_1 x \rrbracket^\gamma$, with some freely picked value x , and a fresh identifier γ , by using VS-FRACAUTH-ALLOC. The rule VS-FRACAUTH-ADD captures that we can update the

authoritative piece $\llbracket \bullet x \rrbracket^\gamma$ in the presence of a permission piece $\llbracket \circ_\pi y \rrbracket^\gamma$ with any partial fraction π . The rule `FRACAUTH-OP` lets us split and rejoin the permission pieces into and from smaller fractions, by similarly splitting the local view that they have recorded. This lets us distribute permission pieces to multiple threads, and later recover the complete record of the changes that they have recorded. Finally, the `FRACAUTH-EQ` rule lets us assert that the value stored in the authoritative piece corresponds to the value recorded by a permission piece with the full fraction.

With this ghost theory in hand we can prove the following functional specification of the `locked_interleaving_example` program:

$$\{\ell \mapsto x\} \text{locked_interleaving_example } \ell \{\ell \mapsto (x + 4)\}$$

We again string together the sub-proofs of the sub-expressions using `HT-BIND`, where the individual sub-proofs are proven as follows:

- Allocate the ghost variables $\llbracket \bullet 0 \rrbracket^\gamma$ and $\llbracket \circ_1 0 \rrbracket^\gamma$ which govern the amount added to the stored value x (initially being 0), using `Vs-CSQ` and `Vs-FRACAUTH-ALLOC`
- Allocate the lock lk , protecting the exclusive reference $\ell \mapsto x$ along with the authoritative piece $\llbracket \bullet 0 \rrbracket^\gamma$, to get the abstract predicate for the lock $P_{lk} \triangleq \text{is_lock } lk (\exists y. \ell \mapsto (x + y) * \llbracket \bullet y \rrbracket^\gamma)$, using `HT-NEWLOCK`
- Duplicate P_{lk} using `LOCK-PERS`, `□-DUP`, and `□-ELIM`
- Split the permission piece $\llbracket \circ_1 0 \rrbracket^\gamma$ into $\llbracket \circ_{\frac{1}{2}} 0 \rrbracket^\gamma * \llbracket \circ_{\frac{1}{2}} 0 \rrbracket^\gamma$, using `FRACAUTH-OP`
- Distribute P_{lk} and one of the permission pieces $\llbracket \circ_{\frac{1}{2}} 0 \rrbracket^\gamma$ to each subproof of the parallel threads with `HT-PAR`
- For each of the parallel threads:
 - Acquire the exclusive reference $\ell \mapsto (x + y)$ and authoritative piece $\llbracket \bullet y \rrbracket^\gamma$ with `HT-ACQUIRE`
 - Deallocate and update the reference with `HT-LOAD` and `HT-STORE`
 - Update the ghost variables to record the value change, from $\llbracket \bullet y \rrbracket^\gamma * \llbracket \circ_{\frac{1}{2}} 0 \rrbracket^\gamma$ to $\llbracket \bullet (y + 2) \rrbracket^\gamma * \llbracket \circ_{\frac{1}{2}} 2 \rrbracket^\gamma$, with `Vs-FRACAUTH-ADD`
 - Release the lock by putting the updated exclusive reference $\ell \mapsto (x + y + 2)$ and the authoritative piece $\llbracket \bullet (y + 2) \rrbracket^\gamma$ (with y instantiated as $y + 2$) back into the lock with `HT-RELEASE`
- Rejoin the permission pieces returned by `HT-PAR` $\llbracket \circ_{\frac{1}{2}} 2 \rrbracket^\gamma * \llbracket \circ_{\frac{1}{2}} 2 \rrbracket^\gamma$ to $\llbracket \circ_1 4 \rrbracket^\gamma$
- Acquire the exclusive reference $\ell \mapsto (x + y)$ and authoritative piece $\llbracket \bullet y \rrbracket^\gamma$ with `HT-ACQUIRE`
- Infer that the value stored in ℓ is $x + 4$ from $\llbracket \bullet y \rrbracket^\gamma * \llbracket \circ_1 4 \rrbracket^\gamma$, using `FRACAUTH-EQ`

We treat the final `acquire` lk as a means of deallocating the lock, to re-obtain the protected resources. We could circumvent the need for this step with another set of specifications for the lock, that keep track of how many times P_{lk} has been duplicated [Hobor et al. 2008; Bizjak et al. 2019]. This would allow us to re-obtain the protected resources by forgetting the lock, in the presence of all of the duplicated

fragments. However, keeping track of the lock fragments adds additional complexity to our proofs, and would not allow us to model the mutex types in our semantic type system Section 4.3.6.

2.3.6 Concluding Remarks

We have shown how Iris can be used to prove high-level specifications of implemented abstractions, such as locks (Section 2.3.4). These specifications are based on abstract predicates, *e.g.*, $\text{is_lock } lk \ R$, constructed such that we can think of them as first-class primitives, akin to $\ell \mapsto v$. As a result, the abstractions seamlessly integrate with each other, and can be nested.

We take a similar approach in Chapter 3, to obtain high-level specifications for the binary channels for message passing presented in Section 2.1.5. The specifications of the channels are based on an abstract predicate that captures exclusive ownership of a channel endpoint, along with the channels compliance with one of our novel dependent separation protocols. The abstract predicate is constructed such that it can be thought of as a first-class citizen of the logic. As a result, we can nest protocols, allowing us to describe the exchange of channels over channels, which we demonstrate in Section 3.2. Additionally, it makes the channels seamlessly integrate with locks, as we can put channels in locks, and vice versa, which we demonstrate in Section 3.4.

We have additionally given an account of various Iris features, such as basic separation logic (Section 2.3.1), persistency (Section 2.3.2), and step-indexing (Section 2.3.3). These Iris features give rise to many of the features of the semantic type system, which we present in Chapter 4. For example, the affine properties of the type system are a direct result of the affine separation logic of Iris, while the copyable types are obtained based on the persistency of the type predicates.

Finally, we have shown how we can use Iris’s ghost state mechanism to prove properties about programs that share exclusive resources via locks (Section 2.3.5). We employ a similar approach to prove the specifications of various programs that combine message passing with locks. In particular, we prove the functional correctness of a load-balancing mapper in Section 3.4. We prove the functional correctness of a map-reduce algorithm in Section 3.5. Finally, we do a manual typing proof of a producer-consumer-based computation client that operates on a single channel endpoint in parallel in Section 4.4.

Actris: Session-Type Based Reasoning in Separation Logic

Abstract Message passing is a useful abstraction for implementing concurrent programs. For real-world systems, however, it is often combined with other programming and concurrency paradigms, such as higher-order functions, mutable state, shared-memory concurrency, and locks. We present **Actris**: a logic for proving functional correctness of programs that use a combination of the aforementioned features. Actris combines the power of modern concurrent separation logics with a first-class protocol mechanism—based on session types—for reasoning about message passing in the presence of other concurrency paradigms. We show that Actris provides a suitable level of abstraction by proving functional correctness of a variety of examples, including a distributed merge sort, a distributed load-balancing mapper, and a variant of the map-reduce model, using concise specifications.

While Actris was already presented in a conference paper (POPL’20), this paper expands the prior presentation significantly. Moreover, it extends Actris to **Actris 2.0** with a notion of *subprotocols*—based on session-type subtyping—that permits additional flexibility when composing channel endpoints, and that takes full advantage of the asynchronous semantics of message passing in Actris. Soundness of Actris 2.0 is proved using a model of its protocol mechanism in the Iris framework. We have mechanised the theory of Actris, together with custom tactics, as well as all examples in the paper, in the Coq proof assistant.

3.1 Introduction

Message-passing programs are ubiquitous in modern computer systems, emphasising the importance of their functional correctness. Programming languages, like Erlang, Elixir, and Go, have built-in primitives that handle spawning of processes and intra-process communication, while other mainstream languages, such as Java, Scala, F#, and C#, have introduced an Actor model [Hewitt et al. 1973] to achieve similar functionality. In both cases the goal remains the same—help design reliable systems, often

with close to constant up-time, using lightweight processes that can be spawned by the hundreds of thousands and that communicate via asynchronous message passing.

While message passing is a useful abstraction, it is not a silver bullet of concurrent programming. In a study of larger Scala projects Tasharofi et al. [2013] write:

We studied 15 large, mature, and actively maintained actor programs written in Scala and found that 80% of them mix the actor model with another concurrency model.

In this study, 12 out of 15 projects did not entirely stick to the Actor model, hinting that even for projects that embrace message passing, low-level concurrency primitives like locks (*i.e.*, mutexes) still have their place. Tu et al. [2019] came to a similar conclusion when studying 6 large and popular Go programs. A suitable solution for reasoning about message-passing programs should thus integrate with other programming and concurrency paradigms.

In this paper we introduce **Actris**—a concurrent separation logic for proving functional correctness of programs that combine message passing with other programming and concurrency paradigms. Actris can be used to reason about programs written in a language that mimics the important features found in aforementioned languages such as higher-order functions, higher-order references, fork-based concurrency, locks, and primitives for asynchronous message passing over channels. The channels of our language are first-class and can be sent as arguments to functions, be sent over other channels (often referred to as delegation), and be stored in references.

Program specifications in Actris are written in an impredicative higher-order concurrent separation logic built on top of the Iris framework [Jung et al. 2015; Krebbers et al. 2017a; Jung et al. 2016, 2018b]. In addition to the usual features of Iris, Actris provides a notion of *dependent separation protocols* to reason about message passing over channels, inspired by the affine variant [Mostrous and Vasconcelos 2014] of binary session types [Honda et al. 1998]. We show that dependent separation protocols integrate seamlessly with other concurrency paradigms, allow delegation of resources, support channel sharing over multiple concurrent threads using locks, and more.

3.1.1 Message Passing in Concurrent Separation Logic

Over the last decade, there has been much work on extensions of concurrent separation logic with reasoning principles for message passing [Oortwijn et al. 2016; Francalanza et al. 2011; Lozes and Villard 2012; Craciun et al. 2015]. These logics typically include some form of mechanism for writing protocol specifications in a high-level manner. Unfortunately, these logics have shortcomings in terms of expressivity. Most importantly, they cannot be used to reason about programs that combine message passing with other programming and concurrency paradigms, such as higher-order functions, fine-grained shared-memory concurrency, and locks.

In a different line of work, researchers have developed expressive extensions of concurrent separation logic that do support proving strong specifications of programs involving some or all combinations of the aforementioned programming and concurrency paradigms. Examples of such logics are TaDA [da Rocha Pinto et al. 2014], iCAP [Svendsen and Birkedal 2014], Iris [Jung et al. 2015], FCSL [Nanevski et al.

2014], and VST [Appel 2014]. However, only a few variants and extensions of these logics address message-passing concurrency.

First, there has been work on the use of separation logic to reason about programs that communicate via message passing over a network. The reasoning principles in such logics are geared towards different programming patterns than the ones used in high-level languages like Erlang, Elixir, Go, and Scala. Namely, on networks all data must be serialised, and packets can be lost or delivered out of order. In high-level languages messages cannot get lost, are ensured to be delivered in order, and are allowed to contain many types of data, including functions, references, and even channel endpoints. Two examples of network logics are *Disel* by Sergey et al. [2018] and *Aneris* by Krogh-Jespersen et al. [2020]. Second, there has been work on the use of separation logic to prove compiler correctness of high-level message-passing languages. Tassarotti et al. [2017] verified a small compiler of a session-typed language into a language where channel buffers are modelled on the heap.

The primary reasoning principle to model the interaction between processes in the aforementioned logics is the notion of a State Transition System (STS). As a simple example, consider the following program, borrowed from Tassarotti et al. [2017]:

$$prog_1 := \text{let } (c, c') := \text{new_chan } () \text{ in fork } \{\text{send } c' \ 42\}; \text{recv } c$$

This program creates two channel endpoints c and c' , forks off a new thread, and sends the number 42 over the channel c' , which is then received by the initiating thread. Modelling the behaviour of this program in an STS requires three states:



The three states model that no message has been sent (**Init**), that a message has been sent but not received (**Sent**), and finally that the message has been sent and received (**Received**). Exactly what this STS represents is made precise by the underlying logic, which determines what constitutes a state and a transition, and how these are related to the channel buffers.

While STSs appear like a flexible and intuitive abstraction to reason about message-passing concurrency, they have their problems:

- Coming up with a good STS that makes the appropriate abstractions is difficult because the STS has to keep track of all possible states that the channel buffers can be in, including all possible interleavings of messages in transit.
- While STSs used for the verification of different modules can be composed at the level of the logic, there is no canonical way of composing them due to their unrestrained structure.
- Finally, STSs are first-order meaning that their states and transitions cannot be indexed by propositions of the underlying logic, which limits what they can express when sending messages containing functions or other channels.

3.1.2 Actris 1.0: Dependent Separation Protocols

Instead of using STSs, Actris extends separation logic with a new notion called *dependent separation protocols*. This notion is inspired by the session type community,

pioneered by Honda et al. [1998], where channel endpoints are given types that describe the expected exchanges. Using binary session types, the channels c and c' in the program $prog_1$ in Section 3.1.1 would have the types $c : ?\mathbb{Z}.\mathbf{end}$ and $c' : !\mathbb{Z}.\mathbf{end}$, where $!T$ and $?T$ denotes that a value of type T is sent or received, respectively. Moreover, the types of the channels c and c' are *duals*—when one does a send the other does a receive, and *vice versa*.

While session types provide a compact way of specifying the behaviour of channels, they can only be used to talk about the *type* of data that is being passed around—not their *payloads*. In this paper, we build on prior work by Bocchi et al. [2010] and Craciun et al. [2015] to attach logical predicates to session types to say more about the payloads, thus vastly extending the expressivity. Concretely, we port session types into separation logic in the form of a construct $c \mapsto prot$, which denotes ownership of a channel c with dependent separation protocol $prot$. Dependent separation protocols $prot$ 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, where finite streams are ultimately terminated by an **end** constructor. Here, v is the value that is being sent or received, P is a separation logic proposition denoting the ownership of the resources being transferred as part of the message, and the variables $\vec{x} : \vec{\tau}$ bind into v , P , and $prot$. The dependent separation protocols for the above example are:

$$c \mapsto ?\langle 42 \rangle\{\mathbf{True}\}.\mathbf{end} \quad \text{and} \quad c' \mapsto !\langle 42 \rangle\{\mathbf{True}\}.\mathbf{end}$$

These protocols state that the endpoint c expects the number 42 to be sent along it, and that the endpoint c' expects to send the number 42. Using this protocol, we can show that $prog_1$ has the specification $\{\mathbf{True}\} prog_1 \{v. v = 42\}$, where v is the result of the evaluation.

Dependent separation protocols $!\vec{x} : \vec{\tau}\langle v \rangle\{P\}.prot$ and $?\vec{x} : \vec{\tau}\langle v \rangle\{P\}.prot$ are *dependent*, meaning that the tail $prot$ can be defined in terms of the previously bound variables $\vec{x} : \vec{\tau}$. A sample program showing the use of such dependency is:

$$\begin{aligned} prog_2 := & \text{let } (c, c') := \text{new_chan } () \text{ in} \\ & \text{fork } \{\text{let } x := \text{recv } c' \text{ in send } c' (x + 2)\}; \\ & \text{send } c \ 40; \text{recv } c \end{aligned}$$

In this program, the main thread sends the number 40 to the forked-off thread, which then adds 2 to it, and sends it back. This program has the same specification as $prog_1$, while we change the dependent separation protocol as follows (we omit the dependent separation protocol for the dual endpoint c'):

$$c \mapsto !(x : \mathbb{Z}) \langle x \rangle\{\mathbf{True}\}.\ ?\langle x + 2 \rangle\{\mathbf{True}\}.\mathbf{end}$$

This protocol states that the second exchanged value is exactly the first with 2 added to it. To do so, it makes use of a dependency on the variable x , which is used to describe the contents of the first message, which the second message then depends on. This variable is bound in the protocol and it is instantiated only when a message is sent. This is different from the logic by Craciun et al. [2015], which does not support dependent protocols. Their logic is limited to protocols analogous to $!\langle x \rangle\{\mathbf{True}\}.\ ?\langle x +$

$2\}\{\text{True}\}.\mathbf{end}$ where x is free, which means the value of x must be known when the protocol is created.

While the prior examples could have been type-checked and verified using the formalisms of Bocchi et al. [2010] and Craciun et al. [2015], the following stateful example cannot:

$$\begin{aligned} \mathit{prog}_3 := & \text{let } (c, c') := \text{new_chan } () \text{ in} \\ & \text{fork } \{\text{let } \ell := \text{recv } c' \text{ in } \ell \leftarrow !\ell + 2; \text{ send } c' ()\}; \\ & \text{let } \ell := \text{ref } 40 \text{ in } \text{send } c \ell; \text{recv } c; !\ell \end{aligned}$$

Here, the main thread stores the value 40 on the heap, and sends a reference ℓ over the channel c to the forked-off thread. The main thread then awaits a signal $()$, notifying that the reference has been updated to 42 by the forked-off thread. This program has the same specification as prog_1 and prog_2 , but the dependent separation protocol is:

$$c \mapsto !(\ell : \text{Loc}) (x : \mathbb{Z}) \langle \ell \rangle \{ \ell \mapsto x \}. ?\langle () \rangle \{ \ell \mapsto (x + 2) \}.\mathbf{end}$$

This protocol denotes that the endpoints first exchange a reference ℓ , as well as a *points-to* connective $\ell \mapsto x$ that describes the ownership and value of the reference ℓ . To perform the exchange c has to give up ownership of the location, while c' acquires it—which is why it can then safely update the received location to 42 before sending the ownership back along with the notification $()$.

The type system by Bocchi et al. [2010] cannot verify this program because it does not support mutable state, while Actris can verify the program because it is a separation logic. The logic by Craciun et al. [2015] cannot verify this program because it does not support dependent protocols, which are crucial here as they make it possible to delay picking the location ℓ used in the protocol until the send operation is performed.

Dependent protocols are also useful to define recursive protocols to reason about programs that use a channel in a loop. Consider the following variant of prog_1 :

$$\begin{aligned} \mathit{prog}_4 := & \text{let } (c, c') := \text{new_chan } () \text{ in} \\ & \text{fork } \{\text{let } \mathit{go} () := (\text{send } c' (\text{recv } c' + 2); \mathit{go} ()) \text{ in } \mathit{go} ()\}; \\ & \text{send } c \ 18; \text{let } x := \text{recv } c \text{ in} \\ & \text{send } c \ 20; \text{let } y := \text{recv } c \text{ in } x + y \end{aligned}$$

The forked-off thread will repeatedly interleave receiving values with sending those values back incremented by two. The program prog_4 has the same specification as before, but now we use the following recursive dependent separation protocol:

$$c \mapsto \mu\mathit{prot}.\ ! (x : \mathbb{Z}) \langle x \rangle \{\text{True}\}.\ ? \langle x + 2 \rangle \{\text{True}\}.\ \mathit{prot}$$

This protocol expresses that it is possible to make repeated exchanges with the forked-off thread to increment a number by 2. The fact that the variable x is bound in the protocol is once again crucial—it allows the use of different numbers for each exchange.

Furthermore, Actris inherently captures some features of conventional session types. One such example is *delegation* of channels as seen in the following program:

```

prog5 := let (c1, c'1) := new_chan () in
  fork {let c := recv c'1 in let y := recv c1 in send c y; send c'1 ()};
  let (c2, c'2) := new_chan () in
  fork {let x := recv c'2 in send c'2 (x + 2)};
  send c1 c2; send c1 40; recv c1; recv c2
    
```

This program uses the channel pair c_2, c'_2 to exchange the number 40 with the second forked-off thread, which adds 2 to it, and sends it back. Contrary to the programs we have seen before, it uses the additional channel pair c_1, c'_1 to delegate the endpoint c_2 to the first forked-off thread, which then sends the number over c_2 . While this program is intricate, the following dependent separation protocols describe the communication concisely:

$$\begin{aligned}
 c_1 &\mapsto ! (c : \mathbf{Chan}) \langle c \rangle \{c \mapsto ! (x : \mathbb{Z}) \langle x \rangle \{\mathbf{True}\}. ? \langle x + 2 \rangle \{\mathbf{True}\}. \mathbf{end}\}. \\
 &\quad ! (y : \mathbb{Z}) \langle y \rangle \{\mathbf{True}\}. ? \langle () \rangle \{c \mapsto ? \langle y + 2 \rangle \{\mathbf{True}\}. \mathbf{end}\}. \mathbf{end}\}. \\
 c_2 &\mapsto ! (x : \mathbb{Z}) \langle x \rangle \{\mathbf{True}\}. ? \langle x + 2 \rangle \{\mathbf{True}\}. \mathbf{end}
 \end{aligned}$$

The first protocol states that the initial value sent must be a channel endpoint with the protocol used in $prog_1$, meaning that the main thread must give up the ownership of the channel endpoint c_2 , thereby delegating it. It then expects a value y to be sent, and finally to receive a notification $()$, along with ownership of the channel c_2 , which has since taken one step by sending y .

Lastly, the dependencies in dependent separation protocols are not limited to first-order data, but can also be used in combination with functions. For example:

```

prog6 := let (c, c') := new_chan () in
  fork {let f := recv c' in send c' (λ(). f() + 2)};
  let ℓ := ref 40 in send c (λ(). !ℓ); recv c ()
    
```

This program exchanges a value to which 2 is added, but postpones the evaluation by wrapping the computation in a closure. The following protocol is used to verify this program:

$$\begin{aligned}
 c &\mapsto ! (P Q : \mathbf{iProp}) (f : \mathbf{Val}) \langle f \rangle \{\{P\} f () \{v. v \in \mathbb{Z} * Q(v)\}\}. \\
 &\quad ? (g : \mathbf{Val}) \langle g \rangle \{\{P\} g () \{v. \exists w. (v = w + 2) * Q(w)\}\}. \mathbf{end}
 \end{aligned}$$

The send constructor $(!)$ does not just bind the function value f , but also the precondition P and postcondition Q of its Hoare triple. In the second message, a Hoare triple is returned that maintains the original pre- and postconditions, but returns an integer of 2 higher. To send the function, the main thread would let $P \triangleq \ell \mapsto 40$ and $Q(v) \triangleq (v = 40)$, and prove $\{P\} (\lambda(). !\ell) () \{Q\}$. This example demonstrates that the state space of dependent separation protocols can be higher-order—it is indexed by the precondition P and postcondition Q of f —which means that they do not have to be agreed upon when creating the protocol, masking the internals of the function from the forked-off thread.

It is worth noting that using dependent recursive protocols it is possible to keep track of a history of what actions have been performed, which, as is shown in Section 3.4, is especially useful when combining channels with locks.

3.1.3 Actris 2.0: Subprotocols

While Actris 1.0’s notion of dependent separation protocols is expressive enough to capture advanced exchanges, as indicated by the examples in the previous section, they are more restrictive than necessary due to their dual nature. The dual nature of dependent separation protocols requires that:

- Sends $(!\vec{x}:\vec{\tau}\langle v\rangle\{P\})$ match up with receives $(?\vec{x}:\vec{\tau}\langle v\rangle\{P\})$, and *vice versa*,
- The logical variables $\vec{x}:\vec{\tau}$ of matched sends and receives are the same, and,
- The propositions P of matched send and receives are the same.

With an asynchronous semantics for message passing, where messages are buffered in both directions, the above notion of duality excludes the verification of certain safe programs. While it is safe for sends (!) to happen before the expected receives (?), duality does not allow that. This is demonstrated by the following safe program:

```

prog7 := let (c, c') := new_chan () in
         fork {send c' 20; send c' (recv c' + 2)};
         send c 20;
         let x := recv c in
         let y := recv c in x + y

```

Here, both threads first send the value 20, and then receive the value of the other thread. After this, they follow a dual behaviour, where the forked-off thread sends a value, which the main thread receives. With asynchronous message passing, this interaction is safe as neither thread blocks when resolving their send, and both messages can be in transit at the same time because channels are buffered. However, with the features of Actris 1.0 presented in the conference version of this paper [Hinrichsen et al. 2020], this program cannot be verified as the two dependent separation protocols of channel endpoints must be strictly dual.

Support for type checking such programs has been studied in the session type community, namely in the context of *asynchronous session subtyping* [Mostrous et al. 2009; Mostrous and Yoshida 2015], in which a subtyping relation $S_1 <: S_2$ is defined, capturing that the session type S_2 can be used in place of S_1 when type checking a program. The relation captures that sends can be swapped ahead of receives $?T.!U.S <: !U.?T.S$. We refer to this as messages being sent *ahead of* the receives.

In this paper, we show that dependent separation protocols are compatible with the idea of asynchronous session subtyping. This gives rise to **Actris 2.0** that supports so-called *subprotocols*. Subprotocols are formalised by a preorder $prot_1 \sqsubseteq prot_2$, which captures (among others) a notion of swapping sends ahead of receives (provided that the send does not depend on the logical variables of the receive). We can then prove that $prog_7$ results in 42 by picking the following dependent separation protocols:

$$\begin{aligned}
c &\mapsto !(x : \mathbb{Z}) \langle x \rangle \{\text{True}\}. ?\langle 20 \rangle \{\text{True}\}. ?\langle x + 2 \rangle \{\text{True}\}. \mathbf{end} \quad \mathbf{and} \\
c' &\mapsto ?(x : \mathbb{Z}) \langle x \rangle \{\text{True}\}. !\langle 20 \rangle \{\text{True}\}. !\langle x + 2 \rangle \{\text{True}\}. \mathbf{end}
\end{aligned}$$

While the main thread satisfies the protocol of c immediately, the forked-off thread does not satisfy the protocol of c' , as it sends the first value before receiving. However, it is possible to weaken the protocol of c' using Actris 2.0's notion of subprotocols:

$$\begin{aligned} & ?(x : \mathbb{Z}) \langle x \rangle \{\text{True}\}. !\langle 20 \rangle \{\text{True}\}. !\langle x + 2 \rangle \{\text{True}\}. \mathbf{end} \\ \sqsubseteq & !\langle 20 \rangle \{\text{True}\}. ?(x : \mathbb{Z}) \langle x \rangle \{\text{True}\}. !\langle x + 2 \rangle \{\text{True}\}. \mathbf{end} \end{aligned}$$

This gives $c' \rightsquigarrow !\langle 20 \rangle \{\text{True}\}. ?(x : \mathbb{Z}) \langle x \rangle \{\text{True}\}. !\langle x + 2 \rangle \{\text{True}\}. \mathbf{end}$. Since the first send (with value 20) is independent of the variable x bound by the receive, the subprotocol relation follows immediately from the swapping property. Note that it is *not* possible to swap the second send (with value $x + 2$) ahead of the receive, as it does in fact depend on variable x bound by the receive.

In addition to allowing the verification of a larger class of programs, Actris 2.0's subprotocols also provide a more extensional approach to reasoning about dependent separation protocols. This is beneficial whenever we want to reuse existing specifications that might use a syntactically different protocol, but that nonetheless logically entail each another. For example, the ordering of logical variables can be changed using the subprotocol relation:

$$!(x : \mathbb{Z})(y : \mathbb{Z}) \langle (x, y) \rangle \{\text{True}\}. \mathit{prot} \sqsubseteq !(y : \mathbb{Z})(x : \mathbb{Z}) \langle (x, y) \rangle \{\text{True}\}. \mathit{prot}$$

Since the subprotocol relation is a first-class logical proposition of Actris 2.0, it also allows the manipulation of separation logic resources, such as moving in ownership. For example, we can show the following *conditional* subprotocol relation:

$$\begin{aligned} \ell'_1 \mapsto 20 \quad - * \\ !(\ell_1, \ell_2 : \text{Loc}) \langle (\ell_1, \ell_2) \rangle \{\ell_1 \mapsto 20 * \ell_2 \mapsto 22\}. \mathit{prot} \sqsubseteq !(\ell_2 : \text{Loc}) \langle (\ell'_1, \ell_2) \rangle \{\ell_2 \mapsto 22\}. \mathit{prot} \end{aligned}$$

Here, we move the ownership of $\ell'_1 \mapsto 20$ into the protocol, to resolve the eventual obligation of sending it, while instantiating the logical variable ℓ_1 with ℓ'_1 .

In addition to the demonstrated features, in the rest of this paper we show that Actris 2.0's subprotocol relation is capable of moving resources from one message to another. This gives rise to a principle similar to *framing*, known from conventional separation logic, but applied to dependent separation protocols. Moreover, inspired by the work of Brandt and Henglein [1998], the subprotocol relation is defined coinductively, allowing us to use the principle of Löb induction to prove subprotocol relations for recursive protocols.

3.1.4 Formal Correspondence to Session Types

Even though Actris's notion of dependent separation protocols is influenced by binary session types, this paper does not provide a formal correspondence between the two systems. However, since Actris is built on top of Iris, it forms a suitable foundation for building logical relation models of type systems. In related work by Hinrichsen et al. [2021c], Actris has been used to define a logical relations model of binary session types, with support for various forms of polymorphism and recursion, asynchronous subtyping, references, and locks/mutexes. Similar to the RustBelt project [Jung et al.

2018a], that work gives rise to an extensible approach for proving type safety, which can be used to manually prove the typing judgements of racy, but safe, programs that cannot be type checked using only the rules of the type system.

3.1.5 Contributions and Outline

This paper introduces **Actris 2.0**: a higher-order impredicative concurrent separation logic built on top of the Iris framework for reasoning about functional correctness of programs with asynchronous message passing that combine higher-order functions, higher-order references, fork-based concurrency, and locks. Concretely, this paper makes the following contributions:

- We introduce *dependent separation protocols* inspired by affine binary session types to model the transfer of resources (including higher-order functions) between channel endpoints. We show that they can be used to handle choice, recursion, and delegation (Section 3.2).
- We introduce *subprotocols* inspired by asynchronous session subtyping. This notion relaxes duality, allowing channels to send messages before receiving others, and gives rise to a more extensional approach to reasoning about dependent separation protocols, providing more flexibility in the design and reuse of protocols. We moreover show how Löb induction is used to reason about recursive subprotocols (Section 3.3).
- We demonstrate the benefits obtained from building Actris on top of Iris by showing how Iris’s support for ghost state and locks can be used to prove functional correctness of programs using manifest sharing, *i.e.*, channel endpoints shared by multiple parties (Section 3.4).
- We provide a case study on Actris and its mechanisation in Coq by proving functional correctness of a variant of the map-reduce model by Dean and Ghemawat [2004] (Section 3.5).
- We give a model of dependent separation protocols in the Iris framework to prove safety (*i.e.*, session fidelity) and postcondition validity of our Hoare triples (Section 3.6).
- We provide a full mechanisation of Actris [Hinrichsen et al. 2021b] using the interactive theorem prover Coq. On top of our Coq mechanisation, we provide custom tactics, which we use to mechanise all examples in the paper (Section 3.7).

3.1.6 Differences from the Conference Version

This paper is an extension of the paper “Actris: Session-type based reasoning in separation logic” presented at the POPL’20 conference [Hinrichsen et al. 2020]. In this paper we present Actris 2.0, which extends Actris 1.0 with the notion of subprotocols. This extension introduces new logical connectives and proof rules, but also involves a significant overhaul of the original model and its Coq mechanisation. We additionally extend the model and mechanisation sections substantially, with additional details,

considerations, and examples, to give a better understanding of how Actris works and how it can be used. Concretely, this paper includes the following extensions compared to the conference version:

- An overview of subprotocols in the introduction (Section 3.1.3).
- A new section on Actris 2.0’s notion of subprotocols (Section 3.3).
- An updated and expanded description of the model of Actris in Iris (Section 3.6).
- An expanded section on the Coq mechanisation, with examples of mechanised proofs using the custom tactics for Actris that we have developed (Section 3.7).

3.2 A Tour of Actris

This section demonstrates the core features of Actris. We first introduce the language (Section 3.2.1) and the logic (Section 3.2.2). We then introduce and iteratively extend a simple distributed merge sort algorithm to demonstrate the main features of Actris (Section 3.2.3–Section 3.2.8). Note that as the point of the sorting algorithms is to showcase the features of Actris, they are intentionally kept simple and no effort has been made to make them efficient (*e.g.*, to avoid spawning threads for small jobs).

3.2.1 The Actris Language

The language used throughout the paper is an untyped functional language with higher-order functions, higher-order mutable references, fork-based concurrency, and primitives for message passing over bidirectional asynchronous channels. The syntax is as follows:

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

We omit the usual operations on pairs, sums, lists, and integers, which are standard. We introduce the following syntactic sugar: lambda abstractions $\lambda x. e$ are defined as $\text{rec } _ \ x := e$, let-bindings $\text{let } x := e_1 \text{ in } e_2$ are defined as $(\lambda x. e_2) \ e_1$, and sequencing $e_1; e_2$ is defined as $\text{let } _ := e_1 \text{ in } e_2$. Here, the underscore $_$ is an anonymous binder, *i.e.*, an arbitrary variable that is fresh in the body of the binding expression.

The language includes the usual operations for heap manipulation. New references can be created using $\text{ref } e$, dereferenced using $!e$, and assigned to using $e_1 \leftarrow e_2$. Concurrency is supported via $\text{fork } \{e\}$, which spawns a new thread e that is executed in the background. The language also supports atomic operations like compare-and-set (CAS), which can be used to implement lock-free data structures and synchronisation primitives, but these are omitted from the syntax.

Message passing is performed over bidirectional channels, which are represented using pairs of buffers (\vec{v}_1, \vec{v}_2) of unbounded size. The $\text{new_chan } ()$ operation creates a new channel whose buffers are empty, and returns a tuple of endpoints (c_1, c_2) . Bidirectionality is obtained by having one endpoint receive from the others send

buffer and *vice versa*. That means, `send ci v` enqueues the value v in its own buffer, *i.e.*, \bar{v}_i , and `recv ci` dequeues a value from the other buffer, *i.e.*, from \bar{v}_2 if $i = 1$ and from \bar{v}_1 if $i = 2$. Message passing is asynchronous, meaning that `send c v` will always reduce, while `recv c` will block as long as the receiving buffer is empty. The exact semantics of the channels will be detailed in Section 3.6.5.

