

Deadlock-Free Asynchronous Message Reordering in Rust with Multiparty Session Types

Zak Cutner

Imperial College London
London, UK

Nobuko Yoshida

Imperial College London
London, UK

Martin Vassor

Imperial College London
London, UK

Abstract

Rust is a modern systems language focused on performance and reliability. Complementing Rust’s promise to provide “fearless concurrency”, developers frequently exploit asynchronous message passing. Unfortunately, sending and receiving messages in an arbitrary order to maximise computation-communication overlap (a popular optimisation in message-passing applications) opens up a Pandora’s box of subtle concurrency bugs.

To guarantee deadlock-freedom by construction, we present RUMPSTEAK: a new Rust framework based on *multiparty session types*. Previous session type implementations in Rust are either built upon synchronous and blocking communication and/or are limited to two-party interactions. Crucially, none support the arbitrary ordering of messages for efficiency.

RUMPSTEAK instead targets asynchronous `async/await` code. Its unique ability is allowing developers to arbitrarily order send/receive messages while preserving deadlock-freedom. For this, RUMPSTEAK incorporates two recent advanced session type theories: (1) k -multiparty compatibility (k -MC), which *globally* verifies the safety of a set of participants, and (2) asynchronous multiparty session subtyping, which *locally* verifies optimisations in the context of a single participant. Specifically, we propose a novel algorithm for asynchronous subtyping that is both sound and decidable.

We first evaluate the performance and expressiveness of RUMPSTEAK against three previous Rust implementations. We discover that RUMPSTEAK is around 1.7–8.6x more efficient and can safely express many more examples by virtue of offering arbitrary ordering of messages. Secondly, we analyse the complexity of our new algorithm and benchmark it against k -MC and a *binary* session subtyping algorithm. We find they are exponentially slower than RUMPSTEAK’s.

CCS Concepts: • Software and its engineering → Development frameworks and environments; Source code generation; • Computer systems organization → Reliability.

Keywords: Rust, Asynchronous Message Passing, Message Reordering, Computation-Communication Overlap, Multiparty Session Types

1 Introduction

Rust is a statically-typed language designed for systems software development. It is rapidly growing in popularity and has been voted “most loved language” over five years of surveys by Stack Overflow [19]. Rust aims to offer the safety of a high-level language without compromising on the performance enjoyed by low-level languages. *Message passing over typed channels* is common in concurrent Rust applications, where (low-level) threads or (high-level) actors communicate efficiently and safely by sending messages containing data.

To improve performance by maximising computation-communication overlap [14, 37, 56], developers often wish to arbitrarily change the order of sending and receiving messages—we will present several examples of this technique, which we refer to as *asynchronous message reordering* (AMR). Our challenge is to remedy communication errors such as deadlocks, which can easily occur in message-passing applications, particularly those that leverage AMR.

To achieve this, we introduce RUMPSTEAK: a framework for efficiently coordinating message-passing processes in Rust using *multiparty session types*. Session types [28, 57] (see [62] for a gentle introduction and [23] for a more exhaustive one) coordinate interactions through *linearly typed channels*, that must be used exactly once, ensuring *protocol compliance* without deadlocks or communication mismatches.

Current state of the art. Since Rust’s *affine type system* is particularly well-suited to session types by statically guaranteeing a linear usage of session channels, there are several previous attempts at implementing session types in Rust [12, 34, 38, 41]. However, their current limitations prevent them from guaranteeing all four of *deadlock-freedom*, *multiparty communication*, *asynchronous execution* and AMR.

Our framework. We motivate the importance of each feature and explain how RUMPSTEAK incorporates this.

Deadlock-freedom. One of the most important properties of concurrent/parallel systems is that their computations

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
PPoPP ’22, April 2–6, 2022, Seoul, Republic of Korea

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9204-4/22/02...\$15.00

<https://doi.org/10.1145/3503221.3508404>

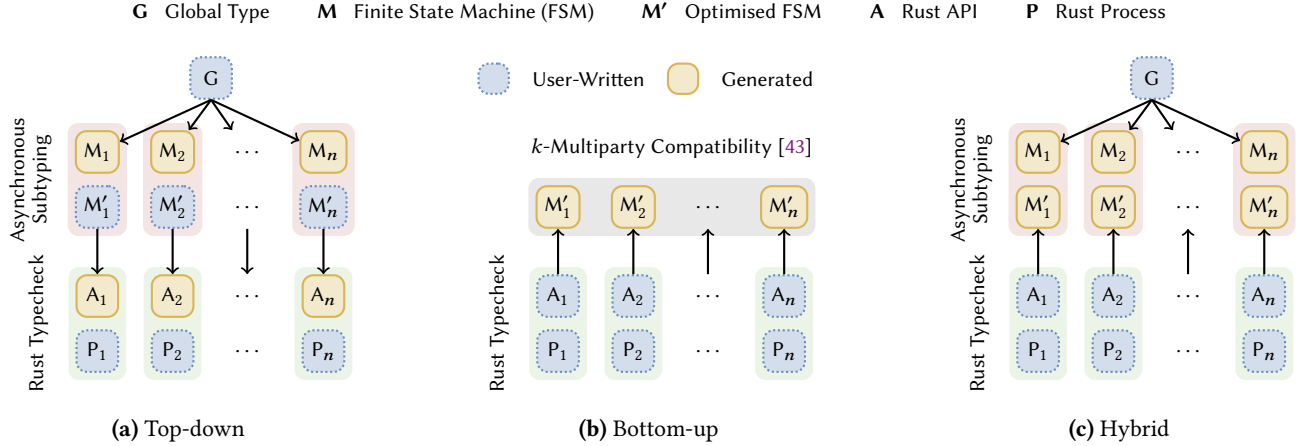


Figure 1. Workflow of the RUMPSTEAK framework (three approaches).

are not blocked. Deadlock-freedom in our context states that the system can always either make progress by exchanging messages or properly terminate.

Multipart communication. Many previous Rust implementations [12, 34, 38] support only *binary* session types, which is limited to two-party communication. On the other hand, the majority of real-world communication protocols consist of more than two participants. RUMPSTEAK therefore uses *multipart* session types (MPST) [29, 30] to ensure deadlock-freedom in protocols with any number of participants (or roles).

Asynchronous execution. Most previous Rust implementations [34, 38, 41] use *synchronous* communication channels. This approach suffers from performance limitations since threads are *blocked* while waiting to receive messages. RUMPSTEAK instead uses *asynchronous* communication, where lightweight asynchronous tasks share a pool of threads. When one task is blocked, another’s work can be scheduled in the meantime to prevent the wasting of computational resources.

Although [12] is also based on asynchronous communication, only RUMPSTEAK closely integrates with Rust’s modern `async/await` syntax, allowing asynchronous programs to be written sequentially. To achieve this, asynchronous functions are annotated with `async`, causing them to return *futures*. Developers can `await` calls to these functions, denoting that execution should continue elsewhere until a result is ready.

Asynchronous message reordering. Our main contribution is offering AMR, which no existing work can provide while preserving deadlock-freedom. To motivate this, we introduce a running example of the *double buffering* protocol [33]. Buffering is frequently used in multimedia applications where a continuous stream of data must be sent from a source (e.g. a graphics card) to a sink (e.g. a CPU). To prevent the source from being blocked while the sink is busy, it writes to a buffer, which is later read by the sink.




Figure 2. Illustration of the double buffering protocol.

We illustrate this effect in Fig. 2. With a single buffer, both the source and sink are constantly blocked as we cannot read and write simultaneously. However, adding a second buffer allows the source and sink to operate on different buffers at once so that (in the best case) they are never blocked and throughput can increase twofold.

As we will see in § 2, the double buffering protocol takes advantage of AMR so it *cannot* be expressed with standard MPST theory [29]; communication with the source and sink are overlapped so both buffers can be accessed at once. The goal of this paper is to provide a method for ensuring that protocols optimised in this way preserve deadlock-freedom.

RUMPSTEAK framework. We achieve these features while allowing three different approaches, summarised in Fig. 1.

In the **top-down approach**, the developer writes a global type, which describes the entire protocol (see § 2 for an example). We project this onto each participant to obtain a local finite state machine (FSM), which describes the protocol from that participant’s perspective. Next, the developer uses AMR to propose optimised FSMs for each participant. We *locally* verify that each optimised FSM is compatible with the projected one using an asynchronous subtyping algorithm. From each optimised FSM, we generate an API so that the developer can write a Rust process implementation. Our

API uses Rust’s type checker to ensure that these processes conform to the protocol, and are therefore deadlock-free.

In the **bottom-up approach**, the developer manually writes an API and process implementation for each participant. We derive the optimised FSMs directly from these APIs and ensure that they are safe using *k*-multiparty compatibility (*k*-MC) [43]. *k*-MC globally verifies that multiple FSMs are compatible with each other without using a global type.

Finally, we propose a **hybrid approach** where the projected FSMs are generated from a global type on one side (as in the top-down approach). On the other side, the developer writes both the APIs and process implementations directly, and we derive the optimised FSMs (as in the bottom-up approach). We then *locally* verify that the optimised FSMs are asynchronous subtypes of the projected FSMs.

To summarise, our subtyping algorithm used in the top-down/hybrid approaches *locally* verifies the correctness of optimisations. This makes it far more scalable than *k*-MC as we will see in § 4. These approaches are also more intuitive to the developer since they use *safety by construction*—it is much easier to write a global type than to determine why a *k*-MC analysis has failed for a complex protocol.

Contribution and outline. In § 2, we give an overview of the design and implementation of RUMPSTEAK. In § 3, we present our new *asynchronous subtyping algorithm* used to check that an optimised FSM is correct. We prove our new algorithm is *sound* (Theorem 7) against the precise MPST asynchronous subtyping theory [25], *terminates* (Theorem 6) and analyse its complexity (Theorem 9). In § 4, we compare the performance and expressiveness of our framework with existing work [7, 12, 38, 41, 43]. We show (1) RUMPSTEAK’s runtime is faster and can express many more asynchronous protocols than other Rust implementations [12, 38, 41]; and (2) our subtyping algorithm is more efficient than existing algorithms [7, 43], confirming our complexity analysis. § 5 discusses related work and concludes. The **supplementary materials** available in the full version [17] contain further examples and proofs of the theorems. We include our source code and benchmarks in a **public GitHub repository** [3].

2 Design and Implementation

This section presents RUMPSTEAK and explains its three workflows: a *top-down approach* using asynchronous subtyping, a *bottom-up approach* using *k*-multiparty compatibility (*k*-MC) [43] and a *hybrid approach* combining the two.

2.1 Top-Down Approach

All three approaches result in the same final application, so we use the example of the top-down workflow (Fig. 1a) to give a general overview of RUMPSTEAK’s implementation.

Two-party protocol. Before going into the details of the top-down approach, we first give an informal insight into

how the protocol is represented at each stage, shown in Fig. 3. For this, we use the example of a simple streaming protocol, later introduced in § 4.

A **global session type** describes the protocol from a global perspective and includes all participants in the protocol. The global type for the streaming protocol (shown below) is read as follows: participant *t* first sends *ready* to participant *s*. Next, *s* replies to *t* with two possible messages: either *value* or *stop*. In the latter case, the protocol terminates (type end). In the former case, the protocol continues with *x*, which is a type variable that is bound by $\mu x.$, i.e. the protocol starts over.

$$G_{ST} = \mu x.t \rightarrow s : \{ready.s \rightarrow t : \{value.x, \quad stop.end\}\}$$

In RUMPSTEAK, developers express global types syntactically with Scribble [55, 63]: a widely used and target-agnostic language for describing multiparty protocols. We show the corresponding Scribble description for the streaming protocol in Fig. 3a.

A **local finite state machine** describes the protocol from the perspective of a single participant. It shows the send and receive actions of that participant, independent of what other participants may be doing in the meantime.

In Fig. 3b, we show the local FSMs for each participant in the streaming protocol, where ! and ? denote *send* and *receive* respectively in session type syntax [62]. For example, the FSM for the source first receives *ready* from the sink, then chooses to either send *value* and start over or simply send *stop* and finish.

The Rust API is an encoding of a local FSM as Rust code. This code then uses Rust’s type checker to confirm that a process implementation, also written in Rust, conforms to the FSM. For example, the encodings for the source and sink are shown in Fig. 3c.

Multiparty protocol. In the remainder of this section, we use a more complex example (the double buffering protocol from § 1), which allows us to illustrate how RUMPSTEAK can verify multiparty protocols. We carefully go through each step of the top-down approach, closely following Fig. 1a.

To guarantee safety in our double buffering example using MPST, the developer defines a global type G_{DB} for the protocol. This protocol has three participants: a source *s*, a kernel *k* (which controls both of the buffers) and a sink *t*. We also show the corresponding Scribble description for the double buffering protocol in Listing 1.

$$G_{DB} = \mu x.k \rightarrow s : \{ready.s \rightarrow k : \{value.t \rightarrow k : \{ready.k \rightarrow t : \{value.x\}\}\}\}$$


Projection. Once we have created a global type, we use a projection algorithm to derive the local FSM for each participant. In RUMPSTEAK, we perform projection using νSCR [2]: a new lightweight and extensible Scribble toolchain implemented in OCaml.

```

1 global protocol Ring(role s,
2   role t) {
3   rec loop {
4     ready() from t to s;
5     choice at s {
6       value() from s to t;
7       continue loop;
8     } or {
9       stop() from s to t;
10    }
11  }
12 }

```

(a) Scribble protocol description



(b) FSMs for the source and sink

```

1 type Source = Receive<T, Ready,
2   Select<T, SourceChoice>>;
3
4 enum SourceChoice {
5   Value(Value, Source),
6   Stop(Stop, End) }
7
8 type Sink = Send<S, Ready,
9   Branch<S, SinkChoice>>;
10
11 enum SinkChoice {
12   Value(Value, Sink),
13   Stop(Stop, End) }

```

(c) Rust API code (simplified excerpt)

Figure 3. Different session type representations used within RUMPSTEAK.

Listing 1. Scribble representation of the double buffering protocol.

```

1 global protocol DoubleBuffering(role s,
2   role k, role t) {
3   rec loop {
4     ready() from k to s;
5     value() from s to k;
6     ready() from t to k;
7     value() from k to t;
8     continue loop;
9   }
10 }

```





(a) Projected FSM for k (M_k)(b) Optimised FSM for k (M'_k)(c) Projected FSM for s (M_s)(d) Projected FSM for t (M_t)

Figure 4. FSMs for roles k, t and s of the double buffering protocol.

We show the projected local FSM for each of the participants in Fig. 4. For instance, the kernel, whose projected type M_k is shown in Fig. 4a, (1) informs the source that it is ready to receive; (2) receives a value from the source; (3) waits for the sink to become ready; (4) sends a value to the sink; and (5) repeats from step 1.

Asynchronous message reordering. Unfortunately, global types cannot represent overlapping communication. Therefore, the projection M_k cannot achieve the simultaneous interactions we saw in Fig. 2. In practice, developers use AMR to produce an optimised FSM for the kernel M'_k (shown in Fig. 4b). Here, the kernel initially sends two *ready* messages to the source. This allows the source to write to the second buffer while the sink is busy reading from the first. Note that the asynchronous queue is effectively acting as the second buffer—the second message from the source waits in the queue while the kernel is processing the first.

Crucially, RUMPSTEAK verifies that M'_k is safe to use in the place of M_k without causing deadlocks or other communication errors. Otherwise, we would be throwing away the safety benefits that come with MPST. To achieve this, we must determine that M'_k is an *asynchronous subtype* of M_k , which we perform using our algorithm presented in § 3.

API design. Once we arrive at an optimised FSM M'_i for each participant, our challenge is to create a Rust API A_i , which uses Rust’s type checker to ensure that a developer-written process P_i conforms to M'_i (a relevant introduction to the Rust programming language can be found in Jung’s thesis [36, Chapter 8]). In the top-down approach, this API is automatically generated from an optimised FSM.

To illustrate, we show the API for the kernel (A_k) in Listing 2, which checks conformance to M'_k . To ensure that our API remains readable by developers and to eliminate extensive boilerplate code, we make use of Rust’s procedural macros [58]. By decorating types with $\#[\dots]$, these macros perform additional compile-time code generation.

Roles (Participants). Each role is represented as a struct, which stores its communication channels with other roles. The struct for the kernel (ls. 1 to 6) contains channels to and from s and t. Developers can in fact use any custom channel that implements Rust’s standard Sink or Stream interfaces [15] for asynchronous sends and receives respectively. This approach minimises the expensive creation of channels

Listing 2. RUMPSTEAK API for the kernel.

```

1  #[derive(Role)]
2  #[message(Label)]
3  struct K(
4      #[route(S)] Channel,
5      #[route(T)] Channel
6  );
7
8  #[derive(Message)]
9  enum Label {
10     Ready(Ready),
11     Value(Value),
12 }
13
14 struct Ready;
15 struct Value(i32);
16
17 #[session]
18 type Kernel = Send<S, Ready, KernelLoop>;
19
20 #[session]
21 struct KernelLoop(Send<S, Ready, Receive<S, Value,
22     Receive<T, Ready, Send<T, Value, Self>>>>);

```

in cases where only bounded or unidirectional channels are required.

Our `#[derive(Role)]` macro generates methods for programmatically retrieving these channels from the struct. Moreover, an optional `#[derive(Roles)]` macro can be applied on a struct containing every role to also generate code for automatically instantiating all roles at once. This will create the necessary combination of channels and assign them to the correct structs in order to reduce human error.

Sessions. Following the approach of MULTICRUSTY [41], we build a set of *generic primitives* to construct a simple API. For instance, the `Send` primitive (l. 21) takes a role, label and continuation as generic parameters. In contrast to the standard approach of creating a type for every state [31], these primitives reduce the number of types required and avoid arbitrary naming of these ‘state types’. For brevity, our API elides two additional parameters used to store channels at runtime, which are reinserted with the `#[session]` macro.

Labels. Internally, RUMPSTEAK sends a `Label` enum (l. 9) over reusable channels to communicate with other participants. Each label is represented as a type (ls. 14 to 15) and our `#[derive(Message)]` macro generates methods for converting to and from the `Label` enum.

Choice. The kernel’s API does not contain choice but we show an example of this below. `Choice` is represented as an enum where each branch contains the label sent/received and a continuation. This design allows processes to pattern match on external choices to determine which label was received. Methods allowing the enum to be used with `Branch` (an *external choice*, l. 3) or `Select` (an *internal choice*) are also generated with the `#[session]` macro. Even so, the enum is

Listing 3. Process implementation for the kernel.

```

1  async fn kernel(role: &mut K) -> Result<> {
2      try_session(role, |s: Kernel<_, _>| async {
3          let mut t = s.send(Ready).await?;
4          loop {
5              let s = t.into_session().send(Ready).await?;
6              let (Value(value), s) = s.receive().await?;
7              let (Ready, s) = s.receive().await?;
8              t = s.send(Value(value)).await?;
9          }
10     }).await
11 }

```

still necessary since Rust’s lack of variadic generics means choice cannot be easily implemented as its own primitive.

```

1  #[session]
2  enum Choice {
3      Continue(Continue, Branch<A, Self>),
4      Stop(Stop, End),
5  }

```

Process implementation. Using the API A_k , the developer writes an implementation for the process P_k ; we show an example in Listing 3. We discuss how our API uses Rust’s type system to check that P_k conforms to the protocol.

Linearity. The linear usage of channels is checked by Rust’s *affine type system*, preventing channels from being used multiple times. When a primitive is executed, it consumes itself, preventing reuse, and returns its continuation.

Developers are prevented from constructing primitives directly using visibility modifiers. They must instead use `try_session` (l. 2). This function accepts a reference to the role being implemented and a closure (or function pointer) to its process implementation. The closure takes the input session type as an argument and returns the terminal type `End`. If a session is discarded, thereby breaking linearity, then the developer will have no `End` to return and Rust’s compiler will complain that the closure does not satisfy its type.

Infinite recursion. At first glance, it seems impossible to implement processes with infinitely recursive types using this approach. Nevertheless, for types without an `End` primitive, such as `Kernel`, we can use an infinite loop (l. 4) to get around this problem. Conveniently, infinite loops are assigned `!`, Rust’s never (or bottom) type, which can be implicitly cast to any other type. This allows the type of the loop to be coerced to `End`, enabling the closure to pass the type checker as before.

Channel reuse. We allow roles to be reused across sessions since the channels they contain are usually expensive to create. Crucially, to prevent communication mismatches between different sessions, we must ensure that the same role is not used in multiple sessions at once. Therefore, `try_session` takes an *exclusive* reference to the role, causing Rust’s borrow checker to enforce this requirement.

2.2 Bottom-Up Approach

In the top-down approach, we generate a Rust API A_i from an optimised FSM M'_i . In the bottom-up approach (Fig. 1b), we do the reverse: we *serialise* each API A_i to obtain an FSM M'_i . Next, we use k -MC on the set of FSMs $M'_{1..n}$. If they are indeed compatible, then the processes $P_{1..n}$, which implement their respective APIs, are free from deadlocks.

k -MC takes the optimised FSMs of all participants and verifies deadlock freedom. In contrast, asynchronous subtyping checks the optimisation of a single participant's FSM in isolation. Therefore, k -MC can be seen as a global analysis of the protocol and asynchronous subtyping as a local analysis of a single participant.

To perform the serialisation of an API to an FSM we provide a Rust function `serialize<S>() -> Fsm` (this is a simplified version). It takes a session type API (such as Kernel from § 2.1) as a generic type parameter S and returns its corresponding FSM. This FSM can be printed in a variety of formats and passed into the k -MC tool for verification.

2.3 Hybrid Approach

Developers may naturally prefer the bottom-up approach since code generation, as used in the top-down approach, can be quite opaque and difficult to understand or debug. Nevertheless, the top-down approach has the advantage of using asynchronous subtyping rather than k -MC—analysing the local types for all participants in the protocol at once is challenging to do scalably (see § 4).

Moreover, when a k -MC analysis fails, it can be difficult to determine how a developer should update a complex protocol to make it free from deadlocks. Safety by construction, as used in the top-down approach, is easier to work with since verification is done locally on each participant.

Therefore, we also propose a third, hybrid approach (Fig. 1c). In this workflow, a global type G is provided by the developer and projected to obtain the FSMs $M_{1..n}$ as before. Rather than the developer proposing the optimised FSMs $M'_{1..n}$ directly, they instead simply write the APIs $A_{1..n}$ (as in the bottom-up approach). These are serialised to $M'_{1..n}$ which can (as in the top-down approach) be checked for safety against $M_{1..n}$ using asynchronous subtyping. In essence, the hybrid approach uses the same theory as the top-down approach, but presents a more developer-friendly interface that uses serialisation rather than code generation.

3 A Sound Asynchronous Multiparty Session Subtyping Algorithm

This section proposes an algorithm for asynchronous multiparty subtyping (§ 3.2), shows its soundness (Theorem 7), termination (Theorem 6) and complexity (Theorem 9), and sketches its implementation.

We begin by formally defining global and local session types. In the remainder of this section, we will omit sorts S from the syntax to simplify the presentation.

Definition 1 (Global and local types).

$$\begin{aligned} S &::= i32 \mid u32 \mid i64 \mid u64 \mid \dots \\ G &::= \text{end} \mid p \rightarrow q : \{\ell_i(S_i).G_i\}_{i \in I} \mid \mu t. G \mid t \\ T &::= \text{end} \mid \oplus_{i \in I} p! \ell_i(S_i).T_i \mid \&_{i \in I} p? \ell_i(S_i).T_i \mid \mu t. T \mid t \end{aligned}$$

$\oplus_{i \in I} p! \ell_i(S_i).T_i$ and $\&_{i \in I} p? \ell_i(S_i).T_i$ represent *internal* and *external* choices respectively where $!$ and $?$ denote send and receive respectively, and all ℓ_i are pairwise distinct.

Defining sound message reordering is non-trivial since it may introduce deadlocks, as shown in the example below.

Example 2 (Correct/incorrect AMR). Consider the following local types

$$T_Q = p? \ell_1. p! \ell_2. \text{end} \quad T_P = q! \ell_1. q? \ell_2. \text{end}$$

which are given by projecting a global type

$$p \rightarrow q : \{\ell_1. q \rightarrow p : \{\ell_2. \text{end}\}\}$$

Reordering the actions of q to become $T'_Q = p! \ell_2. p? \ell_1. \text{end}$, first sending before receiving, retains deadlock-freedom as messages are stored in queues and their reception can be delayed. However, if instead we reorder p 's interactions to become $T'_P = q? \ell_2. q! \ell_1. \text{end}$, we arrive at a deadlock since both processes are simultaneously expecting to receive a message that has not yet been sent.

3.1 Precise Asynchronous Multiparty Session Subtyping

Multiparty session subtyping formulates *refinement* for communication protocols with more than two communicating processes. A process implementing a session type T can be safely used whenever a process implementing one of its supertypes T' is expected. This applies in any context, ensuring that no deadlocks or other communication errors will be introduced. Replacing the implementation of T' with that of the subtype T may allow for more *optimised* communication patterns as we have seen. Ghilezan et al. [25] present the *precise* asynchronous subtyping relation \leq for multiparty session processes. Precision is characterised by both *soundness* (safe process replacement is guaranteed) and *completeness* (any extension of the relation is unsound).

Crucially, asynchronous subtyping supports the optimisation of communications. Under certain conditions, the subtype can anticipate some input/output actions occurring in the supertype, performing them *earlier* than prescribed to achieve the most flexible and precise subtyping. Such reorderings can take two forms:

- R1.** anticipating an input from participant p before a finite number of inputs that are not from p ; or

R2. anticipating an output to participant p before a finite number of inputs (from any participant), and also before other outputs that are not to p .