Throughout the paper, we often use the following syntactic sugar to encapsulate the common behaviour of starting a new process:

$$\text{start } e := \text{let } f := e \text{ in let } (c, c') := \text{new_chan } () \text{ in fork } \{f\ c'\}; c$$

Here, e should evaluate to a function that takes a channel endpoint.

3.2.2 The Actris Logic

Actris is a higher-order impredicative concurrent separation logic with a new notion called *dependent separation protocols* to reason about message-passing concurrency. As we will show in Section 3.6, Actris is built as a library on top of the Iris framework [Jung et al. 2015; Krebbers et al. 2017a; Jung et al. 2016, 2018b] and thus inherits all features of Iris. For the purpose of this section, no prior knowledge of Iris is expected as the majority of Iris’s features are orthogonal to Actris’s. At this point, we are primarily concerned with Iris’s support for nested Hoare triples and guarded recursion, which we need to transfer functions over channels (Section 3.2.4) and to define recursive protocols (Section 3.2.6). An extensive overview of Iris can be found in [Jung et al. 2018b], and a tutorial-style introduction can be found in [Birkedal and Bizjak 2020].

The grammar of Actris and a selection of its rules are displayed in Figures 3.1 and 3.2. The Actris grammar includes the polymorphic lambda-calculus¹ with a number of primitive types and terms operating on these types. Most important is the type `iProp` of propositions and the type `iProto` of dependent separation protocols. The typing judgement is mostly standard and can be derived from the use of meta variables—we use the meta variables P and Q for propositions, the meta variable $prot$ for protocols, the meta variable v for values, and the meta variables t and u for general terms of any type. Apart from that, there is the implicit side-condition that recursive predicates defined using the recursion operator $\mu x : \tau. t$ should be *guarded*. That means, the variable x should appear under a *contractive* term construct. As is usual in logics with guarded recursion [Nakano 2000], the later \triangleright modality is contractive and is used to define recursive predicates. But moreover, as we will demonstrate in Section 3.2.6, the constructors $!\vec{x} : \vec{\tau} \langle v \rangle \{P\}. prot$ and $? \vec{x} : \vec{\tau} \langle v \rangle \{P\}. prot$ of dependent separation protocols are contractive in the tail argument $prot$ to enable the construction of recursive protocols. The rule μ -UNFOLD says that $\mu x : \tau. t$ is in fact a fixpoint of t .

To express program specifications, Actris features Hoare triples $\{P\} e \{v. Q\}$, where P is the precondition and Q the postcondition. The binder v can be used

¹Actris and Iris, which are both formalised as a shallow embedding in Coq, have in fact a predicative Type hierarchy, while propositions `iProp` are impredicative. For brevity’s sake, we omit details about predicativity of Type, as they are standard.

Grammar:

$$\begin{aligned}
 \tau, \sigma &::= x \mid 0 \mid 1 \mid \mathbb{B} \mid \mathbb{N} \mid \mathbb{Z} \mid \text{Type} \mid \forall x : \tau. \sigma \mid \\
 &\quad \text{Loc} \mid \text{Chan} \mid \text{Val} \mid \text{Expr} \mid \text{iProp} \mid \text{iProto} \mid \dots \\
 t, u, P, Q, \text{prot} &::= x \mid \lambda x : \tau. t \mid t(u) \mid t(\tau) \mid \quad (\text{Polymorphic lambda-calculus}) \\
 &\quad \text{True} \mid \text{False} \mid P \wedge Q \mid P \vee Q \mid P \Rightarrow Q \mid \quad (\text{Propositional logic}) \\
 &\quad \forall x : \tau. P \mid \exists x : \tau. P \mid t = u \mid \quad (\text{Higher-order logic with equality}) \\
 &\quad \mu x : \tau. t \mid \triangleright P \mid \quad (\text{Guarded recursion}) \\
 &\quad P * Q \mid P * Q \mid \ell \mapsto v \mid \{P\} e \{v. Q\} \mid \dots \quad (\text{Separation logic}) \\
 &\quad c \mapsto \text{prot} \mid \text{prot}_1 \sqsubseteq \text{prot}_2 \mid \overline{\text{prot}} \mid \text{prot}_1 \cdot \text{prot}_2 \mid \text{end} \\
 &\quad !\vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot} \mid ?\vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot} \mid \dots \quad (\text{Dep. sep. protocols})
 \end{aligned}$$
Ordinary affine separation logic:

$$\begin{array}{c}
 \text{HT-FRAME} \\
 \frac{\{P\} e \{w. Q\}}{\{P * R\} e \{w. Q * R\}} \\
 \text{AFFINE} \quad P * Q \Rightarrow P
 \end{array}
 \quad
 \begin{array}{c}
 \text{HT-VAL} \\
 \frac{}{\{\text{True}\} v \{w. w = v\}}
 \end{array}
 \quad
 \begin{array}{c}
 \text{HT-FORK} \\
 \frac{\{P\} e \{\text{True}\}}{\{P\} \text{fork} \{e\} \{\text{True}\}}
 \end{array}$$

$$\begin{array}{c}
 \text{HT-BIND} \\
 \frac{\{P\} e \{v. Q\} \quad \forall v. \{Q\} K[v] \{w. R\}}{\{P\} K[e] \{w. R\}} \quad K \text{ a call-by-value evaluation context}
 \end{array}$$
Recursion:

$$\begin{array}{c}
 \triangleright\text{-INTRO} \\
 P \Rightarrow \triangleright P
 \end{array}
 \quad
 \begin{array}{c}
 \text{LÖB} \\
 (\triangleright P \Rightarrow P) \Rightarrow P
 \end{array}
 \quad
 \begin{array}{c}
 \mu\text{-UNFOLD} \\
 (\mu x. t) = t[\mu x. t/x]
 \end{array}$$

$$\begin{array}{c}
 \text{HT-REC} \\
 \frac{\{P\} e[v/x][(\text{rec } f \ x := e)/f] \{w. Q\}}{\{\triangleright P\} (\text{rec } f \ x := e) v \{w. Q\}}
 \end{array}$$
Heap manipulation:

$$\begin{array}{c}
 \text{HT-ALLOC} \\
 \{\text{True}\} \text{ref } v \{\ell. \ell \mapsto v\}
 \end{array}
 \quad
 \begin{array}{c}
 \text{HT-LOAD} \\
 \{\ell \mapsto v\} !\ell \{w. (w = v) * \ell \mapsto v\}
 \end{array}
 \quad
 \begin{array}{c}
 \text{HT-STORE} \\
 \{\ell \mapsto v\} \ell \leftarrow w \{\ell \mapsto w\}
 \end{array}$$

Figure 3.1: The grammar and a selection of rules of Actris.

Message passing:

$$\{\text{True}\} \text{new_chan } () \{w. \exists c, c'. (w = (c, c')) * \overline{c \mapsto \text{prot} * c' \mapsto \overline{\text{prot}}}\} \quad (\text{HT-NEW})$$

$$\frac{\{P[\vec{t}/\vec{x}]\} * c \mapsto !\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \text{prot}\} \text{ send } c (v[\vec{t}/\vec{x}]) \{c \mapsto \text{prot}[\vec{t}/\vec{x}]\}}{\quad} \quad (\text{HT-SEND})$$

$$\frac{\{c \mapsto ?\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \text{prot}\} \text{ recv } c \{w. \exists \vec{y}. (w = v[\vec{y}/\vec{x}]) * \overline{c \mapsto \text{prot}[\vec{y}/\vec{x}] * P[\vec{y}/\vec{x}]}\}}{\quad} \quad (\text{HT-RECV})$$

Dependent separation protocol dual and append:

$$\begin{aligned} \overline{!\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \text{prot}} &= ?\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \overline{\text{prot}} \\ \overline{?\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \text{prot}} &= !\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \overline{\text{prot}} \\ \overline{\text{end}} &= \text{end} \\ \overline{\overline{\text{prot}}} &= \text{prot} \\ (!\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \text{prot}_1) \cdot \text{prot}_2 &= !\vec{x}:\vec{\tau}\langle v \rangle\{P\}. (\text{prot}_1 \cdot \text{prot}_2) \\ (?\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \text{prot}_1) \cdot \text{prot}_2 &= ?\vec{x}:\vec{\tau}\langle v \rangle\{P\}. (\text{prot}_1 \cdot \text{prot}_2) \\ \text{prot} \cdot \text{end} &= \text{prot} \\ \text{end} \cdot \text{prot} &= \text{prot} \\ \text{prot}_1 \cdot (\overline{\text{prot}_2 \cdot \text{prot}_3}) &= \overline{(\text{prot}_1 \cdot \text{prot}_2) \cdot \text{prot}_3} \\ \overline{\text{prot}_1 \cdot \text{prot}_2} &= \overline{\text{prot}_1} \cdot \overline{\text{prot}_2} \end{aligned}$$

Figure 3.2: The rules of the Actris Dependent separation protocols.

to talk about the return value of e in the postcondition Q , but is omitted if the result is $()$. Note that Hoare triples are propositions of the logic themselves (*i.e.*, they are of type iProp), so they can be nested to express specifications of higher-order functions. The rules for Hoare triples are mostly standard, but it is worth pointing out the rule HT-REC for recursive functions. This rule has a later modality (\triangleright) in the precondition, which when combined with the LÖB rule allows reasoning about general recursive functions. As usual, the *points-to* connective $\ell \mapsto v$ expresses unique ownership of a location ℓ with value v . Since we consider a garbage collected language, arbitrary separation logic resources can be discarded via the rule AFFINE .

The novel feature of Actris is its support for dependent separation protocols to reason about message-passing programs. This is done using the $c \mapsto \text{prot}$ connective, which expresses unique ownership of a channel endpoint c and states that the endpoint follows the protocol prot . Dependent separation protocols prot are streams of $!\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \text{prot}$ and $?\vec{x}:\vec{\tau}\langle v \rangle\{P\}. \text{prot}$ constructors that are either infinite or finite. The finite streams are ultimately terminated by an **end** constructor. The value v denotes the message that is being sent (!) or received (?), the proposition P denotes the ownership that is transferred along the message, and prot denotes

the protocol that describes the subsequent messages. The logical variables $\vec{x}:\vec{\tau}$ can be used to bind variables in v , P , and $prot$. For example, $!(b : \mathbb{B})(\ell : \text{Loc})(i : \mathbb{N}) \langle (b, \ell) \{ \ell \mapsto i * 10 < i \} \rangle . prot$ expresses that a pair of a boolean and an integer reference whose value is at least 10 is sent. We often omit the proposition $\{P\}$, which simply means it is **True**.

Apart from the constructors for dependent separation protocols, Actris provides two primitive operations. The \overline{prot} connective denotes the *dual* of a protocol. As with conventional session types, it transforms the protocol by changing all sends (!) into receives (?), and *vice versa*. Taking the dual twice thus results in the original protocol. The connective $prot_1 \cdot prot_2$ *appends* the protocols $prot_1$ and $prot_2$, which is achieved by substituting any **end** in $prot_1$ with $prot_2$. Finally, $prot_1 \sqsubseteq prot_2$ states that the protocol $prot_1$ is a *subprotocol* of $prot_2$. The subprotocol relation and its proof rules will be described in Section 3.3.

The rule HT-NEW allow ascribing any protocol to newly created channels using `new_chan` (`()`), obtaining ownership of $c \mapsto prot$ and $c' \mapsto \overline{prot}$ for the respective endpoints. The duality of the protocol guarantees that any receive (?) is matched with a send (!) by the dual endpoint, which is crucial for establishing safety (*i.e.*, session fidelity, see Section 3.6.7).

The rule HT-SEND for `send c w` requires the head of the dependent separation protocol of c to be a send (!) constructor, and the value w that is sent to match up with the ascribed value. To send a message w , we need to give up ownership of $c \mapsto !\vec{x}:\vec{\tau} \langle v \rangle \{P\} . prot$, pick an appropriate instantiation \vec{t} for the variables $\vec{x}:\vec{\tau}$ so that $w = v[\vec{t}/\vec{x}]$, give up ownership of the associated resources $P[\vec{t}/\vec{x}]$, and finally regain ownership of the protocol tail $c \mapsto prot[\vec{t}/\vec{x}]$.

The rule HT-RCV for `recv c` is essentially dual to the rule HT-SEND. We need to give up ownership of $c \mapsto ?\vec{x}:\vec{\tau} \langle v \rangle \{P\} . prot$, and in return acquire 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 $c \mapsto prot[\vec{y}/\vec{x}]$, where \vec{y} are instances of the variables of the protocol.

3.2.3 Basic Protocols

In order to show the basic features of dependent separation protocols, we will prove the functional correctness of a simple distributed merge sort algorithm, whose code is shown in Figure 3.3.

The function `sort_client` takes a comparison function `cmp` and a reference to a linked list l that will be sorted using merge sort. The bulk of the work is done by the `sort_service` function that is parameterised by a channel c over which it receives a reference to the linked list to be sorted. If the list is an empty or singleton list, which is trivially sorted, the function immediately sends back a unit value (`()`) to inform the caller that the work has been completed, and terminates. Otherwise, the list is split into two partitions using the `split` function, which updates the list in-place so that l points to the first partition, and returns a reference l' to the second partition. These partitions are recursively sorted using two newly started instances of `sort_service`. The results of the processes are then requested and merged using the `merge` function, which updates the list in-place so that l points to the merged list. Finally, the unit value (`()`) is sent back along the original channel c .

```

sort_service cmp c :=
  let l := recv c in
  if  $|l| \leq 1$  then send c () else
  let l' := lsplit l in
  let c1 := start sort_service cmp in
  let c2 := start sort_service cmp in
  send c1 l; send c2 l';
  recv c1; recv c2;
  lmerge cmp l l'; send c ()

sort_client cmp l :=
  let c :=
    start sort_service cmp in
  send c l;
  recv c

```

Figure 3.3: A distributed merge sort algorithm (the code for **lmerge** and **lsplit** is standard and thus elided).

In order to verify the correctness of the sorting algorithm we first need a specification for the comparison function *cmp*, which must satisfy the following specification:

$$\begin{aligned}
\mathbf{cmp_spec} \ (I : T \rightarrow \mathbf{Val} \rightarrow \mathbf{iProp}) \ (R : T \rightarrow T \rightarrow \mathbb{B}) \ (cmp : \mathbf{Val}) \triangleq \\
& (\forall x_1 x_2. R \ x_1 \ x_2 \vee R \ x_2 \ x_1) \wedge \\
& (\forall x_1 x_2 v_1 v_2. \{I \ x_1 \ v_1 * I \ x_2 \ v_2\} \ \mathbf{cmp} \ v_1 \ v_2 \ \{r. r = R \ x_1 \ x_2 * I \ x_1 \ v_1 * I \ x_2 \ v_2\})
\end{aligned}$$

Here, R is a decidable total relation on an implicit polymorphic type T , and I is an interpretation predicate that relates language values to elements of type T . While the relation R dictates the ordering, the interpretation predicate I allows for flexibility about what is ordered. Setting I to *e.g.*, $\lambda x v. v \mapsto x$ orders references by what they point to in memory, rather than the memory address itself. To specify how lists are laid out in memory we use the following notation:

$$l \xrightarrow{\text{list}}_I \vec{x} \triangleq \begin{cases} l \mapsto \text{inl} \ () & \text{if } \vec{x} = [] \\ \exists v_1 l_2. l \mapsto \text{inr} \ (v_1, l_2) * I \ x_1 \ v_1 * l_2 \xrightarrow{\text{list}}_I \vec{x}_2 & \text{if } \vec{x} = [x_1] \cdot \vec{x}_2 \end{cases}$$

The channel c adheres to the following dependent separation protocol:

$$\begin{aligned}
\mathbf{sort_prot} \ (I : T \rightarrow \mathbf{Val} \rightarrow \mathbf{iProp}) \ (R : T \rightarrow T \rightarrow \mathbb{B}) \triangleq \\
! (\vec{x} : \mathbf{List} \ T) \ (l : \mathbf{Loc}) \ \langle l \rangle \{ l \xrightarrow{\text{list}}_I \vec{x} \}. \ ? \vec{y} \langle () \rangle \{ l \xrightarrow{\text{list}}_I \vec{y} * \mathbf{sorted_of}_R \ \vec{y} \ \vec{x} \}. \ \mathbf{end}
\end{aligned}$$

The protocol describes the interaction of sending a list reference, and then receiving a unit value $()$ once the list is sorted and the reference is updated to point to the sorted list. The predicate $\mathbf{sorted_of}_R \ \vec{y} \ \vec{x}$ is true iff \vec{y} is a sorted version of \vec{x} with respect to the relation R . The specification of the service and the client is as follows:

$$\begin{array}{ll}
\{ \mathbf{cmp_spec} \ I \ R \ \mathbf{cmp} * c \mapsto \overline{\mathbf{sort_prot} \ I \ R} \cdot \mathbf{prot} \} & \{ \mathbf{cmp_spec} \ I \ R \ \mathbf{cmp} * l \xrightarrow{\text{list}}_I \vec{x} \} \\
\mathbf{sort_service} \ \mathbf{cmp} \ c & \mathbf{sort_client} \ \mathbf{cmp} \ l \\
\{ c \mapsto \mathbf{prot} \} & \{ \exists \vec{y}. \mathbf{sorted_of}_R \ \vec{y} \ \vec{x} * l \xrightarrow{\text{list}}_I \vec{y} \}
\end{array}$$

There are two important things to note about these specifications. First, the protocol $\mathbf{sort_prot}$ is written from the point of view of the client. As such, the precondition for

$$\begin{array}{ll}
 \text{sort_service}_{func} \ c := & \text{sort_client}_{func} \ cmp \ l := \\
 \text{let } cmp := \text{recv } c \text{ in} & \text{let } c := \text{start } \text{sort_service}_{func} \text{ in} \\
 \text{sort_service } cmp \ c & \text{send } c \ cmp; \text{ send } c \ l; \text{ recv } c
 \end{array}$$

Figure 3.4: A version of the sort service that receives the comparison function over the channel.

`sort_service` requires that c follows the dual. Second, the pre- and postcondition of `sort_service` are generalised to have an arbitrary protocol $prot$ appended at the end. It is important to write specifications this way, so they can be embedded in other protocols. We will see examples of that in Section 3.2.6 and Section 3.2.7.

The proofs of these specifications follow from symbolic execution using the rules HT-NEW, HT-SEND, HT-RECV, and the standard separation logic rules.

3.2.4 Transferring Functions

The distributed `sort_service` from the previous section (Figure 3.3) is parametric on a comparison function. To demonstrate Actris’s support for reasoning about functions transferred over channels, we verify the correctness of the program `sort_servicefunc` in Figure 3.4, which receives the comparison function over the channel instead of via a lambda abstraction. To verify this program, we extend the protocol `sort_prot` from Section 3.2.3 as follows:

$$\begin{aligned}
 \text{sort_prot}_{func} &\triangleq !(T : \text{Type}) (I : T \rightarrow \text{Val} \rightarrow \text{iProp}) (R : T \rightarrow T \rightarrow \mathbb{B}) (cmp : \text{Val}) \\
 &\quad \langle cmp \rangle \{ \text{cmp_spec } I \ R \ cmp \}. \text{sort_prot } I \ R
 \end{aligned}$$

The new protocol captures that we first send a comparison function cmp . It includes binders for the polymorphic type T , the interpretation predicate I , and the relation R . The specifications are much the same as before, with the proofs being similar besides the addition of a symbolic execution step to resolve the sending and receiving of the comparison function:

$$\begin{array}{ll}
 \{c \mapsto \overline{\text{sort_prot}_{func}} \cdot prot\} & \{ \text{cmp_spec } I \ R \ cmp * l \xrightarrow{\text{list}}_I \vec{x} \} \\
 \text{sort_service}_{func} \ c & \text{sort_client}_{func} \ cmp \ l \\
 \{c \mapsto prot\} & \{ \exists \vec{y}. l \xrightarrow{\text{list}}_I \vec{y} * \text{sorted_of}_R \vec{y} \vec{x} \}
 \end{array}$$

3.2.5 Choice

Branching communication is commonly modelled using the *choice* session types $\&$ for branching and \oplus for selection. We show that corresponding dependent separation protocols can readily be encoded in Actris. At the level of the programming language, the instructions for choice are encoded by sending and receiving a boolean value that is matched using an if-then-else construct:

$$\begin{aligned}
 &\text{select } e \ e' \triangleq \text{send } e \ e' \\
 &\text{branch } e \text{ with } \mathbf{left} \Rightarrow e_1 \mid \mathbf{right} \Rightarrow e_2 \text{ end} \triangleq \text{if } \text{recv } e \text{ then } e_1 \text{ else } e_2
 \end{aligned}$$

We let `left` := `true` and `right` := `false` to be used together with `select` for readability's sake. Due to the higher-order nature of Actris, the usual protocol specifications for choice from session types can be encoded as regular logical branching within the protocols:

$$\begin{aligned} & prot_1 \{Q_1\} \oplus \{Q_2\} prot_2 \triangleq \\ & \quad ! (b : \mathbb{B}) \langle b \rangle \{ \text{if } b \text{ then } Q_1 \text{ else } Q_2 \}. \text{if } b \text{ then } prot_1 \text{ else } prot_2 \\ & prot_1 \{Q_1\} \& \{Q_2\} prot_2 \triangleq \\ & \quad ? (b : \mathbb{B}) \langle b \rangle \{ \text{if } b \text{ then } Q_1 \text{ else } Q_2 \}. \text{if } b \text{ then } prot_1 \text{ else } prot_2 \end{aligned}$$

We often omit the conditions Q_1 and Q_2 , which simply means that they are `True`. The following rules can be directly derived from the rules HT-SEND and HT-RECV:

$$\text{HT-SELECT} \quad \left\{ \begin{array}{l} c \mapsto prot_1 \{Q_1\} \oplus \{Q_2\} prot_2 * \\ \text{if } b \text{ then } Q_1 \text{ else } Q_2 \end{array} \right\} \text{select } c \ b \ \{ c \mapsto \text{if } b \text{ then } prot_1 \text{ else } prot_2 \}$$

HT-BRANCH

$$\frac{\{P * Q_1 * c \mapsto prot_1\} e_1 \ \{\Phi\} \quad \{P * Q_2 * c \mapsto prot_2\} e_2 \ \{\Phi\}}{\{P * c \mapsto prot_1 \{Q_1\} \& \{Q_2\} prot_2\} \text{branch } c \ \text{with } \mathbf{left} \Rightarrow e_1 \mid \mathbf{right} \Rightarrow e_2 \ \text{end} \ \{\Phi\}}$$

Apart from branching on boolean values, dependent separation protocols can be used to encode choice on any enumeration type (*e.g.*, lists, natural numbers, days of the week, *etc.*). These encodings follow the same scheme.

3.2.6 Recursive Protocols

We will now use choice and recursion to verify the correctness of a sorting service that supports performing multiple sorting jobs in sequence. The code of the sorting service `sort_servicerec` and a possible client `sort_clientrec` are displayed in Figure 3.5. The service `sort_servicerec` contains a loop in which choice is used to either terminate the service, or to sort an individual list using the distributed merge sort algorithm `sort_service` from Section 3.2.3. The client `sort_clientrec` uses the service to sort a nested linked list l of linked lists. It performs this job by starting a single instance of the service at c , and then sequentially sends requests to sort each inner linked list l' in l . Finally, the client selects the terminating branch to end the communication with the service. A protocol for interacting with the sorting service can be defined as follows:

$$\text{sort_prot}_{rec} (I : T \rightarrow \text{Val} \rightarrow \text{iProp}) (R : T \rightarrow T \rightarrow \mathbb{B}) \triangleq \\ \mu(\text{rec} : \text{iProto}). (\text{sort_prot } I \ R \cdot \text{rec}) \oplus \mathbf{end}$$

The protocol uses the choice operator \oplus to specify that the client may either request the service to perform a sorting job, or terminate communication with the service. After the job has been finished the protocol proceeds recursively.

It is important to point out that—as is usual in logics with guarded recursion [Nakano 2000]—the variable x should appear under a *contractive* term construct in the body t of $\mu x : \tau. t$. In our protocol, the recursive variable rec appears under the

```

sort_servicerec cmp c :=
  branch c with
    left ⇒ sort_service cmp c;
           sort_servicerec cmp c
  | right ⇒ ()
  end

sort_clientrec cmp l :=
  let c := start sort_servicerec cmp in
  liter (λl'. select c left;
        send c l'; recv c) l;
  select c right
    
```

Figure 3.5: A recursive version of the sort service that can perform multiple jobs in sequence (the code for the function `liter`, which applies a function to each element of the list, is standard and has been elided).

argument of \oplus , which is defined in terms of $!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$, which, similarly to $?x:\vec{\tau}\langle v\rangle\{P\}.prot$, is contractive in the tail protocol $prot$. The specifications of the service and the client are as follows:

$$\left\{ \begin{array}{l} \text{cmp_spec } I \ R \ \text{cmp} \ * \\ c \mapsto \text{sort_prot}_{rec} \ I \ R \ . \ \text{prot} \end{array} \right\} \quad \left\{ \begin{array}{l} \text{cmp_spec } I \ R \ \text{cmp} \ * \ l \xrightarrow{\text{list}}_J \vec{x} \\ \text{sort_client}_{rec} \ \text{cmp} \ l \\ \exists \vec{y}. |\vec{y}| = |\vec{x}| \ * \ l \xrightarrow{\text{list}}_J \vec{y} \ * \ (\forall i < |\vec{x}|. \text{sorted_of}_R \ \vec{y}_i \ \vec{x}_i) \end{array} \right\}$$

$$\left\{ \begin{array}{l} \text{sort_service}_{rec} \ \text{cmp} \ c \\ c \mapsto \text{prot} \end{array} \right\}$$

We let $J \triangleq \lambda l'. l' \xrightarrow{\text{list}}_I \vec{y}$ to express that l points to a list of lists \vec{x} . The proof of the service follows naturally by symbolic execution using the induction hypothesis (obtained from LÖB), the rules HT-BRANCH and HT-SELECT, and the specification of `sort_service`. Note that we rely on the specification of `sort_service` having an arbitrary protocol as its suffix.

It is worth pointing out that protocols in Actris provide a lot of flexibility. Using just minor changes, we can extend the protocol to support transferring a comparison function over the channel, like the extension made in `sort_clientfunc`, or in a way such that a different comparison function can be used for each sorting job.

3.2.7 Delegation

Delegation is a common feature within communication protocols, and particularly the session-types community—it is the concept of transferring a channel endpoint over a channel. Due to the impredicativity of dependent separation protocols in Actris, reasoning about programs that make use of delegation is readily available. The protocols $!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ and $?x:\vec{\tau}\langle v\rangle\{P\}.prot$ can simply refer to the ownership of protocols $c \mapsto prot'$ in the proposition P .

An example of a program that uses delegation is the `sort_servicedel` variant of the recursive sorting service in Figure 3.6, which allows multiple sorting jobs to be performed in parallel. To enable parallelism, it delegates a new channel c' to an inner sorting service for each sorting job.

The client `sort_clientdel` once again uses the sorting service to sort a nested linked list l of linked lists. The client starts a connection c to the new service, and

```

sort_servicedel cmp c :=
  branch c with
    left ⇒
      let c' :=
        start sort_service cmp in
        send c';
        sort_servicedel cmp c
    | right ⇒ ()
  end

sort_clientdel cmp l =
  let c := start sort_servicedel cmp in
  let k := lnil () in
  liter (λl'. select c left;
        let c' := recv c in
        send c' l'; lcons c' k) l
  send c right;
  liter recv k
    
```

Figure 3.6: A recursive version of the sort service that uses delegation to perform multiple jobs in parallel (the code for the function `lcons`, which adds an element to the head of a list, has been elided).

for each inner list l' , it acquires a delegated channel c' , over which it sends a pointer l' to the inner list that should be sorted. The client keeps track of all channels to delegated services in a linked list k so that it can wait for all of them to finish (using `iter recv`).

A protocol for the delegation service can be defined as follows, denoting that the client can select whether to acquire a connection to a new delegated service or to terminate:

$$\text{sort_prot}_{del} (I : T \rightarrow \text{Val} \rightarrow \text{iProp}) (R : T \rightarrow T \rightarrow \mathbb{B}) \triangleq \\ \mu(\text{rec} : \text{iProto}). (? (c : \text{Chan}) \langle c \rangle \{c \mapsto \text{sort_prot } I \ R\}. \text{rec}) \oplus \text{end}$$

The specifications of the service and the client are as follows:

$$\left\{ \begin{array}{l} \text{cmp_spec } I \ R \ \text{cmp} \ * \\ c \mapsto \text{sort_prot}_{del} \ I \ R \cdot \text{prot} \\ \text{sort_service}_{del} \ \text{cmp} \ c \\ c \mapsto \text{prot} \end{array} \right\} \quad \left\{ \begin{array}{l} \text{cmp_spec } I \ R \ \text{cmp} \ * \ l \ \overset{\text{list}}{\mapsto}_J \ \vec{x} \\ \text{sort_client}_{del} \ \text{cmp} \ l \\ \exists \vec{y}. |\vec{y}| = |\vec{x}| \ * \ l \ \overset{\text{list}}{\mapsto}_J \ \vec{y} \ * \ (\forall i < |\vec{x}|. \text{sorted_of}_R \ \vec{y}_i \ \vec{x}_i) \end{array} \right\}$$

As before, we let $J \triangleq \lambda l' \vec{y}. l' \overset{\text{list}}{\mapsto}_I \vec{y}$ to express that l points to a list of lists \vec{x} . Once again the proofs are straightforward, as they are simply a combination of recursive reasoning combined with the application of Actris's rules for channels.

3.2.8 Dependent Protocols

The protocols we have seen so far have only made limited use of Actris's support for recursion. We now demonstrate Actris's support for dependent protocols, which make it possible to keep track of the history of what messages have been sent and received. We demonstrate this feature by considering a fine-grained version of the distributed merge-sort service. This version `sort_servicefg`, as shown in Figure 3.7, requires the input list to be sent element by element, after which the service sends

```

sort_servicefg cmp c :=
  branch c with
    right ⇒ select c right
  | left ⇒
    let x1 := recv c in
    branch c with
      right ⇒ select c left; send c x1;
        select c right
    | left ⇒
      let x2 := recv c in
      let c1 := start sort_servicefg cmp in
      let c2 := start sort_servicefg cmp in
      select c1 left; send c1 x1;
      select c2 left; send c2 x2;
      splitfg c c1 c2; mergefg cmp c c1 c2
    end
  end
end

splitfg c c1 c2 :=
  branch c with
    right ⇒ select c1 right;
      select c2 right
  | left ⇒
    let x := recv c in
    select c1 left; send c1 x;
    splitfg c c2 c1
  end

mergefg cmp c c1 c2 :=
  branch c1 with
    right ⇒ assert false
  | left ⇒
    let x := recv c1 in
    mergeauxfg cmp c x c1 c2
  end
end

mergeauxfg cmp c x c1 c2 :=
  branch c2 with
    right ⇒ select c left;
      send c x1;
      transfer c1 c
  | left ⇒
    let y := recv c2 in
    if cmp x y then
      select c left; send c x;
      mergeauxfg cmp c y c2 c1
    else
      select c left; send c y;
      mergeauxfg cmp c x c1 c2
    end
  end

sort_clientfg cmp l :=
  let c :=
    start sort_servicefg cmp in
  send.all c l; recv.all c l
  
```

Figure 3.7: A fine-grained version of the sort service that transfers elements one by one (the code for the functions **transfer**, **send_all**, and **recv_all** has been elided).

the sorted list back in the same fashion. We use choice to indicate whether the whole list has been sent (**right**) or another element remains to be sent (**left**).

The structure of **sort_service**_{fg} is somewhat similar to the coarse-grained merge-sort algorithm that we have seen before. The base cases of the empty or the singleton list are handled initially. This is achieved by waiting for at least two values before starting the recursive sub-services *c*₁ and *c*₂. In the base cases the values are sent back immediately, as they are trivially sorted. The inductive case is handled by starting two sub-services at *c*₁ and *c*₂ that are sent the two initially received elements, respectively, after which the **split**_{fg} function is used to receive and forward the remaining values to the sub-services alternately. Once the **right** flag is received, the **split**_{fg} function terminates, and the algorithm moves to the second phase in which the **merge**_{fg} function merges the stream of values returned by the sub-services and forwards them to the parent service.

The merge_{fg} function initially acquires the first value x from the first sub-service, which it uses in the recursive call as the current largest value. The recursive function merge_{fg}^{aux} recursively requests a value y from the sub-service of which the current largest value was not acquired from. It then compares x and y using the comparison function cmp , and forwards the smallest element. This is repeated until the **right** flag is received from either sub-service, after which the remaining values of the other are forwarded to the parent service using the **transfer** function.

The interface of the client sort_client_{fg} for this service is similar to the previous ones. It takes a reference to a linked list, which is then sorted. It performs this task by sending the elements of the linked list to the sort service using the **send_all** function (which modifies the list in-place by removing the sent elements), and puts the received values back into the linked list using the **recv_all** function (which also modifies the list in-place). A suitable protocol for proving functional correctness of the fine-grained sorting service is as follows:

$$\begin{aligned}
 & \text{sort_prot}_{fg} (I : T \rightarrow \text{Val} \rightarrow \text{iProp}) (R : T \rightarrow T \rightarrow \mathbb{B}) \triangleq \text{sort_prot}_{fg}^{head} I R [] \\
 & \text{sort_prot}_{fg}^{head} (I : T \rightarrow \text{Val} \rightarrow \text{iProp}) (R : T \rightarrow T \rightarrow \mathbb{B}) \triangleq \mu(\text{rec} : \text{List } T \rightarrow \text{iProto}). \\
 & \lambda \vec{x}. (! (x : T) (v : \text{Val}) \langle v \rangle \{I x v\}. \text{rec } (\vec{x} \cdot [x]) \} \oplus \text{sort_prot}_{fg}^{tail} I R \vec{x} []) \\
 & \text{sort_prot}_{fg}^{tail} (I : T \rightarrow \text{Val} \rightarrow \text{iProp}) (R : T \rightarrow T \rightarrow \mathbb{B}) \triangleq \\
 & \mu(\text{rec} : \text{List } T \rightarrow \text{List } T \rightarrow \text{iProto}). \lambda \vec{x} \vec{y}. \\
 & (? (y : T) (v : \text{Val}) \langle v \rangle \{(\forall i < |\vec{y}|. R \vec{y}_i y) * I y v\}. \text{rec } \vec{x} (\vec{y} \cdot [y]) \} \ \&_{\{\vec{x} \equiv_p \vec{y}\}} \text{end}
 \end{aligned}$$

The protocol is split into two phases $\text{sort_prot}_{fg}^{head}$ and $\text{sort_prot}_{fg}^{tail}$, mimicking the behaviour of the program. The $\text{sort_prot}_{fg}^{head}$ phase is indexed by the values \vec{x} that have been sent so far. The protocol describes that one can either send another value and proceed recursively, or stop, which moves the protocol to the next phase.