To denote such reorderings, two kinds of finite sequences of inputs/outputs are defined by [25], where $p \neq q$. $\mathcal{A}^{(p)}$ is a sequence containing receives from participants apart from p ; $\mathcal{B}^{(p)}$ is a sequence containing receives from *any* participant ($r = p$ is allowed) and sends to participants apart from p .

$$\mathcal{A}^{(p)} ::= q?l \mid q!l. \mathcal{A}^{(p)} \quad \mathcal{B}^{(p)} ::= r?l \mid q!l \mid r?l. \mathcal{B}^{(p)} \mid q!l. \mathcal{B}^{(p)}$$

The *tree refinement relation* \lesssim is defined coinductively on infinite session type trees that contain only single-inputs (SI) and single-outputs (SO). We leave the formal definition of \lesssim in [25] and shall explain its essence with our sound algorithm. Based on \lesssim , the subtyping relation \leq for all types (including internal and external choice) is given as

$$\frac{\forall U \in \llbracket T \rrbracket_{\text{so}} \quad \forall V' \in \llbracket T' \rrbracket_{\text{si}} \quad \exists W \in \llbracket U \rrbracket_{\text{si}} \quad \exists W' \in \llbracket V' \rrbracket_{\text{so}} \quad W \lesssim W'}{T \leq T'}$$

where $\llbracket T \rrbracket_{\text{so}}$ (resp. $\llbracket T \rrbracket_{\text{si}}$) is the minimal set of trees containing only single *outputs* (resp. *inputs*) of the session type tree T . Using existential quantifiers for $\llbracket U \rrbracket_{\text{si}}$ and $\llbracket V' \rrbracket_{\text{so}}$ allows external choices to be added and internal choices to be removed (see the Appendix B.2.1 for examples of \leq).

3.2 Our Algorithm

The precise asynchronous subtyping in [25] is *undecidable*—even when limited to two participants, as proven in [42]. This subsection introduces our *practical* algorithm which is *sound* and *terminates*, through the use of a bound on recursion.

Prefixes. We first define reduction rules for finite single-input and single-output (SISO) session type *prefixes*.

Definition 3 (Prefix reduction). Let us define the syntax of *prefixes* as $\pi, \rho ::= \epsilon \mid p!l(S) \mid p?l(S) \mid \pi_1.\pi_2$ where “.” denotes the concatenation operator. We define the reduction $\pi \rightarrow \pi'$ as the smallest relation that ensures

$$\begin{aligned} [\rightarrow\text{I}] \quad & \langle p?l.\pi \parallel p?l.\pi' \rangle \rightarrow \langle \pi \parallel \pi' \rangle \\ [\rightarrow\text{O}] \quad & \langle p!l.\pi \parallel p!l.\pi' \rangle \rightarrow \langle \pi \parallel \pi' \rangle \\ [\rightarrow\mathcal{A}] \quad & \langle p?l.\pi \parallel \mathcal{A}^{(p)}.p?l.\pi' \rangle \rightarrow \langle \pi \parallel \mathcal{A}^{(p)}.\pi' \rangle \\ [\rightarrow\mathcal{B}] \quad & \langle p!l.\pi \parallel \mathcal{B}^{(p)}.p!l.\pi' \rangle \rightarrow \langle \pi \parallel \mathcal{B}^{(p)}.\pi' \rangle \end{aligned}$$

$[\rightarrow\text{I}]$ and $[\rightarrow\text{O}]$ erase the top input and output prefixes respectively; $[\rightarrow\mathcal{A}]$ formalises **R1** from § 3.1 (permuting the input $p?l$) while $[\rightarrow\mathcal{B}]$ represents **R2** (permuting the output $p!l$).

Example 4 (Prefix reduction). In Example 2, We can consider both T_Q and T_P as prefixes since they already contain no choice. The (safe) reordering T'_Q can be achieved using $[\rightarrow\mathcal{B}]$, where $\mathcal{B}^{(p)} = p?l_1$:

$$\langle p!l_2.p?l_1.\text{end} \parallel p?l_1.p!l_2.\text{end} \rangle \rightarrow \langle p?l_1.\text{end} \parallel p?l_1.\text{end} \rangle$$

On the other hand, the (unsafe) reordering T'_P cannot be achieved with $[\rightarrow\mathcal{A}]$, since $\mathcal{A}^{(q)} = q!l_1$ violates the definition of $\mathcal{A}^{(q)}$.

$$\begin{aligned} & \frac{\epsilon; \emptyset \vdash_k \langle \epsilon, T, n \rangle \leq \langle \epsilon, T', n \rangle \quad k, n \in \mathbb{N}}{T \leq T'} \text{ [INIT]} \\ & \frac{}{\rho; \Sigma \vdash_k \langle \epsilon, \text{end}, n \rangle \leq \langle \epsilon, \text{end}, n' \rangle} \text{ [END]} \\ & \frac{\Sigma [\langle \pi \parallel T \rangle \leq \langle \pi' \parallel T' \rangle] = \rho \quad \text{act}(\rho') \supseteq \text{act}(\pi')}{\rho.\rho'; \Sigma \vdash_k \langle \pi, T, n \rangle \leq \langle \pi', T', n' \rangle} \text{ [ASM]} \\ & \frac{\langle \pi_1 \parallel \pi'_1 \rangle \rightarrow \langle \pi_2 \parallel \pi'_2 \rangle \quad \rho; \Sigma \vdash_k \langle \pi_2, T, n \rangle \leq \langle \pi'_2, T', n' \rangle}{\rho; \Sigma \vdash_k \langle \pi_1, T, n \rangle \leq \langle \pi'_1, T', n' \rangle} \text{ [SUB]} \\ & \frac{\forall i \in I. \forall j \in J. \rho.p!l_i; \Sigma \vdash_k \langle \pi.p!l_i, T_i, n \rangle \leq \langle \pi'.q?l_j, T'_j, n' \rangle}{\rho; \Sigma \vdash_k \langle \pi, \oplus_{i \in I} p!l_i, T_i, n \rangle \leq \langle \pi', \&_{j \in J} q?l_j, T'_j, n' \rangle} \text{ [OI]} \\ & \frac{\forall i \in I. \exists j \in J. \rho.p!l_i; \Sigma \vdash_k \langle \pi.p!l_i, T_i, n \rangle \leq \langle \pi'.q!l_j, T'_j, n' \rangle}{\rho; \Sigma \vdash_k \langle \pi, \oplus_{i \in I} p!l_i, T_i, n \rangle \leq \langle \pi', \oplus_{j \in J} q!l_j, T'_j, n' \rangle} \text{ [OO]} \\ & \frac{\forall j \in J. \exists i \in I. \rho.p?l_i; \Sigma \vdash_k \langle \pi.p?l_i, T_i, n \rangle \leq \langle \pi'.q?l_j, T'_j, n' \rangle}{\rho; \Sigma \vdash_k \langle \pi, \&_{i \in I} p?l_i, T_i, n \rangle \leq \langle \pi', \&_{j \in J} q?l_j, T'_j, n' \rangle} \text{ [II]} \\ & \frac{\exists i \in I. \exists j \in J. \rho.p?l_i; \Sigma \vdash_k \langle \pi.p?l_i, T_i, n \rangle \leq \langle \pi'.q!l_j, T'_j, n' \rangle}{\rho; \Sigma \vdash_k \langle \pi, \&_{i \in I} p?l_i, T_i, n \rangle \leq \langle \pi', \oplus_{j \in J} q!l_j, T'_j, n' \rangle} \text{ [IO]} \\ & \frac{\Sigma' = \Sigma [\langle \pi \parallel \mu t.T \rangle \leq \langle \pi' \parallel T' \rangle \mapsto \rho] \quad \rho; \Sigma' \vdash_k \langle \pi, T[\mu t.T/t], n-1 \rangle \leq \langle \pi', T', n' \rangle \quad n > 0}{\rho; \Sigma \vdash_k \langle \pi, \mu t.T, n \rangle \leq \langle \pi', T', n' \rangle} \text{ [\mu L]} \\ & \frac{\Sigma' = \Sigma [\langle \pi \parallel T \rangle \leq \langle \pi' \parallel \mu t.T' \rangle \mapsto \rho] \quad \rho; \Sigma' \vdash_k \langle \pi, T, n \rangle \leq \langle \pi', T'[\mu t.T'/t], n'-1 \rangle \quad n' > 0}{\rho; \Sigma \vdash_k \langle \pi, T, n \rangle \leq \langle \pi', \mu t.T', n' \rangle} \text{ [\mu R]} \\ & \frac{\rho; \Sigma \vdash_{k-1} \langle \pi, T, n \rangle \leq \langle \pi', T', n' \rangle \quad \rho; \Sigma \vdash_{k-1} \langle \pi', T', n' \rangle \leq \langle \pi'', T'', n'' \rangle \quad k > 0}{\rho; \Sigma \vdash_k \langle \pi, T, n \rangle \leq \langle \pi'', T'', n'' \rangle} \text{ [TRA]} \end{aligned}$$

Figure 5. Asynchronous subtyping algorithm rules.

Algorithm. We define the rules for our asynchronous subtyping algorithm in Fig. 5, following the style of [22]. Our rules use the function $\text{act}(W)$, the set of input and output actions of π such that $\text{act}(\epsilon) = \emptyset$, $\text{act}(p?l.\pi) = \{p?\} \cup \text{act}(\pi)$ and $\text{act}(p!l.\pi) = \{p!\} \cup \text{act}(\pi)$.

Our algorithm operates on triples of $\langle \pi, T, n \rangle$, where π is a session prefix and n is a bound on the number of recursions to unroll. We keep track of ρ , a prefix containing all actions in the subtype seen so far, and Σ , a set of subtyping assumptions. Each assumption in Σ is associated with the value of ρ as it was at the time of the assumption. An additional bound k is included to limit applications of the $[\text{TRA}]$ rule. We use our algorithm to check whether T is a subtype of T' by beginning with the $[\text{INIT}]$ rule. If a proof derivation can be found then we conclude that $T \leq T'$. If not, then either $T \not\leq T'$ or $T \leq T'$ but this cannot be shown by our algorithm since \leq is undecidable (hence our algorithm cannot be complete).

Our algorithm works as follows: **(1)** If both prefixes π and π' are empty and $T = T' = \text{end}$, then we have nothing

left to check and we terminate with success ([END]). (2) If $\langle \pi \parallel T \rangle \leq \langle \pi' \parallel T' \rangle$ is already in our set of assumptions, then we perform a check on π' to ensure that no actions have been forgotten by the subtype (see the Appendix B.3 for more detail). If this check passes, then we terminate with success ([ASM]). (3) We attempt to reduce the pair of prefixes $\langle \pi \parallel \pi' \rangle$. If we can, then the algorithm repeats from (1) ([SUB]). (4) If not, we try to pop one action from the start of both T and T' and push them to the end of π and π' respectively. If this is possible, we repeat from (1) ([OI,OO,II,IO]). Note that the quantifiers permit subtyping for internal and external choices. For example, [OO] says T is a subtype of T' if it has a subset of T' 's internal choices (defined with $\forall i \in I. \exists j \in J$). (5) Otherwise, we attempt to unroll recursion in T (resp. T'), decrement the bound n (resp. n') by one and repeat from (1). If the bound is already zero, we instead terminate ([$\mu L, \mu R$]).

Double buffering example. We show the execution of our algorithm to check the optimised type from § 2 for the kernel in the double buffering protocol (i.e. $T \leq T'$).

$$\begin{array}{c}
T = s!ready.T' \quad T' = \mu x.s!ready.s?copy.t?ready.t!copy.x \\
\\
\frac{\text{act}(\pi_1.\pi_2.\pi_3.\pi_4) \supseteq \text{act}(\epsilon)}{(\div) \quad \rho_5; \Sigma_3 \vdash \langle \epsilon, T', 1 \rangle \leq \langle \epsilon, T'', 0 \rangle} \text{ [ASM]} \\
\frac{(\star) \quad \rho_5; \Sigma_3 \vdash \langle \pi_4, T', 1 \rangle \leq \langle \pi_4, T'', 0 \rangle}{(\dagger) \quad \rho_5; \Sigma_3 \vdash \langle \pi_3.\pi_4, T', 1 \rangle \leq \langle \pi_3.\pi_4, T'', 0 \rangle} \text{ [SUB]} \\
\frac{(\ddagger) \quad \rho_5; \Sigma_3 \vdash \langle \pi_2.\pi_3.\pi_4, T', 1 \rangle \leq \langle \pi_2.\pi_3.\pi_4, T'', 0 \rangle}{(\dagger) \quad \rho_5; \Sigma_3 \vdash \langle \pi_2.\pi_3.\pi_4, T', 1 \rangle \leq \langle \pi_2.\pi_3.\pi_4, T'', 0 \rangle} \text{ [SUB]} \\
\frac{\rho_5; \Sigma_3 \vdash \langle \pi_1.\pi_2.\pi_3.\pi_4, T', 1 \rangle \leq \langle \pi_2.\pi_3.\pi_4.\pi_1, T'', 0 \rangle}{\rho_4; \Sigma_3 \vdash \langle \pi_1.\pi_2.\pi_3, t!copy.T', 1 \rangle \leq \langle \pi_2.\pi_3.\pi_4, s!ready.T'', 0 \rangle} \text{ [OO]} \\
\frac{\rho_4; \Sigma_2 \vdash \langle \pi_1.\pi_2.\pi_3, t!copy.T', 1 \rangle \leq \langle \pi_2.\pi_3.\pi_4, T', 1 \rangle}{\rho_3; \Sigma_2 \vdash \langle \pi_1.\pi_2, t?ready.t!copy.T', 1 \rangle \leq \langle \pi_2.\pi_3, t!copy.T', 1 \rangle} \text{ [IO]} \\
\frac{\rho_2; \Sigma_2 \vdash \langle \pi_1, T'', 1 \rangle \leq \langle \pi_2, t?ready.t!copy.T', 1 \rangle}{\rho_1; \Sigma_2 \vdash \langle \epsilon, s!ready.T'', 1 \rangle \leq \langle \epsilon, T'', 1 \rangle} \text{ [OI]} \\
\frac{(\ast) \quad \rho_1; \Sigma_1 \vdash \langle \epsilon, T', 2 \rangle \leq \langle \epsilon, T'', 1 \rangle}{\rho_1; \Sigma_1 \vdash \langle \pi_1, T', 2 \rangle \leq \langle \pi_1, T'', 1 \rangle} \text{ [SUB]} \\
\frac{\rho_1; \Sigma_1 \vdash \langle \pi_1, T', 2 \rangle \leq \langle \pi_1, T'', 1 \rangle}{\epsilon; \Sigma_1 \vdash \langle \epsilon, T, 2 \rangle \leq \langle \epsilon, s!ready.T'', 1 \rangle} \text{ [OO]} \\
\frac{\epsilon; \Sigma_1 \vdash \langle \epsilon, T, 2 \rangle \leq \langle \epsilon, s!ready.T'', 1 \rangle}{\epsilon; \emptyset \vdash \langle \epsilon, T, 2 \rangle \leq \langle \epsilon, T', 2 \rangle} \text{ [MR]} \\
\frac{\epsilon; \emptyset \vdash \langle \epsilon, T, 2 \rangle \leq \langle \epsilon, T', 2 \rangle}{T \leq T'} \text{ [INIT]}
\end{array}$$

$$T'' = s?copy.t?ready.t!copy.T'$$

$$\begin{array}{l}
\pi_1 = s!ready \quad \pi_2 = s?copy \quad \pi_3 = t?ready \quad \pi_4 = t!copy \\
\rho_1 = \pi_1 \quad \rho_2 = \rho_1.\pi_1 \quad \rho_3 = \rho_2.\pi_2 \quad \rho_4 = \rho_3.\pi_3 \quad \rho_5 = \rho_4.\pi_4
\end{array}$$

$$\Sigma_1 = [\langle \epsilon \parallel T \rangle \leq \langle \epsilon \parallel T' \rangle \mapsto \epsilon]$$

$$\Sigma_2 = \Sigma_1 [\langle \epsilon \parallel T' \rangle \leq \langle \epsilon \parallel T'' \rangle \mapsto s!ready]$$

$$\Sigma_3 = \Sigma_2 [\langle \pi_1.\pi_2.\pi_3 \parallel t!copy.T' \rangle \leq \langle \pi_2.\pi_3.\pi_4 \parallel T' \rangle \mapsto \rho_4]$$

$$(\dagger) = \langle \pi_1.\pi_2.\pi_3.\pi_4 \parallel \pi_2.\pi_3.\pi_4.\pi_1 \rangle \rightarrow \langle \pi_2.\pi_3.\pi_4 \parallel \pi_2.\pi_3.\pi_4 \rangle \text{ [}\mathcal{A}\text{]}$$

$$(\ddagger) = \langle \pi_2.\pi_3.\pi_4 \parallel \pi_2.\pi_3.\pi_4 \rangle \rightarrow \langle \pi_3.\pi_4 \parallel \pi_3.\pi_4 \rangle \text{ [}\uparrow\text{I]}$$

$$(\star) = \langle \pi_3.\pi_4 \parallel \pi_3.\pi_4 \rangle \rightarrow \langle \pi_4 \parallel \pi_4 \rangle \text{ [}\uparrow\text{I]}$$

$$(\ast) = \langle \pi_1 \parallel \pi_1 \rangle \rightarrow \langle \epsilon \parallel \epsilon \rangle \text{ [}\uparrow\text{O]} \quad (\div) = \langle \pi_4 \parallel \pi_4 \rangle \rightarrow \langle \epsilon \parallel \epsilon \rangle \text{ [}\uparrow\text{O]}$$

See also the Appendix B.4 for the ring and alternating bit protocols, which include internal and external choices.

Properties. We prove the correctness and complexity of our algorithm. See the Appendix B for the proofs. We define the size of prefixes as $|\epsilon| = 0$ and $|\rho?l.\pi| = |\rho!l.\pi| = 1 + |\pi|$.

Lemma 5. *Given finite prefixes π and π' , $\langle \pi \parallel \pi' \rangle$ can be reduced only a finite number of times.*

Theorem 6 (Termination). *Our subtyping algorithm always eventually terminates.*

Theorem 7 (Soundness). *Our subtyping algorithm is sound.*

Lemma 8. *Given finite prefixes π and π' , the time complexity of reducing $\langle \pi \parallel \pi' \rangle$ is $O(\min(|\pi|, |\pi'|))$.*

Theorem 9 (Complexity). *Consider T and T' as (possibly infinite) trees $\mathcal{T}(T)$ and $\mathcal{T}(T')$ with asymptotic branching factors b and b' respectively [21, 39]. Our algorithm has time complexity $O(n \min(b, b')^n)$ and space complexity $O(n \min(b, b'))$ in the worst case to determine if $T \leq T'$ with bound n .*

Algorithm implementation. In practice, we make some minor alterations to the algorithm when implementing it in RUMPSTEAK. As outlined in § 2, RUMPSTEAK acts on FSMs rather than local types so we modify our bounds-checking accordingly. We also represent prefixes as lazily-removable lists for greater memory efficiency—this additionally allows a slightly simplified termination condition in the case of [ASM]. Finally, we provide some opportunities for the algorithm to “short circuit” in cases where we can tell early that a subtype is not valid. See the Appendix B.5 for more details.

4 Evaluation

In this section, we evaluate how RUMPSTEAK performs with respect to existing tools. First, in § 4.1, we evaluate the runtime performance of programs written in RUMPSTEAK versus the same benchmarks implemented using other Rust session type tools. In the spirit of Rust’s emphasis on efficiency, this runtime performance is of particular significance for developers. Secondly, in § 4.2, we evaluate how RUMPSTEAK’s verification of message reordering (our subtyping algorithm) scales compared to existing verification tools. Although not a runtime cost, subtyping is known to be a computationally challenging problem that often scales poorly.

4.1 Session-Based Rust Implementations

We compare RUMPSTEAK’s runtime against three other session type implementations in Rust: (1) SESH [38], a synchronous implementation of binary session types; (2) FERRITE [12], an implementation of shared binary session types [5] supporting asynchronous execution; and (3) MULTICRUSTY [41], a synchronous MPST implementation based on SESH.

We perform a series of benchmarks shown in Fig. 6. We execute these using a 16-core AMD Opteron™ 6200 Series

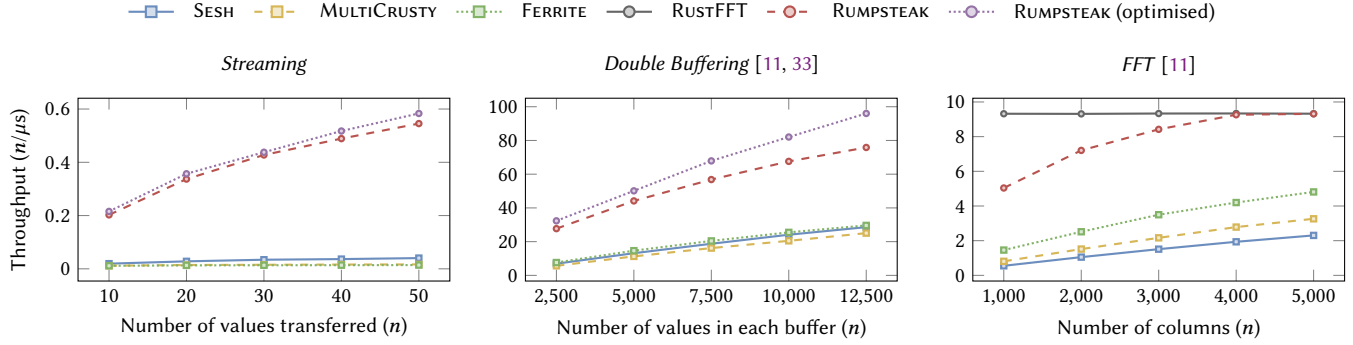


Figure 6. Benchmarking RUMPSTEAK’s runtime performance against previous work in Rust (raw data in the Appendix C.1).

CPU @ 2.6GHz with hyperthreading, 128GB of RAM, Ubuntu 18.04.5 LTS and Rust Nightly 2021-07-06. We use version 0.3.5 of the Criterion.rs library [27] to perform microbenchmarking and a *multi-threaded* asynchronous runtime from version 1.11.0 of the Tokio library [60].

Streaming. This protocol has two participants, a source s and a sink t , with a combination of recursion and choice. As stated previously, its global type G_{ST} is given as

$$G_{ST} = \mu x.t \rightarrow s : \{ \text{ready}.s \rightarrow t : \{ \text{value}.x, \text{ stop}.end \} \}$$

We benchmark the protocol by varying the number of *values* that are sent before the exchange *stops*. We use AMR to create an optimised version for RUMPSTEAK: if the source knows it will send at least n values before *stopping*, then these messages can be *unrolled* and sent all at once—only receiving *ready* from the sink after sending all n values. For benchmarking, we unroll the first 5 values in the optimised version.

Our results show RUMPSTEAK reaches around 14.5x the throughput of other implementations. At first, it is limited by channel creation overheads, the cost of which becomes less significant as more messages are sent. The throughput levels off as message passing overheads become the bottleneck. The optimised version eases this overhead, as the source is less frequently blocked on receiving *ready* from the sink—this effect could be increased by unrolling more messages.

SESH and MULTICRUSTY have much lower throughputs since they use more expensive synchronous communication. These implementations also create a new channel for each interaction. This causes their throughput to stay constant as there are fewer one-time costs to spread than in RUMPSTEAK. Interestingly, FERRITE performs similarly to SESH and MULTICRUSTY, despite employing asynchronous execution like RUMPSTEAK. This is because its design requires that the source and sink are implemented using *recursion* rather than *iteration*. FERRITE also requires particularly strong safety requirements at compile time. Therefore, the sink’s output buffer must be guarded with a mutex instead of being accessed directly.

Double Buffering. We benchmark our running example by performing only two iterations. This allows for both of the kernel’s buffers to be filled and, importantly, for the protocol to terminate. SESH and FERRITE do not support MPST so their implementations use binary session types between pairs of participants. This approach does *not* provide the same safety as MPST and so we cannot verify deadlock-freedom as we can in MULTICRUSTY and RUMPSTEAK (see Table 1).

We parameterise the size of the buffers and measure the throughput. Performance is similar to that of the stream benchmark; RUMPSTEAK’s throughput reaches around 3.2x that of the others. Moreover, we confirm the intuition described in § 1. Increasing the size of the buffers does generally improve throughput. However, it remains limited when only a single buffer can be used at once. Fig. 6 shows that optimising the kernel as described in § 2 increases the throughput since the source and sink operate on different buffers at once.

FFT. The fast Fourier transform (FFT) is an algorithm for computing the discrete Fourier transform of a vector. We take $n \times 8$ matrices, where the FFT is computed on each column to produce a new, transformed matrix, and implement the Cooley-Tukey FFT algorithm which divides the problem into—in our case—eight. Each of these problems can be solved independently by different processes, which communicate with one another using message passing. The exchanges between participants are described in [11, Fig. 7].

We have chosen to use FFT as a benchmark as it is a standard problem with numerous implementations, and we can therefore compare implementations based on MPST with actual high-optimised implementations.

We implement a concurrent version of the algorithm in SESH, FERRITE, MULTICRUSTY and RUMPSTEAK, which uses eight processes. Each process works on an array of n inputs, representing one column. Arithmetic operations are performed in a pairwise fashion on two of these arrays.

We use the same approach as in the double buffering protocol to write the binary implementations for SESH and FERRITE. Interestingly, to represent this as a sequence of binary

Table 1. Expressiveness of RUMPSTEAK compared to previous work.

Protocol	n	C	R	IR	AMR	SESH	FERRITE	MULTICRUSTY	RUMPSTEAK	k -MC	SOUNDBINARY
Two Adder [2]	2	✓	✓			✓	✓	✓	✓	✓	✓
Three Adder	3					x	x	✓	✓	✓	x
Streaming	2	✓	✓			✓	✓	✓	✓	✓	✓
Optimised Streaming	2	✓	✓		✓	x	x	x	✓	✓	✓
Ring [11]	3		✓	✓		x	x	✓	✓	✓	x
Optimised Ring [11]	3		✓	✓	✓	x	x	x	✓	✓	x
Ring With Choice [11]	3	✓	✓	✓		x	x	✓	✓	✓	x
Optimised Ring With Choice [11]	3	✓	✓	✓	✓	x	x	x	✓	✓	x
Double Buffering [11]	3		✓	✓		x	x	✓	✓	✓	x
Optimised Double Buffering [11, 33]	3		✓	✓	✓	x	x	x	✓	✓	x
Alternating Bit [1, 43]	2	✓	✓	✓		x	x	x	✓	✓	✓
Elevator [6, 43]	3	✓	✓	✓	✓	x	x	x	✓	✓	x
FFT [11]	8					x	x	✓	✓	✓	x
Optimised FFT [11]	8				✓	x	x	x	✓	✓	x
Authentication [48]	3	✓				x	x	✓	✓	✓	x
Client-Server Log [41]	3	✓	✓	✓		x	x	✓	✓	✓	x
Hospital [7]	2	✓	✓	✓	✓	x	x	x	x	x	✓

n Number of participants C Choice R Recursion IR Infinite recursion AMR Asynchronous message reordering
 ✓ Expressible x Expressible using endpoint types (but without deadlock-freedom guarantee) x Not expressible

sessions, we require additional synchronisation of all participants at each stage of the protocol. While we can perform AMR to RUMPSTEAK’s version, in practice we found that this does not have much effect since the protocol is already heavily synchronised so cannot progress any faster.