The $\text{sort_prot}_{fg}^{tail}$ phase is dependent on the list of values \vec{x} received in the first phase, and the list of values \vec{y} returned so far. The condition $(\forall i < |\vec{y}|. R \vec{y}_i y)$ states that the received element is larger than any of the elements that have previously been returned, which maintains the invariant that the stream of received elements is sorted. When the **right** flag is received $\vec{x} \equiv_p \vec{y}$ shows that the received values \vec{y} are a permutation of the ones \vec{x} that were sent, making sure that all of the sent elements have been accounted for.

The top-level specification of the service and client are similar to the specifications of the coarse-grained version of distributed merge sort:

$$\begin{array}{ll}
 \{ \text{cmp_spec } I R \text{ cmp} * c \mapsto \overline{\text{sort_prot}_{fg}} \cdot \text{prot} \} & \{ \text{cmp_spec } I R \text{ cmp} * l \stackrel{\text{list}}{\mapsto}_I \vec{x} \} \\
 \text{sort_prot}_{fg} c & \text{sort_client}_{fg} \text{ cmp } l \\
 \{ c \mapsto \text{prot} \} & \{ \exists \vec{y}. l \stackrel{\text{list}}{\mapsto}_I \vec{y} * \text{sorted_of}_R \vec{y} \vec{x} \}
 \end{array}$$

Proving these specifications requires one to pick appropriate specifications for the auxiliary functions to capture the required invariants with regard to sorting. After having picked these specifications, the parts of the proofs that involve communication are mostly straightforward, but require a number of trivial auxiliary results about sorting and permutations.

3.3 Subprotocols

This section describes **Actris 2.0**, which extends Actris 1.0—as presented in the conference version of this paper [Hinrichsen et al. 2020]—with *subprotocols*, inspired by asynchronous subtyping of session types. Subprotocols have two key features. First, they expose the asynchronous nature of channels in the Actris logic by relaxing the requirements of duality, thereby making it possible to prove functional correctness of a larger class of programs. Second, they give rise to a more extensional approach to reasoning about dependent separation protocols, as we can work up to the subprotocol relation and not equality, thereby providing more flexibility in the design and reuse of protocols.

We first introduce the subprotocol relation and its proof rules (Section 3.3.1). We then show how subprotocols can be employed to prove a mapper service, which handles requests one at a time, while its client may send multiple requests up front (Section 3.3.2). Next, we verify a list reversal service whose protocol involves a minimal specification of lists, which we then reuse through subprotocols to obtain a protocol with a more generic specification of lists (Section 3.3.3). Finally, we show that the subprotocol relation is coinductive, and thereby when combined with LÖB induction, can be used to reason about recursive protocols (Section 3.3.4).

3.3.1 The Subprotocol Relation

The dependent separation protocol of a channel is picked upon channel creation (using the rule HT-NEW), which then determines how the channel endpoints should interact. To ensure safe communication, Actris adapts the notion of duality from session types, which requires every send (!) of one endpoint to be paired with a receive (?) for the other endpoint, and *vice versa*. However, strictly working with a channel’s protocol and its dual is more restrictive than necessary, as either endpoint is agnostic to some variance from the protocol by the other endpoint. We capture the flexibility of safe variances via a new notion—the *subprotocol relation*:

$$prot_1 \sqsubseteq prot_2$$

This relation describes that protocol $prot_1$ is *stronger* than $prot_2$, or conversely, that protocol $prot_2$ is *weaker* than $prot_1$. More specifically, this means that $prot_2$ can be used *in place of* $prot_1$ whenever such a protocol is expected during the verification. This property is captured by the following monotonicity rule for channel ownership:

$$\frac{c \rightsquigarrow prot_1 \quad prot_1 \sqsubseteq prot_2}{c \rightsquigarrow prot_2}$$

The subprotocol relation is inspired by asynchronous subtyping for session types [Mostrous et al. 2009; Mostrous and Yoshida 2015], which allows (1) sending subtypes (contravariance), (2) receiving supertypes (covariance), and (3) swapping sends ahead of receives. These variations are sound, as (1) the originally expected type that is to be sent can be derived from the subtype, (2) the originally expected type that is to be received can be derived from the supertype, and (3) sends do not block because

channels are buffered in both directions. These variances, along with the swapping property, can be generalised to dependent separation protocols using the following proof rules:

$$\frac{\sqsubseteq\text{-SEND-MONO}' \quad \forall \vec{x}:\vec{\tau}. P_2 \multimap P_1 \quad \forall \vec{x}:\vec{\tau}. \text{prot}_1 \sqsubseteq \text{prot}_2}{!\vec{x}:\vec{\tau} \langle v \rangle \{P_1\}. \text{prot}_1 \sqsubseteq !\vec{x}:\vec{\tau} \langle v \rangle \{P_2\}. \text{prot}_2}$$

$$\frac{\sqsubseteq\text{-RECV-MONO}' \quad \forall \vec{x}:\vec{\tau}. P_1 \multimap P_2 \quad \forall \vec{x}:\vec{\tau}. \text{prot}_1 \sqsubseteq \text{prot}_2}{?\vec{x}:\vec{\tau} \langle v \rangle \{P_1\}. \text{prot}_1 \sqsubseteq ?\vec{x}:\vec{\tau} \langle v \rangle \{P_2\}. \text{prot}_2}$$

$$\frac{\sqsubseteq\text{-SWAP}' \quad ?\vec{x}:\vec{\tau} \langle v \rangle \{P\}. !\vec{y}:\vec{\sigma} \langle w \rangle \{Q\}. \text{prot} \sqsubseteq !\vec{y}:\vec{\sigma} \langle w \rangle \{Q\}. ?\vec{x}:\vec{\tau} \langle v \rangle \{P\}. \text{prot}}$$

The rules $\sqsubseteq\text{-SEND-MONO}'$ and $\sqsubseteq\text{-RECV-MONO}'$ use *separation implication* $P \multimap Q$ —which states that ownership of Q can be obtained by giving up ownership of P —to mimic the contra- and covariance of session subtyping. The rule $\sqsubseteq\text{-SWAP}'$ states that sends can be swapped ahead of receives. To be well-formed, this rule has the implicit side condition that $\vec{x}:\vec{\tau}$ does not bind into w and Q , and that $\vec{y}:\vec{\sigma}$ does not bind into v and P .

With this set of rules we can make subprotocol derivations such as the following:

$$\begin{array}{ll} ?(i : \mathbb{Z}) \langle i \rangle \{i < 42\}. !(j : \mathbb{Z}) \langle j \rangle \{j > 42\}. \text{prot} & \sqsubseteq\text{-SEND-MONO}' \\ \sqsubseteq ?(i : \mathbb{Z}) \langle i \rangle \{i < 42\}. !(j : \mathbb{Z}) \langle j \rangle \{j > 50\}. \text{prot} & \sqsubseteq\text{-RECV-MONO}' \\ \sqsubseteq ?(i : \mathbb{Z}) \langle i \rangle \{i < 40\}. !(j : \mathbb{Z}) \langle j \rangle \{j > 50\}. \text{prot} & \sqsubseteq\text{-SWAP}' \\ \sqsubseteq !(j : \mathbb{Z}) \langle j \rangle \{j > 50\}. ?(i : \mathbb{Z}) \langle i \rangle \{i < 40\}. \text{prot} & \end{array}$$

Here, we strengthen the proposition of the send (by increasing the bound from $j > 42$ to $j > 50$), weaken the proposition of the receive (by reducing the bound from $i < 42$ to $i < 40$), while also swapping the send ahead of the receive.

While the aforementioned rules cover the intuition behind Actris's subprotocol relation, Actris's actual rules for subprotocols provide a number of additional features.

1. They can manipulate the logical variables $\vec{x}:\vec{\tau}$ that appear in protocols.
2. They can transfer ownership of resources in and out of messages.
3. They can reason about recursive protocols defined using Löb induction.

The key idea to obtain these features is to consider the subprotocol relation $\text{prot}_1 \sqsubseteq \text{prot}_2$ as a first-class logical connective. That is, $\text{prot}_1 \sqsubseteq \text{prot}_2$ is an Iris proposition that can be combined freely with the logical connectives of Iris and Actris (e.g., separating conjunction, separating implication, higher-order quantification). Since $\text{prot}_1 \sqsubseteq \text{prot}_2$ is an Iris proposition, the proof rules are in fact (separating) implications in Iris. For readability, we use inference rules to denote each rule ($P_1 * \dots * P_n \multimap Q$) as:

$$\frac{P_1 \quad \dots \quad P_n}{Q}$$

Logical variable manipulation and resource transfer:

$$\frac{\text{\(\(\subseteq\)-SEND-OUT}}{\forall \vec{x} : \vec{\tau}. P \multimap (\text{\(\subseteq\)} prot_1 \subseteq !\langle v \rangle. prot_2)} \quad prot_1 \neq \text{end}}{\text{\(\subseteq\)} prot_1 \subseteq !\vec{x} : \vec{\tau} \langle v \rangle \{P\}. prot_2}$$

$$\frac{\text{\(\subseteq\)-SEND-IN}}{P[\vec{t}/\vec{x}]}{\text{\(\subseteq\)} !\vec{x} : \vec{\tau} \langle v \rangle \{P\}. prot \subseteq !\langle v[\vec{t}/\vec{x}] \rangle. prot[\vec{t}/\vec{x}]}$$

$$\frac{\text{\(\subseteq\)-RECV-OUT}}{\forall \vec{x} : \vec{\tau}. P \multimap (\text{\(\subseteq\)} \langle v \rangle. prot_1 \subseteq prot_2)} \quad prot_2 \neq \text{end}}{\text{\(\subseteq\)} ?\vec{x} : \vec{\tau} \langle v \rangle \{P\}. prot_1 \subseteq prot_2}$$

$$\frac{\text{\(\subseteq\)-RECV-IN}}{P[\vec{t}/\vec{x}]}{\text{\(\subseteq\)} \langle v[\vec{t}/\vec{x}] \rangle. prot[\vec{t}/\vec{x}] \subseteq ?\vec{x} : \vec{\tau} \langle v \rangle \{P\}. prot}$$

Monotonicity and swapping:

$$\frac{\text{\(\subseteq\)-SEND-MONO}}{\triangleright(\text{\(\subseteq\)} prot_1 \subseteq prot_2)}{\text{\(\subseteq\)} !\langle v \rangle. prot_1 \subseteq !\langle v \rangle. prot_2}$$

$$\frac{\text{\(\subseteq\)-RECV-MONO}}{\triangleright(\text{\(\subseteq\)} prot_1 \subseteq prot_2)}{\text{\(\subseteq\)} ?\langle v \rangle. prot_1 \subseteq ?\langle v \rangle. prot_2}$$

$$\frac{\text{\(\subseteq\)-SWAP}}{?\langle v \rangle. !\langle w \rangle. prot \subseteq !\langle w \rangle. ?\langle v \rangle. prot}$$

Reflexivity and transitivity:

$$\frac{\text{\(\subseteq\)-REFL}}{prot \subseteq prot} \quad \frac{\text{\(\subseteq\)-TRANS}}{prot_1 \subseteq prot_2 \quad prot_2 \subseteq prot_3}{prot_1 \subseteq prot_3}$$

Dual and append:

$$\frac{\text{\(\subseteq\)-DUAL}}{prot_2 \subseteq prot_1}{prot_1 \subseteq prot_2} \quad \frac{\text{\(\subseteq\)-APPEND}}{prot_1 \subseteq prot_2 \quad prot_3 \subseteq prot_4}{prot_1 \cdot prot_3 \subseteq prot_2 \cdot prot_4}$$

Channel ownership:

$$\frac{\text{\(\subseteq\)-CHAN-MONO}}{c \multimap prot_1 \quad prot_1 \subseteq prot_2}{c \multimap prot_2}$$

Figure 3.8: The rules of Actris 2.0 for subprotocols.

The full set of rules for subprotocols is shown in Figure 3.8. The first four rules account for logical variable manipulation and resource transfer: Rules \sqsubseteq -SEND-OUT and \sqsubseteq -RECV-OUT generalise over the logical variables $\vec{x} : \vec{\tau}$ and transfer ownership of P out of the weaker sending protocol $!\vec{x} : \vec{\tau} \langle v \rangle \{P\}.prot$, and stronger receiving protocol $?\vec{x} : \vec{\tau} \langle v \rangle \{P\}.prot$, respectively. Rule \sqsubseteq -SEND-IN weakens a sending protocol $!\vec{x} : \vec{\tau} \langle v \rangle \{P\}.prot$ by instantiating the logical variables $\vec{x} : \vec{\tau}$ and transferring ownership of $P[\vec{t}/\vec{x}]$ into the protocol. Dually, the rule \sqsubseteq -RECV-IN strengthens a receiving protocol $?\vec{x} : \vec{\tau} \langle v \rangle \{P\}.prot$ by instantiating the logical variables $\vec{x} : \vec{\tau}$ and transferring ownership of $P[\vec{t}/\vec{x}]$ into the protocol.

To demonstrate the intuition behind these rules consider the following proof of the subprotocol relation presented in Section 3.1.3, where we transfer ownership of $\ell_1 \mapsto 20$ into a protocol, while instantiating the logical variable ℓ_1 accordingly:

$$\begin{array}{c}
 \frac{}{\ell'_1 \mapsto 20 * \ell_2 \mapsto 22 * *} \\
 \frac{}{!(\ell_1, \ell_2 : \text{Loc}) \langle (\ell_1, \ell_2) \rangle \{ \ell_1 \mapsto 20 * \ell_2 \mapsto 22 \}.prot \sqsubseteq ! \langle (\ell'_1, \ell_2) \rangle .prot} \sqsubseteq\text{-SEND-OUT} \\
 \hline
 \frac{}{\ell'_1 \mapsto 20 * *} \sqsubseteq\text{-SEND-IN} \\
 \frac{}{!(\ell_1, \ell_2 : \text{Loc}) \langle (\ell_1, \ell_2) \rangle \{ \ell_1 \mapsto 20 * \ell_2 \mapsto 22 \}.prot \sqsubseteq} \\
 \frac{}{!(\ell_2 : \text{Loc}) \langle (\ell'_1, \ell_2) \rangle \{ \ell_2 \mapsto 22 \}.prot}
 \end{array}$$

We first use rule \sqsubseteq -SEND-OUT to generalise over the logical variable ℓ_2 and transfer ownership of $\ell_2 \mapsto 22$ out of the weaker protocol (*i.e.*, the send on the RHS), and then use \sqsubseteq -SEND-IN to instantiate the logical variables ℓ'_1 and ℓ_2 and transfer ownership of $\ell'_1 \mapsto 20$ and $\ell_2 \mapsto 22$ into the stronger protocol (*i.e.*, the send on the LHS).

The rules for monotonicity (\sqsubseteq -SEND-MONO and \sqsubseteq -RECV-MONO) and swapping (\sqsubseteq -SWAP) in Figure 3.8 differ in two aspects from the rules for monotonicity (\sqsubseteq -SEND-MONO' and \sqsubseteq -RECV-MONO') and swapping (\sqsubseteq -SWAP') that we have seen in the beginning of this section. First, the actual rules only apply to protocols whose head does not have logical variables $\vec{x} : \vec{\tau}$ and resources P , *i.e.*, protocols of the shape $!\langle v \rangle .prot$ or $?\langle v \rangle .prot$, instead of those of the shape $!\vec{x} : \vec{\tau} \langle v \rangle \{P\}.prot$ or $?\vec{x} : \vec{\tau} \langle v \rangle \{P\}.prot$. While this restriction might seem to make the rules more restrictive, the more general rules for monotonicity (\sqsubseteq -SEND-MONO' and \sqsubseteq -RECV-MONO') and swapping (\sqsubseteq -SWAP') are derivable from these simpler rules. This is done using the rules for logical variable manipulation and resource transfer. Second, the actual rules for monotonicity have a later modality (\triangleright) in their premise. The later modality makes these rules stronger (by \triangleright -INTRO we have that P entails $\triangleright P$), and thereby internalizes its coinductive nature into the Actris logic so LÖB induction can be used to prove subprotocol relations for recursive protocols (Section 3.3.4).

The remaining rules in Figure 3.8 express that the subprotocol relation is reflexive (\sqsubseteq -REFL) and transitive (\sqsubseteq -TRANS), as well as that the dual operation is anti-monotone (\sqsubseteq -DUAL) and the append operation is monotone (\sqsubseteq -APPEND).

Let us consider the following subprotocol relation to provide some further insight into the expressivity of our rules, where logical variables are omitted for simplicity:

$$!\langle v \rangle \{P\}. ?\langle w \rangle \{Q\}. prot \sqsubseteq !\langle v \rangle \{P * R\}. ?\langle w \rangle \{Q * R\}. prot$$

Here we extend the protocol $!\langle v \rangle\{P\}.?\langle w \rangle\{Q\}.prot$ with a so-called *frame* R , which describes resources that must be sent along with the originally expected resources P , and which are reacquired along with the resources Q that are sent back. The above subprotocol relation mimics the frame rule of separation logic (HT-FRAME), which makes it possible to apply specifications while maintaining a *frame* of resources R :

$$\frac{\{P\}e\{w.Q\}}{\{P * R\}e\{w.Q * R\}}$$

The frame-like subprotocol relation is proven as follows:

$$\frac{}{Q * R \multimap \quad ?\langle w \rangle.prot \sqsubseteq ?\langle w \rangle\{Q * R\}.prot} \quad \sqsubseteq\text{-RECV-IN}$$

$$\frac{}{R \multimap \quad ?\langle w \rangle\{Q\}.prot \sqsubseteq ?\langle w \rangle\{Q * R\}.prot} \quad \sqsubseteq\text{-RECV-OUT}$$

$$\frac{}{R \multimap \quad !\langle v \rangle.?\langle w \rangle\{Q\}.prot \sqsubseteq !\langle v \rangle.?\langle w \rangle\{Q * R\}.prot} \quad \sqsubseteq\text{-SEND-MONO}$$

$$\frac{}{P * R \multimap !\langle v \rangle\{P\}.?\langle w \rangle\{Q\}.prot \sqsubseteq !\langle v \rangle.?\langle w \rangle\{Q * R\}.prot} \quad \sqsubseteq\text{-SEND-IN}$$

$$\frac{}{!\langle v \rangle\{P\}.?\langle w \rangle\{Q\}.prot \sqsubseteq !\langle v \rangle\{P * R\}.?\langle w \rangle\{Q * R\}.prot} \quad \sqsubseteq\text{-SEND-OUT}$$

We use rule $\sqsubseteq\text{-SEND-OUT}$ to transfer P and the frame R out of the weaker protocol (*i.e.*, the send on the RHS), and then use rule $\sqsubseteq\text{-SEND-IN}$ to transfer P into the stronger protocol (*i.e.*, the send on the LHS), leaving us with a context in which we still own the frame R . We then use rule $\sqsubseteq\text{-SEND-MONO}$ to proceed with the receiving part of the protocol in a dual fashion—we use rule $\sqsubseteq\text{-RECV-OUT}$ to transfer out Q of the stronger protocol (*i.e.*, the receive on the LHS), and use rule $\sqsubseteq\text{-RECV-IN}$ to transfer Q and the frame R into the weaker protocol (*i.e.*, the receive on the RHS).

3.3.2 Swapping

Subprotocols make it possible to verify message-passing programs whose order of sends and receives does not match up w.r.t. duality. As an example of such a program, let us consider the mapper service and client in Figure 3.9. The service `mapper_service` is a loop, which iteratively receives an element, maps a function over that element, and sends the resulting value back. Conversely, the client `mapper_client` sends all of the elements of the list l up front, and only requests the mapped results back once all elements have been sent. Since the former interleaves the sends and receives, while the latter does not, the dependent separation protocols for the service and client cannot be dual of each other. However, the communication between the service and client is in fact safe as messages are buffered. We now show that using subprotocols we can prove that this is indeed the case. We define the protocol based on the interleaved communication:

$$\text{mapper_prot } (I_T : T \rightarrow \text{Val} \rightarrow \text{iProp}) (I_U : U \rightarrow \text{Val} \rightarrow \text{iProp}) (f : T \rightarrow U) \triangleq$$

$$\mu(\text{rec} : \text{iProto}).$$

$$(!\langle x : T \rangle (v : \text{Val}) \langle v \rangle\{I_T x v\}.?\langle w : \text{Val} \rangle \langle w \rangle\{I_U (f x) w\}.rec) \oplus \text{end}$$

The protocol is parameterised by representation predicates I_T and I_U that relate language-level values to elements of type T and U in the Iris/Actris logic, and a

```

mapper_service f_v c :=
  branch c with
    left  => send (f_v (recv c)) ;
           mapper_service f_v c
    | right => ()
  end

mapper_client f_v l :=
  let c := start (mapper_service f_v) in
  let n := |l| in
  send_all c l; recvN c l n;
  select c right;

```

Figure 3.9: A mapper service whose verification relies on swapping (the code for the functions `send_all` and `recvN` has been elided).

function $f : T \rightarrow U$ in Iris/Actris that specifies the behaviour of the language-level function f_v . The connection between f and f_v is formalised as:

$$\text{f_spec } (I_T : T \rightarrow \text{Val} \rightarrow \text{iProp}) (I_U : U \rightarrow \text{Val} \rightarrow \text{iProp}) (f : T \rightarrow U) (f_v : \text{Val}) \triangleq \forall x v. \{I_T x v\} f_v v \{w. I_U (f x) w\}$$

Since `mapper_prot` describes an interleaved sequence of transactions, `mapper_service` can be readily verified against the protocol $\overline{\text{mapper_prot}}$ using just the symbolic execution rules from Section 3.2. However, to verify `mapper_client` against the protocol `mapper_prot`, we need to weaken the protocol using Actris’s rules for subprotocols. Given a list of n elements, the subprotocol relation (together with an intermediate step) that describes this weakening is:

$$\begin{aligned} & \text{mapper_prot } I_T I_U f \\ \sqsubseteq & \text{!}(\text{left}).!(x_1 : T) (v_1 : \text{Val}) \langle v_1 \rangle \{I_T x_1 v_1\}. \quad n \text{ times } \mu\text{-UNFOLD and} \\ & ?(y_1 : U) \langle y_1 \rangle \{I_U (f x_1) y_1\}. \cdots \quad \text{weaken } \oplus \text{ into !}(\text{left}) \\ & \text{!}(\text{left}).!(x_n : T) (v_n : \text{Val}) \langle v_n \rangle \{I_T x_n v_n\}. \\ & ?(y_n : U) \langle y_n \rangle \{I_U (f x_n) y_n\}. \\ & \text{mapper_prot } I_T I_U f \\ \sqsubseteq & \text{!}(\text{left}).!(x_1 : T) (v_1 : \text{Val}) \langle v_1 \rangle \{I_T x_1 v_1\}. \cdots \quad n \text{ times } \sqsubseteq\text{-SWAP}' \\ & \text{!}(\text{left}).!(x_n : T) (v_n : \text{Val}) \langle v_n \rangle \{I_T x_n v_n\}. \\ & ?(y_1 : U) \langle y_1 \rangle \{I_U (f x_1) y_1\}. \cdots \\ & ?(y_n : U) \langle y_n \rangle \{I_U (f x_n) y_n\}. \\ & \text{mapper_prot } I_T I_U f \end{aligned}$$

Both steps are proven by induction on n . In the first step, we unfold the recursive protocol n times using $\mu\text{-UNFOLD}$ and the derived rule $(\text{prot}_1 \oplus \text{prot}_2) \sqsubseteq \text{!}(\text{left}). \text{prot}_1$ to weaken the choices. Recall from Section 3.2.5 that \oplus is defined in terms of the send protocol (!), allowing us to prove $(\text{prot}_1 \oplus \text{prot}_2) \sqsubseteq \text{!}(\text{left}). \text{prot}_1$ using $\sqsubseteq\text{-SEND-OUT}$ and $\sqsubseteq\text{-SEND-IN}$. The second step involves swapping all sends ahead of the receives using the rule $\sqsubseteq\text{-SWAP}'$.

The weakened protocol that we have obtained follows the behaviour of the client, making its verification straightforward using Actris’s rules for symbolic execution.

Concretely, we prove the following specifications for the service and the client:

$$\left\{ \begin{array}{l} \text{f_spec } I_T I_U f f_V * \\ c \mapsto \overline{\text{mapper_prot } I_T I_U f \cdot \text{prot}} \\ \text{mapper_service } f_V c \\ \{c \mapsto \text{prot}\} \end{array} \right\} \quad \left\{ \begin{array}{l} \text{f_spec } I_T I_U f f_V * l \mapsto_{I_T} \vec{x} \\ \text{mapper_client } f_V l \\ \{l \mapsto_{I_U} \text{map } f \vec{x}\} \end{array} \right\}$$

3.3.3 Minimal Protocols

An essential feature of separation logic is the ability to assign strong and minimal specifications to libraries, so that each library can be verified once against its specification, which in turn can be used to verify as many client programs as possible. To achieve a similar goal for message-passing programs we would like to assign strong and minimal protocols to services, so that each service can be verified once against its protocol, which in turn can be used to verify as many clients as possible. One of the key ingredients of separation logic to allow for such specifications is the frame rule (HT-FRAME). In Section 3.3.1 we showed that subprotocols allow framing in protocols. In this section we give a detailed example of framing in protocols by considering the service `list_rev_service` in Figure 3.10, which receives a linked list, reverses it, and sends it back.

To specify this service, we could use a protocol similar to the sorting service in Section 3.2.3, defined in terms of the representation predicate $l \mapsto_{I_T} \vec{x}$ for linked lists:

$$\text{list_rev_prot}_{I_T} \triangleq ! (l : \text{Loc}) (\vec{x} : \text{List } T) \langle l \rangle \{ l \mapsto_{I_T} \vec{x} \}. ? \langle () \rangle \{ l \mapsto_{I_T} \text{reverse } \vec{x} \}. \text{end}$$

Although it is possible to verify the service against the protocol $\overline{\text{list_rev_prot}_{I_T}}$, that approach is not quite satisfactory. Unlike the sorting service, the reversal service does not access the list elements, but only changes the structure of the list. Hence, there is no need to keep track of the ownership of the elements through the predicate I_T . A self-contained and minimal protocol for this service would instead be the following:

$$\text{list_rev_prot} \triangleq ! (l : \text{Loc}) (\vec{v} : \text{List Val}) \langle l \rangle \{ l \mapsto \vec{v} \}. ? \langle () \rangle \{ l \mapsto \text{reverse } \vec{v} \}. \text{end}$$

Here, $l \mapsto \vec{v}$ is a version of the list representation predicate that does not keep track of the resources of the elements, but only describes the structure of the list. It is defined as:

$$l \mapsto \vec{v} \triangleq \begin{cases} l \mapsto \text{inl } () & \text{if } \vec{v} = [] \\ \exists l_2. l \mapsto \text{inr } (v_1, l_2) * l_2 \mapsto \vec{v}_2 & \text{if } \vec{v} = [v_1] \cdot \vec{v}_2 \end{cases}$$

Once we have verified the service against the minimal protocol, a client might still want to interact with the list reversal service through the protocol $\text{list_rev_prot}_{I_T}$. This can be achieved by proving the subprotocol relation $\text{list_rev_prot} \sqsubseteq \text{list_rev_prot}_{I_T}$. To do so, we first establish a relation between the two list representation predicates:

$$l \mapsto_{I_T} \vec{x} ** (\exists \vec{v}. l \mapsto \vec{v} * \bigstar_{(x,v) \in (\vec{x}, \vec{v})} . I_T x v) \quad (\text{LIST-REL})$$

```

list_rev_service c :=
  let l := recv c in lreverse l; send ()

list_rev_client l :=
  let c := start list_rev_service in
  send c l; recv c

```

Figure 3.10: A list reversing service (the code for the function `lreverse` has been elided).

Here, $\star_{(x,v) \in (\vec{x}, \vec{v})}$ is the pair-wise iterated separation conjunction over two lists of equal length, and $\ast\ast$ is a bi-directional separation implication. The above result thus states that $l \overset{\text{list}}{\mapsto}_{I_T} \vec{x}$ can be split into two parts, ownership of the links of the list $l \overset{\text{list}}{\mapsto} \vec{v}$, and a range of interpretation predicates I_T for each element of the list, and *vice versa*. With this result at hand, the proof of the desired subprotocol relation is carried out as follows:

$$\begin{aligned}
& \text{list_rev_prot} \\
&= !(l : \text{Loc})(\vec{v} : \text{List Val}) \langle l \rangle \{l \overset{\text{list}}{\mapsto} \vec{v}\}. ?\langle () \rangle \{l \overset{\text{list}}{\mapsto} \text{reverse } \vec{v}\}. \text{end} \\
&\sqsubseteq !(l : \text{Loc})(\vec{v} : \text{List Val})(\vec{x} : \text{List } T) \langle l \rangle \left\{ l \overset{\text{list}}{\mapsto} \vec{v} \ast \star_{(x,v) \in (\vec{x}, \vec{v})} . I_T x v \right\}. \\
&\quad ?\langle () \rangle \left\{ l \overset{\text{list}}{\mapsto} (\text{reverse } \vec{v}) \ast \star_{(x,v) \in (\vec{x}, \vec{v})} . I_T x v \right\}. \text{end} \\
&\sqsubseteq !(l : \text{Loc})(\vec{x} : \text{List } T) \langle l \rangle \{l \overset{\text{list}}{\mapsto}_{I_T} \vec{x}\}. ?\langle () \rangle \{l \overset{\text{list}}{\mapsto}_{I_T} \text{reverse } \vec{x}\}. \text{end} \\
&= \text{list_rev_prot}_{I_T}
\end{aligned}$$

We first frame the interpretation predicates owned by the list $\star_{(x,v) \in (\vec{x}, \vec{v})} . I_T x v$, using an approach similar to the frame example in Section 3.3.1, and then use LIST-REL to combine it with $l \overset{\text{list}}{\mapsto} \vec{v}$ and $l \overset{\text{list}}{\mapsto} \text{reverse } \vec{v}$ for the sending and receiving step, to turn them into $l \overset{\text{list}}{\mapsto}_{I_T} \vec{x}$ and $l \overset{\text{list}}{\mapsto}_{I_T} \text{reverse } \vec{x}$, respectively. Note that the logical variable \vec{v} is changed into \vec{x} , using the subprotocol rules for logical variable manipulation. With this subprotocol relation at hand, it is possible to prove the following specifications for the service and client:

$$\begin{array}{cc}
\{c \mapsto \overline{\text{list_rev_prot}} \cdot \text{prot}\} & \{l \overset{\text{list}}{\mapsto}_{I_T} \vec{x}\} \\
\text{list_rev_service } c & \text{list_rev_client } l \\
\{c \mapsto \text{prot}\} & \{l \overset{\text{list}}{\mapsto}_{I_T} \text{reverse } \vec{x}\}
\end{array}$$

3.3.4 Subprotocols and Recursion

We conclude this section by showing how subprotocol relations involving recursive protocols can be proved using LÖB induction. Recall from Section 3.2 that the principle of LÖB induction is as follows:

$$(\triangleright P \Rightarrow P) \Rightarrow P$$

By letting P to be $\text{prot}_1 \sqsubseteq \text{prot}_2$, this means that we can prove $\text{prot}_1 \sqsubseteq \text{prot}_2$ under the assumption of the induction hypothesis $\triangleright(\text{prot}_1 \sqsubseteq \text{prot}_2)$. The later modality (\triangleright)

ensures that we do not immediately use the induction hypothesis, but first apply the monotonicity rule for send (\sqsubseteq -SEND-MONO) or receive (\sqsubseteq -RECV-MONO), which is done typically after unfolding the recursion operator using μ -UNFOLD. The monotonicity rules \sqsubseteq -SEND-MONO or \sqsubseteq -RECV-MONO contain a later modality (\triangleright) in their premise, which makes it possible to strip off the later of the induction hypotheses (by monotonicity of \triangleright).