We benchmark the protocol by varying the number of columns in the matrix. We also compare these implementations with the most downloaded open-source Rust FFT implementation, RUSTFFT [61]. RUSTFFT does not use concurrency and computes the FFT of a matrix by iteratively performing the Cooley-Tukey algorithm on each column.

FERRITE performs better than before since its FFT implementation does not suffer from the limitations explained previously. Like RUMPSTEAK, it benefits from the use of asynchronous execution, although the additional synchronisation from representing the problem as a set of binary sessions causes RUMPSTEAK’s throughput to remain around 1.9x greater.

Most excitingly, RUMPSTEAK achieves RUSTFFT’s throughput for large matrices. RUSTFFT’s implementation is highly tuned for low-level efficiency and does not incur any overheads associated with message passing. Moreover, asynchronous executors are generally designed for higher-level tasks such as networking; MPI-based communication [44] would more likely be used to parallelise this problem in practice. Therefore, we think it impressive that even a basic parallel implementation using RUMPSTEAK can reach the same performance as a state-of-the-art sequential implementation.

Expressiveness. Table 1 discusses the expressiveness of RUMPSTEAK compared with SESH, FERRITE and MULTICRUSTY. Since SESH and FERRITE support only binary session types,

they are unable to guarantee deadlock-freedom in protocols with more than two participants. MULTICRUSTY has greater expressiveness than SESH and FERRITE since it implements MPST but, unlike RUMPSTEAK, still cannot ensure deadlock-freedom for protocols optimised using AMR. In addition, many optimisations, like the ones we have benchmarked, break duality between pairs of participants so are not expressible at all by SESH, FERRITE and MULTICRUSTY. Meanwhile, our powerful API and new subtyping algorithm allow RUMPSTEAK to express many examples using AMR.

4.2 Verifying Asynchronous Message Reordering

We perform a second set of benchmarks, shown in Fig. 7, to evaluate RUMPSTEAK’s asynchronous subtyping algorithm. We compare it against (1) SOUNDBINARY [7], a sound subtyping algorithm defined for *binary* session types only; and (2) k -MC [43], an algorithm for directly checking the compatibility of a set of FSMs without the need for a global type. Compatibility is checked up to a bound k on the size of each process’ asynchronous queue. We note that neither SOUNDBINARY nor k -MC provide a runtime framework like RUMPSTEAK does, they are used only to verify the AMR.

We benchmark with the same machine as before. SOUNDBINARY and k -MC are written in Haskell so we must run each tool’s binary rather than simply timing Rust functions. We, therefore, provide a command-line for RUMPSTEAK’s subtyping algorithm so it is comparable with these other tools. RUMPSTEAK’s binary is compiled with Rust 1.54.0 and we use Hyperfine [51] to compare the execution time for each tool.

Streaming (from § 4.1). We vary n , the number of *values* we unroll. Using SOUNDBINARY and RUMPSTEAK, we check

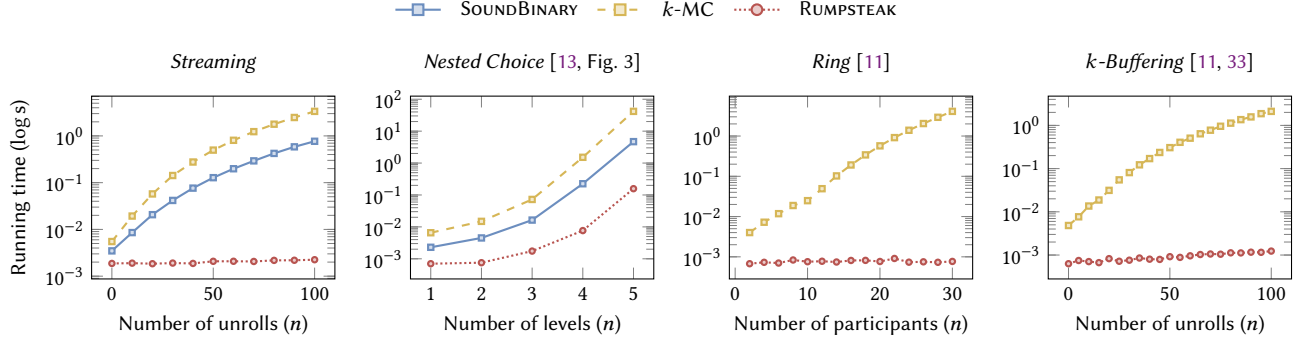


Figure 7. Benchmarking RUMPSTEAK’s subtyping performance against previous work (raw data in the Appendix C.2).

that the optimised source is a subtype of its projected version and, using k -MC, we check that the optimised source is compatible with the sink. Our results show that RUMPSTEAK scales significantly better than both SOUNDBINARY and k -MC. While all three implementations can verify up to 100 unrolls in under a second, the execution time taken by SOUNDBINARY and k -MC increases exponentially while RUMPSTEAK’s remains mostly flat. This is consistent with our complexity analysis in Theorem 9, since unrolling more values does not increase the branching factor of the subtype.

Nested choice. We next consider a protocol from Chen et al. [13, Fig. 3] containing nested choice. We perform this nesting up to a parameterised number of levels n to increase the complexity. Specifically, we check that $T_n \leq T'_n$ where

$$\begin{aligned}
 T_0 &= T'_0 = \text{end} \\
 T_{n+1} &= !m.(?r.T_n \ \& \ ?s.T_n \ \& \ ?u.T_n) \oplus !p.(?r.T_n \ \& \ ?s.T_n) \\
 T'_{n+1} &= ?r.(!m.T'_n \ \oplus \ !p.T'_n \ \oplus \ !q.T'_n) \ \& \ ?s.(!m.T'_n \ \oplus \ !p.T'_n)
 \end{aligned}$$

We find that RUMPSTEAK again performs more scalably than SOUNDBINARY and k -MC. Here, RUMPSTEAK’s improved efficiency is significant—for five levels, k -MC takes around 40s while RUMPSTEAK requires only a fraction of a second. Nonetheless, we observe that RUMPSTEAK also exhibits exponential performance, in contrast to our stream benchmark. This is because our algorithm bounds recursion but not choice, causing an explosion in the paths to visit as the number of nested choices increases. This is also explained by our complexity analysis in Theorem 9 since a greater number of nested choices increases the depth to which we must explore.

Ring [11]. We consider a ring protocol with a parameterised number of participants. Each participant (aside from the first, which initiates the protocol) receives a value from its preceding neighbour then sends a value to its succeeding neighbour. Providing that this receive is not dependent on the send, we can use AMR to send before receiving.

We benchmark verifying this optimisation using both k -MC and RUMPSTEAK. In this case, we cannot use SOUNDBINARY since the protocol is multiparty. While RUMPSTEAK can verify each participant’s subtype individually, k -MC must consider the entire protocol at once. Unsurprisingly, k -MC is

significantly less scalable as its running time grows exponentially, whereas RUMPSTEAK’s performance remains mostly constant.

k -buffering. We extend the double buffering protocol to a parameterised number of buffers. We benchmark verifying the optimisation to the kernel discussed previously. Crucially, with a greater number of buffers, we can unroll a greater number of ready messages. We again compare only k -MC and RUMPSTEAK since the protocol is multiparty. Similarly to other benchmarks, RUMPSTEAK scales more efficiently than k -MC.

Expressiveness. Table 1 discusses the expressiveness of RUMPSTEAK’s algorithm against SOUNDBINARY and k -MC (note again that SOUNDBINARY and k -MC only verify AMR and do not provide language implementations like RUMPSTEAK). As we discussed, SOUNDBINARY supports only two parties so cannot express many of the case studies. However, it can express some unbounded binary protocols that RUMPSTEAK and k -MC cannot, such as the Hospital example [7, § 1]¹. Unsurprisingly, k -MC’s expressiveness in the table coincides with RUMPSTEAK’s as k -MC with $k=\infty$ is equivalent to the liveness property induced by asynchronous multiparty subtyping [25, Theorem 7.3]. On the other hand, k -MC can verify a wider syntax of local FSMs than those corresponding to Definition 1.

5 Conclusion and Related Work

There are a vast number of studies on session types, some of which are implemented in programming languages [4] and tools [24]. The top-down approach using Scribble [2, 55, 63] we presented in this paper has been implemented for a number of other programming languages such as Java [31, 32, 40], Go [10], TypeScript [45], Scala [54], MPI-C [49], Erlang [47], Python [18], F# [46], F★ [64] and Actor DSL [26]. Several implementations use an EFSM-based approach to generate

¹Notice, for the case of RUMPSTEAK, that we can manually write the endpoints, and the framework checks the conformance to the protocol. However, RUMPSTEAK cannot verify the *deadlock-freedom* property of the protocol, hence the *amber cross*.

APIs from Scribble for target programming languages such as [10, 31, 32, 45, 46, 54, 64] to ensure correctness by construction, but none are integrated with AMR like RUMPSTEAK.

RUMPSTEAK provides three approaches for gaining *efficiency* and ensuring *deadlock-freedom* in message-passing Rust applications: (1) the *top-down* approach for ensuring *correctness/safety by construction* and maximising *asynchrony* with local analysis; (2) the *bottom-up* approach using global analysis on a set of FSMs; and (3) the *hybrid* approach, which combines inference with local analysis. All three approaches are backed by RUMPSTEAK’s API (described in § 2), which uses Rust’s affine type system to ensure protocol compliance.

Note that RUMPSTEAK is the first framework (for any programming language) to enable the bottom-up (by integrating with k -MC [43]) and the hybrid approaches.

Rust session type implementations. We closely compared the performance and expressiveness of RUMPSTEAK with three existing works in § 4: (1) SESH [38], a synchronous implementation of binary session types; (2) FERRITE [12], an implementation of shared binary session types [5] supporting asynchronous execution; and (3) MULTICRUSTY [41], a synchronous MPST implementation based on SESH.

Of these previous implementations, only MULTICRUSTY can also support MPST with deadlock-freedom. However, to represent a multiparty session, MULTICRUSTY requires defining a tuple of binary sessions for each role as well as their order of use. This leads to a more complex and less intuitive API than RUMPSTEAK’s. Our API requires fewer definitions and is also both more expressive and performant (see § 4).

RUMPSTEAK and FERRITE support asynchronous execution, which is more efficient than the synchronous and blocking communication used by SESH and MULTICRUSTY. However, only RUMPSTEAK supports Rust’s idiomatic `async/await` syntax. FERRITE instead requires nesting sequential communication, which is more verbose and less efficient. Iteration cannot be easily expressed (see § 4) and it requires stricter compile-time concurrency guarantees than RUMPSTEAK. We do not compare directly against [34] (another synchronous implementation of binary session types) as this uses an older edition of Rust and has several limitations already discussed in [38, 41].

RUMPSTEAK can perform AMR, which is not possible in any of these existing implementations. Combined with its straightforward design and use of asynchronous execution, we found that it is therefore much more efficient than all previous work. It has around 14.5x, 3.2x and 1.9x greater throughput than the next fastest implementation in the stream, double buffering and FFT protocols respectively. RUMPSTEAK can also compete with a popular Rust implementation in the FFT benchmark, achieving the same throughput for large input sizes.

Recent work [20] builds a DSL to offer protocol conformance for Rust, supported by `typestates`. They do not yet

explore combining the DSL with communication channels like RUMPSTEAK and [5, 34, 38, 41]. If this is achieved, it would be interesting to integrate session type tooling with their DSL and compare this with RUMPSTEAK.

Verification of asynchronous subtyping. Interest in AMR has grown recently, both in theory and practice. An extension of the Iris framework, Actris [35], formalises a variant of binary asynchronous subtyping, which is in turn implemented in the Coq proof assistant. The asynchronous session subtyping defined in [13, 25] is *precise* but was shown to be undecidable, even for binary sessions [9, 42]. Hence, in general, checking $M'_i \leq M_i$ is undecidable. Various limited classes of session types for which $M'_i \leq M_i$ is decidable [7, 8, 11, 42] are proposed but they are not applicable to our use cases since (1) the relations in [7, 9, 42] are *binary* and the same limitations do not work for multiparty sessions; and (2) the relation in [11, Def. 6.1] does not handle subtyping across unrolling recursions, e.g. the relation is inapplicable to the double buffering algorithm [33] (see [11, Remark 8.1]).

Our new algorithm for asynchronous optimisation is *terminating* (see Theorem 6), *sound* (see Theorem 7) and capable of verifying optimisations in a range of classical examples (see § 4). It is also more performant than a global analysis based on k -MC, and an existing two-party sound algorithm [7] (see § 4). This is because our algorithm is implemented efficiently (see § 3.2) and can often execute in linear rather than exponential time (see Theorem 9).