Our approach for proving subprotocol relations using LöB induction is similar to that of Brandt and Henglein [1998] for proving subtyping relations for recursive types using coinduction. Brandt and Henglein [1998] however have a syntactic restriction on proofs to ensure that the induction hypothesis is not used immediately (*i.e.*, is used in a *contractive* fashion), while we use the later modality (\triangleright) of Iris to achieve that. To demonstrate how our approach works, we prove $prot_1 \sqsubseteq prot_2$, where:

$$\begin{aligned} prot_1 &\triangleq \mu(rec : iProto). (list_rev_prot \cdot rec) \oplus \mathbf{end} \\ prot_2 &\triangleq \mu(rec : iProto). (list_rev_prot_{I_T} \cdot rec) \oplus \mathbf{end} \end{aligned}$$

Here, $list_rev_prot$ and $list_rev_prot_{I_T}$ are the protocols from Section 3.3.3, for which we already proved $list_rev_prot \sqsubseteq list_rev_prot_{I_T}$. The proof is as follows:

$$\begin{array}{c} \frac{}{prot_1 \sqsubseteq prot_2 \text{ } \ast} \qquad \frac{}{prot_1 \sqsubseteq prot_2} \qquad (*) \\ \frac{prot_1 \sqsubseteq prot_2 \text{ } \ast \qquad list_rev_prot \cdot prot_1 \sqsubseteq list_rev_prot_{I_T} \cdot prot_2}{\triangleright(prot_1 \sqsubseteq prot_2) \text{ } \ast \qquad \triangleright(list_rev_prot \cdot prot_1 \sqsubseteq list_rev_prot_{I_T} \cdot prot_2)} \triangleright \text{ mono} \\ \frac{\triangleright(prot_1 \sqsubseteq prot_2) \text{ } \ast \qquad \triangleright(list_rev_prot \cdot prot_1 \sqsubseteq list_rev_prot_{I_T} \cdot prot_2)}{\triangleright(prot_1 \sqsubseteq prot_2) \text{ } \ast \qquad (list_rev_prot \cdot prot_1) \oplus \mathbf{end} \sqsubseteq (list_rev_prot_{I_T} \cdot prot_2) \oplus \mathbf{end}} \oplus \text{ mono} \\ \frac{}{\triangleright(prot_1 \sqsubseteq prot_2) \text{ } \ast \qquad (list_rev_prot \cdot prot_1) \oplus \mathbf{end} \sqsubseteq (list_rev_prot_{I_T} \cdot prot_2) \oplus \mathbf{end}} \mu\text{-UNFOLD} \\ \frac{\triangleright(prot_1 \sqsubseteq prot_2) \text{ } \ast \qquad prot_1 \sqsubseteq prot_2}{prot_1 \sqsubseteq prot_2} \text{ L\"OB} \end{array}$$

The proof starts with rule LÖB and unfolding the recursive types. We then proceed with monotonicity of \oplus , *i.e.*, $\triangleright(prot_1 \sqsubseteq prot_2 \wedge prot_3 \sqsubseteq prot_4) \text{ } \ast (prot_1 \oplus prot_3) \sqsubseteq (prot_2 \oplus prot_4)$, which itself follows from \sqsubseteq -SEND-MONO since selection (\oplus) is defined in terms of send (!). In step (*), we use \sqsubseteq -APPEND and $list_rev_prot \sqsubseteq list_rev_prot_{I_T}$, which we proved in Section 3.3.3.

While the protocols in the prior examples are similar in structure, our approach scales to protocols for which that is not the case. For example, consider the subprotocol relation $prot_1 \sqsubseteq prot_2$, where:

$$\begin{aligned} prot_1 &\triangleq \mu(rec : iProto). ! (x : \mathbb{Z}) \langle x \rangle. ? \langle x + 2 \rangle. rec \\ prot_2 &\triangleq \mu(rec : iProto). ! (x : \mathbb{Z}) \langle x \rangle. ! (y : \mathbb{Z}) \langle y \rangle. ? \langle x + 2 \rangle. ? \langle y + 2 \rangle. rec \end{aligned}$$

Intuitively, these protocols are related, as we can unfold the body of $prot_1$ twice, the body of $prot_2$ once, and swap the second receive over the first send.

The proof is as follows:

$prot_1 \sqsubseteq prot_2 \text{ -*}$	$prot_1 \sqsubseteq prot_2$		
$prot_1 \sqsubseteq prot_2 \text{ -*}$	$? \langle x+2 \rangle. ? \langle y+2 \rangle. prot_1 \sqsubseteq$	$? \langle x+2 \rangle. ? \langle y+2 \rangle. prot_2$	$\sqsubseteq\text{-RECV-MONO}$
$prot_1 \sqsubseteq prot_2 \text{ -*}$	$! y \langle y \rangle. ? \langle x+2 \rangle. ? \langle y+2 \rangle. prot_1 \sqsubseteq$	$! y \langle y \rangle. ? \langle x+2 \rangle. ? \langle y+2 \rangle. prot_2$	$\sqsubseteq\text{-SEND-MONO}'$
$prot_1 \sqsubseteq prot_2 \text{ -*}$	$? \langle x+2 \rangle. ! y \langle y \rangle. ? \langle y+2 \rangle. prot_1 \sqsubseteq$	$! y \langle y \rangle. ? \langle x+2 \rangle. ? \langle y+2 \rangle. prot_2$	$\sqsubseteq\text{-SWAP}'$
$\triangleright(prot_1 \sqsubseteq prot_2) \text{ -*}$	$\triangleright(? \langle x+2 \rangle. ! y \langle y \rangle. ? \langle y+2 \rangle. prot_1 \sqsubseteq$	$! y \langle y \rangle. ? \langle x+2 \rangle. ? \langle y+2 \rangle. prot_2)$	$\triangleright\text{-mono}$
$\triangleright(prot_1 \sqsubseteq prot_2) \text{ -*}$	$! x \langle x \rangle. ? \langle x+2 \rangle. ! y \langle y \rangle. ? \langle y+2 \rangle. prot_1 \sqsubseteq$	$! x \langle x \rangle. ! y \langle y \rangle. ? \langle x+2 \rangle. ? \langle y+2 \rangle. prot_2$	$\sqsubseteq\text{-SEND-MONO}'$
$\triangleright(prot_1 \sqsubseteq prot_2) \text{ -*}$	$prot_1 \sqsubseteq prot_2$	$prot_1 \sqsubseteq prot_2$	$\mu\text{-UNFOLD}$
		$prot_1 \sqsubseteq prot_2$	L\"OB

After we use $\sqsubseteq\text{-SEND-MONO}'$ for the first time, we strip off the later of the induction hypothesis $\triangleright(prot_1 \sqsubseteq prot_2)$. Subsequently, when we use $\sqsubseteq\text{-SEND-MONO}'$ and $\sqsubseteq\text{-RECV-MONO}$, there are no more lateres to strip. We therefore introduce the later implicitly using $\triangleright\text{-INTRO}$ after applying the appropriate monotonicity rule (not shown).

3.4 Manifest Sharing via Locks

Since dependent separation protocols and the connective $c \mapsto prot$ for ownership of protocols are first-class objects of the Actris logic, they can be used like any other logical connective. This means that protocols can be combined with any other mechanism that Actris inherits from Iris. In particular, they can be combined with Iris's generic invariant and ghost state mechanism, and can be used in combination with Iris's abstractions for reasoning about other concurrency connectives like locks, barriers, lock-free data structures, *etc.*

In this section we demonstrate how dependent separation protocols can be combined with lock-based concurrency. This combination allows us to prove functional correctness of programs that make use of the notion of *manifest sharing* [Balzer and Pfenning 2017; Balzer et al. 2019], where channel endpoints are shared between multiple parties. Instead of having to extend Actris, we make use of locks and ghost state that Actris inherits from Iris. We present the basic idea with an introductory example of sharing a channel endpoint between two parties (Section 3.4.1). We then consider a more challenging example of a distributed load-balancing mapper (Section 3.4.2).

3.4.1 Locks and Ghost State

Using the language from Section 3.2.1 it is possible to implement locks using a spin lock, ticket lock, or a more sophisticated implementation. For the purpose of this

$\{R\} \text{new_lock } () \{lk. \text{is_lock } lk \ R\}$	(HT-NEW-LOCK)
$\{\text{is_lock } lk \ R\} \text{acquire } lk \ \{R\}$	(HT-ACQUIRE)
$\{\text{is_lock } lk \ R * R\} \text{release } lk \ \{\text{True}\}$	(HT-RELEASE)
$\text{is_lock } lk \ R \multimap \text{is_lock } lk \ R * \text{is_lock } lk \ R$	(LOCK-DUP)

Figure 3.11: The rules Actris inherits from Iris for locks.

```

prog_lock := let c := start (λc. let lk := new_lock () in
                             fork {acquire lk; send c 21; release lk};
                             acquire lk; send c 21; release lk) in
    recv c + recv c
    
```

Figure 3.12: A sample program that combines locks and channels to achieve manifest sharing.

paper, we abstract over the concrete implementation and assume that we have operations `new_lock`, `acquire` and `release` that satisfy the common separation logic specifications for locks as shown in Figure 3.11.

The `new_lock ()` operation creates a new lock, which can be thought of as a mutex. The operation `acquire lk` will atomically take the lock or block in the case the lock is already taken, and `release lk` releases the lock so that it may be acquired by other threads. The specifications in Figure 3.11 make use of the representation predicate `is_lock lk R`, which expresses that a lock `lk` guards the resources described by the proposition `R`. When creating a new lock one has to give up ownership of `R`, and in turn, obtains the representation predicate `is_lock lk R` (HT-NEW-LOCK). The representation predicate can then be freely duplicated so it can be shared between multiple threads (LOCK-DUP). When entering a critical section using `acquire lk`, a thread gets exclusive ownership of `R` (HT-ACQUIRE), which has to be given up when releasing the lock using `release lk` (HT-RELEASE). The resources `R` that are protected by the lock are therefore invariant in-between any of the critical sections.

To show how locks can be used, consider the program in Figure 3.12, which uses a lock to share a channel endpoint between two threads, that each send 21 to the main thread. The following dependent protocol, where n denotes the number of messages that should be exchanged, captures the expected interaction from the point of view of the main thread:

$$\text{lock_prot} \triangleq \mu(\text{rec} : \mathbb{N} \rightarrow \text{iProto}). \lambda n. \text{if } (n = 0) \text{ then } \mathbf{end} \text{ else } ?\langle 21 \rangle. \text{rec } (n - 1)$$

Since $c \mapsto \overline{\text{lock_prot } n}$ is an exclusive resource, we need a lock to share it between the threads that send 21. For this we will use the following lock invariant:

$$\text{is_lock } lk \ (\exists n. \text{auth}_\gamma n * c \mapsto \overline{\text{lock_prot } n})$$

$$\begin{array}{ll}
\text{True} \Rightarrow \exists \gamma. \text{auth}_\gamma 0 & (\text{AUTH-INIT}) \\
\text{auth}_\gamma n \Rightarrow \text{contrib}_\gamma * \text{auth}_\gamma (1 + n) & (\text{AUTH-ALLOC}) \\
\text{auth}_\gamma (1 + n) * \text{contrib}_\gamma \Rightarrow \text{auth}_\gamma n & (\text{AUTH-DEALLOC}) \\
\text{auth}_\gamma n * \text{contrib}_\gamma \multimap n > 0 & (\text{AUTH-CONTRIB-POS})
\end{array}$$

Figure 3.13: The authoritative contribution ghost theory.

The natural number n is existentially quantified since it changes over time depending on the values that are sent. To tie the number n to the number of contributions made by the threads that share the channel endpoint, we make use of the connectives $\text{auth}_\gamma n$ and contrib_γ , which are defined using Iris’s “ghost theory” mechanism for “user-defined” ghost state [Jung et al. 2015, 2018b].

The $\text{auth}_\gamma n$ fragment can be thought of as an authority that keeps track of the number of ongoing contributions n , while each contrib_γ is a token that witnesses that a contribution is still in progress. These concepts are made precise by the rules in Figure 3.13. The rule AUTH-INIT expresses that an authority $\text{auth}_\gamma 0$ can always be created, which is given some fresh ghost identifier γ . Using the rules AUTH-ALLOC and AUTH-DEALLOC, one can allocate and deallocate tokens contrib_γ as long as the count n of ongoing contributions in $\text{auth}_\gamma n$ is updated accordingly. The rule AUTH-CONTRIB-POS expresses that ownership of a token contrib_γ implies that the count n of $\text{auth}_\gamma n$ must be positive.

Most of the rules in Figure 3.13 involve the logical connective \Rightarrow of a so-called *view shift*. The view shift connective, which Actris inherits from Iris, can be thought of as a “ghost update”, which is made precise by the structural rules VS-CSQ and VS-FRAME rules, that establish the connection between \Rightarrow and the Hoare triples of the logic:

$$\begin{array}{c}
\text{VS-CSQ} \\
\frac{P \Rightarrow P' \quad \{P'\} e \{v. Q'\} \quad \forall v. Q' \Rightarrow Q}{\{P\} e \{v. Q\}}
\end{array}
\qquad
\begin{array}{c}
\text{VS-FRAME} \\
\frac{P \Rightarrow Q}{P * R \Rightarrow Q * R}
\end{array}$$

With the ghost state in place, we can now state suitable specifications for the program. The specification of the top-level program is shown on the right, while the left Hoare triple shows the auxiliary specification of both threads that send the integer 21:

$$\begin{array}{ll}
\{ \text{contrib}_\gamma * \text{is.lock } lk \ (\exists n. \text{auth}_\gamma n * c \multimap \overline{\text{lock.prot } n}) \} & \{ \text{True} \} \\
\text{acquire } lk; \text{ send } c \ 21; \text{ release } lk & \text{prog.lock} \\
\{ \text{True} \} & \{ v. v = 42 \}
\end{array}$$

To establish the initial lock invariant, we use the rules AUTH-INIT and AUTH-ALLOC to create the authority $\text{auth}_\gamma 2$ and two contrib_γ tokens. The contrib_γ tokens play a crucial role in the proofs of the sending threads to establish that the existentially quantified variable n is positive (using AUTH-CONTRIB-POS). Knowing $n > 0$, these

```

par_mapper_worker fv lk c :=
  acquire lk; select c left;
  branch c with
    right ⇒ release lk
  | left ⇒ let x := recv c in release lk;      (* acquire work *)
          let y := fv x in                    (* map it *)
          acquire lk;
          select c right; send c y;          (* send it back *)
          release lk;
          par_mapper_worker fv lk c
end

```

Figure 3.14: A worker of the distributed mapper service.

threads can establish that the protocol $\overline{\text{lock_prot } n}$ has not terminated yet (*i.e.*, is not **end**). This is needed to use the rule HT-SEND to prove the correctness of sending 21, and thereby advancing the protocol from $\overline{\text{lock_prot } n}$ to $\text{lock_prot } (n - 1)$. Subsequently, the sending threads can deallocate the token contrib_γ (using AUTH-DEALLOC) to decrement the n of $\text{auth}_\gamma n$ accordingly to restore the lock invariant.

3.4.2 A Distributed Load-Balancing Mapper

This section demonstrates a more interesting use of manifest sharing. We show how Actris can be used to verify functional correctness of a distributed load-balancing mapper that maps a function f_v over a list. Our distributed mapper consists of one client that distributes the work, and a number of workers that perform the function f_v on individual elements of the list. To enable communication between the client and the workers, we make use of a single channel. One endpoint is used by the client to distribute the work between the workers, while the other endpoint is shared between all workers to request and return work from the client. The implementation of the workers, which can be found in Figure 3.14, consists of a loop over three phases:

1. The worker notifies the client that it wants to perform work (via `select c left`), after which it is then notified (via `branch`) whether there is more work or all elements have been mapped. If there is more work, the worker receives an element x that needs to be mapped. Otherwise, the worker will terminate.
2. The worker maps the function f_v on x .
3. The worker notifies the client that it wants to send back a result (by using `select c right`), and subsequently sends back the result y of mapping f_v on x .

The first and last phases are in a critical section guarded by a lock lk since they involve interaction over a shared channel endpoint. As the sharing behaviour is encapsulated by the worker, we omit the code of the client for brevity's sake.²

²The entire code is present in the accompanied Coq development [Hinrichsen et al. 2021b].

$$\begin{array}{ll}
 \text{True} \Rightarrow \exists \gamma. \text{auth}_\gamma 0 \emptyset & (\text{AUTHM-INIT}) \\
 \text{auth}_\gamma n X \Rightarrow \text{auth}_\gamma (1 + n) X * \text{contrib}_\gamma \emptyset & (\text{AUTHM-ALLOC}) \\
 \text{auth}_\gamma n X * \text{contrib}_\gamma \emptyset \Rightarrow \text{auth}_\gamma (n - 1) X & (\text{AUTHM-DEALLOC}) \\
 \text{auth}_\gamma n X * \text{contrib}_\gamma Y \Rightarrow \text{auth}_\gamma n (X \uplus Z) * \text{contrib}_\gamma (Y \uplus Z) & (\text{AUTHM-ADD}) \\
 Z \subseteq Y * & \\
 \text{auth}_\gamma n X * \text{contrib}_\gamma Y \Rightarrow \text{auth}_\gamma n (X \setminus Z) * \text{contrib}_\gamma (Y \setminus Z) & (\text{AUTHM-REMOVE}) \\
 \text{auth}_\gamma n X * \text{contrib}_\gamma Y * n > 0 * Y \subseteq X & (\text{AUTHM-CONTRIB-AGREE}) \\
 \text{auth}_\gamma 1 X * \text{contrib}_\gamma Y * Y = X & (\text{AUTHM-CONTRIB-AGREE1})
 \end{array}$$

Figure 3.15: The authoritative contribution ghost theory extended with multisets.

A protocol that describes the interaction from the client's point of view is as follows:

$$\begin{array}{l}
 \text{par_mapper_prot } (I_T : T \rightarrow \text{Val} \rightarrow \text{iProp}) (I_U : U \rightarrow \text{Val} \rightarrow \text{iProp}) (f : T \rightarrow \text{List } U) \triangleq \\
 \mu(\text{rec} : \mathbb{N} \rightarrow \text{MultiSet } T \rightarrow \text{iProto}). \lambda n X. \\
 \text{if } n = 0 \text{ then } \mathbf{end} \ \mathbf{else} \\
 (! (x : T) (v : \text{Val}) \langle v \rangle \{I_T x v\}. \text{rec } n (X \uplus \{x\}) \oplus \text{rec } (n - 1) X \\
 \quad \{(n=1) \Rightarrow (X=\emptyset)\} \& \{\text{True}\} \\
 ?(x : T) (\ell : \text{Loc}) \langle \ell \rangle \{x \in X * \ell \xrightarrow{\text{hs}}_{I_U} (f x)\}. \text{rec } n (X \setminus \{x\})
 \end{array}$$

Similarly to `mapper_prot` from Section 3.3.2, the protocol is parameterised by representation predicates I_T and I_U , and a function $f : T \rightarrow \text{List } U$ in Iris/Actris logic, related through the `f_spec` specification. Similar to the protocol `lock_prot` from Section 3.4.1, the protocol `par_mapper_prot` is indexed by the number of remaining workers n . On top of that, it carries a multiset X describing the values currently being processed by all the workers. The multiset X is used to make sure that the returned results are in fact the result of mapping the function f . The condition $(n = 1) \Rightarrow (X = \emptyset)$ on the branching operator ($\&$) expresses that the last worker may only request more work if there are no ongoing jobs.

To accommodate sharing of the channel endpoint between all workers using a lock invariant, we extend the authoritative contribution ghost theory from Section 3.4.1. We do this by adding multisets X and Y to the connectives $\text{auth}_\gamma n X$ and $\text{contrib}_\gamma Y$. These multisets keep track of the values held by the workers. The rules for the ghost theory extended with multisets are shown in Figure 3.15. The rules `AUTHM-INIT`, `AUTHM-ALLOC` and `AUTHM-DEALLOC` are straightforward generalisations of the ones we have seen before. The new rules `AUTHM-ADD` and `AUTHM-REMOVE` determine that the multiset Y of $\text{contrib}_\gamma Y$ can be updated as long as it is done in accordance with the multiset X of $\text{auth}_\gamma n X$. Finally, the `AUTHM-CONTRIB-AGREE` rule expresses that the multiset Y of $\text{contrib}_\gamma Y$ must be a subset of the multiset X of $\text{auth}_\gamma n X$, while the stricter rule `AUTHM-CONTRIB-AGREE1` asserts equality between X and Y when only one contribution remains.

The specifications of `par_mapper_worker` and a top-level client `par_mapper_client` that uses n workers to map f_v over the linked list ℓ are as follows:

$$\left\{ \begin{array}{l} \text{f_spec } I_T \ I_U \ f \ f_v \ * \ \text{contrib}_\gamma \ \emptyset \ * \\ \text{is_lock } lk \ \left(\frac{\exists n \ X. \ \text{auth}_\gamma \ n \ X \ *}{c \mapsto \text{par_mapper_prot } I_T \ I_U \ f \ n \ X} \right) \end{array} \right\} \quad \left\{ \begin{array}{l} \text{f_spec } I_T \ I_U \ f \ f_v \ * \\ 0 < n \ * \ \ell \mapsto_{I_T} \vec{x} \end{array} \right\}$$

$$\text{par_mapper_worker } f_v \ lk \ c \quad \text{par_mapper_client } n \ f_v \ \ell$$

$$\{\text{True}\} \quad \left\{ \begin{array}{l} \exists \vec{y}. \ \vec{y} \equiv_p \text{flatMap } f \ \vec{x} \ * \\ \ell \mapsto_{I_U} \vec{y} \end{array} \right\}$$

The lock invariant and specification of `par_mapper_worker` are similar to those used in the simple example in Section 3.4.1. The specification of `par_mapper_client` $n \ f_v \ \ell$ simply states that the resulting linked list points to a permutation of performing the map at the level of the logic. To specify that, we make use of `flatMap` : $(T \rightarrow \text{List } U) \rightarrow (\text{List } T \rightarrow \text{List } U)$, whose definition is standard.

The proof of the client involves allocating the channel with the mapper protocol `par_mapper_prot`, with the initial number of workers n . Subsequently, we use the rules `AUTHM-INIT` and `AUTHM-ALLOC` to create the authority `authγ n ∅` and n tokens `contribγ ∅`, which allow us to establish the lock invariant and to distribute the tokens among the mappers. The proof of the mapper proceeds as usual. After acquiring the lock, the mapper obtains ownership of the lock invariant. Since the worker owns the token `contribγ ∅`, it knows that the number of remaining workers n is positive, which allows it to conclude that the protocol has not terminated (*i.e.*, is not `end`). After using the rules for channels, the rules `AUTHM-ADD` and `AUTHM-REMOVE` are used to update the authority, which is needed to reestablish the lock invariant so the lock can be released.

3.5 Case Study: Map-Reduce

As a means of demonstrating the use of Actris for verifying more realistic programs, we present a proof of functional correctness of a simple distributed load-balancing implementation of the map-reduce model by Dean and Ghemawat [2004].

Since Actris is not concerned with distributed systems over networks, we consider a version of map-reduce that distributes the work over forked-off threads on a single machine. This means that we do not consider mechanics like handling the failure, restarting, and rescheduling of nodes that a version that operates on a network has to consider.

In order to implement and verify our map-reduce version we make use of the implementation and verification of the fine-grained distributed merge sort algorithm (Section 3.2.8) and the distributed load-balancing mapper (Section 3.4.2). As such, our map-reduce implementation is mostly a suitable client that glues together communication with these services. The purpose of this section is to give a high-level description of the implementation. The actual code and proofs can be found in the accompanied Coq development [Hinrichsen et al. 2021b].

3.5.1 A Functional Specification of Map-Reduce

The purpose of the map-reduce model is to transform an input set of type `List T` into an output set of type `List V` using two functions f (often called “map”) and g (often called “reduce”):

$$f : T \rightarrow \text{List } (K * U) \quad g : (K * \text{List } U) \rightarrow \text{List } V$$

An implementation of map-reduce performs the transformation in three steps:

1. First, the function f is applied to each element of the input set. This results in lists of key/value pairs which are then flattened using a `flatMap` operation (an operation that takes a list of lists and appends all nested lists):

$$\text{flatMap } f \quad : \quad \text{List } T \rightarrow \text{List } (K * U)$$

2. Second, the resulting lists of key/value pairs are grouped together by their key (this step is often called “shuffling”):

$$\text{group} \quad : \quad \text{List } (K * U) \rightarrow \text{List } (K * \text{List } U)$$

3. Finally, the grouped key/value pairs are passed on to the g function, after which the results are flattened to aggregate the results. This again is done using a `flatMap` operation:

$$\text{flatMap } g \quad : \quad \text{List } (K * \text{List } U) \rightarrow \text{List } V$$

The complete functionality of map-reduce is equivalent to applying the following `map_reduce` function on the entire data set:

$$\text{map_reduce} \quad : \quad \text{List } T \rightarrow \text{List } V \quad \triangleq \quad (\text{flatMap } g) \circ \text{group} \circ (\text{flatMap } f)$$

A standard instance of map-reduce is counting word occurrences, where we let $T \triangleq K \triangleq \text{String}$ and $U \triangleq \mathbb{N}$ and $V \triangleq \text{String} * \mathbb{N}$ with:

$$\begin{aligned} f : \text{String} &\rightarrow \text{List } (\text{String} * \mathbb{N}) \triangleq \lambda x. [(x, 1)] \\ g : (\text{String} * \text{List } \mathbb{N}) &\rightarrow \text{List } (\text{String} * \mathbb{N}) \triangleq \lambda (k, \vec{n}). [(k, \sum_{i < |\vec{n}|} \vec{n}_i)] \end{aligned}$$

3.5.2 Implementation of Map-Reduce

The general distributed model of map-reduce is achieved by distributing the phases of mapping, shuffling, and reducing, over a number of worker nodes (*e.g.*, nodes of a cluster or individual CPUs). To perform the computation in a distributed way, there is some work involved in coordinating the jobs over these worker nodes, which is usually done as follows:

1. Split the input data into chunks and delegate these chunks to the mapper nodes, that each apply the “map” function f to their given data in parallel.
2. Collect the complete set of mapped results and “shuffle” them, *i.e.*, group them by key. The grouping is often implemented with a distributed sorting algorithm.

3. Split the shuffled data into chunks and delegate these chunks to the reducer nodes that each apply the “reduce” function g to their given data in parallel.
4. Collect and aggregate the complete set of result of the reducers.

Our variant of the map-reduce model is defined as a function `map_reducev n m fv gv l`, which coordinates the work for performing map-reduce on a linked list l between n mappers performing the “map” function f_v and m workers performing the “reduce” function g_v . To make the implementation more interesting, we prevent storing intermediate values locally by forwarding/returning them immediately as they are available/requested. The global structure is as follows:

1. Start n instances of the load-balancing `par_mapper_worker` from Section 3.4 with the f_v function. Additionally start an instance of `sort_servicefg` from Section 3.2, parameterised by a concrete comparison function on the keys, corresponding to $\lambda(k_1, -) (k_2, -). k_1 < k_2$. Note that the type of keys are restricted to be \mathbb{Z} for brevity’s sake.
2. Perform a loop that handles communication with the mappers. If a mapper requests work, pop a value from the input list. If a mapper returns work, forward it to the sorting service. This process is repeated until all inputs have been mapped and forwarded.
3. Start m instances of the `par_mapper_worker`, parameterised by g_v .
4. Perform a loop that handles communication with the mappers. If a mapper requests work, group elements returned by the sort service. If a mapper returns work, aggregate the returned value in a the linked list. Grouped elements are created by requesting and aggregating elements from the sorter until the key changes.

The aggregated linked list then contains the fully mapped input set upon completion.

3.5.3 Functional Correctness of Map-Reduce

The specification of the map-reduce program is as follows:

$$\left\{ 0 < n * 0 < m * \text{f_spec } I_T \text{ } I_{\mathbb{Z}*U} f f_v * \text{f_spec } I_{\mathbb{Z}*List} U \text{ } I_V g g_v * l \xrightarrow{\text{init}}_{I_T} \vec{x} \right\}$$

$$\text{map_reduce}_v n m f_v g_v l$$

$$\left\{ \exists \vec{z}. \vec{z} \equiv_p \text{map_reduce } f g \vec{x} * l \xrightarrow{\text{init}}_{I_V} \vec{z} \right\}$$

The `f_spec` predicates (as introduced in Section 3.3.2) establish a connection between the functions f and g on the logical level and the functions f_v and g_v in the language. These make use of the various interpretation predicates I_T , $I_{\mathbb{Z}*U}$, $I_{\mathbb{Z}*List} U$, and I_V for the types in question. Lastly, the $l \xrightarrow{\text{init}}_{I_T} \vec{x}$ predicate determines that the input is a linked list of the initial type T . The postcondition asserts that the result \vec{z} is a permutation of the original linked list \vec{x} applied to the functional specification `map_reduce` of map-reduce from Section 3.5.1.

3.6 The Model of Actris

We prove the adequacy theorem of Actris—given a Hoare triple $\{\text{True}\} e \{\phi\}$ that is derivable in Actris, we prove that e cannot get stuck (*i.e.*, safety), and if e terminates, the resulting value v satisfies ϕ (*i.e.*, postcondition validity). To do that, we construct a model of Actris as a shallow embedding in the Iris framework [Jung et al. 2015; Krebbers et al. 2017a; Jung et al. 2016, 2018b]. This means that the type `iProto` of dependent separation protocols, the subprotocol relation $\text{prot}_1 \sqsubseteq \text{prot}_2$, and the connective $c \multimap \text{prot}$ for the channel ownership, are definitions in Iris, and the Actris proof rules are lemmas about these definitions in Iris. Adequacy of Actris is then simply a consequence of Iris’s adequacy theorem.

In this section we describe the relevant aspects of the model of Actris. We model the type `iProto` of dependent separation protocols as the solution of a recursive domain equation, and describe how the operators for dual and composition are defined (Section 3.6.1). We then define the subprotocol relation $\text{prot}_1 \sqsubseteq \text{prot}_2$ and prove its proof rules as lemmas (Section 3.6.2). To connect protocols to the endpoint channel buffers in the semantics we define the *protocol consistency relation*, which ensures that a pair of protocols is consistent with the messages in their associated buffers (Section 3.6.3). On top of the protocol consistency relation, we define the *Actris ghost theory* for dependent separation protocols (Section 3.6.4), which forms the key ingredient for defining the connective $c \multimap \text{prot}$ for channel ownership (Section 3.6.6) that links protocols to the semantics of channels (Section 3.6.5). We then show how adequacy follows from the embedding in Iris (Section 3.6.7). Finally, we show how to solve the recursive domain equation for the type `iProto` of dependent separation protocols (Section 3.6.8).

3.6.1 The Model of Dependent Separation Protocols

To construct a model of dependent separation protocols, we first need to determine what they mean semantically. The challenging part involves the constructors $!\vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}$ and $? \vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}$, whose (higher-order and impredicative) logical variables $\vec{x} : \vec{\tau}$ bind into the communicated value v , the transferred resources P , and the tail protocol prot . We model these constructors as predicates over the communicated value and the tail protocol. To describe the transferred resources P , we model these protocols as Iris predicates (functions to `iProp`) instead of meta-level predicates (functions to `Prop`). This gives rise to the following recursive domain equation:

$$\begin{aligned} \text{action} &::= \text{send} \mid \text{recv} \\ \text{iProto} &\cong 1 + (\text{action} \times (\text{Val} \rightarrow \blacktriangleright \text{iProto} \rightarrow \text{iProp})) \end{aligned}$$

The left part of the sum type (the unit type 1) indicates that the protocol has terminated, while the right part describes a message that is exchanged, expressed as an Iris predicate. Since the recursive occurrence of `iProto` in the predicate appears in negative position, we guard it using Iris’s *type-level later* (\blacktriangleright) operator (whose only constructor is $\text{next} : T \rightarrow \blacktriangleright T$).

The exact way the solution is constructed is detailed in Section 3.6.8. For now, we assume a solution exists, and define the dependent separation protocols constructors as follows:

$$\begin{aligned} \mathbf{end} &\triangleq \mathbf{inj}_1 () \\ !\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot &\triangleq \mathbf{inj}_2 (\mathbf{send}, \lambda w prot'. \exists \vec{x}:\vec{\tau}. (v = w) * P * (prot' = \mathbf{next} prot)) \\ ?\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot &\triangleq \mathbf{inj}_2 (\mathbf{recv}, \lambda w prot'. \exists \vec{x}:\vec{\tau}. (v = w) * P * (prot' = \mathbf{next} prot)) \end{aligned}$$

The definitions of $!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ and $?\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ make use of the (higher-order and impredicative) existential quantifiers of Iris to constrain the actual message w and tail $prot'$ so that they agree with the message v and tail $prot$ prescribed by the protocol.

Recursive Protocols Iris's guarded recursion operator $\mu x. t$ requires the recursion variable x to appear under a *contractive* term construct in t . Hence, to use Iris's recursion operator to construct recursive protocols, it is essential that the protocols $!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ and $?\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ are contractive in the tail $prot$. To show why this is the case, let us first define what it means for a function $f : T \rightarrow U$ to be contractive:

$$\forall x, y. \triangleright(x = y) \Rightarrow f x = f y$$

Examples of contractive functions are the later modality $\triangleright : \mathbf{iProp} \rightarrow \mathbf{iProp}$ and the constructor $\mathbf{next} : T \rightarrow \blacktriangleright T$. The protocols $!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ and $?\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ are defined so that $prot$ appears below a \mathbf{next} , and hence we can prove that they are contractive in $prot$.

Operations With these definitions at hand, the dual $\overline{(_)}$ and append $(_ \cdot _)$ operations are defined using Iris's guarded recursion operator $(\mu x. t)$:

$$\begin{aligned} \overline{\mathbf{send}} &\triangleq \mathbf{recv} \\ \overline{\mathbf{recv}} &\triangleq \mathbf{send} \\ \overline{(_)} &\triangleq \mu rec. \lambda prot. \begin{cases} \mathbf{inj}_1 () & \text{if } prot = \mathbf{inj}_1 () \\ \mathbf{inj}_2 (\overline{a}, \lambda w prot'. \exists prot''. \\ \quad \Phi w (\mathbf{next} prot'') * \\ \quad prot' = \mathbf{next} (rec prot'')) & \text{if } prot = \mathbf{inj}_2 (a, \Phi) \end{cases} \\ (_ \cdot prot_2) &\triangleq \mu rec. \lambda prot_1. \begin{cases} prot_2 & \text{if } prot_1 = \mathbf{inj}_1 () \\ \mathbf{inj}_2 (a, \lambda w prot'. \exists prot''. \\ \quad \Phi w (\mathbf{next} prot'') * \\ \quad prot' = \mathbf{next} (rec prot'')) & \text{if } prot_1 = \mathbf{inj}_2 (a, \Phi) \end{cases} \end{aligned}$$

The base cases of both definitions are as expected. In the recursive cases, we construct a new predicate, given the original predicate Φ . In these new predicates, we quantify over an original tail protocol $prot''$ such that $\Phi w (\mathbf{next} prot'')$ holds, and unify the new tail protocol $prot'$ with the result of the recursive call $rec prot''$.

The equational rules for dual $\overline{(_)}$ and append $(_ \cdot _)$ from Figure 3.1 are proven as lemmas in Iris using LöB induction. This is possible as the recursive call $rec\ prot''$ appears below a `next` constructor—since the `next` constructor is contractive, we can strip-off the later from the induction hypothesis when proving the equality for the tail.

Difference from the Conference Version In the conference version of this paper [Hinrichsen et al. 2020], we described two versions of the recursive domain equation for dependent separation protocols: an “ideal” version (as used in this paper), where `iProto` appears in negative position, and an “alternative” version, where `iProto` appears in positive position. At that time, we were unable to construct a solution of the “ideal” version, so we used the “alternative” version. In Section 3.6.8 we show how we are now able to solve the “ideal” version.