Deadlock detection. There is also recent progress [52, 53] on static deadlock detection in Rust by tracking locks and unlocks in Rust’s intermediate representation (MIR) [59]. This targets shared memory, rather than message-passing applications, and does not attempt safety by construction. In future work, we could also analyse the MIR directly in RUMPSTEAK, replacing the use of our API. However, it is difficult to identify concurrency primitives in the MIR (this would be made harder with `async/await`); [52] currently supports only three concurrent data structures from Rust’s ecosystem.

Acknowledgments

We thank (in alphabetical order) Julia Gabet, Paul Kelly, Roland Kuhn, Nicolas Lagaillardie, Julien Lange, and Fangyi Zhou, for their valuable proofreading, feedback and advice on early versions of this work. We also thank the PPOPP reviewers for the helpful comments in their reviews.

The work is supported by EPSRC EP/T006544/1, EP/K011715/1, EP/K034413/1, EP/L00058X/1, EP/N027833/1, EP/N028201/1, EP/T014709/1, EP/V000462/1, and NCSS/EPSRC VeTSS.

A Artifact

A.1 Content of the artifact

The artifact [16] contains the following files:

```
.
|-- Artifact.md
|-- Artifact.pdf
|-- Dockerfile
|-- gen_fig_6.sh
|-- gen_fig_7.sh
`-- getting_started.sh
```

Note that an internet connection is required to use the artifact.

A.2 Getting started guide (Docker artifact)

In this subsection, we will run the benchmarks used to produce Figs. 6 and 7 of the paper in a Docker container. At the end of this section, you should be have all the plots of those figures.

Note that the benchmark results may not match those in Figs. 6 and 7 when run in a Docker container or a machine with different specification than used in the paper. See the subsection on *Claims supported or not by the artifact* for more discussion of this.

To run the getting started guide, we assume you are running a Linux machine with Docker and Gnuplot installed, as well as standard Unix tools (`tail`, `awk`, `cut`, etc.).

This *Getting started* subsection is fully automated: extract the archive and run the `getting_started.sh` script. *The script takes a long time to run all the benchmarks. On the author's laptop, it took approximately 2 hours and 15 minutes.* When the script finishes, you should have a few `*.png` plots which correspond to the plots shown in Figs. 6 and 7 of the paper.

```
1 $ cd /path/to/extracted/artifact/
2 $ ./getting_started.sh
```

Once run, to clean-up your system, in addition to removing the archive folder, you should remove the docker image (the following command assumes you don't have other docker images/containers on your system):

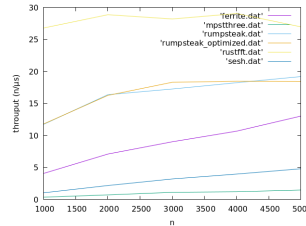
```
1 $ docker rmi rumpsteak_tool
2 $ docker system prune
```

A.2.1 Output. The `getting_started.sh` generates figures similar to the one used in the paper².

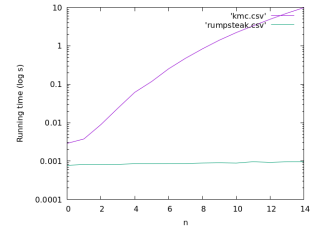
We show in Fig. 8 the figures `fft.png` and `ring.png` that are the equivalent of Fig. 6 (subfigure FFT) and Fig. 7 (subfigure Ring) in the paper.

A.2.2 Analysing the raw data.

²They are not exactly the same in the sense that they are not produced by the same tool (Tikz vs. Gnuplot), but the data they use is produced similarly.



(a) Figure `fig_6_fft.png` generated by the artifact, which corresponds to Example FFT in Fig. 6 in the paper.



(b) Figure `fig_7_ring.png` generated by the artifact, which corresponds to Example Ring in Fig. 7 in the paper.

Figure 8. Two examples of the figures generated. Note that those figure are obtain from a different run of the artifact, in a lower-end computer, which explains the differences with figures presented in the paper.

```
.
|-- double_buffering
|   |-- ...
|-- fft
|   |-- ...
|-- report
|   `-- index.html
`-- stream
    |-- ...
```

Figure 9. The files created by the Criterion benchmark (in `results/criterion/`). A detailed report is available by viewing `index.html` in a browser.

Runtime benchmarks. The results used to generate Fig. 6 are generated with `Criterion`. `Criterion` produces the tree of files shown in Fig. 9. The report can be viewed in a web browser.

Subtyping benchmarks. `Hyperfine` generates a CSV file (located at `results/data/*/*.csv`) for the results of each benchmark. For each row, this file contains

1. `command`: the actual command that was run;
2. `mean`: the mean time to execute the command;
3. `stddev`: the standard deviation of executing the command;
4. `median`: the median time to execute the command;
5. `user`: the mean *user* time to run the command;
6. `system`: the mean *system* time to run the command;
7. `min`: the minimum time to run the command;
8. `max`: the maximum time to run the command; and
9. `parameter_n`: the value of the parameter used in the command

where all times are given in seconds. Since we do not care about execution time spent in the kernel (all relevant computation is done in user space) we take only the timings

in the user column. We provide a script to plot the mean user execution time against the parameter value for each tool.

A.3 Claims supported or not by the artifact

A.3.1 Claims supported by the artifact.

- Results presented in Fig. 7 are fairly stable across different machines, even when run in Docker. One should expect to obtain similar plots on their machine.

A.3.2 Claims partially supported by the artifact.

- Results presented in Fig. 6 are quite dependent on the number of cores provided by the machine and on other programs being run on the machine simultaneously. The results presented in the paper are run on a 16-core AMD Opteron 6200 Series at 2.6GHz with hyperthreading and 128GB of RAM.

Attempts run on lower-end processors or from within a Docker container may not provide similar results (in particular, the `rustfft` implementation is highly optimised for sequential execution and outperforms approaches based on message passing on processors with fewer cores). Nonetheless, the performance of Rumpsteak vs. other MPST implementations would remain mostly comparable.

A.3.3 Claims not supported by the artifact.

- NuScr is an external tool not contributed by the authors, and therefore not part of the claims of the paper.

B Multiparty Asynchronous Subtyping

In the main paper, we mentioned a few definitions from [25] that we omitted due to space constraints. We explain these in the first section of this appendix. In the later sections, we provide examples of the rules shown in the paper, as well as proofs of the theorems stated. Finally, we provide more details on its implementation.

B.1 Synchronous session subtyping

We first give the rules for *synchronous session subtyping* given by Chen et al. [13] in Fig. A.10. The relation \leq : on sorts is defined as the least reflexive binary relation such that $\text{nat} \leq \text{int}$ [25].

$$\frac{}{\text{end} \leq \text{end}} \text{[SUB-END]}$$

$$\frac{\forall i \in I \ S_i \leq S'_i \ T_i \leq T'_i}{\&_{i \in I \cup J} p?l_i(S_i).T_i \leq \&_{i \in I \cup J} p?l_i(S'_i).T'_i} \text{[SUB-BRA]}$$

$$\frac{\forall i \in I \ S_i \leq S_i \ T_i \leq T'_i}{\oplus_{i \in I} p!l_i(S_i).T_i \leq \oplus_{i \in I \cup J} p!l_i(S'_i).T'_i} \text{[SUB-SEL]}$$

Figure A.10. Subtyping rules for synchronous session types.

B.2 Precise Asynchronous Multiparty Subtyping

Asynchronous subtyping is more complex as it allows the order of operations to be swapped for efficiency. The *tree refinement relation* \lesssim is defined coinductively on session types that have only single-inputs (SI) and single-outputs (SO). It is specified for type trees, which are possibly infinite trees representing a session type. An example of a type tree is given in Fig. A.13 and the tree refinement relation by Ghilezan et al. [25] is given in Fig. A.11. The function $\text{act}(W)$, the set of input and output actions in a tree W , is defined in Fig. A.12.

Single-input and single-output types are session types which do *not* include branching, i.e. a type generated from the grammar (Ghilezan et al. [25]) $T ::= \text{end} \mid p?l.T \mid p!l.T$.

Remark. As mentioned in [25], checking the set of actions within $\text{[REF-}\mathcal{A}\text{]}$ and $\text{[REF-}\mathcal{B}\text{]}$ is important. If this were not included, then unsound recursive subtypes that “forget” some interactions would be allowed. Ghilezan et al. [25] give the following example of a potential subtype that forgets to input an initial l' message. If the $\text{[REF-}\mathcal{A}\text{]}$ rule were allowed to be used then T would incorrectly be a subtype of T' .

$$T = \mathcal{T}(\mu t.p?l.t) \quad T' = q?l'.T = q?l'.\mathcal{T}(\mu t.p?l.t)$$

$$\frac{T \leq q?l'.T'}{\frac{}{T = p?l.T \leq q?l'.p?l.T = T'} \text{[REF-}\mathcal{A}\text{]}}$$

B.2.1 Examples of asynchronous subtyping.

Ring protocol. We show an example of the subtyping rules for a ring protocol with choice. The projected and optimised local types are given by T' and T respectively.

$$T' = \mu t.a?add.c! \left\{ \begin{array}{l} \text{add.t} \\ \text{sub.t} \end{array} \right\} \quad T = \mu t.c! \left\{ \begin{array}{l} \text{add.a?add.t} \\ \text{sub.a?add.t} \end{array} \right\}$$

Considering the tree types $T = \mathcal{T}(T)$ and $T' = \mathcal{T}(T')$, we must show that $T \leq T'$ using the tree refinement definition from the main paper in order to prove that $T \leq T'$.

$$\forall U \in \llbracket T \rrbracket_{\text{so}} : \forall V' \in \llbracket T' \rrbracket_{\text{si}} : \exists W \in \llbracket U \rrbracket_{\text{si}} : \exists W' \in \llbracket V' \rrbracket_{\text{so}} : W \lesssim W'$$

However, in our case, since T and T' are already SI trees, we can express the definition more simply using only SO tree transformations.

$$\forall W \in \llbracket T \rrbracket_{\text{so}} : \exists W' \in \llbracket T' \rrbracket_{\text{so}} : W \lesssim W'$$

We define the SO sets for each tree coinductively and use coinduction to show that $W \lesssim W'$ in all cases.

$$\pi_1 = c!add.a?add \quad \pi_2 = c!sub.a?add \quad \pi_3 = a?add$$

$$\forall \pi_1.W \in \llbracket T \rrbracket_{\text{so}} : W \in \llbracket T \rrbracket_{\text{so}} \quad \forall \pi_2.W \in \llbracket T \rrbracket_{\text{so}} : W \in \llbracket T \rrbracket_{\text{so}}$$

$$\forall \pi_3.\pi_1.W' \in \llbracket T' \rrbracket_{\text{so}} : \pi_3.W' \in \llbracket T' \rrbracket_{\text{so}}$$

$$\forall \pi_3.\pi_2.W' \in \llbracket T' \rrbracket_{\text{so}} : \pi_3.W' \in \llbracket T' \rrbracket_{\text{so}}$$

1. Using the coinductive hypothesis $\pi_1.W \lesssim \pi_3.\pi_1.W'$, we show that $W \lesssim \pi_3.W'$.

$$\begin{array}{c}
\frac{}{\text{end} \lesssim \text{end}} \text{[REF-END]} \quad \frac{S' \leq S \quad W \lesssim W'}{p?l(S).W \lesssim p?l(S').W'} \text{[REF-IN]} \quad \frac{S \leq S' \quad W \lesssim W'}{p!l(S).W \lesssim p!l(S').W'} \text{[REF-OUT]} \\
\frac{S' \leq S \quad W \lesssim \mathcal{A}^{(p)}.W' \quad \text{act}(W) = \text{act}(\mathcal{A}^{(p)}.W')}{p?l(S).W \lesssim \mathcal{A}^{(p)}.p?l(S').W'} \text{[REF-A]} \\
\frac{S \leq S' \quad W \lesssim \mathcal{B}^{(p)}.W' \quad \text{act}(W) = \text{act}(\mathcal{B}^{(p)}.W')}{p!l(S).W \lesssim \mathcal{B}^{(p)}.p!l(S').W'} \text{[REF-B]}
\end{array}$$

Figure A.11. Tree refinement relation rules for asynchronous session type trees.

$$\begin{array}{l}
\text{act}(\text{end}) = \emptyset \quad \text{act}(p?l(S).W) = \{p?\} \cup \text{act}(W) \\
\text{act}(p!l(S).W) = \{p!\} \cup \text{act}(W)
\end{array}$$

Figure A.12. Definition of the function $\text{act}(W)$ on a tree W .




Figure A.13. An example of a session type and its corresponding type tree.

$$\frac{W \lesssim a? \text{add}. W' = \pi_3. W'}{a? \text{add}. W \lesssim a? \text{add}. a? \text{add}. W'} \text{[REF-IN]} \\
\frac{}{c! \text{add}. a? \text{add}. W \lesssim a? \text{add}. c! \text{add}. a? \text{add}. W'} \text{[REF-B]}$$

2. Using the coinductive hypothesis $\pi_2. W \lesssim \pi_3. \pi_2. W'$, we show that $W \lesssim \pi_3. W'$.

$$\frac{W \lesssim a? \text{add}. W' = \pi_3. W'}{a? \text{add}. W \lesssim a? \text{add}. a? \text{add}. W'} \text{[REF-IN]} \\
\frac{}{c! \text{sub}. a? \text{add}. W \lesssim a? \text{add}. c! \text{sub}. a? \text{add}. W'} \text{[REF-B]}$$

Double buffering protocol. We also show the optimisation to the double buffering protocol discussed in § 2. The kernel sends two *ready* messages at once, allowing the source to fill up both buffers sooner.

$$\begin{array}{l}
T = s! \text{ready}. T' = s! \text{ready}. \mu x. s! \text{ready}. s? \text{copy}. t? \text{ready}. t! \text{copy}. x \\
T' = \mu x. s! \text{ready}. s? \text{copy}. t? \text{ready}. t! \text{copy}. x
\end{array}$$

As before, we consider the tree types $T = \mathcal{T}(T)$ and $T' = \mathcal{T}(T')$ which are already both SI and SO trees. Therefore, to prove that $T \leq T'$ we need only to show that $T \lesssim T'$.

$$\frac{}{T' \leq W} \text{[REF-IN]} \\
\frac{}{t! \text{copy}. T' \leq t! \text{copy}. W} \text{[REF-IN]} \\
\frac{}{t? \text{ready}. t! \text{copy}. T' \leq t? \text{ready}. t! \text{copy}. W} \text{[REF-IN]} \\
\frac{}{s? \text{copy}. t? \text{ready}. t! \text{copy}. T' \leq s? \text{copy}. t? \text{ready}. t! \text{copy}. W} \text{[REF-IN]} \\
\frac{}{s! \text{ready}. W \leq s? \text{copy}. t? \text{ready}. t! \text{copy}. s! \text{ready}. W} \text{[REF-B]} \\
\frac{}{T = s! \text{ready}. T' \leq s! \text{ready}. W = T'} \text{[REF-OUT]} \\
W = s? \text{copy}. t? \text{ready}. t! \text{copy}. T'$$

B.3 Proofs for the Subtyping Algorithm

Lemma 3. Given finite prefixes π and π' , $\langle \pi \parallel \pi' \rangle$ can be reduced only a finite number of times.

Proof. We prove this in a similar style to [22, Lemma 10], using a well-founded relation. We first consider the relation $R = \{(\langle \pi_1 \parallel \pi'_1 \rangle, \langle \pi_2 \parallel \pi'_2 \rangle) \mid \langle \pi_2 \parallel \pi'_2 \rangle \rightarrow \langle \pi_1 \parallel \pi'_1 \rangle\}$. We define the function $\text{terms}(\pi)$, which returns the number of terms in π , such that

$$\begin{array}{l}
\text{terms}(\epsilon) = 0 \\
\text{terms}(p!l(S)) = 1 \quad \text{terms}(p?l(S)) = 1 \\
\text{terms}(\pi_1. \pi_2) = \text{terms}(\pi_1) + \text{terms}(\pi_2)
\end{array}$$

and we define $\text{terms}(\langle \pi \parallel \pi' \rangle)$ such that $\text{terms}(\langle \pi \parallel \pi' \rangle) = \langle \text{terms}(\pi) \parallel \text{terms}(\pi') \rangle$. We also define the lexicographical ordering

$$\begin{array}{l}
\langle m \parallel n \rangle < \langle m' \parallel n' \rangle \\
\iff m < m' \text{ or } (m = m' \text{ and } n < n')
\end{array}$$

We next show that reducing a pair of prefixes decrements the number of terms in the pair by induction over our reduction rules.

$$\begin{array}{l}
(\langle \pi_1 \parallel \pi'_1 \rangle, \langle \pi_2 \parallel \pi'_2 \rangle) \in R \\
\implies \text{terms}(\langle \pi_1 \parallel \pi'_1 \rangle) < \text{terms}(\langle \pi_2 \parallel \pi'_2 \rangle) \quad (1)
\end{array}$$

- Case [→I]. We have that $\pi_2 = p?l(S). \pi_1$ and $\pi'_2 = p?l(S'). \pi'_1$.

Since $\langle \pi_2 \parallel \pi'_2 \rangle \rightarrow \langle \pi_1 \parallel \pi'_1 \rangle$, we have that, by definition of R , $(\langle \pi_1 \parallel \pi'_1 \rangle, \langle \pi_2 \parallel \pi'_2 \rangle) \in R$.

As we know that $\text{terms}(\pi_2) = \text{terms}(\pi_1) + 1$ and $\text{terms}(\pi'_2) = \text{terms}(\pi'_1) + 1$, we also have that

$$\text{terms}(\langle \pi_1 \parallel \pi'_1 \rangle) < \text{terms}(\langle \pi_2 \parallel \pi'_2 \rangle)$$

as required.

- Case $[\rightarrow 0]$. We have that $\pi_2 = p!\ell(S).\pi_1$ and $\pi'_2 = p!\ell(S').\pi'_1$.

Since $\langle \pi_2 \parallel \pi'_2 \rangle \rightarrow \langle \pi_1 \parallel \pi'_1 \rangle$, we have that, by definition of R , $(\langle \pi_1 \parallel \pi'_1 \rangle, \langle \pi_2 \parallel \pi'_2 \rangle) \in R$.

As we know that $\text{terms}(\pi_2) = \text{terms}(\pi_1) + 1$ and $\text{terms}(\pi'_2) = \text{terms}(\pi'_1) + 1$, we also have that

$$\text{terms}(\langle \pi_1 \parallel \pi'_1 \rangle) < \text{terms}(\langle \pi_2 \parallel \pi'_2 \rangle)$$

as required.

- Case $[\rightarrow \mathcal{A}]$. We have that $\pi_2 = p?\ell(S).\pi_1$ and $\pi'_2 = \mathcal{A}^{(p)}.p?\ell(S').\pi'_1$.

Since $\langle \pi_2 \parallel \pi'_2 \rangle \rightarrow \langle \pi_1 \parallel \mathcal{A}^{(p)}. \pi'_1 \rangle$, by definition of R , we have that $(\langle \pi_1 \parallel \mathcal{A}^{(p)}. \pi'_1 \rangle, \langle \pi_2 \parallel \pi'_2 \rangle) \in R$.

As we know $\text{terms}(\pi_2) = \text{terms}(\pi_1) + 1$ and $\text{terms}(\pi'_2) = \text{terms}(\mathcal{A}^{(p)}. \pi'_1) + 1$, we have that $\text{terms}(\langle \pi_1 \parallel \mathcal{A}^{(p)}. \pi'_1 \rangle) < \text{terms}(\langle \pi_2 \parallel \pi'_2 \rangle)$ as required.

- Case $[\rightarrow \mathcal{B}]$. We have that $\pi_2 = p!\ell(S).\pi_1$ and $\pi'_2 = \mathcal{B}^{(p)}.p!\ell(S').\pi'_1$.

Since $\langle \pi_2 \parallel \pi'_2 \rangle \rightarrow \langle \pi_1 \parallel \mathcal{B}^{(p)}. \pi'_1 \rangle$, by definition of R , we have that $(\langle \pi_1 \parallel \mathcal{B}^{(p)}. \pi'_1 \rangle, \langle \pi_2 \parallel \pi'_2 \rangle) \in R$.

As we know $\text{terms}(\pi_2) = \text{terms}(\pi_1) + 1$ and $\text{terms}(\pi'_2) = \text{terms}(\mathcal{B}^{(p)}. \pi'_1) + 1$, we have that $\text{terms}(\langle \pi_1 \parallel \mathcal{B}^{(p)}. \pi'_1 \rangle) < \text{terms}(\langle \pi_2 \parallel \pi'_2 \rangle)$ as required.

The ordering of terms of pairs is well-founded since both components are bounded from below by 0. Therefore, from Eq. (1), R is also well-founded and so pairs of prefixes cannot be reduced ad infinitum. \square

Theorem 4 (Termination). *Our subtyping algorithm always eventually terminates.*

Proof. We prove termination by arguing that each of our subtyping algorithm rules can be run only a finite number of times.

- $[\text{END}]$ and $[\text{ASM}]$ can each be run only once since they are terminating rules.
- $[\text{SUB}]$ can be run only a finite number of times since a pair of prefixes can be reduced only a finite number of times, as proven in Lemma 3.
- $[\text{OI}]$, $[\text{OO}]$, $[\text{II}]$ and $[\text{IO}]$ can be run only a finite number of times since (1) the number of terms in T and T' is finite and (2) recursion must be explicitly unrolled with $[\mu\text{L}]$ or $[\mu\text{R}]$, which themselves can only be run a finite number of times.

- $[\mu\text{L}]$ and $[\mu\text{R}]$ can be run only a finite number of times since (1) the bounds n and n' are finite; (2) each execution of the rule decrements n or n' respectively; and (3) no rule allows n or n' to be incremented.
- $[\text{TRA}]$ can be run only a finite number of times since (1) the bound k is finite; (2) each execution of the rule decrements k ; and (3) no rule allows k to be incremented. \square

Theorem 5 (Soundness). *Our subtyping algorithm is sound.*

Proof. To prove that our algorithm is sound, we must show that each rule in the precise subtyping by [25] is matched by a rule in our algorithm.

- The subtyping relation rule is matched by $[\text{OI}]$, $[\text{OO}]$, $[\text{II}]$ and $[\text{IO}]$.
- $[\text{REF-END}]$ is matched by $[\text{END}]$. Since both prefixes are empty and $T = T' = \text{end}$, we trivially have that $\pi.T \leq \pi'.T'$.
- The coinductive behaviour of the refinement relation rules is matched by $[\text{ASM}]$. From our map of assumptions Σ , we have that $\pi.T \leq \pi'.T'$.
- $[\text{REF-IN}]$ and $[\text{REF-OUT}]$ are matched by $[\rightarrow \text{I}]$ and $[\rightarrow \text{O}]$ respectively, which can be applied with $[\text{SUB}]$.
- $[\text{REF-}\mathcal{A}]$ and $[\text{REF-}\mathcal{B}]$ are matched by $[\rightarrow \mathcal{A}]$ and $[\rightarrow \mathcal{B}]$ respectively, which can be applied with $[\text{SUB}]$. To ensure that actions are not forgotten when using coinduction, the refinement relation rules additionally require that the actions of the resulting type trees are equal. Our reduction rules emit this check since they deal only with finite sequences. We instead prevent actions from being forgotten in our $[\text{ASM}]$ rule by ensuring that each reordered action in the supertype (which is contained in π') is encountered in the recursive part of the subtype (ρ') by using the subset relation. Ghilezan et al. [25] use the example of checking $T = \mu t.p?\ell.t \leq q?\ell'.T = T'$ to demonstrate why comparing the actions is necessary—we show how our algorithm also correctly rejects this subtype in Fig. A.14 (with a bound of 2 for brevity).

We cannot apply $[\text{ASM}]$ as the final rule since

$$\text{act}(p?\ell.\text{end}) = \{p?\} \not\supseteq \{q?\} = \text{act}(q?\ell'.\text{end}).$$

This subset check spots that the $q?$ action was not present in the recursive part of the supposed subtype. Otherwise, our algorithm would incorrectly conclude that $T \leq T'$.

It is straightforward to also argue that our algorithm preserves reflexivity. We have that $\langle \epsilon, T, n \rangle \leq \langle \epsilon, T, n \rangle$, providing n is sufficiently large to ensure that each recursion can be visited at least once. The prefixes of both sides will always be identical so will reduce to $\langle \epsilon \parallel \epsilon \rangle$ using $[\rightarrow \text{I}]$ or $[\rightarrow \text{O}]$. These

$$\begin{array}{c}
\Sigma_1 [\langle p?l \parallel T \rangle \leq \langle q?l' \parallel T \rangle] = p?l \text{ act}(p?l.\text{end}) \not\leq \text{act}(q?l'.\text{end}) \quad [\text{ASM}] \\
\frac{p?l.p?l; \Sigma_3 \vdash \langle p?l, T, 0 \rangle \leq \langle q?l', T, 1 \rangle}{\frac{p?l.p?l; \Sigma_3 \vdash \langle p?l.p?l, T, 0 \rangle \leq \langle q?l'.p?l, T, 1 \rangle}{\frac{p?l; \Sigma_3 \vdash \langle p?l, p?l.T, 0 \rangle \leq \langle q?l', p?l.T, 1 \rangle}{\frac{p?l; \Sigma_2 \vdash \langle p?l, p?l.T, 1 \rangle \leq \langle q?l', T, 2 \rangle}{\frac{p?l; \Sigma_1 \vdash \langle p?l, T, 1 \rangle \leq \langle q?l', T, 2 \rangle}{\frac{\epsilon; \Sigma_1 \vdash \langle \epsilon, p?l.T, 1 \rangle \leq \langle \epsilon, q?l'.T, 2 \rangle}{\epsilon; \emptyset \vdash \langle \epsilon, T, 2 \rangle \leq \langle \epsilon, q?l'.T, 2 \rangle}}}}}}}} \quad [\text{SUB}] \quad [\text{II}] \quad [\mu\text{R}] \quad [\mu\text{L}] \quad [\text{II}] \quad [\mu\text{L}] \\
\Sigma_1 = [\langle \epsilon \parallel T \rangle \leq \langle \epsilon \parallel T' \rangle \mapsto \epsilon] \quad \Sigma_2 = \Sigma_1 [\langle p?l \parallel T \rangle \leq \langle q?l' \parallel T \rangle \mapsto p?l] \\
\Sigma_3 = \Sigma_2 [\langle p?l \parallel p?l.T \rangle \leq \langle q?l' \parallel T \rangle \mapsto p?l]
\end{array}$$

Figure A.14. Demonstration of how our algorithm correctly prevents actions from being forgotten.

reductions will be applied using [SUB] until either end is encountered and [END] can be applied or the algorithm loops, which allows the application of [ASM]. \square

Lemma 6. *Given finite prefixes π and π' , the time complexity of reducing $\langle \pi \parallel \pi' \rangle$ is $O(\min(|\pi|, |\pi'|))$.*

Proof. We prove this quite simply by induction over our reduction rules.