In the conference version of this paper, the proposition P appeared under a later modality in the definitions of the protocols $!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ and $?x:\vec{\tau}\langle v\rangle\{P\}.prot$, making these protocols contractive in P . This choice was motivated by the ability to construct recursive protocols like $\mu rec.!(c:\text{Chan})\langle c\rangle\{c \mapsto prot\}.prot'$, where the payload refers to the recursion variable rec . In the current version (without the later modality) we can still construct such protocols, because $c \mapsto prot$ is contractive in $prot$. We removed the later modality because it is incompatible with the rules $\sqsubseteq\text{-SEND-OUT}$ and $\sqsubseteq\text{-RECV-OUT}$ for subprotocols.

3.6.2 The Model of the Subprotocol Relation

We now model the subprotocol relation $prot_1 \sqsubseteq prot_2$ from Section 3.3. For legibility, we present it in the style of an inference system through its constructors, whereas it is formally defined using Iris’s guarded recursion operator $(\mu x.t)$:

$$\text{inj}_1() \sqsubseteq \text{inj}_1()$$

$$\frac{\forall v, prot_2. \Phi_2 v (\text{next } prot_2) \multimap \exists prot_1. \Phi_1 v (\text{next } prot_1) * \triangleright(prot_1 \sqsubseteq prot_2)}{\text{inj}_2(\text{send}, \Phi_1) \sqsubseteq \text{inj}_2(\text{send}, \Phi_2)}$$

$$\frac{\forall v, prot_1. \Phi_1 v (\text{next } prot_1) \multimap \exists prot_2. \Phi_2 v (\text{next } prot_2) * \triangleright(prot_1 \sqsubseteq prot_2)}{\text{inj}_2(\text{recv}, \Phi_1) \sqsubseteq \text{inj}_2(\text{recv}, \Phi_2)}$$

$$\frac{\forall v_1, v_2, prot_1, prot_2. (\Phi_1 v_1 (\text{next } prot_1) * \Phi_2 v_2 (\text{next } prot_2)) \multimap \exists prot. \triangleright(prot_1 \sqsubseteq !\langle v_2 \rangle. prot) * \triangleright(?\langle v_1 \rangle. prot \sqsubseteq prot_2)}{\text{inj}_2(\text{recv}, \Phi_1) \sqsubseteq \text{inj}_2(\text{send}, \Phi_2)}$$

To be a well-formed guarded recursion definition, every recursive occurrence of \sqsubseteq is guarded by later modality (\triangleright) . Aside from the later being required for well-formedness, these later make it possible to reason about the subprotocol relation using LöB induction; both to prove the subprotocol rules from Figure 3.8 as lemmas, and for Actris users to reason about recursive protocols as shown in Section 3.3.4. The relation is defined in a syntax-directed fashion (*i.e.*, there are no overlapping rules), and therefore all constructors need to be defined so that they are closed under monotonicity and transitivity.

The first constructor states that terminating protocols ($\mathbf{end} \triangleq \mathbf{inj}_1 ()$) are related. The other constructors concern the protocols $!\vec{x} : \vec{\tau} \langle v \rangle \{P\}.prot$ and $? \vec{x} : \vec{\tau} \langle v \rangle \{P\}.prot$, which are modelled as $\mathbf{inj}_2(\mathbf{send}, \Phi)$ and $\mathbf{inj}_2(\mathbf{recv}, \Phi)$, where $\Phi : \mathbf{Val} \rightarrow \mathbf{iProto} \rightarrow \mathbf{iProp}$ is a predicate over the communicated value and tail protocol. While the actual constructors are somewhat intimidating because they are defined in terms of these predicates in the model, they essentially correspond to the following high-level versions:

$$\frac{\forall \vec{y} : \vec{\sigma}. P_2 * \exists \vec{x} : \vec{\tau}. (v_1 = v_2) * P_1 * \triangleright (prot_1 \sqsubseteq prot_2)}{!\vec{x} : \vec{\tau} \langle v_1 \rangle \{P_1\}. prot_1 \sqsubseteq ! \vec{y} : \vec{\sigma} \langle v_2 \rangle \{P_2\}. prot_2}$$

$$\frac{\forall \vec{x} : \vec{\tau}. P_1 * \exists \vec{y} : \vec{\sigma}. (v_1 = v_2) * P_2 * \triangleright (prot_1 \sqsubseteq prot_2)}{? \vec{x} : \vec{\tau} \langle v_1 \rangle \{P_1\}. prot_1 \sqsubseteq ? \vec{y} : \vec{\sigma} \langle v_2 \rangle \{P_2\}. prot_2}$$

$$\frac{\forall \vec{x} : \vec{\tau}, \vec{y} : \vec{\sigma}. (P_1 * P_2) * \exists prot. \triangleright (prot_1 \sqsubseteq ! \langle v_2 \rangle. prot) * \triangleright (? \langle v_1 \rangle. prot \sqsubseteq prot_2)}{? \vec{x} : \vec{\tau} \langle v_1 \rangle \{P_1\}. prot_1 \sqsubseteq ! \vec{y} : \vec{\sigma} \langle v_2 \rangle \{P_2\}. prot_2}$$

To obtain syntax-directed rules, the first rule combines \sqsubseteq -SEND-OUT, \sqsubseteq -SEND-IN, and \sqsubseteq -SEND-MONO, and dually, the second rule combines \sqsubseteq -RECV-OUT, \sqsubseteq -RECV-IN, and \sqsubseteq -RECV-MONO. The third rule combines \sqsubseteq -RECV-OUT, \sqsubseteq -SEND-OUT and \sqsubseteq -SWAP and bakes in transitivity, instead of asserting that $prot_1$ and $prot_2$ are equal to $!\langle v_2 \rangle. prot$ and $? \langle v_1 \rangle. prot$, respectively.

The rules from the beginning of this section are defined by generalising the high-level rules to arbitrary predicates. For example, the rule $\mathbf{inj}_2(\mathbf{send}, \Phi_1) \sqsubseteq \mathbf{inj}_2(\mathbf{send}, \Phi_2)$ requires that for any value v and tail protocol $prot_2$ that satisfy the predicate Φ_2 , there is a stronger tail protocol $prot_1$ (i.e., where $prot_1 \sqsubseteq prot_2$), so that the same value v and stronger tail protocol $prot_1$ that satisfy the predicate Φ_1 .

The rules in Figure 3.8 on page 62 are proven as lemmas. Those for logical variable and resource manipulation (\sqsubseteq -SEND-OUT, \sqsubseteq -SEND-IN, \sqsubseteq -RECV-OUT and \sqsubseteq -RECV-IN) monotonicity (\sqsubseteq -SEND-MONO and \sqsubseteq -RECV-MONO), and swapping (\sqsubseteq -SWAP) follow almost immediately from the definition, whereas those for reflexivity (\sqsubseteq -REFL), transitivity (\sqsubseteq -TRANS), and the dual and append operator (\sqsubseteq -DUAL and \sqsubseteq -APPEND) are proven using LÖB induction.

3.6.3 Protocol Consistency

To connect dependent separation protocols to the semantics of channels in Section 3.6.6, we define the *protocol consistency relation* $\mathbf{prot_consistent} \vec{v}_1 \vec{v}_2 prot_1 prot_2$, which expresses that protocols $prot_1$ and $prot_2$ are *consistent* w.r.t. channel buffers containing values \vec{v}_1 and \vec{v}_2 . The consistency relation is defined as:

$$\mathbf{prot_consistent} \vec{v}_1 \vec{v}_2 prot_1 prot_2 \triangleq \exists prot. (\overline{? \langle \vec{v}_{2.1} \rangle. \dots ? \langle \vec{v}_{2.|\vec{v}_2|} \rangle. prot} \sqsubseteq prot_1) * (? \langle \vec{v}_{1.1} \rangle. \dots ? \langle \vec{v}_{1.|\vec{v}_1|} \rangle. \overline{prot} \sqsubseteq prot_2)$$

Intuitively, $\mathbf{prot_consistent} \vec{v}_1 \vec{v}_2 prot_1 prot_2$ ensures that for all messages \vec{v}_1 in transit from the endpoint described by $prot_1$ to the endpoint described by $prot_2$, the protocol

$prot_2$ is expecting to receive these message in order (and *vice versa* for \vec{v}_2), after which the remaining protocols $prot$ and \overline{prot} are dual. To account for weakening we close the relation under subprotocols (by using \sqsubseteq instead of equality), which additionally captures ownership of the resources associated with the messages \vec{v}_1 and \vec{v}_2 .

Closure under the subprotocol relation gives us that $\text{prot_consistent } \vec{v}_1 \vec{v}_2 prot_1 prot_2$ and $prot_1 \sqsubseteq prot'_1$ implies $\text{prot_consistent } \vec{v}_1 \vec{v}_2 prot'_1 prot_2$, and ensures that the consistency relation enjoys the following rules corresponding to creating a channel, sending a message, and receiving a message:

$$\begin{aligned} & \text{prot_consistent } [] [] prot \overline{prot} \\ & \text{prot_consistent } \vec{v}_1 \vec{v}_2 (!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot_1) prot_2 * P[\vec{t}/\vec{x}] \multimap \\ & \quad \triangleright^{\vec{v}_2} (\text{prot_consistent } (\vec{v}_1 \cdot [v[\vec{t}/\vec{x}]]) \vec{v}_2 prot_1 prot_2) \\ & \text{prot_consistent } \vec{v}_1 ([w] \cdot \vec{v}_2) (? \vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot_1) prot_2 \multimap \\ & \quad \exists \vec{y}. (w = v[\vec{y}/\vec{x}]) * P[\vec{y}/\vec{x}] * \triangleright (\text{prot_consistent } \vec{v}_1 \vec{v}_2 prot_1 prot_2) \end{aligned}$$

The first rule states that dual protocols are consistent w.r.t. a pair of empty buffers. The second rule states that a protocol $!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot_1$ can be advanced to $prot_1$ by giving up ownership of $P[\vec{t}/\vec{x}]$ and enqueueing the value $v[\vec{t}/\vec{x}]$ in the buffer \vec{v}_1 . Dually, the third rule states that given a protocol $? \vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot_1$ and a buffer that contains value w as its head, we learn that w is equal to $v[\vec{y}/\vec{x}]$, and that we can obtain ownership of $P[\vec{y}/\vec{x}]$ by advancing the protocol to $prot_1$ and dequeuing the value w from the buffer. Since the relation is symmetric, *i.e.*, if $\text{prot_consistent } \vec{v}_1 \vec{v}_2 prot_1 prot_2$ then $\text{prot_consistent } \vec{v}_2 \vec{v}_1 prot_2 prot_1$, we obtain similar rules for the protocol $prot_2$ on the right-hand side.

The last two rules are proved by inversion on the definition of the subprotocol relation (\sqsubseteq). Since the subprotocol relation (\sqsubseteq) is defined using guarded recursion, we obtain a later modality (\triangleright) for each inversion. To prove the first rule, we need to perform a number of inversions equal to the size of the buffer \vec{v}_2 , whereas for the second rule we need to perform just a single inversion. The later modalities will be eliminated through physical program steps in the semantics of channels in Section 3.6.5.

3.6.4 The Actris Ghost Theory

To provide a general interface for adopting Actris's reasoning principles for arbitrary message-passing languages, we employ a standard ghost theory approach to compartmentalise channel ownership.

We use an approach similar to the ghost theory for contributions that we used in Section 3.4. The authority $\text{prot_ctx } (\gamma_1, \gamma_2) \vec{v}_1 \vec{v}_2$ governs the global state of the buffers \vec{v}_1 and \vec{v}_2 . The tokens $\text{prot_own}_l (\gamma_1, \gamma_2) prot_l$ and $\text{prot_own}_r (\gamma_1, \gamma_2) prot_r$ provide local views of that state, being that the protocols $prot_l$ and $prot_r$ are consistent with the buffers. As we will see in Section 3.6.6, the authority can then be shared through a lock, while the tokens can be distributed to individual threads. The ghost connectives are identified by the shared ghost identifiers γ_1 and γ_2 for the protocols $prot_l$ and $prot_r$, respectively.

$$\begin{aligned}
 \text{True} &\Rightarrow \exists \gamma. (\gamma \mapsto_{\bullet} \text{prot}) * (\gamma \mapsto_{\circ} \text{prot}) && \text{(HO-GHOST-ALLOC)} \\
 (\gamma \mapsto_{\bullet} \text{prot}) * (\gamma \mapsto_{\circ} \text{prot}') &\Rightarrow \triangleright(\text{prot} = \text{prot}') && \text{(HO-GHOST-AGREE)} \\
 (\gamma \mapsto_{\bullet} \text{prot}) * (\gamma \mapsto_{\circ} \text{prot}') &\Rightarrow (\gamma \mapsto_{\bullet} \text{prot}'') * (\gamma \mapsto_{\circ} \text{prot}'') && \text{(HO-GHOST-UPDATE)}
 \end{aligned}$$

Figure 3.16: Higher-order ghost variables in Iris.

To define the connectives of the Actris ghost theory we use Iris’s higher-order ghost variables, whose rules are shown in Figure 3.16. Higher-order ghost variables come in pairs $\gamma \mapsto_{\bullet} \text{prot}$ and $\gamma \mapsto_{\circ} \text{prot}$, which always hold the same protocol prot . They can be allocated together (HO-GHOST-ALLOC), are always required to hold the same protocol (HO-GHOST-AGREE), and can only be updated together (HO-GHOST-UPDATE). The subtle part of the higher-order ghost variables is that they involve ownership of a protocol of type `iProto`, which is defined in terms of Iris propositions `iProp`. Due to the dependency on `iProp`, which is covered in detail in Section 3.6.8, the rule HO-GHOST-AGREE only gives the equality between the protocols under a later modality (\triangleright). The Actris ghost theory connectives are then defined as:

$$\begin{aligned}
 \text{prot_ctx } (\gamma_1, \gamma_2) \vec{v}_1 \vec{v}_2 &\triangleq \exists \text{prot}_1, \text{prot}_2. \gamma_1 \mapsto_{\bullet} \text{prot}_1 * \gamma_2 \mapsto_{\bullet} \text{prot}_2 * \\
 &\quad \triangleright \text{prot_consistent } \vec{v}_1 \vec{v}_2 \text{ prot}_1 \text{ prot}_2 \\
 \text{prot_own}_l (\gamma_1, \gamma_2) \text{ prot}_l &\triangleq \exists \text{prot}'_l. \gamma_1 \mapsto_{\circ} \text{prot}'_l * \triangleright(\text{prot}'_l \sqsubseteq \text{prot}_l) \\
 \text{prot_own}_r (\gamma_1, \gamma_2) \text{ prot}_r &\triangleq \exists \text{prot}'_r. \gamma_2 \mapsto_{\circ} \text{prot}'_r * \triangleright(\text{prot}'_r \sqsubseteq \text{prot}_r)
 \end{aligned}$$

The authority $\text{prot_ctx } (\gamma_1, \gamma_2) \vec{v}_1 \vec{v}_2$ asserts that the buffers \vec{v}_1 and \vec{v}_2 are consistent with respect to the protocols prot_1 and prot_2 (via $\text{prot_consistent } \vec{v}_1 \vec{v}_2 \text{ prot}_1 \text{ prot}_2$). It also asserts the higher-order authoritative ownership $\gamma_1 \mapsto_{\bullet} \text{prot}_1$ and $\gamma_2 \mapsto_{\bullet} \text{prot}_2$ of both protocols. The tokens assert higher-order fragmented ownership $\gamma_1 \mapsto_{\circ} \text{prot}'_l$ and $\gamma_2 \mapsto_{\circ} \text{prot}'_r$ of protocols prot'_l and prot'_r that are weaker than the protocol arguments prot_l and prot_r (via $\text{prot}'_l \sqsubseteq \text{prot}_l$ and $\text{prot}'_r \sqsubseteq \text{prot}_r$). Explicit weakening under the subprotocol relation may seem redundant, as weakening is already accounted for in prot_consistent . However, it allows us to weaken the protocols of the tokens without the presence of the authority as shown by the rules $\text{PROTO-}\sqsubseteq\text{-L}$ and $\text{PROTO-}\sqsubseteq\text{-R}$ in Figure 3.17. The later modality (\triangleright) makes sure that $\text{prot_own}_l (\gamma_1, \gamma_2) \text{ prot}$ and $\text{prot_own}_r (\gamma_1, \gamma_2) \text{ prot}$ are contractive in prot . For readability we condense the ghost state identifiers (γ_1, γ_2) into a single identifier γ from now on.

With these definitions at hand, we prove the rules of the ghost theory presented in Figure 3.17. The rule PROTO-ALLOC corresponds to allocation of a buffer pair, the rules PROTO-SEND-L and PROTO-SEND-R correspond to sending a message, and the rules PROTO-RECV-L and PROTO-RECV-R correspond to receiving a message. These are proved through a combination of the rules for higher-order ghost state from Figure 3.16, and the rules for the protocol consistency relation prot_consistent from Section 3.6.3.

$$\begin{array}{l}
\text{True} \Rightarrow \exists \gamma. \text{prot_ctx} \gamma \ [] \ [] * \text{prot_own}_l \gamma \ \text{prot} * \text{prot_own}_r \gamma \ \overline{\text{prot}} \quad (\text{PROTO-ALLOC}) \\
\text{prot_ctx} \gamma \ \vec{v}_1 \ \vec{v}_2 * \text{prot_own}_l \gamma \ (!\ \vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}) * P[\vec{t}/\vec{x}] \Rightarrow \quad (\text{PROTO-SEND-L}) \\
\triangleright^{|\vec{v}_2|} (\text{prot_ctx} \gamma \ (\vec{v}_1 \cdot [v[\vec{t}/\vec{x}]]) \ \vec{v}_2) * \text{prot_own}_l \gamma \ (\text{prot}[\vec{t}/\vec{x}]) \\
\text{prot_ctx} \gamma \ \vec{v}_1 \ \vec{v}_2 * \text{prot_own}_r \gamma \ (!\ \vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}) * P[\vec{t}/\vec{x}] \Rightarrow \quad (\text{PROTO-SEND-R}) \\
\triangleright^{|\vec{v}_1|} ((\text{prot_ctx} \gamma \ \vec{v}_1 \ (\vec{v}_2 \cdot [v[\vec{t}/\vec{x}]]) * \text{prot_own}_r \gamma \ (\text{prot}[\vec{t}/\vec{x}])) \\
\text{prot_ctx} \gamma \ \vec{v}_1 \ ([w] \cdot \vec{v}_2) * \text{prot_own}_l \gamma \ (? \vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}) \Rightarrow \quad (\text{PROTO-RECV-L}) \\
\triangleright \exists \vec{y}. (w = v[\vec{y}/\vec{x}]) * P[\vec{y}/\vec{x}] * \text{prot_ctx} \gamma \ \vec{v}_1 \ \vec{v}_2 * \text{prot_own}_l \gamma \ \text{prot} \\
\text{prot_ctx} \gamma \ ([w] \cdot \vec{v}_1) \ \vec{v}_2 * \text{prot_own}_r \gamma \ (? \vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}) \Rightarrow \quad (\text{PROTO-RECV-R}) \\
\triangleright \exists \vec{y}. (w = v[\vec{y}/\vec{x}]) * P[\vec{y}/\vec{x}] * \text{prot_ctx} \gamma \ \vec{v}_1 \ \vec{v}_2 * \text{prot_own}_r \gamma \ \text{prot} \\
\text{prot_own}_l \gamma \ \text{prot} * \text{prot} \sqsubseteq \text{prot}' \multimap \text{prot_own}_l \gamma \ \text{prot}' \quad (\text{PROTO-}\sqsubseteq\text{-L}) \\
\text{prot_own}_r \gamma \ \text{prot} * \text{prot} \sqsubseteq \text{prot}' \multimap \text{prot_own}_r \gamma \ \text{prot}' \quad (\text{PROTO-}\sqsubseteq\text{-R})
\end{array}$$

Figure 3.17: The Actris ghost theory.

3.6.5 Semantics of Channels

Since Iris is parametric in the programming language that is used, there are various approaches to extend Iris with support for channels:

- Instantiate Iris with a language that has native support for channels. This approach was carried out in the original Iris paper [Jung et al. 2015] and by Tassarotti et al. [2017].
- Instantiate Iris with a language that has low-level concurrency primitives, but no native support for channels, and implement channels as a library in that language. This approach was carried out by Bizjak et al. [2019] for a lock-free implementation of channels.

In this paper we take the second approach. We use HeapLang, the default language shipped with Iris, and implement bidirectional channels using a pair of mutable linked lists protected by a lock. Although this implementation is not efficient, contrary to *e.g.*, the implementation by Bizjak et al. [2019], it has the benefit that it gives a clear declarative semantics that corresponds exactly to the intuitive semantics of channels described in Section 3.2.1.

Our implementation of bidirectional channels in HeapLang is displayed in Figure 3.18. New channels are created by the `new_chan` function, which allocates two empty mutable linked lists l and r using `lnil` (`()`), along with a lock lk using `new_lock` (`()`), and returns the tuples (l, r, lk) and (r, l, lk) , where the order of the linked lists l and r determines the side of the endpoints. We refer to the list in the left position as the endpoint's own buffer, and the list in the right position as the other endpoint's buffer.

```

new_chan () := let (l, r, lk) := (lnil (), lnil (), new_lock ()) in
                ((l, r, lk), (r, l, lk))

send c v := let (l, r, lk) := c in
             acquire lk;
             lsnoc l v; skipN |r|;
             release lk

try_recv c := let (l, r, lk) := c in
              acquire lk;
              let ret := (if (lisnil r) then (inj1 ()) else (inj2 (lpop l))) in
              release lk; ret

recv c := match (try_recv c) with
          inj1 () ⇒ recv c
        | inj2 v ⇒ v
        end
    
```

Figure 3.18: Implementation of bidirectional channels in HeapLang.

Values are sent over a channel endpoint (l, r, lk) using the `send` function. This function operates in an atomic fashion by first acquiring the lock via `acquire lk`, thereby entering the critical section, after which the value is enqueued (*i.e.*, appended to the end) of the endpoint’s own buffer using `lsnoc l v`. The `skipN |r|` instruction is a no-op that is inserted to aid the proof. We come back to the reason why this instruction is needed in Section 3.6.6.

Values are received over a channel endpoint (l, r, lk) using the `send` function, which performs a loop that repeatedly calls the helper function `try_recv`. This helper function attempts to receive a value atomically, and fails if there is no value in the other endpoint’s buffer. The function `try_recv` acquires the lock with `acquire lk`, and then checks whether the other endpoint’s buffer is empty using `lisnil r`. If it is empty, nothing is returned (*i.e.*, `inj1 ()`), while otherwise the value is dequeued and returned (*i.e.*, `inj2 (lpop l)`).

3.6.6 The Model of Channel Ownership

To link the physical contents of the bidirectional channel c to the Actris ghost theory we define the channel ownership connective as follows:

$$c \mapsto \text{prot} \triangleq \exists \gamma, l, r, lk. \left(\begin{array}{l} (c = (l, r, lk) * \text{prot_own}_l \gamma \text{ prot}) \vee \\ (c = (r, l, lk) * \text{prot_own}_r \gamma \text{ prot}) \end{array} \right) * \\
 \text{is_lock } lk \ (\exists \vec{v}_1 \vec{v}_2. l \xrightarrow{\text{list}} \vec{v}_1 * r \xrightarrow{\text{list}} \vec{v}_2 * \text{prot_ctx } \gamma \vec{v}_1 \vec{v}_2)$$

The predicate states that the referenced channel endpoint c is either the left (l, r, lk) or the right (r, l, lk) side of a channel, and that we have exclusive ownership of the ghost token `prot_ownl γ prot` or `prot_ownr γ prot` for the corresponding side. Iris’s lock

representation predicate `is_lock` (previously presented in Section 3.4) is used to make sharing of the buffers possible. The lock invariant is governed by lock lk , and carries the ownership $l \mapsto \bar{v}_1$ and $r \mapsto \bar{v}_2$ of the mutable linked lists containing the channel buffers, as well as `prot_ctx γ \bar{v}_1 \bar{v}_2` , which asserts protocol consistency of the buffers with respect to the protocols.

With the definition of the channel endpoint ownership along with the ghost theory and lock rules we then prove the channel rules HT-NEW, HT-SEND and HT-RCV from Figure 3.1. The proofs are carried out through symbolic execution to the point where the critical section is entered, after which the rules of the Actris ghost theory (Figure 3.17) are used to allocate or update the ghost state appropriately so that it matches the physical channel buffers.

The Need for Skip Instructions The rules `PROTO-SEND-L` and `PROTO-SEND-R` from Figure 3.17 contain a number of later modalities (\triangleright) proportional to the other endpoint’s buffer in their premise. As explained in Section 3.6.3 these later modalities are the consequence of having to perform a number of inversions on the subprotocol relation, which is defined using guarded recursion, and thus contains a later modality for each recursive unfolding.

To eliminate these later modalities, we instrument the code of the `send` function with the `skipN $|r|$` instruction, which performs a number of skips equal to the size of the other endpoint’s buffer r . The `skipN` instruction has the following specification:

$$\{\triangleright^n P\} \text{skipN } n \{P\}$$

Instrumentation with skip instructions appears more often in work on step-indexing, see *e.g.*, [Svendsen et al. 2016; Giarrusso et al. 2020]. It is needed due to the fact that current step-indexed logics like Iris unify physical/program steps and logical steps, *i.e.*, for each physical/program step at most one later can be eliminated from the hypotheses. Svendsen et al. [2016] proposed a more liberal version of step-indexed, called *transfinite step-indexing*, to avoid this problem. However, transfinite step-indexing is not available in Iris.

3.6.7 Adequacy of Actris

Having constructed the model of Actris in Iris, we now obtain the following main result:

Theorem 1 (Adequacy of Actris) *Let ϕ be a first-order predicate over values and suppose the Hoare triple $\{\text{True}\} e \{\phi\}$ is derivable in Actris, then:*

- **(Safety):** *The program e will not get stuck.*
- **(Postcondition validity):** *If the main thread of e terminates with a value v , then the postcondition ϕv holds at the meta-level.*

Since Actris is an internal logic embedded in Iris, the proof is an immediate consequence of Iris’s adequacy theorem [Krebbbers et al. 2017a; Jung et al. 2018b]. Finally, note that safety implies *session fidelity*—any message that is received has in fact been sent.

3.6.8 Solving the Recursive Domain Equation for Protocols

Recall the recursive domain equation for dependent separation protocols from Section 3.6.1:

$$\text{iProto} \cong 1 + (\text{action} \times (\text{Val} \rightarrow \blacktriangleright \text{iProto} \rightarrow \text{iProp}))$$

This recursive domain equation shows that iProto depends on the type iProp of Iris propositions. To use types that depend on iProp as part of higher-order ghost state in Iris, such types need to be bi-functorial in iProp . Hence, this means that to construct iProto , in a way that it can be used in combination with the higher-order ghost variables in Figure 3.16, we need to solve the following recursive domain equation:

$$\text{iProto}(X^-, X^+) \cong 1 + (\text{action} \times (\text{Val} \rightarrow \blacktriangleright \text{iProto}(X^+, X^-) \rightarrow X^+))$$

Since the recursive occurrence of iProto appears in negative position, the polarity needs to be inverted for iProto to be bi-functorial.

The version of Iris’s recursive domain solver based on [America and Rutten 1989; Birkedal et al. 2010] as mechanised in Iris’s Coq development is not readily able to construct a solution of $\text{iProto}(X^-, X^+)$. Concretely, the solver can only construct solutions of non-parameterised recursive domain equations. While a general construction for solving such recursive domain equations exists [Birkedal et al. 2012, § 7], that construction has not been mechanised in Coq. We circumvent this shortcoming by solving the following recursive domain equation instead, in which we unfold the recursion once by hand:

$$\begin{aligned} \text{iProto}_2(X^-, X^+) &\cong \\ &1 + (\text{action} \times (\text{Val} \rightarrow \blacktriangleright (1 + (\text{action} \times (\text{Val} \rightarrow \blacktriangleright \text{iProto}_2(X^-, X^+) \rightarrow X^-)))) \rightarrow X^+) \end{aligned}$$

Here, the polarity in the recursive occurrence is fixed, allowing us to solve the equation $\text{iProto}_2(X^-, X^+)$ using Iris’s existing recursive domain equation solver. This is sufficient because a solution of $\text{iProto}_2(X^-, X^+)$ is isomorphic to a solution of $\text{iProto}(X^-, X^+)$.

3.7 Coq Mechanisation

The definition of the Actris logic, its model, and the proofs of all examples in this paper have been fully mechanised using the Coq proof assistant [Coq Development Team 2021]. In this section we will elaborate on the mechanisation effort (Section 3.7.1), and go through the full proof of a message-passing program (Section 3.7.2) and a sub-protocol relation (Section 3.7.3) showcasing the tactics for Actris. During the section we display proofs and proof states taken directly from the Coq mechanisation, which differs in notation from the paper as shown in Figure 3.19.

3.7.1 Mechanisation Effort

The mechanisation of Actris is built on top of the mechanisation of Iris [Krebbbers et al. 2017a; Jung et al. 2016, 2018b]. To carry out proofs in separation logic, we use the MoSeL Proof Mode (formerly Iris Proof Mode) [Krebbbers et al. 2017b, 2018],

	Paper	Coq Mechanisation
Send	$!x_1 \dots x_n \langle v \rangle \{P\}. prot$	<code><! x_1 .. x_n> MSG v {{ P }}; prot</code>
Receive	$?x_1 \dots x_n \langle v \rangle \{P\}. prot$	<code><? x_1 .. x_n> MSG v {{ P }}; prot</code>
End	end	<code>END</code>
Dual	$prot$	<code>iProto_dual prot</code>
Literals	$()$, 5 , $true$	<code>#()</code> , <code>#5</code> , <code>#true</code>
Logical variables	$x, y, z, -$	<code>"x"</code> , <code>"y"</code> , <code>"z"</code> , <code><></code>
Types	$1, \mathbb{N}, \mathbb{Z}$	<code>()</code> , <code>nat</code> , <code>Z</code>

Figure 3.19: The difference in syntax between the paper and Coq source code.

which provides an embedded proof assistant for separation logic in Coq. Building Actris on top of the Iris and MoSeL framework in Coq has a number of tangible advantages:

- By defining channels on top of HeapLang, the default concurrent language shipped with Iris, we do not have to define a full programming language semantics, and can reuse all of the program libraries and Coq machinery, including the tactics for symbolic execution of non message-passing programs.
- Since Actris is mechanised as an Iris library that provides support for the `iProto` type, the subprotocol relation $prot_1 \sqsubseteq prot_2$, the $c \mapsto prot$ connective, the various operations on protocols, and the proof rules as lemmas, we get all of the features of Iris for free, such as the ghost state mechanisms for reasoning about concurrency.
- When proving the Actris proof rules, we can make use of the MoSeL Proof Mode to carry out proofs directly using separation logic, thus reasoning at a high-level of abstraction.
- We can make use of the extendable nature of the MoSeL Proof Mode to define custom tactics for symbolic execution of message-passing programs.

These advantages made it possible to mechanise Actris, along with the examples of the paper, with a small Coq development of a total size of about 5000 lines of code (comments and whitespace included). The line count of the different components are shown in Figure 3.20.

3.7.2 Tactic Support for Session Type-Based Reasoning

To carry out interactive Actris's proof using symbolic execution, we follow the methodology described in the original Iris Proof Mode paper [Krebbers et al. 2017b], which in particular means that the logic in Coq is presented in weakest precondition style rather than using Hoare triples. For handling `send` or `recv` we define the following tactics:

```
wp_send (t1 .. tn) with "[H1 .. Hn]" and wp_recv (y1 .. yn) as "H".
```

These tactics roughly perform the following actions:

Component	Sections	~LOC
The Actris model	Section 3.6.1–Section 3.6.4	1500
Channel implementation and proof rules	Section 3.6.5–Section 3.6.6	350
Tactics for symbolic execution	Section 3.7.2	500
Utilities (linked lists, permutations, <i>etc.</i>)	n.a.	450
Authoritative contribution ghost theory	Section 3.4	150
Recursive domain equation theory solver	Section 3.6.8	100
Examples:		
• Basic examples	Section 3.1 and Section 3.4.1	400
• Coarse-grained distributed merge sort	Section 3.2.3–Section 3.2.7	250
• Fine-grained distributed merge sort	Section 3.2.8	300
• Mapper with swapping	Section 3.3.2	400
• List reversal	Section 3.3.3	100
• Distributed mapper	Section 3.4.2	200
• Distributed map-reduce	Section 3.5	300
Total		5000

Figure 3.20: Overview of the Actris Coq mechanisation.

- Find a `send` or `recv` in evaluation position of the program under consideration.
- Find a corresponding $c \mapsto \text{prot}$ hypothesis in the separation logic context.
- Normalise the protocol prot using the rules for duals, composition, recursion, and swapping so it has a $!\vec{x}:\vec{\tau}\langle v \rangle\{P\}.\text{prot}$ or $? \vec{x}:\vec{\tau}\langle v \rangle\{P\}.\text{prot}$ construct in its head position.
- In case of `wp_send`, instantiate the variables $\vec{x}:\vec{\tau}$ using the terms $(\tau_1 \dots \tau_n)$, and create a goal for the proposition P with the hypotheses $[\#1 \dots \#n]$. Hypotheses prefixed with $\$$ will automatically be consumed to resolve a subgoal of P if possible. In case the terms $(\tau_1 \dots \tau_n)$ are omitted, an attempt is made to determine these using unification.
- In case of `wp_recv`, introduce the variables $\vec{x}:\vec{\tau}$ into the context by naming them $(y_1 \dots y_n)$, and create a hypothesis $\#$ for P .

The implementation of these tactics follows the approach by [Krebbers et al. 2017b]. The protocol normalisation is implemented via logic programming with type classes.

As an example we will go through a proof of the following program:

```

prog_ref_swap_loop :=  $\lambda()$ . let  $c := \text{start}$  (rec go  $c'$  := let  $\ell := \text{recv } c'$  in
                                                     $\ell \leftarrow !\ell + 2;$ 
                                                    send  $c' ()$ ; go  $c'$ ) in

    let  $\ell_1 := \text{ref } 18$  in
    let  $\ell_2 := \text{ref } 20$  in
    send  $c \ell_1$ ; send  $c \ell_2$ ;
    recv  $c$ ; recv  $c$ ;
     $!\ell_1 + !\ell_2$ 
    
```

```

1 Lemma prog_ref_swap_loop_spec :
2   ∀ Φ, Φ #42 -* WP prog_ref_swap_loop #() {Φ}.
3 Proof.
4   iIntros (Φ) "HΦ". wp_lam.
5   wp_apply (start_chan_spec prot_ref_loop); iIntros (c) "Hc".
6   - iLöb as "IH". wp_lam.
7     wp_recv (1 x) as "H1". wp_load. wp_store. wp_send with "[H1]".
8     do 2 wp_pure _. by iApply "IH".
9   - wp_alloc l1 as "H11". wp_alloc l2 as "H12".
10    wp_send with "[H11]". wp_send with "[H12]".
11    wp_recv as "H11". wp_recv as "H12".
12    wp_load. wp_load.
13    wp_pures. by iApply "HΦ".
14 Qed.

```

Figure 3.21: Proof of message-passing program

Here, the forked-off thread acts as a service that recursively receives locations, adds 2 to their stored number, and then sends back a flag indicating that the location has been updated. The main thread, acting like a client, first allocates two new references, to 18 and 20, respectively, which are both sent to the service after which the update flags are received. It finally dereferences the updated locations, and adds their values together, thus returning 42. To verify this program, we use the following recursive protocol:

$$\text{prot_ref_loop} \triangleq \mu(\text{rec} : \text{iProto}). !(\ell : \text{Loc})(x : \mathbb{Z}) \langle \ell \rangle \{ \ell \mapsto x \}. ?\langle () \rangle \{ \ell \mapsto x + 2 \}. \text{rec}$$

The service (in the forked-off thread) follows the (dual of) the protocol exactly, while the main thread follows a weakened version, where the recursion is unfolded twice, after which the second send has been swapped in front of the first receive, allowing it to first send both values before receiving:

$$\begin{aligned} \text{prot_ref_loop} \sqsubseteq & !(\ell_1 : \text{Loc})(x_1 : \mathbb{Z}) \langle \ell_1 \rangle \{ \ell_1 \mapsto x_1 \}. \\ & !(\ell_2 : \text{Loc})(x_2 : \mathbb{Z}) \langle \ell_2 \rangle \{ \ell_2 \mapsto x_2 \}. \\ & ?\langle () \rangle \{ \ell_1 \mapsto (x_1 + 2) \}. \\ & ?\langle () \rangle \{ \ell_2 \mapsto (x_2 + 2) \}. \text{prot_ref_loop} \end{aligned}$$

The full Coq proof of the program is shown in Figure 3.21. The proved lemma is logically equivalent to the specification $\{\text{True}\} \text{prog_ref_swap_loop} () \{v.v = 42\}$, but is presented in weakest precondition style as is common in Iris in Coq. We start the proof on line 4 by introducing the postcondition continuation Φ , and the hypothesis $H\Phi$: $\Phi \#42$, and then evaluate the lambda expression with wp_lam . On line 5 we apply the specification start_chan_spec for start_ by picking the expected protocol prot_ref_loop . This leaves us with two subgoals, separated by bullets -: one for the forked-off thread, and one for the main thread.

In the proof of the recursively-defined forked-off thread, we use iLöb as "IH" for LÖB induction on line 6. This leaves us with the following intermediate proof state:

```

"IH" : ▷ (c ↦ iProto_dual prot_ref_loop -*
  WP (rec: "go" "c'" :=
    let: "l" := recv "c'" in "l" <- ! "l" + #2;;
    send "c'" #()); "go" "c'") c {{ _, True }})
-----□
"Hc" : c ↦ iProto_dual prot_ref_loop
-----*
WP (rec: "go" "c'" :=
  let: "l" := recv "c'" in "l" <- ! "l" + #2;;
  send "c'" #()); "go" "c'") c {{ _, True }}
    
```

We now resolve the application of c to the recursive function with wp_lam . This lets us strip the later from the Löb induction hypothesis, as the program has taken a step. For brevity's sake we refer to the recursive program as prog_rec , in the following proof states.

```

"IH" : c ↦ iProto_dual prot_ref_loop -* WP prog_rec c {{ _, True }}
-----□
"Hc" : c ↦ iProto_dual prot_ref_loop
-----*
WP let: "l" := recv c in "l" <- ! "l" + #2;;
  send c #(); prog_rec c {{ _, True }}
    
```

On line 7 we resolve the proof of the body of the recursive function. So far, the proof only used Iris's standard tactics, we now use the Actris tactic for receive wp_recv (1 x) as "H1", to resolve the receive in evaluation position, introducing the received logical variables l and x , along with the predicate of the protocol $l \mapsto \#x$ naming it H1. To do so, the protocol is normalised, unfolding the recursive definition once, as well as resolving the dualisation of the head, turning it into a receive as expected. This leads to the following proof state:

```

"IH" : c ↦ iProto_dual prot_ref_loop -* WP prog_rec c {{ _, True }}
-----□
"H1" : l ↦ #x
"Hc" : c ↦ iProto_dual (<?> MSG #() {{ l ↦ #(x + 2) }}; prot_ref_loop)
-----*
WP let: "l" := #l in "l" <- ! "l" + #2;; send c #(); prog_rec c {{ _, True }}
    
```

We then use wp_load and wp_store to resolve the dereferencing and updating of the location:

```

"IH" : c ↦ iProto_dual prot_ref_loop -* WP prog_rec c {{ _, True }}
-----□
"H1" : l ↦ #(x + 2)
"Hc" : c ↦ iProto_dual (<?> MSG #() {{ l ↦ #(x + 2) }}; prot_ref_loop)
-----*
WP send c #(); prog_rec c {{ _, True }}
    
```

We then use Actris's tactic wp_send with "[$\$H1$]" to resolve the send operation in evaluation position, by giving up the ownership of "H1". Again, the protocol is automatically normalised by resolving the dualisation of the receive (?) to obtain the send (!) as expected.

We finally close the proof of the forked-off thread on line 8. We first take two pure evaluation steps revolving the sequencing of operations with `do 2 wp_pure _` to reach the recursive call. This results in the proof state:

```
"IH" : c ↦ iProto_dual prot_ref_loop -* WP prog_rec c {{ _, True }}
-----□
"Hc" : c ↦ iProto_dual prot_ref_loop
-----*
WP prog_rec c {{ _, True }}
```

We then use `by iApply "IH"` to close the proof by using the LÖB induction hypothesis.

The proof of the main thread follows similarly. On line 9 we use `wp_alloc 11` as "`H11`" and `wp_alloc 12` as "`H12`", to resolve the allocations of the new locations, binding the logical variables of the locations to 11 and 12, and adding hypotheses "`H11`" and "`H12`" for ownership of these locations to the separation logic proof context. The proof state is then:

```
"HΦ" : Φ #42
"Hc" : c ↦ prot_ref_loop
"H11" : 11 ↦ #18
"H12" : 12 ↦ #20
-----*
WP send c #11;; send c #12;; recv c;; recv c;; ! #11 + ! #12 {{ v, Φ v }}
```

On line 10, we then resolve the first send operation with Actris's tactic `wp_send` with "`[$H11]`", by giving up ownership of the location 11. Here, the protocol is normalised by unfolding the recursive definition, after which the head symbol is a send (!) as expected. The resulting proof state is as follows:

```
"HΦ" : Φ #42
"H12" : 12 ↦ #20
"Hc" : c ↦ (<?> MSG #() {{ 11 ↦ #(18 + 2) }}); prot_ref_loop
-----*
WP send c #12;; recv c;; recv c;; ! #11 + ! #12 {{ v, Φ v }}
```

To resolve the second send operation, we need to weaken the protocol using swapping (rule \sqsubseteq -SWAP'), which is taken care of automatically by Actris's tactic `wp_send` with "`[$H12]`". The normalisation detects that the protocol has a receive (?) as a head symbol, and therefore attempts swapping. To do so it steps ahead of the receive (?), and unfolds the recursive definition, which results in a send (!) as the first symbol after the head. It then detects that there are no dependencies between the two, and can thus apply the swapping rule \sqsubseteq -SWAP', moving the send (!) ahead of the receive (?). With the head symbol now being a send (!), the symbolic execution continues as normal, resulting in the proof state:

```
"HΦ" : Φ #42
"Hc" : c ↦ (<?> MSG #() {{ 11 ↦ #(18 + 2) }};
          <?> MSG #() {{ 12 ↦ #(20 + 2) }}); prot_ref_loop
-----*
WP recv c;; recv c;; ! #11 + ! #12 {{ v, Φ v }}
```

On line 11 we then proceed as expected with `wp_recv` as "`H11`" and `wp_recv` as "`H12`", to resolve the receive operations, giving us back the updated point-to resources:

```

"HΦ" : Φ #42
"H11" : l1 ↦ #(18 + 2)
"H12" : l2 ↦ #(20 + 2)
"Hc" : c ↦ prot_ref_loop
-----*
WP ! #l1 + ! #l2 {{ v, Φ v }}

```

At line 12 we then continue by using `wp_load` twice to dereference the reacquired and updated locations, and then use trivial symbolic execution to resolve the remaining computations. On line 13 we finally close the proof with the continuation hypothesis by `iApply "HΦ"`.

3.7.3 Tactic Support for Subprotocols

While the Actris tactics automatically apply the subprotocol rules during symbolic execution, as shown in Section 3.7.2, we sometimes want to prove subprotocol relations as explicit lemmas. We have tactic support for such proofs as well, which is integrated with the existing MoSeL tactics `iIntros`, `iExists`, `iFrame`, `iModIntro`, and `iSplitL/iSplitR` by automatically using the subprotocol rules to turn the goal into a goal where the regular Iris tactics apply.

- `iIntros (x1 .. xn) "H1 .. Hm"` transforms the subprotocol goal to begin with n universal quantification and m implications, using the rules \sqsubseteq -SEND-IN and \sqsubseteq -RECV-IN, and then introduces the quantifiers (naming them $x1 .. xn$) into the Coq context, and the hypotheses (naming them $H1 .. Hm$) into the separation logic context.
- `iExists (t1 .. tn)` transforms the subprotocol goal to start with n existential quantifiers, using the \sqsubseteq -SEND-OUT and \sqsubseteq -RECV-OUT rules, and then instantiates these quantifiers with the terms $t1 .. tn$ specified by the pattern.
- `iFrame "H"` transforms the subprotocol goal into a separating conjunction between the payload predicates of the head symbols of either protocol, using the rules \sqsubseteq -SEND-OUT and \sqsubseteq -RECV-OUT, and then tries to solve the payload predicate subgoal using "H".
- `iModIntro` transforms the subprotocol goal into a goal starting with a later modality (\triangleright), using the rules \sqsubseteq -SEND-MONO and \sqsubseteq -RECV-MONO, and then introduces that later by stripping off a later from any hypothesis in the separation logic context.
- `iSplitL/iSplitR "H1 .. Hn"` transforms the subprotocol goal into a separating conjunction between the payload predicates of the head symbols of either protocol, using the rules \sqsubseteq -SEND-OUT and \sqsubseteq -RECV-OUT, and then creates two subgoals. For `iSplitL` the left subgoal is given the hypotheses $H1 .. Hn$ from the separation logic context, while the right subgoal is given any remaining hypotheses, and *vice versa* for `iSplitR`.

The extensions of these tactics are implemented by defining custom type class instances that hook into the existing MoSeL tactics as described in [Krebbbers et al. 2017b].

```

1 Lemma list_rev_subprot :
2   ⊢ (<! (l : loc) (vs : list val)> MSG #1 {{ llist l vs }};
3     <?> MSG #() {{ llist internal_eq l (reverse vs) }}; END) ⊑
4     (<! (l : loc) (xs : list T)> MSG #1 {{ llistI IT l xs }};
5     <?> MSG #() {{ llistI IT l (reverse xs) }}; END).
6 Proof.
7   iIntros (l xs) "H1".
8   iDestruct (Hlr with "H1") as (vs) "[H1 HIT]".
9   iExists l, vs. iFrame "H1".
10  iModIntro. iIntros "H1".
11  iSplitL.
12  { rewrite big_sepL2_reverse_2. iApply Hlr.
13    iExists (reverse vs). iFrame "H1 HIT". }
14  done.
15 Qed.

```

Figure 3.22: Proof of subprotocol relation

To demonstrate these tactics, we will go through a proof of the subprotocol relation for the list reversing service presented in Section 3.3.3:

$$\begin{aligned} & ! (l : \text{Loc}) (\vec{v} : \text{List Val}) \langle l \rangle \{ l \xrightarrow{\text{int}} \vec{v} \}. ? \langle () \rangle \{ l \xrightarrow{\text{int}} \text{reverse } \vec{v} \}. \mathbf{end} \\ & \sqsubseteq ! (l : \text{Loc}) (\vec{x} : \text{List } T) \langle l \rangle \{ l \xrightarrow{\text{int}}_{IT} \vec{x} \}. ? \langle () \rangle \{ l \xrightarrow{\text{int}}_{IT} \text{reverse } \vec{x} \}. \mathbf{end} \end{aligned}$$

Recall that the following conversion between the list representation predicate with payload $l \xrightarrow{\text{int}}_{IT} \vec{x}$ and one without payload $l \xrightarrow{\text{int}} \vec{v}$ holds:

$$\text{Hlr} \quad : \quad l \xrightarrow{\text{int}}_{IT} \vec{x} \mathbf{**} (\exists \vec{v}. l \xrightarrow{\text{int}} \vec{v} \mathbf{*}_{(x,v) \in (\vec{x}, \vec{v})} \cdot IT \ x \ v)$$

The full Coq proof is shown in Figure 3.22. On line 7 we start the proof by introducing the logical variables l , xs and the payload $llistI \ IT \ l \ xs$ of the weaker protocol with the tactic `iIntros (l xs) "H1"`. This tactic will implicitly apply the rule `⊑-SEND-IN`, so the goal starts with a universal quantification $\forall (l : \text{loc}) (xs : \text{list } T). \text{llistI } IT \ l \ xs \mathbf{-*} \dots$, which is then introduced based on the regular Iris introduction pattern. This gives us:

```

"H1" : llistI IT l vs
-----*
(<! (l : loc) (vs : list val)> MSG #1 {{ llist l vs }};
 <?> MSG #() {{ llist l (reverse vs) }}; END) ⊑
(<!> MSG #1; <?> MSG #() {{ llistI IT l (reverse xs) }}; END)

```

To obtain the payload predicate expected by the stronger protocol, we use the lemma `Hlr`, to derive `llist l vs` and `[* list] x;v ∈ xs;vs, IT x v` from `llistI l xs` with the tactic `iDestruct (Hlr with "H1") as (vs) "[H1 HIT]"` on line 8. The resulting proof state is:

```

"H1" : llist l vs
"HIT" : [* list] x;v ∈ xs;vs, IT x v
-----*
(<! (l : loc) (vs : list val)> MSG #1 {{ llist l vs }};
 <?> MSG #() {{ llist l (reverse vs) }}; END) ⊑
(<!> MSG #1; <?> MSG #() {{ llistI IT l (reverse xs) }}; END)

```

At line 9 we instantiate the logical variables of the stronger protocol with the logical variables `l` and `vs` using `iExists l, vs`. This will implicitly apply the rule \sqsubseteq -SEND-OUT, which makes the goal start with $\exists(l : \text{loc}) (vs : \text{list val})$, so the existentials can be instantiated as usual. To resolve the payload predicate obligation `l1ist l vs`, we use `iFrame "H1"`. This uses the \sqsubseteq -SEND-OUT to turn the goal into `l1ist l vs * ...`, where the left subgoal is resolved using "H1". We then have the following remaining proof state:

```
"HIT" : [* list] x;v ∈ xs;vs, IT x v
-----*
(<!> MSG #1; <?> MSG #() {{ l1ist l (reverse vs) }}; END) ⊆
(<!> MSG #1; <?> MSG #() {{ l1istI IT l (reverse xs) }}; END)
```

As the head symbols of both protocols are sends (!) with no logical variables or payload predicates, we use `iModIntro` on line 10, which first applies \sqsubseteq -SEND-MONO to step over the sends, and then introduces the later modality (\triangleright). This gives us the proof state:

```
"HIT" : [* list] x;v ∈ xs;vs, IT x v
-----*
(<?> MSG #() {{ l1ist l (reverse vs) }}; END) ⊆
(<?> MSG #() {{ l1istI IT l (reverse xs) }}; END)
```

On line 10, similarly to before, we use `iIntros "H1"`, to introduce the payload predicate, but this time we do it for the stronger protocol, as dictated by \sqsubseteq -RECV-IN:

```
"HIT" : [* list] x;v ∈ xs;vs, IT x v
"H1" : l1ist l (reverse vs)
-----*
(<?> MSG #() ; END) ⊆
(<?> MSG #() {{ l1istI IT l (reverse xs) }}; END)
```

To resolve the payload predicate of the weaker protocol, we use `iSplitL "H1 HIT"` on line 11, that first use \sqsubseteq -RECV-OUT, to turn the goal into `l1istI IT l (reverse xs) * ...`, and then use the goal splitting pattern of Iris, to give us two subgoals, where we use the hypotheses "H1" and "HIT" in the left subgoal. The first subgoal is then:

```
"HIT" : [* list] x;v ∈ xs;vs, IT x v
"H1" : l1ist l (reverse vs)
-----*
l1istI IT l (reverse xs)
```

On line 12, we first use the lemma `H1r` in the right-to-left direction, and then rewrite the hypothesis "HIT" using a lemma from the Iris library with `rewrite big_sepL2_reverse_2`. We do this to obtain `[*list] x;v ∈ reverse xs;reverse vs, IT x v`, in order to match the proof goal. This gives the proof obligation:

```
"HIT" : [* list] x;v ∈ reverse xs;reverse vs, IT x v
"H1" : l1ist l (reverse vs)
-----*
∃ vs : list val, l1ist l vs * ([* list] x;v ∈ reverse xs;vs, IT x v)
```

We finally close the proof on line 13 with `iExists (reverse vs)`, followed by `iFrame "H1 HIT"`, as the goal matches the hypotheses exactly, when picking `reverse vs` as the existential quantification. We then move on to the second subgoal:

```
-----*
(<?> MSG #(); END)  $\sqsubseteq$  (<?> MSG #(); END)
```

We resolve this subgoal, on line 14, with the tactic `done`, which tries to close the proof, by automatically applying \sqsubseteq -REFL.

3.8 Related Work

As Actris combines results from both the separation logic and session types community, there is an abundance of related work. This section briefly elaborates on the relation to message passing in separation logic (Section 3.8.1) and process calculi (Section 3.8.2), session types (Section 3.8.3), endpoint sharing (Section 3.8.4), and verification efforts of map-reduce (Section 3.8.5).

3.8.1 Message Passing and Separation Logic

Lozes and Villard [2012] proposed a logic, based on previous work by Villard et al. [2009], to reason about programs written in a small imperative language with message passing using channels similar to ours. Messages are labelled, and protocols are handled with a combination of finite-state automata (FSA) with correspondingly labelled transitions and predicates associated with each state of the automata. This combination is similar to, but less general than, STSs in Iris. Their language does not support higher-order functions or delegation, but since their language is restricted to structured concurrency (*i.e.*, not fork-based) and their logic is linear (*i.e.*, not affine), they ensure that all resources like channels and memory are properly deallocated.

Craciun et al. [2015] introduced “session logic”, a variant of separation logic that includes predicates for protocol specifications similar to ours. This work includes support for mutable state, ownership transfer via message-passing, delegation through higher-order channels, choice using a special type of disjunction operator on the protocol level, and a sketch of an approach to verify deadlock freedom of programs. Combined, these features allow them to verify interesting and non-trivial message-passing programs. Their logic as a whole is not higher-order, which means that sending functions over channels is not possible. Moreover, their logic does not support protocol-level logical variables that can connect the transferred message with the tail protocol. It is therefore not possible to model dependent protocols like we do in Actris. Their work includes a notion of subtyping as weakening and strengthening of the payload predicates, however they do not consider swapping, and do not allow manipulation of resources (or binders by construction) as a part of their relation. There also exists no support for other concurrency primitives such as locks, which by extension means that manifest sharing is not possible. In Actris we get this for free by building on top of Iris, and reusing its ghost state mechanism. Their work has not been mechanised in a proof assistant, but example programs can be checked using the HIP/SLEEK verifier.

The original Iris [Jung et al. 2015] includes a small message-passing language with channels that do not preserve message order. It was included to demonstrate that

Iris is flexible enough to handle other concurrency models than standard shared-memory concurrency. Since the Hoare-triples for send and receive only reason about the entire channel buffer, protocol reasoning must be done via STSs or other forms of ghost state.

Hamin and Jacobs [2019] take an orthogonal direction and use separation logic to prove deadlock freedom of programs that communicate via message passing using a custom logic tailored to this purpose. They did not provide abstractions akin to our session-type based protocols. Instead one has to reason using invariants and ghost state explicitly.

Mansky et al. [2017] take yet another direction and verify the functional correctness of a message-passing system written in C using the VST framework in Coq [Appel 2014]. While they do not verify message-passing programs like we do, they do verify that the implementation of their message-passing system is resilient to faulty behaviour in the presence of malicious senders and receivers.

Tassarotti et al. [2017] prove correctness and termination preservation of a compiler from a simple language with session types to a functional language with mutable state, where the channels are implemented using references on the heap. This work is also done in Iris. The session types they consider are more like standard session types, which cannot express functional properties of messages, but only their types.

The Diesel logic by Sergey et al. [2018] and the Aneris logic by Krogh-Jespersen et al. [2020] can be used to reason about message-passing programs that work on network sockets. Channels can only be used to send strings, are not order preserving, and messages can be dropped but not duplicated. Since only strings are sent over channels complex data (such as functions) must be marshalled and unmarshalled in order to be sent over the network. Both Diesel and Aneris therefore address a different problem than we do.

The SteelCore framework by [Swamy et al. 2020] is an extensible concurrent separation logic based in F^* , which has been used to encode unidirectional synchronous channels. The channels are encoded as a shallow embedding, which is tied together with ghost state to follow a protocol, using a trace of the communicated messages. Their protocols are defined as a dependent sequence of value obligations with associated separation logic predicates, dictating what can be sent over the channel, including the transfer of ownership. Their channels are first-class and can also be transferred, effectively achieving delegation. They have postulated that bidirectional asynchronous communication is possible, but have not yet done that. Finally, their protocols do not include higher-order protocol-level logical variables, or any notion of subtyping.

3.8.2 Separation Logic and Process Calculi

Another approach is to verify message-passing programs written in some dialect of process calculus. We focus on related work that combines process calculus with separation logic. Neither of the approaches below support delegation or concurrency paradigms other than message passing.

Francalanza et al. [2011] use separation logic to verify programs written in a CCS-like language. Channels model memory location, which has the effect that their

input-actions behave a lot like our updates of mutable state with variable substitutions updating the state. As a proof of concept they prove the correctness of an in-place quick-sort algorithm.

Oortwijn et al. [2016] use separation logic and the mCRL2 process calculus to model communication protocols. The logic itself operates on a high level of abstraction and deals exclusively with intraprocess communication where a fractional separation logic is used to distribute channel resources to concurrent threads. Protocols are extracted from code, but there is no formal connection between the specification logic and the underlying language.

3.8.3 Session Types

Seminal work on linear type systems for the pi calculus by Kobayashi et al. [1996] led to the creation of binary session types by Honda et al. [1998].

Bocchi et al. [2010] pushed the boundaries of what can be verified with multi-party session types while staying within a decidable fragment of first-order logic. They use first-order predicates to describe properties of values being sent and received. Decidability is maintained by imposing restrictions on these predicates, such as ensuring that nothing is sent that will be invalidated down the line. The constraints on the logic do, however, limit what programs can be verified. The work includes standard subtyping on communicated values and on choices, but no notion of swapping sends ahead of receives.

Later work by Dardha et al. [2012] helped merge the linear type systems of Kobayashi with Honda’s session types, which facilitated the incorporation of session types in mainstream programming languages like Java [Hu et al. 2010], Go [Lange et al. 2018], and OCaml [Padovani 2017; Imai et al. 2019]. These works focus on adding session-typed support for the Actor model in existing languages, but do not target functional correctness.

Thiemann and Vasconcelos [2020] introduced label dependent session types, where tails of protocols can depend on the communicated message, which allows for encoding of the choice operators using send and receive. This is similar to our encoding of the choice operators in terms of Actris’s dependent send and receive (Section 3.2.5).

Actris’s subprotocol relation is inspired by the notion of session subtyping, for which seminal work was carried out by Gay and Hole [2005]. Mostrous et al. [2009] extended session subtyping to multiparty asynchronous session types, and as part of that, introduced the notion of swapping sends ahead of receives for independent channels. Mostrous and Yoshida [2015] later considered swapping over the same channel in the context of binary session types. Our subprotocol relation is most closely related to the work of Mostrous and Yoshida [2015], although they define subtyping as a simulation on infinite trees, using so-called asynchronous contexts, whereas we define it using Iris’s support for guarded recursion. It should be noted that the work by Gay and Hole [2005] differs from the work by Mostrous et al. [2009] and Mostrous and Yoshida [2015] in the orientation of the subtyping relation, as discussed by Gay [2016]. Our subprotocol relation uses the orientation of Gay and Hole [2005].

Session subtyping for recursive type systems is universally carried out as a type simulation on infinite trees [Gay and Hole 2005; Mostrous et al. 2009; Mostrous and Yoshida 2015], which complicates subtyping under the recursive operator. Gay et al. [2020] provide further insights on this problem, although they investigate duality rather than subtyping. To reason about recursive types, Brandt and Henglein [1998] present a coinductive formulation of subtyping (which they apply to regular type systems, rather than session types). We use a similar coinductive formulation, but instead of ordinary coinduction, we use Iris’s support for guarded recursion, which lets us prove subtyping relations of recursive protocols using LöB induction.

3.8.4 Endpoint Sharing

One of the key features of session types is that endpoints are owned by a single process. While these endpoints can be delegated (*i.e.*, transferred from one process to another), they typically cannot be shared (*i.e.*, be accessed by multiple processes concurrently). However, as we demonstrate in Section 3.4, sharing channels endpoints is often desirable, and possible in Actris.

In the pi calculus community there has been prior work on endpoint sharing, *e.g.*, by Atkey et al. [2016]; Kobayashi [2006]; Padovani [2014]. The latest contribution in this line of work is by Balzer and Pfenning [2017]; Balzer et al. [2019], who developed a type system based on session types with support for manifest sharing. Manifest sharing is the notion of sharing a channel endpoint between multiple processes using a lock-like structure to ensure mutual exclusion. Their key idea to ensure mutual exclusion using a type system is to use adjoint modalities to connect two classes of types: types that are linear, and thus denote unique channel ownership, and types that are unrestricted, and thus can be shared. The approach to endpoint sharing in Actris is different: dependent separation protocols do not include a built-in notion for endpoint sharing, but can be combined with Iris’s general-purpose mechanisms for sharing, like locks.

3.8.5 Verification of Map-Reduce

To our knowledge the only verification related to the map-reduce model [Dean and Ghemawat 2004] is by Ono et al. [2011], who made two mechanisations in Coq. The first took a functional model of map-reduce and verified a few specific mappers and reducers, extracted these to Haskell, and ran them using Hadoop Streaming. The second did the same by annotating Java mappers and reducers using JML and proving them correct using the Krakatoa tool [Marché et al. 2004], using a combination of SAT-solvers and the Coq proof assistant. While they worked on verifying specific mappers and reducers, our case study focuses on verifying the communication of a map-reduce model that can later be parameterised with concrete mappers and reducers.

3.9 Conclusion and Future Work

In this paper, we have given a comprehensive account of the Actris logic, which incorporates a protocol mechanism based on session types into concurrent separation logic to enable functional correctness proofs of programs that combine message-passing with other programming and concurrency paradigms. Considering the rich literature on session types and concurrent separation logic, we expect there to be many promising directions for future work.

One of the most prominent extensions of binary session types is multi-party session types [Honda et al. 2008], often called choreographies, which allow concise specifications of message transfers between more than two parties. It would be interesting to explore a multi-party version of dependent separation protocols, similar to the multi-party version of session logic by Costea et al. [2018], to allow Actris to more readily verify programs that make use of multi-party communication.

In addition to safety (*i.e.*, session fidelity), conventional session type systems guarantee properties like deadlock and resource-leak freedom. Since Actris is an extension of concurrent separation logic that supports reasoning about several concurrency primitives and not only message passing, ensuring deadlock freedom is hard. The only prior work in this direction that we are aware of is by Hamin and Jacobs [2019] and Craciun et al. [2015], but it is not immediately obvious how to integrate that with Iris or Actris. Resource-leak freedom has been studied in Iron, an extension of Iris by Bizjak et al. [2019], which makes it possible to prove resource-leak freedom of non-structured fork-based concurrent programs. It would be interesting to build dependent separation protocols on top of Iron instead of Iris.

Semantic Session Typing

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.

4.1 Introduction

Session types [Honda et al. 1998] 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 [Gay 2008], asynchronous subtyping [Mostrous et al. 2009], and sharing via locks [Balzer and Pfenning 2017]. 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^o.
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 [Appel and McAllester 2001; Ahmed 2004; Ahmed et al. 2010], using *logical relations* defined in terms of a program logic [Appel et al. 2007; Dreyer et al. 2009, 2019].

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 [Jung et al. 2015, 2016; Krebbers et al. 2017a; Jung et al. 2018b; Krebbers et al. 2017b, 2018] 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 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 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 [Jung et al. 2015, 2016; Krebbers et al. 2017a; Jung et al. 2018b] has proved to be sufficiently expressive to define semantic type systems for *e.g.*, Rust [Jung et al. 2018a, 2021] and Scala [Giarrusso et al. 2020], due to its state-of-the-art built-in support for *e.g.*, resource ownership, recursion, polymorphism, and concurrency. In addition, we make

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.

use of the Actris framework for message passing in Iris [Hinrichsen et al. 2020, 2021a]. 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 [Krebbers et al. 2017b, 2018].

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 (Section 4.2). 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 (Section 4.3).
- We demonstrate the benefit of our type system being semantic by integrating the manual verification of safe but not conventionally type-checkable programs (Section 4.4).
- We provide insight on the benefits of a semantic type system in regards to mechanisation efforts (Section 4.5). All of our results are mechanised in the Coq proof assistant and can be found in [Hinrichsen et al. 2021b].

4.2 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 (Section 4.2.1). We provide a brief overview of Iris (Section 4.2.2), and then present a lightweight affine type system (Section 4.2.3) as the core upon which we built our session-type system (Section 4.2.4). Our affine type system is inspired by RustBelt [Jung et al. 2018a, 2021], but drops Rust-specific features like borrowing and lifetimes to focus on session types.

4.2.1 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 $\text{rec } x := e$, and $\text{let } x := e_1 \text{ in } e_2$ for $(\lambda x. e_2) e_1$, and $e_1; e_2$ for $\text{let } _ := e_1 \text{ in } e_2$. Message passing is given an asynchronous semantics: $\text{new_chan } ()$ returns a pair (c_1, c_2) of channel endpoints that operate on buffers (\vec{v}_1, \vec{v}_2) that are initially empty, $\text{send } c_i w$ enqueues message w in \vec{v}_i , while $\text{recv } c_i$ blocks until a message w is available in $\vec{v}_{(\text{if } i=2 \text{ then } 1 \text{ else } 2)}$, and then dequeues and returns w . Mutable references ℓ are allocated with $\text{ref } e$, 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 .

4.2.2 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 $P, Q \in \text{iProp}$ 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 Section 4.2.3 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 (Section 4.2.4) by means of the connective $(c \mapsto -)$ 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 (Section 4.4).

To define recursive types semantically (Section 4.3.3), 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: $(\triangleright P) * \text{wp } e \{ \Phi \} \dashv\text{* } \text{wp } e \{ w. P * \Phi w \}$ 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”, 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. [2018b] for an extensive account of the Iris model and proof rules.

4.2.3 Term Types

The definitions of our semantic type system are shown in Figure 4.1. Types Type_k are indexed by kinds; \star for *term types*, and \blacklozenge for *session types*. Meta-variables $A, B \in \text{Type}_\star$ are used for term types, $S \in \text{Type}_{\blacklozenge}$ for session types, and $K \in \text{Type}_k$ for types of any kind. Term types Type_\star are defined as Iris predicates over values, and session types $\text{Type}_{\blacklozenge}$ are defined as dependent separation protocols of Actris (Section 4.2.4).

Type Formers The ground types (unit type $\mathbf{1}$, Boolean type \mathbf{B} , and integer type \mathbf{Z}) are defined through membership of the corresponding set ($\{\{\}\}$, \mathbb{B} , and \mathbb{Z}). 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 ($\dashv\text{*}$) 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 Section 4.3.3.

Typing Judgement As is common in substructural type systems with operations that perform strong updates, we use a typing judgement $\Gamma \Vdash e : A \dashv\text{=} \Gamma'$ (defined in Figure 4.1) 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 .

Term types:

$$\begin{aligned}
\text{Type}_\star &\triangleq \text{Val} \rightarrow \text{iProp} \\
\text{any} &\triangleq \lambda w. \text{True} \\
\mathbf{1} &\triangleq \lambda w. w \in \{()\} \\
\mathbf{B} &\triangleq \lambda w. w \in \mathbb{B} \\
\mathbf{Z} &\triangleq \lambda w. w \in \mathbb{Z} \\
\text{ref}_{\text{uniq}} A &\triangleq \lambda w. \exists v. w \in \text{Loc} * (w \mapsto v) * \triangleright(A v) \\
A_1 \times A_2 &\triangleq \lambda w. \exists v_1, v_2. w = (v_1, v_2) * \triangleright(A_1 v_1) * \triangleright(A_2 v_2) \\
A_1 + A_2 &\triangleq \lambda w. \exists v. (w = \text{inl } v * \triangleright(A_1 v)) \vee (w = \text{inr } v * \triangleright(A_2 v)) \\
A \multimap B &\triangleq \lambda w. \forall v. \triangleright(A v) \multimap \text{wp } (w v) \{B\} \\
\text{chan } S &\triangleq \lambda w. w \mapsto S
\end{aligned}$$

Session Types:

$$\begin{aligned}
\text{Type}_\blacklozenge &\triangleq \text{iProto} \\
\text{end} &\triangleq \text{end} \\
!A. S &\triangleq !(v : \text{Val}) \langle v \rangle \{A v\}. S \\
?A. S &\triangleq ?(v : \text{Val}) \langle v \rangle \{A v\}. S \\
\oplus\{\vec{S}\} &\triangleq !(l : \mathbb{Z}) \langle l \rangle \left\{ l \in \text{dom}(\vec{S}) \right\}. \vec{S}(l) \\
\&\{\vec{S}\} &\triangleq ?(l : \mathbb{Z}) \langle l \rangle \left\{ l \in \text{dom}(\vec{S}) \right\}. \vec{S}(l)
\end{aligned}$$

Judgements:

$$\begin{aligned}
\Gamma \vDash \sigma &\triangleq \star_{(x,A) \in \Gamma}. A(\sigma(x)) \\
\Gamma \vDash e : A \ni \Gamma' &\triangleq \forall \sigma. (\Gamma \vDash \sigma) \multimap \text{wp } e[\sigma] \{v. A v * (\Gamma' \vDash \sigma)\}
\end{aligned}$$

Figure 4.1: Typing judgements and type formers of the semantic type system.

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 $\star_{(x,A) \in \Gamma}$ 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 $\Gamma \vDash e : A \ni \Gamma'$ 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 post-condition 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' .

Selection of Iris's proof rules for weakest preconditions:

$$\begin{array}{l}
\Phi v \text{ * wp } v \{ \Phi \} \quad (\text{WP-VAL}) \\
\ell \mapsto v \text{ * wp } !\ell \{ w. (v = w) * (\ell \mapsto v) \} \quad (\text{WP-LOAD}) \\
\text{wp } e_1 \{ v. \text{wp } e_2[v/x] \{ \Phi \} \} \text{ * wp } (\text{let } x := e_1 \text{ in } e_2) \{ \Phi \} \quad (\text{WP-LET}) \\
\text{wp } e_1 \{ \Phi_1 \} \text{ * wp } e_2 \{ \Phi_2 \} \text{ * wp } (e_1 \parallel e_2) \left\{ w. \exists w_1, w_2. (w = (w_1, w_2)) * \Phi_1 w_1 * \Phi_2 w_2 \right\} \quad (\text{WP-PAR})
\end{array}$$

Selection of semantic typing rules:

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

Figure 4.2: A selection of Iris's proof rules and semantic typing rules.

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 $\Gamma \vDash e : A \vDash \Gamma'$, 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 4.2.

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 $\Gamma, x : \mathbf{any}$ as ownership of A has been moved out. For substructural type systems this is crucial as in expressions such as $\text{let } x := y \text{ in } e$, 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 \mathbf{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 $[] \vDash e : A \vDash \Gamma$, 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 [Krebbbers et al. 2017a; Jung et al. 2018b]. 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.

4.2.4 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 $\vec{S} : \mathbb{Z} \xrightarrow{\text{fin}} \text{Type}_\blacklozenge$ 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* [Hinrichsen et al. 2020], 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 $!\vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}$ and $?\vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}$ 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 $!(\ell : \text{Loc}) (i : \mathbb{Z}) \langle \ell \rangle \{\ell \mapsto i * 10 \leq i\}. ?(\langle () \rangle \{\ell \mapsto (i + 1)\}). \text{end}$ 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 $c \mapsto \text{prot}$ denotes ownership of a channel c with a dependent separation protocol prot . The Actris proof rules are shown in Figure 4.3. The rule for **new_chan** $()$ allows ascribing any protocol to a new channel, obtaining ownership of $c \mapsto \text{prot}$ and $c' \mapsto \overline{\text{prot}}$ 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 $c \mapsto !\vec{x} : \vec{\tau} \langle v \rangle \{P\}. \text{prot}$, pick an appropriate 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 4.1. Since session types ($\text{Type}_\blacklozenge$) are defined as dependent separation protocols iProto , the channel type **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 $(?)$

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

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

Session Typing Rules The session typing rules are shown in Figure 4.3. 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:

$$\begin{aligned} \text{select } x \ l &\triangleq \text{send } x \ l \\ \text{branch } x \ \text{with } &\triangleq \text{let } i := \text{recv } x \ \text{in} \\ l_1 \Rightarrow e_1 \mid \dots \mid & \quad \text{if } i = l_1 \ \text{then } e_1 \ \text{else } \dots \\ l_n \Rightarrow e_n & \quad \text{if } i = l_n \ \text{then } e_n \ \text{else } () \ () \end{aligned}$$

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 definition of `select` (\oplus) and `branch` ($\&$).

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 Section 4.2.3 remains applicable without change.

4.3 Extending the Type System

We demonstrate the extensibility of our approach to session types by adding term- and session-level subtyping (Section 4.3.1 and Section 4.3.7), copyable types (Section 4.3.2), term- and session-level equi-recursive types (Section 4.3.3), term- and session-level polymorphism (Section 4.3.4 and Section 4.3.5), and locks/mutexes (Section 4.3.6). While we only present a small representative selection of rules associated with each extension, all rules can be found in Section 4.A.

4.3.1 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:

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

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 $\Gamma \vDash \sigma$ for the context Γ , then we can give those up to obtain the resources $\Gamma' \vDash \sigma$ for Γ' . With these definitions at hand, we prove the usual subsumption rule as a lemma:

$$\frac{\Gamma_1 <:\text{ctx} \Gamma'_1 \quad \Gamma'_1 \vDash e : A \doteq \Gamma'_2 \quad A <: B \quad \Gamma'_2 <:\text{ctx} \Gamma_2}{\Gamma_1 \vDash e : B \doteq \Gamma_2}$$

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:

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

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 section 4.3.2 and section 4.3.7.

4.3.2 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 section 4.2.3, 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 **B**, **Z**, or **Z * 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**:

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

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 **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 (**1**, **B**, **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 (\square) of Iris, where $\square P$ 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, $\square P$ is always persistent, allowing the proposition P to be freely duplicated using the rule $\square P \multimap (\square P * P)$. 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 [Jung et al. 2018a, 2021] 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.

4.3.3 Equi-Recursive Term and Session Types

We extend our type system with equi-recursive types using Iris’s fixpoint operator. Recall from Section 4.2.2 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 $!\vec{x} : \vec{\tau} \langle v \rangle \{ P \}$, *prot* and $?\vec{x} : \vec{\tau} \langle v \rangle \{ P \}$, *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 [Gay et al. 2020], such as $\mu(X : \blacklozenge). !(\text{chan } X). X$, 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).

4.3.4 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 \blacklozenge). 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 $\forall(X : \blacklozenge). \mathbf{chan} (!\mathbf{Z}.X) \multimap \mathbf{chan} X$ (for describing a function that reads an integer from a channel with an arbitrary tail X). Universal and existential types are defined as follows:

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

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 $X : \mathbf{Type}_k$. 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* [Wright 1995] 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 Section 4.4.A.

4.3.5 Polymorphism in Session Types

A more interesting extension is polymorphism in session types [Gay 2008]. An example is the following type, which describes the interaction with a polymorphic computation service:

$$\mathbf{compute_type} \triangleq \mu(\mathit{rec} : \blacklozenge). \oplus \{ \mathbf{cont} : !_{(X:\star)} (\mathbf{1} \multimap X). ?X. \mathit{rec}, \mathbf{stop} : \mathbf{end} \}$$

The service can be used by sending computation requests $\mathbf{1} \multimap X$, 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:

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

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 $!\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ and $?\vec{x}:\vec{\tau}\langle v\rangle\{P\}.prot$ 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:

$$\frac{\Gamma \vDash e : A[\vec{K}/\vec{X}] \ni \Gamma', x : \mathbf{chan} (!\vec{x}:\vec{k}.A.S)}{\Gamma \vDash \mathbf{send} \ x \ e : \mathbf{1} \ni \Gamma', x : \mathbf{chan} \ S[\vec{K}/\vec{X}]}$$

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

The second rule requires the result y of \mathbf{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
    | left  $\Rightarrow$  let f := recv c in send (f ()) ; go c
    | right  $\Rightarrow$  ()
  end

```

We can prove the typing judgement $\Gamma \vDash \mathbf{compute_service} : \mathbf{chan} \ \overline{\mathbf{compute_type}} \rightarrow () \ni \Gamma$ using only the typing rules of our semantic type system. In section 4.4.2 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.

4.3.6 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 [2017] 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 Section 4.3.5 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 4.4 and include the type former `mutex`, which signifies that the mutex is acquired.

Iris's proof rules for locks:

$$\begin{aligned}
& \text{is_lock } lk \ R \ \multimap \ \Box(\text{is_lock } lk \ R) \\
& \quad R \ \multimap \ \text{wp } \text{new_lock } () \ \{lk. \text{is_lock } lk \ R\} \\
& \text{is_lock } lk \ R \ \multimap \ \text{wp } \text{acquire } lk \ \{R\} \\
& \text{is_lock } lk \ R \ * \ R \ \multimap \ \text{wp } \text{release } lk \ \{\text{True}\}
\end{aligned}$$

Semantic typing rules for mutexes:

$$\begin{aligned}
& \text{copyable } (\text{mutex } A) \\
& \Gamma \Vdash \text{newmutex } : A \rightarrow \text{mutex } A \Vdash \Gamma \\
& \Gamma, x : \text{mutex } A \Vdash \text{acquiremutex } x : A \Vdash \Gamma, x : \overline{\text{mutex}} A \\
& \frac{\Gamma \Vdash e : A \Vdash \Gamma', x : \overline{\text{mutex}} A}{\Gamma \Vdash \text{releasemutex } x \ e : \mathbf{1} \Vdash \Gamma', x : \text{mutex } A}
\end{aligned}$$

Figure 4.4: Iris's proof rules for locks and semantic typing rules for mutexes.

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:

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

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 4.4 and make use of the representation predicate `is_lock 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 `is_lock 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:

$$\begin{aligned}
\text{mutex } A &\triangleq \lambda w. \exists lk, \ell. (w = (lk, \ell)) * \text{is_lock } lk \ (\exists v. (\ell \mapsto v) * \triangleright(A \ v)) \\
\overline{\text{mutex}} A &\triangleq \lambda w. \exists lk, \ell. (w = (lk, \ell)) * \text{is_lock } lk \ (\exists v. (\ell \mapsto v) * \triangleright(A \ v)) * (\ell \mapsto -)
\end{aligned}$$

The `mutex` type former states that its values are pairs of locks and boxed values. The `mutex` type former additionally asserts ownership of the reference, implying that the lock has been acquired. The typing rules for mutexes as shown in Figure 4.4 are proven as lemmas.

4.3.7 Session-Level Subtyping

Session-level subtyping $S <: T$, originally presented by Gay and Hole [2005], relates session subtypes S with session supertypes 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:

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

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 ($\&$) [Mostrous et al. 2009]². For example, using swapping, a client of the computation service from Section 4.3.5, 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* [Hinrichsen et al. 2021a], for which the rules are shown in Figure 4.5. 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 -SEND-OUT) and (2) left-hand side receiving protocols (\sqsubseteq -RECV-OUT), and (3) move in binders and propositions of right-hand side sending protocols (\sqsubseteq -SEND-IN) and (4) left-hand side receiving protocols (\sqsubseteq -RECV-IN). Rule \sqsubseteq -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 -SEND-MONO and \sqsubseteq -RECV-MONO account for the monotonicity of the subprotocol relation in the tails,

²Discrepancies in the direction between the swapping rules of Mostrous et al. [2009] and us will be discussed in Section 4.6.

Actris's proof rules for subprotocols:

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

Semantic subtyping rules for session polymorphism:

$$\begin{array}{c}
\frac{S_1 < : !A. S_2}{S_1 < : !_{(\vec{x}:\vec{k})} A. S_2} \qquad \frac{?A. S_1 < : S_2}{?_{(\vec{x}:\vec{k})} A. S_1 < : S_2} \\
!_{(\vec{x}:\vec{k})} A. S < : !A[\vec{K}/\vec{X}]. S[\vec{K}/\vec{X}] \qquad ?A[\vec{K}/\vec{X}]. S[\vec{K}/\vec{X}] < : ?_{(\vec{x}:\vec{k})} A. S
\end{array}$$

Figure 4.5: A selection of the Actris's proof rules for subprotocols and semantic session subtyping rules. $\sqsubseteq\text{-SEND-OUT}$ and $\sqsubseteq\text{-RECV-OUT}$ only hold if $\tau \in \vec{\tau}$ is inhabited

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 4.5 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 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 4.5. 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 $?A[\vec{K}/\vec{X}]. S[\vec{K}/\vec{X}] <: ?(\vec{X}: \vec{k}) A. S$.

4.4 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. [2018a, 2021], 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 (Section 4.4.1), and then show a more realistic example by proving the typing judgement of a parallel client of the computation service from Section 4.3.5 that uses a producer/consumer pattern (Section 4.4.2).

4.4.1 Receiving in Parallel

Consider the example from Section 4.1 (where the locks have been made explicit):

$$\begin{array}{ll} \text{threadprog} \triangleq & \text{lockprog} \triangleq \\ \lambda c. lk. \text{acquire } lk; & \lambda c. \text{let } lk := \text{new_lock } () \text{ in} \\ \quad \text{let } x := \text{recv } c \text{ in} & \quad (\text{threadprog } c \text{ } lk \parallel \text{threadprog } c \text{ } lk) \\ \quad \text{release } lk; & \\ x & \end{array}$$

We want to prove the following typing judgement:

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

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:

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

The proof of above obligation is carried out using Iris's support for fractional permissions $\llbracket \frac{q}{i} \rrbracket^\gamma$ 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:

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

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

$$(\text{is_lock } lk \text{ (chaninv } c) * \llbracket \frac{1}{2} \rrbracket^\gamma) \multimap \text{wp threadprog } c \text{ lk } \{v. v \in \mathbb{Z}\}$$

First, the lock invariant is obtained by acquiring the lock. The channel can then either be in state (i) or (ii), as having 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 $\llbracket \frac{1}{2} \rrbracket^\gamma$ 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 $\llbracket \frac{1}{1} \rrbracket^\gamma$ in state (iii).

4.4.2 A Parallel Computation Client

In section 4.3.5 we considered the session type `compute_type` for a client of a polymorphic recursive computation service. We now consider a client `compute_client`,

<pre> compute_client \triangleq $\lambda l c.$ let $n := \text{llength } l$ in let $ctr := \text{ref } 0$ in let $l' := \text{lnil } ()$ in let $lk := \text{new_lock } ()$ in (produce $l \ ctr \ lk \ c$ consume $l' \ n \ ctr \ lk \ c$); l' </pre>	<pre> produce \triangleq rec go $l \ ctr \ lk \ c :=$ if lisnil l then acquire lk; select c stop; release lk else acquire lk; select c cont; send c (lpop l); $ctr \leftarrow !ctr + 1$ release lk; go $l \ ctr \ lk \ c$ </pre>	<pre> consume \triangleq rec go $l \ n \ ctr \ lk \ c :=$ if $n = 0$ then $()$ else acquire lk; if $!ctr = 0$ then release lk; go $l \ n \ ctr \ lk \ c$ else let $x := \text{recv } c$ in $ctr \leftarrow !ctr - 1$; release lk; go $l \ (n - 1) \ ctr \ lk \ c$; lcons $x \ l$ </pre>
---	--	---

Figure 4.6: 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).

shown in Figure 4.6, 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:

$$\Gamma \vDash \text{compute_client} : \text{list } (\mathbf{1} \multimap A) \multimap \text{chan compute_type} \multimap \text{list } A \approx \Gamma$$

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

The client `compute_client` operates on a channel endpoint c , where the computation service has the other endpoint. The client creates a shared counter ctr 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 `produce` and consumer `consume` in parallel, which both race for the lock lk to access the channel c and counter ctr . 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 ctr thereafter. The consumer `consume` adds the results one-by-one to the list l' by receiving them on the channel c , and decreasing the shared counter ctr thereafter. When both the producer and consumer terminate, the client returns the list l' that then contains the results.

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

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.

rules. We first use the weakest precondition rule for `llength`:

$$\mathbf{list} \ A \ l \ * \ \mathbf{wp} \ \mathbf{llength} \ l \ \{n. n = |\vec{v}| * l \xrightarrow{\text{list}}_A \vec{v}\}$$

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

$$l \xrightarrow{\text{list}}_A \vec{v} \triangleq \begin{cases} l \mapsto \text{inl} \ () & \text{if } \vec{v} = [] \\ \exists l_2. l \mapsto \text{inr} \ (v_1, l_2) * A \ v_1 * l_2 \xrightarrow{\text{list}}_A \vec{v}_2 & \text{if } \vec{v} = [v_1] \cdot \vec{v}_2 \end{cases}$$

The remainder of the proof is similar to the proof of the parallel receive in Section 4.4.1—we establish a lock invariant `chaninvpc ctr c A` to share the counter `ctr` and the channel `c` between the producer and consumer, and use a fractional permission $\{\underline{1}\}_1^\gamma$ to determine the state of the shared channel `c`:

$$\begin{aligned} \mathbf{chaninv}_{pc} \ \mathit{ctr} \ c \ A \ \triangleq \ \exists n. \ \mathit{ctr} \ \mapsto \ n \ * \\ (c \mapsto ((?A)^n \cdot \mathbf{compute_type}) \vee \quad (i) \\ (c \mapsto ((?A)^n \cdot \mathbf{end}) * \{\underline{1}\}_1^\gamma) \quad (ii) \end{aligned}$$

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 `ctr`. 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 $\{\underline{1}\}_1^\gamma$, which must be released before closing the channel.

The proof is carried out by allocating the fractional permission $\{\underline{1}\}_1^\gamma$ (with a fresh identifier γ), after which the weakest precondition rules for parallel composition (see Figure 4.2), the producer `produce`, and consumer `consume` are used:

$$\begin{aligned} \mathbf{is_lock} \ lk \ (\mathbf{chaninv}_{pc} \ \mathit{ctr} \ c \ A) * \{\underline{1}\}_1^\gamma * l \xrightarrow{\text{list}}_{(\mathbf{1} \multimap A)} \vec{v} \ * \\ \mathbf{wp} \ \mathbf{produce} \ l \ \mathit{ctr} \ lk \ c \ \{l \xrightarrow{\text{list}}_{(\mathbf{1} \multimap A)} []\} \\ \mathbf{is_lock} \ lk \ (\mathbf{chaninv}_{pc} \ \mathit{ctr} \ c \ A) * l \xrightarrow{\text{list}}_A [] \ * \\ \mathbf{wp} \ \mathbf{consume} \ l \ n \ \mathit{ctr} \ lk \ c \ \{\exists \vec{w}. |\vec{w}| = n * l \xrightarrow{\text{list}}_A \vec{w}\} \end{aligned}$$

The proof of `produce` proceeds as follows. Owning $\{\underline{1}\}_1^\gamma$ 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:

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

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

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

If the list is non-empty, we resolve a computation step (by selecting the `cont` branch and sending a computation with type $\mathbf{1} \multimap A$) resulting in $c \mapsto (?A)^{n+1} \cdot \text{compute_type}$. After incrementing the shared counter ctr , 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 $\overline{\mathbf{1}}_1$.

The proof of `consume` proceeds as follows. We only perform a receive operation when the shared counter ctr 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 ctr we can reestablish the lock invariant.

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

4.5 Mechanisation in Coq

In this paper, we have used what is often called the “foundational approach” to semantic type safety [Appel and McAllester 2001; Ahmed 2004; Ahmed et al. 2010]. 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 hides unnecessary details related to the embedding of separation logic in Coq [Krebbers et al. 2017b, 2018].

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 → Type :=
  | Ltty : (val → iProp Σ) → lty Σ tty_kind
  | Lsty : iProto Σ → lty Σ sty_kind.
Notation ltty Σ := (lty Σ tty_kind).
Notation lsty Σ := (lty Σ sty_kind).
```

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.

Typing contexts are represented as association lists:

```
Inductive ctx_item  $\Sigma$  := CtxItem {
  ctx_item_name : string;
  ctx_item_type : lttty  $\Sigma$  }.
Notation ctx  $\Sigma$  := (list (ctx_item  $\Sigma$ )).
```

The semantic term typing judgement is defined as:

```
(* lttty_car: lttty  $\Sigma$   $\rightarrow$  (val  $\rightarrow$  iProp  $\Sigma$ ) is the
   inverse of Lttty *)
Definition ltttyped ( $\Gamma$ 1  $\Gamma$ 2 : ctx  $\Sigma$ )
  (e : expr) (A : lttty  $\Sigma$ ) : iProp  $\Sigma$  :=
  ■  $\forall$  vs, ctx_ltttyped vs  $\Gamma$ 1 -*
    WP subst_map vs e {{ v, lttty_car A v * ctx_ltttyped vs  $\Gamma$ 2 }}.
Notation " $\Gamma$ 1  $\models$  e : A  $\models$   $\Gamma$ 2" :=
  (ltttyped  $\Gamma$ 1  $\Gamma$ 2 e A) : bi_scope.
Notation " $\Gamma$ 1  $\models$  e : A  $\models$   $\Gamma$ 2" :=
  ( $\vdash$  ltttyped  $\Gamma$ 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

The typing judgement is identical to the definition in Section 4.2.3, 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 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* (■), which ensures that it does not capture any separation logic resources.⁶ We define two notations so 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 ltttyped_let  $\Gamma$ 1  $\Gamma$ 2  $\Gamma$ 3 x e1 e2 A1 A2 :
  ( $\Gamma$ 1  $\models$  e1 : A1  $\models$   $\Gamma$ 2) -*
  (ctx_cons x A1  $\Gamma$ 2  $\models$  e2 : A2  $\models$   $\Gamma$ 3) -*
  ( $\Gamma$ 1  $\models$  (let: x := e1 in e2) : A2  $\models$  ctx_filter_eq x  $\Gamma$ 2 ++ ctx_filter_ne x  $\Gamma$ 3).
```

The typing rule of `let` shows the handling of shadowing of variables: `ctx_cons x A1 Γ 2` removes all bindings of `x` from `Γ 2` before adding the new binding, and `ctx_filter_eq x Γ 2` makes sure that potentially overshadowed variables are preserved. Dealing with

⁵Since `HeapLang`'s operational semantics is defined on closed terms, the use of strings does not cause issues with variable capture. See also [Pierce et al. 2020, Section STL] for a discussion on the use of strings for variables.

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

shadowing in the proof is trivial due to some general-purpose lemmas for $\Gamma \models \sigma$ (`ctx_ltyped` Γ `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 : lty  $\Sigma$  k  $\rightarrow$  ltt  $\Sigma$ ) : ltt  $\Sigma$  :=
  Ltty ( $\lambda$  w,  $\forall$  X, WP w #() {{ ltt_car (C X) }}).
Notation " $\forall$  X, C" := (lty_forall ( $\lambda$  X, C)) : lty_scope.

Lemma ltyped_tlam  $\Gamma$ 1  $\Gamma$ 2  $\Gamma'$  e k (C : lty  $\Sigma$  k  $\rightarrow$  ltt  $\Sigma$ ) :
  ( $\forall$  K,  $\Gamma$ 1  $\models$  e : C K  $\Rightarrow$  []) -*
  ( $\Gamma$ 1 ++  $\Gamma$ 2  $\models$  ( $\lambda$ : <>, e) : ( $\forall$  X, C X)  $\Rightarrow$   $\Gamma$ 2).
Lemma ltyped_tapp  $\Gamma$   $\Gamma$ 2 e k (C : lty  $\Sigma$  k  $\rightarrow$  ltt  $\Sigma$ ) K :
  ( $\Gamma$   $\models$  e : ( $\forall$  X, C X)  $\Rightarrow$   $\Gamma$ 2) -*
  ( $\Gamma$   $\models$  e #() : C K  $\Rightarrow$   $\Gamma$ 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 [Pfenning and Elliott 1988], 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 : gmap Z (lty  $\Sigma$ ) ) : lty  $\Sigma$  :=
  Lsty (<a@(i: Z)> MSG #x {{  $\ulcorner$ is_Some (Ss !! i)  $\urcorner$  }}; lsty_car (Ss !!! i)).
Notation lty_select := (lty_choice Send).

Lemma ltyped_select  $\Gamma$  x i S Ss :
   $\Gamma$  !! x = Some (chan (lty_select Ss))  $\rightarrow$ 
  Ss !! i = Some S  $\rightarrow$ 
   $\Gamma$   $\models$  select x #i : ()  $\Rightarrow$  env_cons x (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 `<a@($\vec{x}:\vec{\tau}$)> MSG v {{ P }}`; `prot` expands to `Actris's ! $\vec{x}:\vec{\tau}$ (v){P}`. `prot` or `? $\vec{x}:\vec{\tau}$ (v){P}`. `prot` depending on the action `a`. The definition uses the finite map library `gmap` of `std++` [The `Coq-std++` Team 2020], 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 `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 uses `Actris's` proof rules, and is 6 lines of code.

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

```

Lemma ltyped_safety e  $\sigma$  es  $\sigma'$  e' :
  ( $\exists$  A, []  $\models$  e : A  $\Rightarrow$  [])  $\rightarrow$ 
  rtc erased_step ([e],  $\sigma$ ) (es,  $\sigma'$ )  $\rightarrow$  e'  $\in$  es  $\rightarrow$ 
  is_Some (to_val e')  $\vee$  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.

4.6 Related Work

Session Types Seminal work on subtyping for binary recursive session types for a synchronous pi-calculus was done by Gay and Hole [2005]. Mostrous et al. [2009] expand on this work by adding support for multi-party asynchronous recursive session types, and later for higher-order process calculi [Mostrous and Yoshida 2015]. These two works present the session subtyping relation with inverted orientations, inverting the sub- and supertypes, which has been discussed by Gay [2016]. Our semantic session subtyping relation uses the same orientation as Gay and Hole. Mostrous et al. [2009] also present an output-input swapping rule, which inspired our swapping rule in Section 4.3.7, 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. [2017], precisely because of the swapping rule.

Gay [2008] 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. [2012] expanded on this work by adding subtyping and delegation, while still only conjecturing that recursion was a possible extension. Caires et al. [2013] 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 [2020] 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 [2017] and Balzer et al. [2019] 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 [Krebbers et al. 2017b; Jung et al. 2018a; Giarrusso et al. 2020], program refinement [Krebbers et al. 2017b; Krogh-Jespersen et al. 2017; Tassarotti et al. 2017; Timany et al. 2018; Frumin et al. 2018], robust safety [Swasey et al. 2017], and non-interference [Frumin et al. 2020]. The most immediately related work in this area is the RustBelt project [Jung et al. 2018a], 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. [2020] 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 [2010], Wadler [2012], Carbone et al. [2017], and Dardha and Gay [2018]. As part of this line of work, Perez *et al.* used logical relations to prove termination [Pérez et al. 2012] and confluence [Pérez et al. 2014] 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. [2017] 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. [2020, 2021a] 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. [2020] 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. [2020] 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 [2019] 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. [2020] 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.

4.7 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.

4.A The Complete 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 [Hinrichsen et al. 2021b].

In particular, the appendix shows the type and judgement definitions in Figure 4.7, the typing rules in Figures 4.8 and 4.9, and the subtyping rules in Figures 4.11 to 4.14.

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 (Section 4.A.1). We furthermore describe kinded subtyping and type equivalence (Section 4.A.2), shared reference types (Section 4.A.3), and discuss the internal versions of all judgements of the type system (Section 4.A.4).

We omit the typing rule for polymorphic sends from Section 4.3.5 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).

4.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), *e.g.*:

$$\text{TY-REFUNIQLoad-Move} \\ \Gamma, x : \mathbf{ref}_{\text{uniq}} A \Vdash !x : A \ni \Gamma, x : \mathbf{ref}_{\text{uniq}} \mathbf{any}$$

$$\text{TY-REFUNIQLoad-Copy} \\ \frac{\text{copyable } A}{\Gamma, x : \mathbf{ref}_{\text{uniq}} A \Vdash !x : A \ni \Gamma, x : \mathbf{ref}_{\text{uniq}} A}$$

The full version of our type system unifies these rules as a 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 `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 (\square). 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>.

4.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.

4.A.3 Shared References

We also have an additional type former `refshr`, 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 `refshr` 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 `refuniq`.

4.A.4 Internal Judgements

In Section 4.5 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 (■). 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:

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

Session Types:

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

Judgements:

$$\begin{aligned}
\Gamma \vDash \sigma &\triangleq \star_{(x,A) \in \Gamma}. A(\sigma(x)) \\
\Gamma \vDash e : A \vDash \Gamma' &\triangleq \blacksquare(\forall \sigma. (\Gamma \vDash \sigma) \multimap \text{wp } e[\sigma] \{v. A v * (\Gamma' \vDash \sigma)\}) \\
A <: B &\triangleq \blacksquare(\forall v. A v \multimap B v) \\
S <: T &\triangleq \blacksquare(S \sqsubseteq T) \\
K <:> L &\triangleq K <: L \wedge L <: K \\
\Gamma <:\text{ctx} \Gamma' &\triangleq \blacksquare(\forall \sigma. (\Gamma \vDash \sigma) \multimap (\Gamma' \vDash \sigma)) \\
\text{copyable } A &\triangleq A <: \text{copy } A
\end{aligned}$$

Figure 4.7: Typing judgements and type formers.

Basics:

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

Figure 4.8: Term typing rules.

Product and Sums:

$$\frac{\text{TY-PAIR} \quad \Gamma \vDash e_2 : A_2 \ni \Gamma' \quad \Gamma' \vDash e_1 : A_1 \ni \Gamma''}{\Gamma \vDash (e_1, e_2) : A_1 \times A_2 \ni \Gamma''}$$

$$\frac{\text{TY-INL} \quad \Gamma \vDash e : A \ni \Gamma'}{\Gamma \vDash \text{inl } e : A + B \ni \Gamma''}$$

$$\frac{\text{TY-INR} \quad \Gamma \vDash e : B \ni \Gamma'}{\Gamma \vDash \text{inr } e : A + B \ni \Gamma''}$$

$$\text{TY-FST} \quad \Gamma, x : A_1 \times A_2 \vDash \text{fst } x : A_1 \ni \Gamma, x : \text{uncopy } A_1 \times A_2$$

$$\text{TY-SND} \quad \Gamma, x : A_1 \times A_2 \vDash \text{snd } x : A_2 \ni \Gamma, x : A_1 \times \text{uncopy } A_2$$

$$\frac{\text{TY-CASE} \quad \Gamma \vDash e_1 : A + B \ni \Gamma' \quad \Gamma' \vDash e_2 : A \multimap C \ni \Gamma'' \quad \Gamma' \vDash e_3 : B \multimap C \ni \Gamma''}{\Gamma \vDash \text{case } e_1 \ e_2 \ e_3 : C \ni \Gamma''}$$

Polymorphism:

$$\frac{\text{TY-TLAM} \quad \Gamma \vDash e : A \ni \Gamma'' \quad X \notin FV(\Gamma, \Gamma')}{\Gamma \cdot \Gamma' \vDash \lambda _. e : \forall X. A \ni \Gamma''}$$

$$\frac{\text{TY-TAPP} \quad \Gamma \vDash e : \forall X. A \ni \Gamma'}{\Gamma \vDash e () : A[K/X] \ni \Gamma''}$$

$$\frac{\text{TY-PACK} \quad \Gamma \vDash e : A[K/X] \ni \Gamma'}{\Gamma \vDash e : \exists X. A \ni \Gamma''}$$

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

References:

$$\frac{\text{TY-TOREFSHR} \quad \Gamma \vDash e : \text{ref}_{\text{uniq}}(\text{copy } A) \ni \Gamma'}{\Gamma \vDash e : \text{ref}_{\text{shr}} A \ni \Gamma''}$$

$$\frac{\text{TY-REFSHRLOAD} \quad \Gamma \vDash e : \text{ref}_{\text{shr}} A \ni \Gamma'}{\Gamma \vDash !e : A \ni \Gamma''}$$

$$\frac{\text{TY-REFSHRSTORE} \quad \Gamma \vDash e_2 : \text{copy } A \ni \Gamma' \quad \Gamma' \vDash e_1 : \text{ref}_{\text{shr}} A \ni \Gamma''}{\Gamma \vDash e_1 \leftarrow e_2 : \mathbf{1} \ni \Gamma''}$$

$$\frac{\text{TY-REFUNIQUALLOC} \quad \Gamma \vDash e : A \ni \Gamma'}{\Gamma \vDash \text{ref } e : \text{ref}_{\text{uniq}} A \ni \Gamma''}$$

$$\frac{\text{TY-REFUNIQLFREE} \quad \Gamma \vDash e : \text{ref}_{\text{uniq}} A \ni \Gamma'}{\Gamma \vDash \text{free } e : \mathbf{1} \ni \Gamma''}$$

$$\frac{\text{TY-REFUNIQLSTORE} \quad \Gamma \vDash e : B \ni \Gamma', x : \text{ref}_{\text{uniq}} A}{\Gamma \vDash x \leftarrow e : \mathbf{1} \ni \Gamma', x : \text{ref}_{\text{uniq}} B}$$

$$\text{TY-REFUNIQLLOAD} \quad \Gamma, x : \text{ref}_{\text{uniq}} A \vDash !x : A \ni \Gamma, x : \text{ref}_{\text{uniq}}(\text{uncopy } A)$$

Figure 4.9: Term typing rules (cont.)

Channels:

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

$$\begin{array}{c}
 \text{TY-CHANRECV} \\
 \Gamma, x : \text{chan } (?A.S) \models \text{recv } x : A \ni \Gamma, x : \text{chan } S
 \end{array}$$

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

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

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

Locks:

$$\begin{array}{c}
 \text{TY-MUTEXALLOC} \\
 \Gamma \models \text{newmutex} : A \rightarrow \text{mutex } A \ni \Gamma
 \end{array}$$

$$\begin{array}{c}
 \text{TY-MUTEXACQUIRE} \\
 \Gamma, x : \text{mutex } A \models \text{acquiremutex } x : A \ni \Gamma, x : \overline{\text{mutex}} A
 \end{array}$$

$$\begin{array}{c}
 \text{TY-MUTEXRELEASE} \\
 \Gamma \models e : A \ni \Gamma', x : \overline{\text{mutex}} A \\
 \hline
 \Gamma \models \text{releasemutex } x \ e : \mathbf{1} \ni \Gamma', x : \text{mutex } A
 \end{array}$$

Figure 4.10: Term typing rules (cont.)

Subtyping Properties:

$$\begin{array}{c}
 \text{SUBTY-REFL} \\
 \frac{}{K <: K}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-TRANS} \\
 \frac{K <: L \quad L <: M}{K <: M}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-BI} \\
 \frac{K <: L \quad L <: K}{K <:> L}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-BI-REFL} \\
 \frac{}{K <:> K}
 \end{array}$$

$$\begin{array}{c}
 \text{SUBTY-BI-TRANS} \\
 \frac{K <:> L \quad L <:> M}{K <:> M}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-BI-TRANS-LEFT} \\
 \frac{K <:> L \quad L <: M}{K <: M}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-BI-TRANS-RIGHT} \\
 \frac{K <: L \quad L <:> M}{K <: M}
 \end{array}$$

$$\begin{array}{c}
 \text{SUBTY-BI-SYM} \\
 \frac{L <:> K}{K <:> L}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-REC-UNFOLD} \\
 \frac{}{\mu X. K <:> K(\mu X. K)}
 \end{array}$$

Term Subtyping:

$$\begin{array}{c}
 \text{SUBTY-ANY} \\
 \frac{}{A <: \text{any}}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-LOLLI} \\
 \frac{C <: A \quad B <: D}{A \multimap B <: C \multimap D}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-ARR} \\
 \frac{C <: A \quad B <: D}{A \rightarrow B <: C \rightarrow D}
 \end{array}$$

$$\begin{array}{c}
 \text{SUBTY-PRODUCT} \\
 \frac{A <: C \quad B <: D}{A \times B <: C \times D}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-SUM} \\
 \frac{A <: C \quad B <: D}{A + B <: C + D}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-FORALL} \\
 \frac{\forall X. (A <: B)}{\forall X. A <: \forall X. B}
 \end{array}$$

$$\begin{array}{c}
 \text{SUBTY-EXIST} \\
 \frac{\forall X. (A <: B)}{\exists X. A <: \exists X. B}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-EXIST-ELIM} \\
 \frac{}{A[K/X] <: \exists X. A}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-REF-UNIQU} \\
 \frac{A <: B}{\text{ref}_{\text{uniq}} A <: \text{ref}_{\text{uniq}} B}
 \end{array}$$

$$\begin{array}{c}
 \text{SUBTY-REF-SHR} \\
 \frac{A <:> B}{\text{ref}_{\text{shr}} A <: \text{ref}_{\text{shr}} B}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-CHAN} \\
 \frac{S <: T}{\text{chan } S <: \text{chan } T}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SUBTY-MUTEX} \\
 \frac{A <:> B}{\text{mutex } A <: \text{mutex } B}
 \end{array}$$

$$\begin{array}{c}
 \text{SUBTY-MUTEXGUARD} \\
 \frac{A <:> B}{\text{mutex } A <: \text{mutex } B}
 \end{array}$$

Figure 4.11: Subtyping rules.

Copyable Types:		
$\frac{\text{SUBTY-COPY}}{A <: B}$ $\frac{}{\text{copy } A <: \text{copy } B}$	$\frac{\text{SUBTY-COPY-INTRO}}{\text{copyable } A}$ $\frac{}{A <: \text{copy } A}$	SUBTY-COPY-ELIM $\text{copy } A <: A$
$\frac{\text{SUBTY-UNCOPY}}{A <: B}$ $\frac{}{\text{uncopy } A <: \text{uncopy } B}$	$\text{SUBTY-UNCOPY-INTRO}$ $A <: \text{uncopy } A$	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
$\text{SUBTY-COPYABLE-UNIT}$ $\text{copyable } \mathbf{1}$	$\text{SUBTY-COPYABLE-BOOL}$ $\text{copyable } \mathbf{B}$	$\text{SUBTY-COPYABLE-INT}$ $\text{copyable } \mathbf{Z}$
$\frac{\text{SUBTY-COPYABLE-PRODUCT}}{\text{copyable } A \quad \text{copyable } B}$ $\frac{}{\text{copyable } (A \times B)}$	$\frac{\text{SUBTY-COPYABLE-SUM}}{\text{copyable } A \quad \text{copyable } B}$ $\frac{}{\text{copyable } (A + B)}$	
$\frac{\text{SUBTY-COPYABLE-EXISTS}}{\forall X. \text{copyable } A}$ $\frac{}{\text{copyable } (\exists X. A)}$	$\text{SUBTY-COPYABLE-REFSHR}$ $\text{copyable } (\text{ref}_{\text{shr}} X)$	
$\text{SUBTY-COPYABLE-MUTEX}$ $\text{copyable } (\text{mutex } X)$		
Context Subtyping:		
$\frac{\text{CTX-PERMUTE}}{\Gamma' \text{ is a permutation of } \Gamma}$ $\frac{}{\Gamma <:_{\text{ctx}} \Gamma'}$	CTX-REFL $\Gamma <:_{\text{ctx}} \Gamma$	$\frac{\text{CTX-TRANS}}{\Gamma_1 <:_{\text{ctx}} \Gamma_2 \quad \Gamma_2 <:_{\text{ctx}} \Gamma_3}$ $\frac{}{\Gamma_1 <:_{\text{ctx}} \Gamma_3}$
CTX-NIL $\Gamma <:_{\text{ctx}} []$	$\frac{\text{CTX-CONS}}{A <: B \quad \Gamma <:_{\text{ctx}} \Gamma'}$ $\frac{}{x : A, \Gamma <:_{\text{ctx}} x : B, \Gamma'}$	$\frac{\text{CTX-APP}}{\Gamma_1 <:_{\text{ctx}} \Gamma_2 \quad \Gamma'_1 <:_{\text{ctx}} \Gamma'_2}$ $\frac{}{\Gamma_1 \cdot \Gamma'_1 <:_{\text{ctx}} \Gamma_2 \cdot \Gamma'_2}$
CTX-COPY $x : A <:_{\text{ctx}} x : A, x : \text{uncopy } A$	$\frac{\text{CTX-COPYABLE}}{\text{copyable } A}$ $\frac{}{x : A <:_{\text{ctx}} x : A, x : A}$	

Figure 4.12: Subtyping rules (cont.)

Session Subtyping:

$$\begin{array}{c}
\text{SUBTY-SEND} \\
\frac{B <: A \quad S <: T}{!A. S <: !B. T}
\end{array}
\quad
\begin{array}{c}
\text{SUBTY-RECV} \\
\frac{A <: B \quad S <: T}{?A. S <: ?B. T}
\end{array}
\quad
\begin{array}{c}
\text{SUBTY-SEND-IN} \\
\frac{}{!(\vec{X}:\vec{k}) A. S <: !A[\vec{K}/\vec{X}]. S[\vec{K}/\vec{X}]}
\end{array}$$

$$\begin{array}{c}
\text{SUBTY-RECV-IN} \\
\frac{}{?A[\vec{K}/\vec{X}]. S[\vec{K}/\vec{X}] <: ?_{(\vec{X}:\vec{k})} A. S}
\end{array}
\quad
\begin{array}{c}
\text{SUBTY-SEND-OUT} \\
\frac{S <: !A. T}{S <: !_{(\vec{X}:\vec{k})} A. T}
\end{array}
\quad
\begin{array}{c}
\text{SUBTY-RECV-OUT} \\
\frac{?A. S <: T}{?_{(\vec{X}:\vec{k})} A. S <: T}
\end{array}$$

$$\begin{array}{c}
\text{SUBTY-SELECT} \\
\frac{\forall i. \vec{S}_i <: \vec{T}_i}{\oplus\{\vec{l}_i : \vec{S}_i\}_{i \in \vec{i}} <: \oplus\{\vec{l}_i : \vec{T}_i\}_{i \in \vec{i}}}
\end{array}
\quad
\begin{array}{c}
\text{SUBTY-SELECT-SUBSETEQ} \\
\frac{\vec{j} \subseteq \vec{i}}{\oplus\{\vec{l}_i : \vec{S}_i\}_{i \in \vec{i}} <: \oplus\{\vec{l}_j : \vec{S}_j\}_{j \in \vec{j}}}
\end{array}$$

$$\begin{array}{c}
\text{SUBTY-BRANCH} \\
\frac{\forall i. \vec{S}_i <: \vec{T}_i}{\&\{\vec{l}_i : \vec{S}_i\}_{i \in \vec{i}} <: \&\{\vec{l}_i : \vec{T}_i\}_{i \in \vec{i}}}
\end{array}
\quad
\begin{array}{c}
\text{SUBTY-BRANCH-SUBSETEQ} \\
\frac{\vec{i} \subseteq \vec{j}}{\&\{\vec{l}_i : \vec{S}_i\}_{i \in \vec{i}} <: \&\{\vec{l}_j : \vec{S}_j\}_{j \in \vec{j}}}
\end{array}$$

$$\begin{array}{c}
\text{SUBTY-SWAP-RECV-SEND} \\
?A. !B. S <: !B. ?A. S
\end{array}$$

$$\begin{array}{c}
\text{SUBTY-SWAP-BRANCH-SEND} \\
\&\{l_1 : !A. S_1, \dots, l_n : !A. S_n\} <: !A. \&\{l_1 : S_1, \dots, l_n : S_n\}
\end{array}$$

$$\begin{array}{c}
\text{SUBTY-SWAP-RECV-SELECT} \\
?A. \oplus\{l_1 : S_1, \dots, l_n : S_n\} <: \oplus\{l_1 : ?A. S_1, \dots, l_n : ?A. S_n\}
\end{array}$$

$$\begin{array}{c}
\text{SUBTY-SWAP-BRANCH-SELECT} \\
\&\{l_1 : \oplus\{l'_1 : S_{(1,1)}, \dots, l'_m : S_{(1,m)}\}, \dots, \\
\quad \dots, \\
l_n : \oplus\{l'_1 : S_{(n,1)}, \dots, l'_m : S_{(n,m)}\}\} <: \oplus\{l'_1 : \&\{l_1 : S_{(1,1)}, \dots, l_n : S_{(n,1)}\}, \\
\quad \dots, \\
l'_m : \&\{l_1 : S_{(n,1)}, \dots, l_n : S_{(n,m)}\}\}
\end{array}$$

Figure 4.13: Subtyping rules (cont.)

Duality Subtyping:

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

Append Subtyping:

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

Figure 4.14: Subtyping rules (cont.)

Bibliography

- Amal Ahmed. 2004. *Semantics of types for mutable state*. Ph.D. Dissertation. Princeton University. <https://dl.acm.org/doi/10.5555/1037736>
- 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). <https://doi.org/10.1145/1709093.1709094>
- Akka. 2021. The Akka Project. (2021). <https://akka.io/>
- Akka.NET. 2021. The Akka.NET Project. (2021). <https://getakka.net/>
- Pierre America and Jan J. M. M. Rutten. 1989. Solving Reflexive Domain Equations in a Category of Complete Metric Spaces. *JCSS* 39, 3 (1989). [https://doi.org/10.1016/0022-0000\(89\)90027-5](https://doi.org/10.1016/0022-0000(89)90027-5)
- Andrew W. Appel. 2011. Verified Software Toolchain. In *ESOP (LNCS, Vol. 7226)*. https://doi.org/10.1007/978-3-642-19718-5_1
- Andrew W. Appel. 2014. *Program Logics - for Certified Compilers*. Cambridge University Press. <http://www.cambridge.org/de/academic/subjects/computer-science/programming-languages-and-applied-logic/program-logics-certified-compilers?format=HB>
- Andrew W. Appel and David A. McAllester. 2001. An indexed model of recursive types for foundational proof-carrying code. *TOPLAS* 23, 5 (2001). <https://doi.org/10.1145/504709.504712>
- 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*. <https://doi.org/10.1145/1190216.1190235>
- Robert Atkey, Sam Lindley, and J. Garrett Morris. 2016. Conflation Confers Concurrency. In *Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*. https://doi.org/10.1007/978-3-319-30936-1_2
- Mehdi Bagherzadeh, Nicholas Fireman, Anas Shawesh, and Raffi Khatchadourian. 2020. Actor concurrency bugs: a comprehensive study on symptoms, root causes, API usages, and differences. *PACMPL* 4, OOPSLA (2020). <https://doi.org/10.1145/3428282>

- Stephanie Balzer and Frank Pfenning. 2017. Manifest Sharing with Session Types. *PACMPL* 1, ICFP (2017). <https://doi.org/10.1145/3110281>
- Stephanie Balzer, Bernardo Toninho, and Frank Pfenning. 2019. Manifest Deadlock-Freedom for Shared Session Types. In *ESOP (LNCS, Vol. 11423)*. https://doi.org/10.1007/978-3-030-17184-1_22
- Bedrock Systems A/S. 2021. The Bedrock Systems Project. (2021). <https://bedrocksystems.com/>
- Lennart Beringer, Adam Petcher, Katherine Q. Ye, and Andrew W. Appel. 2015. Verified Correctness and Security of OpenSSL HMAC. <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/beringer>
- Lars Birkedal and Aleš Bizjak. 2020. Lecture Notes on Iris: Higher-Order Concurrent Separation Logic. <https://iris-project.org/tutorial-material.html>.
- Lars Birkedal, Rasmus Ejlers Møgelberg, Jan Schwinghammer, and Kristian Støvring. 2012. First steps in synthetic guarded domain theory: step-indexing in the topos of trees. *LMCS* 8, 4 (2012). [https://doi.org/10.2168/LMCS-8\(4:1\)2012](https://doi.org/10.2168/LMCS-8(4:1)2012)
- Lars Birkedal, Kristian Støvring, and Jacob Thamsborg. 2010. The category-theoretic solution of recursive metric-space equations. *TCS* 411, 47 (2010). <https://doi.org/10.1016/j.tcs.2010.07.010>
- Ales Bizjak, Daniel Gratzer, Robbert Krebbers, and Lars Birkedal. 2019. Iron: managing obligations in higher-order concurrent separation logic. *PACMPL* 3, POPL (2019). <https://doi.org/10.1145/3290378>
- Laura Bocchi, Kohei Honda, Emilio Tuosto, and Nobuko Yoshida. 2010. A Theory of Design-by-Contract for Distributed Multiparty Interactions. In *CONCUR*. https://doi.org/10.1007/978-3-642-15375-4_12
- Michael Brandt and Fritz Henglein. 1998. Coinductive Axiomatization of Recursive Type Equality and Subtyping. *Fundamenta Informaticae* 33, 4 (1998). <https://doi.org/10.3233/FI-1998-33401>
- Mario Bravetti, Marco Carbone, Julien Lange, Nobuko Yoshida, and Gianluigi Zavattaro. 2021. A Sound Algorithm for Asynchronous Session Subtyping and its Implementation. *LMCS* 17, 1 (2021). <https://lmcs.episciences.org/7238>
- Mario Bravetti, Marco Carbone, and Gianluigi Zavattaro. 2017. Undecidability of asynchronous session subtyping. *Information and Computation* 256 (2017). <https://doi.org/10.1016/j.ic.2017.07.010>
- Stephen D. Brookes. 2004. A Semantics for Concurrent Separation Logic. In *CONCUR (LNCS, Vol. 3170)*. https://doi.org/10.1007/978-3-540-28644-8_2
- Luis Caires, Jorge A. Pérez, Frank Pfenning, and Bernardo Toninho. 2013. Behavioral Polymorphism and Parametricity in Session-Based Communication. In *ESOP (LNCS, Vol. 7792)*. https://doi.org/10.1007/978-3-642-37036-6_19

- Luís Caires and Frank Pfenning. 2010. Session Types as Intuitionistic Linear Propositions. In *CONCUR (LNCS, Vol. 6269)*. https://doi.org/10.1007/978-3-642-15375-4_16
- Marco Carbone, Fabrizio Montesi, Carsten Schürmann, and Nobuko Yoshida. 2017. Multiparty session types as coherence proofs. *Acta Informatica* 54, 3 (2017). <https://doi.org/10.1007/s00236-016-0285-y>
- David Castro, Francisco Ferreira, and Nobuko Yoshida. 2020. EMTST: Engineering the Meta-theory of Session Types. In *TACAS (LNCS, Vol. 12079)*. https://doi.org/10.1007/978-3-030-45237-7_17
- Ernie Cohen, Eyad Alkassar, Vladimir Boyarinov, Markus Dahlweid, Ulan Degenbaev, Mark A. Hillebrand, Bruno Langenstein, Dirk Leinenbach, Michal Moskal, Steven Obua, Wolfgang J. Paul, Hristo Pentchev, Elena Petrova, Thomas Santen, Norbert Schirmer, Sabine Schmaltz, Wolfram Schulte, Andrey Shadrin, Stephan Tobies, Alexandra Tsyban, and Sergey Tverdyshev. 2009. Invariants, Modularity, and Rights. In *PSI*. https://doi.org/10.1007/978-3-642-11486-1_4
- Coq Development Team. 2021. The Coq Proof Assistant. (2021). <https://coq.inria.fr>
- Andreea Costea, Wei-Ngan Chin, Shengchao Qin, and Florin Craciun. 2018. Automated Modular Verification for Relaxed Communication Protocols. In *APLAS*. https://doi.org/10.1007/978-3-030-02768-1_16
- Florin Craciun, Tibor Kiss, and Andreea Costea. 2015. Towards a Session Logic for Communication Protocols. In *ICECCS*. <https://doi.org/10.1109/ICECCS.2015.33>
- Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. 2014. TaDA: A Logic for Time and Data Abstraction. In *ECOOP*. https://doi.org/10.1007/978-3-662-44202-9_9
- Ornela Dardha and Simon J. Gay. 2018. A New Linear Logic for Deadlock-Free Session-Typed Processes. In *FOSSACS (LNCS, Vol. 10803)*. https://doi.org/10.1007/978-3-319-89366-2_5
- Ornela Dardha, Elena Giachino, and Davide Sangiorgi. 2012. Session Types Revisited. In *PPDP*. https://doi.org/10.1007/978-3-030-17184-1_22
- Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. 2015. *The Lean Theorem Prover (System Description)*. Cham. https://doi.org/10.1007/978-3-319-21401-6_26
- Jeffrey Dean and Sanjay Ghemawat. 2004. MapReduce: Simplified Data Processing on Large Clusters. In *OSDI*. <http://www.usenix.org/events/osdi04/tech/dean.html>

- Derek Dreyer, Amal Ahmed, and Lars Birkedal. 2009. Logical Step-Indexed Logical Relations. In *LICS*. <https://doi.org/10.1109/LICS.2009.34>
- Derek Dreyer, Georg Neis, Andreas Rossberg, and Lars Birkedal. 2010. A relational modal logic for higher-order stateful ADTs. In *POPL*. <https://doi.org/10.1145/1706299.1706323>
- 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/>.
- Adrian Francalanza, Julian Rathke, and Vladimiro Sassone. 2011. Permission-Based Separation Logic for Message-Passing Concurrency. *LMCS* 7, 3 (2011). [https://doi.org/10.2168/LMCS-7\(3:7\)2011](https://doi.org/10.2168/LMCS-7(3:7)2011)
- Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2018. ReLoC: A Mechanised Relational Logic for Fine-Grained Concurrency. In *LICS*. <https://doi.org/10.1145/3209108.3209174>
- Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2020. Compositional Non-Interference for Fine-Grained Concurrent Programs. To appear in S&P'21.
- Simon J. Gay. 2008. Bounded polymorphism in session types. *MSCS* 18, 5 (2008). <https://doi.org/10.1017/S0960129508006944>
- Simon J. Gay. 2016. Subtyping Supports Safe Session Substitution. In *Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*. https://doi.org/10.1007/978-3-319-30936-1_5
- Simon J. Gay and Malcolm Hole. 2005. Subtyping for session types in the pi calculus. *Acta Informatica* 42, 2-3 (2005). <https://doi.org/10.1007/s00236-005-0177-z>
- Simon J. Gay, Peter Thiemann, and Vasco T. Vasconcelos. 2020. Duality of Session Types: The Final Cut. In *PLACES (EPTCS, Vol. 314)*. <https://doi.org/10.4204/EPTCS.314.3>
- 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). <https://doi.org/10.1145/3408996>
- Liang Gu, Alexander Vaynberg, Bryan Ford, Zhong Shao, and David Costanzo. 2011. CertiKOS: a certified kernel for secure cloud computing. In *APSys '11 Asia Pacific Workshop on Systems, Shanghai, China, July 11-12, 2011*. <https://doi.org/10.1145/2103799.2103803>
- Jafar Hamin and Bart Jacobs. 2019. Transferring Obligations Through Synchronizations. In *ECOOP*. <https://doi.org/10.4230/LIPIcs.ECOOP.2019.19>

-
- Carl Hewitt, Peter Boehler Bishop, and Richard Steiger. 1973. A Universal Modular ACTOR Formalism for Artificial Intelligence. In *IJCAI*. <http://ijcai.org/Proceedings/73/Papers/027B.pdf>
- Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. 2020. Actris: Session-type based reasoning in separation logic. *PACMPL* 4, POPL (2020). <https://doi.org/10.1145/3371074>
- Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. 2021a. Actris 2.0: Asynchronous session-type based reasoning in separation logic. (2021). Manuscript under review.
- Jonas Kastberg Hinrichsen, Daniël Louwrik, Robbert Krebbers, and Jesper Bengtson. 2021b. 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>.
- Jonas Kastberg Hinrichsen, Daniël Louwrik, Robbert Krebbers, and Jesper Bengtson. 2021c. Machine-checked semantic session typing. In *CPP*. <https://doi.org/10.1145/3437992.3439914>
- C. A. R. Hoare. 1969. An Axiomatic Basis for Computer Programming. *Commun. ACM* 12, 10 (1969). <https://doi.org/10.1145/363235.363259>
- Aquinas Hobor, Andrew Appel, and Francesco Nardelli. 2008. Oracle Semantics for Concurrent Separation Logic. In *Programming Languages and Systems*. Lecture Notes in Computer Science, Vol. 4960. Springer Berlin / Heidelberg. http://dx.doi.org/10.1007/978-3-540-78739-6_27
- Kohei Honda. 1993. Types for Dyadic Interaction. In *CONCUR*. https://doi.org/10.1007/3-540-57208-2_35
- Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Discipline for Structured Communication-Based Programming. In *ESOP (LNCS, Vol. 1381)*. <https://doi.org/10.1007/BFb0053567>
- Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008. Multiparty asynchronous session types. In *POPL*. <https://doi.org/10.1145/1328438.1328472>
- Raymond Hu, Dimitrios Kouzapas, Olivier Pernet, Nobuko Yoshida, and Kohei Honda. 2010. Type-Safe Eventful Sessions in Java. In *ECOOP*. https://doi.org/10.1007/978-3-642-14107-2_16
- Keigo Imai, Nobuko Yoshida, and Shoji Yuen. 2019. Session-OCaml: A session-based library with polarities and lenses. *Science of Computer Programming* 172 (2019). <https://doi.org/10.1016/j.scico.2018.08.005>
- Iris Development Team. 2021. The Mechanisation of Iris. (2021). <https://gitlab.mpi-sws.org/iris/iris/>

- Samin S. Ishtiaq and Peter W. O’Hearn. 2001. BI as an Assertion Language for Mutable Data Structures. In *POPL*. <https://dl.acm.org/doi/10.1145/373243.375719>
- Jacques-Henri Jourdan and Robbert Krebbers. 2018. Iris Tutorial at POPL. Available online at <https://gitlab.mpi-sws.org/iris/tutorial-popl18>.
- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018a. RustBelt: Securing the Foundations of the Rust Programming Language. *PACMPL* 2, POPL (2018). <https://doi.org/10.1145/3158154>
- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2021. Safe systems programming in Rust: The promise and the challenge. To appear in CACM.
- Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-Order Ghost State. In *ICFP*. <https://doi.org/10.1145/2951913>
- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018b. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *JFP* 28 (2018). <https://doi.org/10.1017/S0956796818000151>
- 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*. <https://doi.org/10.1145/2676726.2676980>
- Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. 2009. seL4: formal verification of an OS kernel. In *SOSP*. <https://doi.org/10.1145/1629575.1629596>
- Naoki Kobayashi. 2006. A New Type System for Deadlock-Free Processes. In *CONCUR (LNCS, Vol. 4137)*. https://doi.org/10.1007/11817949_16
- Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. 1996. Linearity and the Pi-Calculus. In *POPL*. <https://doi.org/10.1145/237721.237804>
- 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). <https://doi.org/10.1145/3236772>
- Robbert Krebbers, Ralf Jung, Ales Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017a. The Essence of Higher-Order Concurrent Separation Logic. In *ESOP (LNCS, Vol. 10201)*. https://doi.org/10.1007/978-3-662-54434-1_26
- Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017b. Interactive Proofs in Higher-Order Concurrent Separation Logic. In *POPL*. <https://doi.org/10.1145/3093333.3009855>

- Neelakantan R. Krishnaswami, Aaron Turon, Derek Dreyer, and Deepak Garg. 2012. Superficially substructural types. In *ICFP*. <https://doi.org/10.1145/2364527.2364536>
- Morten Krogh-Jespersen, Kasper Svendsen, and Lars Birkedal. 2017. A relational model of types-and-effects in higher-order concurrent separation logic. In *POPL*. <https://doi.org/10.1145/3093333.3009877>
- Morten Krogh-Jespersen, Amin Timany, Marit Edna Ohlenbusch, Simon Oddershede Gregersen, and Lars Birkedal. 2020. Aneris: A Mechanised Logic for Modular Reasoning about Distributed Systems. In *ESOP*. https://doi.org/10.1007/978-3-030-44914-8_13
- Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: a verified implementation of ML. In *POPL*. <https://doi.org/10.1145/2535838.2535841>
- Leslie Lamport. 1977. Proving the Correctness of Multiprocess Programs. *IEEE TSE* 3, 2 (1977). <https://doi.org/10.1109/TSE.1977.229904>
- Julien Lange, Nicholas Ng, Bernardo Toninho, and Nobuko Yoshida. 2018. A static verification framework for message passing in Go using behavioural types. In *ICSE*. <https://doi.org/10.1145/3180155.3180157>
- Xavier Leroy. 2006. Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In *POPL*. <https://doi.org/10.1145/1111037.1111042>
- Étienne Lozes and Jules Villard. 2012. Shared Contract-Obedient Endpoints. In *ICE*. <https://doi.org/10.4204/EPTCS.104.3>
- William Mansky, Andrew W. Appel, and Aleksey Nogin. 2017. A verified messaging system. *PACMPL* 1, OOPSLA (2017). <https://doi.org/10.1145/3133911>
- Claude Marché, Christine Paulin-Mohring, and Xavier Urbain. 2004. The KRAKA-TOA tool for certification of JAVA/JAVACARD programs annotated in JML. *JLAMP* 58, 1-2 (2004). <https://doi.org/10.1016/j.jlap.2003.07.006>
- Robin Milner. 1978. A Theory of Type Polymorphism in Programming. *JCSS* 17, 3 (1978). [https://doi.org/10.1016/0022-0000\(78\)90014-4](https://doi.org/10.1016/0022-0000(78)90014-4)
- Dimitris Mostrous and Vasco Thudichum Vasconcelos. 2014. Affine Sessions. In *COORDINATION*. https://doi.org/10.1007/978-3-662-43376-8_8
- Dimitris Mostrous and Nobuko Yoshida. 2015. Session typing and asynchronous subtyping for the higher-order π -calculus. *Information and Computation* 241 (2015). <https://doi.org/10.1016/j.ic.2015.02.002>
- Dimitris Mostrous, Nobuko Yoshida, and Kohei Honda. 2009. Global Principal Typing in Partially Commutative Asynchronous Sessions. In *ESOP (LNCS, Vol. 5502)*. https://doi.org/10.1007/978-3-642-00590-9_23

- Hiroshi Nakano. 2000. A Modality for Recursion. In *LICS*. <https://doi.org/10.1109/LICS.2000.855774>
- Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. 2014. Communicating State Transition Systems for Fine-Grained Concurrent Resources. In *ESOP*. https://doi.org/10.1007/978-3-642-54833-8_16
- Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. 2002. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*. LNCS, Vol. 2283. <https://doi.org/10.1007/3-540-45949-9>
- Peter W. O’Hearn. 2004. Resources, Concurrency and Local Reasoning. In *CONCUR*. https://doi.org/10.1007/978-3-540-28644-8_4
- Kosuke Ono, Yoichi Hirai, Yoshinori Tanabe, Natsuko Noda, and Masami Hagiya. 2011. Using Coq in Specification and Program Extraction of Hadoop MapReduce Applications. In *SEFM*. https://doi.org/10.1007/978-3-642-24690-6_24
- Wytse Oortwijn, Stefan Blom, and Marieke Huisman. 2016. Future-based Static Analysis of Message Passing Programs. In *PLACES*. <https://doi.org/10.4204/EPTCS.211.7>
- Luca Padovani. 2014. Deadlock and lock freedom in the linear π -calculus. In *CSL-LICS*. <https://doi.org/10.1145/2603088.2603116>
- Luca Padovani. 2017. A simple library implementation of binary sessions. *JFP* 27 (2017). <https://doi.org/10.1017/S0956796816000289>
- Jorge A. Pérez, Luís Caires, Frank Pfenning, and Bernardo Toninho. 2012. Linear Logical Relations for Session-Based Concurrency. In *ESOP (LNCS, Vol. 7211)*. https://doi.org/10.1007/978-3-642-28869-2_27
- 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). <https://doi.org/10.1016/j.ic.2014.08.001>
- Frank Pfenning and Conal Elliott. 1988. Higher-Order Abstract Syntax. In *PLDI*. <https://doi.org/10.1145/53990.54010>
- Benjamin C. Pierce et al. 2020. Programming Language Foundations. <https://softwarefoundations.cis.upenn.edu/plf-current/index.html>
- John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *LICS*. <https://doi.org/10.1109/LICS.2002.1029817>
- Arjen Rouvoet, Casper Bach Poulsen, Robbert Krebbers, and Eelco Visser. 2020. Intrinsically-typed definitional interpreters for linear, session-typed languages. In *CPP*. ACM. <https://doi.org/10.1145/3372885.3373818>

- Ilya Sergey, James R. Wilcox, and Zachary Tatlock. 2018. Programming and proving with distributed protocols. *PACMPL* 2, POPL (2018). <https://doi.org/10.1145/3158116>
- Kasper Svendsen and Lars Birkedal. 2014. Impredicative Concurrent Abstract Predicates. In *ESOP*. https://doi.org/10.1007/978-3-642-54833-8_9
- Kasper Svendsen, Lars Birkedal, and Matthew J. Parkinson. 2010. Verifying Generics and Delegates. In *ECOOP*. https://doi.org/10.1007/978-3-642-14107-2_9
- Kasper Svendsen, Filip Sieczkowski, and Lars Birkedal. 2016. Transfinite Step-Indexing: Decoupling Concrete and Logical Steps. In *ESOP*. https://doi.org/10.1007/978-3-662-49498-1_28
- Nikhil Swamy, Aseem Rastogi, Aymeric Fromherz, Denis Merigoux, Danel Ahman, and Guido Martínez. 2020. SteelCore: an extensible concurrent separation logic for effectful dependently typed programs. *PACMPL* 4, ICFP (2020). <https://doi.org/10.1145/3409003>
- David Swasey, Deepak Garg, and Derek Dreyer. 2017. Robust and compositional verification of object capability patterns. *PACMPL* 1, OOPSLA (2017). <https://doi.org/10.1145/3133913>
- Samira Tasharofi, Peter Dinges, and Ralph E. Johnson. 2013. Why Do Scala Developers Mix the Actor Model with other Concurrency Models?. In *ECOOP*. https://doi.org/10.1007/978-3-642-39038-8_13
- Joseph Tassarotti, Ralf Jung, and Robert Harper. 2017. A Higher-Order Logic for Concurrent Termination-Preserving Refinement. In *ESOP (LNCS, Vol. 10201)*. https://doi.org/10.1007/978-3-662-54434-1_34
- The CertiK Team. 2021. The CertiK Project. (2021). <https://certik.io/>
- The Coq-std++ Team. 2020. An extended “standard library” for Coq. Available online at <https://gitlab.mpi-sws.org/iris/stdpp>.
- The Erlang Team. 2021. The Erlang Language Project. (2021). <https://www.erlang.org/>
- The Go Team. 2021. The Go Language Project. (2021). <https://golang.org/>
- Peter Thiemann. 2019. Intrinsically-Typed Mechanized Semantics for Session Types. In *PPDP*. <https://doi.org/10.1145/3354166.3354184>
- Peter Thiemann and Vasco T. Vasconcelos. 2020. Label-dependent session types. *PACMPL* 4, POPL (2020). <https://doi.org/10.1145/3371135>
- 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). <https://doi.org/10.1145/3158152>

- Tengfei Tu, Xiaoyu Liu, Linhai Song, and Yiyang Zhang. 2019. Understanding Real-World Concurrency Bugs in Go. In *ASPLOS*. <https://doi.org/10.1145/3297858.3304069>
- Jules Villard, Étienne Lozes, and Cristiano Calcagno. 2009. Proving Copyless Message Passing. In *APLAS*. https://doi.org/10.1007/978-3-642-10672-9_15
- Philip Wadler. 2012. Propositions as sessions. In *ICFP*. <https://doi.org/10.1145/2364527.2364568>
- Andrew K. Wright. 1995. Simple Imperative Polymorphism. *Lisp and Symbolic Computation* 8, 4 (1995). <https://dl.acm.org/doi/10.1007/BF01018828>
- Andrew K. Wright and Matthias Felleisen. 1994. A Syntactic Approach to Type Soundness. *Information and Computation* 115, 1 (1994). <https://doi.org/10.1006/inco.1994.1093>