- Inductive case $\langle p?l.\pi \parallel p?l.\pi' \rangle$. Then, we can perform a [→I] reduction to get $\langle \pi \parallel \pi' \rangle$. By the inductive hypothesis, the complexity of reducing $\langle \pi \parallel \pi' \rangle$ is $O(\min(|\pi|, |\pi'|))$. Therefore, since we require one additional reduction step, the complexity of reducing $\langle p?l.\pi \parallel p?l.\pi' \rangle$ is $O(\min(|\pi|, |\pi'|)+1) = O(\min(|\pi|+1, |\pi'|+1)) = O(\min(|p?l.\pi|, |p?l.\pi'|))$ as required.
- Inductive case $\langle p!l.\pi \parallel p!l.\pi' \rangle$. The proof is the same as for the $\langle p?l.\pi \parallel p?l.\pi' \rangle$ case, except a [→O] reduction is applied.
- Inductive case $\langle p?l.\pi \parallel \mathcal{A}^{(p)}.p?l.\pi' \rangle$. Then, we can perform a [→A] reduction to get $\langle \pi \parallel \mathcal{A}^{(p)}.\pi' \rangle$. By the inductive hypothesis, the complexity of reducing $\langle \pi \parallel \mathcal{A}^{(p)}.\pi' \rangle$ is $O(\min(|\pi|, |\mathcal{A}^{(p)}.\pi'|))$. Therefore, since we require one additional reduction step, the complexity of reducing $\langle p?l.\pi \parallel \mathcal{A}^{(p)}.p?l.\pi' \rangle$ is $O(\min(|\pi|, |\mathcal{A}^{(p)}.\pi'|)+1) = O(\min(|\pi|+1, |\mathcal{A}^{(p)}.\pi'|+1)) = O(\min(|p?l.\pi|, |\mathcal{A}^{(p)}.p?l.\pi'|))$ as required.
- Inductive case $\langle p!l.\pi \parallel \mathcal{B}^{(p)}.p!l.\pi' \rangle$. The proof is the same as for the $\langle p?l.\pi \parallel \mathcal{A}^{(p)}.p?l.\pi' \rangle$ case, except a [→B] reduction is applied.
- Base case $\langle \pi \parallel \pi' \rangle$ where $\langle \pi \parallel \pi' \rangle$ cannot be reduced. Since $\forall \pi. |\pi| > 0$, the complexity of reducing $\langle \pi \parallel \pi' \rangle$ is $O(0) = O(\min(|\pi|, |\pi'|))$ as required. \square

Theorem 7 (Complexity). *Consider T and T' as (possibly infinite) trees $\mathcal{T}(T)$ and $\mathcal{T}(T')$ with asymptotic branching factors b and b' respectively [21, 39]. Our algorithm has time complexity $O(n \min(b, b')^n)$ and space complexity $O(n \min(b, b'))$ in the worst case to determine if $T \leq T'$ with bound n .*

Proof. Let us consider for now only the left tree $\mathcal{T}(T)$ which has asymptotic branching factor b . In the worst case, the number of nodes we have to explore in the tree is

$$1 + b + b^2 + \dots + b^{n-1} = \sum_{i=0}^{n-1} b^i = \frac{b^n - 1}{b - 1}$$

Therefore, exploring the left tree up to a depth of n has time complexity $O(b^n)$. Similarly, exploring the right tree $\mathcal{T}(T')$ to depth n has complexity $O(b'^n)$. Note that in the worst case we cannot reduce any prefixes as we go along and must therefore do the entire reduction at the end of each exploration path.

Suppose at the end of some exploration path we have the pair of prefixes $\langle \pi \parallel \pi' \rangle$. Since $|\pi| = |\pi'|$ (our algorithm does not add an uneven number of terms to either side of the prefix pair), the time complexity of reducing this pair is $O(\min(|\pi|, |\pi'|)) = O(|\pi|)$ from Lemma 8.

At the end of our exploration, we will have b^{n-1} prefixes, each with size n . Therefore, the complexity of reducing all the pairs of prefixes is $O(nb^{n-1})$ so the total complexity for the exploration and reduction is $O(b^n + nb^{n-1}) = O(nb^n)$.

Our algorithm stops when the bound n is reached in either of the two trees so the overall time complexity of exploration is $O(\min(nb^n, nb'^n)) = O(n \min(b, b')^n)$.

Considering space complexity, it is clear that the greatest amount of memory will be required at the end of an exploration path (when the prefixes are greatest in length). For the left tree, at this point, we will need to store (1) a prefix of length n (since we are considering the worst case); and (2) the other b siblings to visit at $n-1$ levels.

Therefore, the total space complexity for the left tree is $O(n + b(n-1)) = O(nb)$ and the space complexity for exploring both trees is $O(\min(nb, nb')) = O(n \min(b, b'))$. \square

B.4 Algorithm Examples

Ring protocol. We again use the ring protocol with choice and show that our algorithm can successfully check the

optimisation to b. The derivation trees are shown in Fig. A.15.

$$T = \mu t.c! \left\{ \begin{array}{l} \text{add.a?add.t} \\ \text{sub.a?add.t} \end{array} \right\} \quad T' = \mu t.a?add.c! \left\{ \begin{array}{l} \text{add.t} \\ \text{sub.t} \end{array} \right\}$$

Alternating bit protocol. We consider the alternating bit protocol [1]. We construct a global type G for the protocol such that when projected onto the receiver, its local type matches the protocol specification.

$$G = \mu t.s \rightarrow r : \left\{ d0.r \rightarrow s : \left\{ \begin{array}{l} a0.\mu u.s \rightarrow r : \left\{ d1.r \rightarrow s : \left\{ \begin{array}{l} a0.u \\ a1.t \end{array} \right\} \right\} \\ a1.t \end{array} \right\} \right\} \quad \text{transitions.len() > 0} \implies !\text{transitions}[0].0$$

$$G \upharpoonright s = \mu t.r!d0.r? \left\{ \begin{array}{l} a0.\mu x.r!d1.r? \left\{ \begin{array}{l} a0.x \\ a1.t \end{array} \right\} \\ a1.t \end{array} \right\}$$

$$G \upharpoonright r = \mu t.s?d0.s! \left\{ \begin{array}{l} a0.\mu x.s?d1.s! \left\{ \begin{array}{l} a0.x \\ a1.t \end{array} \right\} \\ a1.t \end{array} \right\}$$

$$T = \mu t.s? \left\{ \begin{array}{l} d0.s!a0.t \\ d1.s!a1.t \end{array} \right\} \quad T' = G \upharpoonright r$$

We then use our subtyping algorithm to confirm that the type given by the protocol specification for the receiver [1] is a subtype of its projected version. In this derivation, we omit some exploration paths for brevity. The derivation tree is in Fig. A.16.

B.5 Implementation of the Algorithm

In practice, we implement our asynchronous subtyping algorithm on FSMs M and M' rather than local types T and T' . We discuss the practical considerations behind some of our implementation decisions and explain why these are equivalent to the theory presented in § 3.1.

Prefixes. We define prefixes somewhat differently in Rust to avoid copying memory where possible. A prefix is a struct containing three elements:

1. A list of lazy-removable transitions which make up the prefix. A boolean for each element indicates whether the corresponding transition has been lazily removed. A transition is either $p!l(S)$ or $p?l(S)$, which is identical to a prefix term in the theory.
2. A start index, which indicates that the first start elements in transitions should be ignored as they have been lazily removed.
3. A list of indexes of elements that have been lazily removed by setting their boolean to true.

```
1 struct Prefix {
2     transitions: Vec<(bool, Transition)>,
3     start: usize,
4     removed: Vec<usize>,
5 }
```

Elements can be lazily removed either by incrementing `start` or by setting the element's boolean to `true` and adding its index to `removed`. We favour the first option so as to maintain the invariant

where the tuple indexing syntax $(x, y).0$ will evaluate to x , the first element of the tuple. To ensure that this invariant holds, we must advance `start` as far as possible when removing a transition at the head of the prefix.

We also give the option of storing snapshots to previous versions of a prefix. A snapshot stores (1) the size of the transitions list; (2) the value of the `start` field; and (3) the size of the removed list, all taken at the time of the snapshot.

```
1 struct Snapshot {
2     size: usize,
3     start: usize,
4     removed: usize,
5 }
```

We can easily revert a prefix to a previous snapshot by (1) finding the elements of `removed` that have been added since the snapshot; (2) setting the boolean to false for each of these elements to restore them; (3) truncating transitions to its previous size; (4) restoring `start` to its previous value; and (5) truncating `removed` to its previous size.

Visitor. We use the visitor pattern [50] to traverse a pair of FSMs M and M' . In our visitor, we store (1) the fsm's we are traversing; (2) a matrix of history (as we will see, this is equivalent to the assumptions map Σ in the theory); and (3) a pair of prefixes, as in the theory.

```
1 struct SubtypeVisitor {
2     fsm's: Pair<Fsm>,
3     history: Matrix<Previous>,
4     prefixes: Pair<Prefix>,
5 }
```

The history matrix stores a value for each combination of states in M and M' (it effectively has the type $|M| \times |M'| \rightarrow \text{Previous}$). Each of these values stores a `Previous` struct containing the number of visits this combination of states has remaining and optionally (if it has been visited before) a pair of snapshots taken during the last visit to this combination.

$$\begin{array}{c}
\frac{(\star) \rho_3; \Sigma_2 \vdash \langle c!add.a?add, T, 0 \rangle \leq \langle a?add.c!sub, T', 0 \rangle \quad (\dagger) \rho_4; \Sigma_2 \vdash \langle c!sub.a?add, T, 0 \rangle \leq \langle a?add.c!add, T', 0 \rangle}{\rho_1; \Sigma_2 \vdash \langle c!add, a?add, T, 0 \rangle \leq \langle a?add, c! \left\{ \begin{array}{l} add.T' \\ sub.T' \end{array} \right\}, 0 \rangle} \text{[IO]} \quad \frac{\quad}{\rho_2; \Sigma_2 \vdash \langle c!sub, a?add, T, 0 \rangle \leq \langle a?add, c! \left\{ \begin{array}{l} add.T' \\ sub.T' \end{array} \right\}, 0 \rangle} \text{[IO]} \\
\hline
\frac{\epsilon; \Sigma_2 \vdash \langle \epsilon, c! \left\{ \begin{array}{l} add.a?add.T \\ sub.a?add.T \end{array} \right\}, 0 \rangle \leq \langle \epsilon, a?add.c! \left\{ \begin{array}{l} add.T' \\ sub.T' \end{array} \right\}, 0 \rangle}{\epsilon; \Sigma_1 \vdash \langle \epsilon, c! \left\{ \begin{array}{l} add.a?add.T \\ sub.a?add.T \end{array} \right\}, 0 \rangle \leq \langle \epsilon, T', 1 \rangle} \text{[}\mu\text{R]} \\
\hline
\frac{\quad}{\epsilon; \emptyset \vdash \langle \epsilon, T, 1 \rangle \leq \langle \epsilon, T', 1 \rangle} \text{[}\mu\text{L]} \\
(\star) = \frac{\frac{\langle c!add.a?add \parallel a?add.c!add \rangle \rightarrow \langle a?add \parallel a?add \rangle \quad \frac{\langle a?add \parallel a?add \rangle \rightarrow \langle \epsilon \parallel \epsilon \rangle \quad \frac{\text{act}(\rho_3.\text{end}) \supseteq \text{act}(\text{end})}{\rho_3; \Sigma_2 \vdash \langle \epsilon, T, 0 \rangle \leq \langle \epsilon, T', 0 \rangle} \text{[ASM]}}{\rho_3; \Sigma_2 \vdash \langle a?add, T, 0 \rangle \leq \langle a?add, T', 0 \rangle} \text{[SUB]}}{\rho_3; \Sigma_2 \vdash \langle c!add.a?add, T, 0 \rangle \leq \langle a?add.c!add, T', 0 \rangle} \text{[}\mathcal{B}\text{]} \\
(\dagger) = \frac{\frac{\langle c!sub.a?add \parallel a?add.c!sub \rangle \rightarrow \langle a?add \parallel a?add \rangle \quad \frac{\langle a?add \parallel a?add \rangle \rightarrow \langle \epsilon \parallel \epsilon \rangle \quad \frac{\text{act}(\rho_4.\text{end}) \supseteq \text{act}(\text{end})}{\rho_4; \Sigma_2 \vdash \langle \epsilon, T, 0 \rangle \leq \langle \epsilon, T', 0 \rangle} \text{[ASM]}}{\rho_4; \Sigma_2 \vdash \langle a?add, T, 0 \rangle \leq \langle a?add, T', 0 \rangle} \text{[SUB]}}{\rho_4; \Sigma_2 \vdash \langle c!sub.a?add, T, 0 \rangle \leq \langle a?add.c!sub, T', 0 \rangle} \text{[}\mathcal{B}\text{]} \\
\rho_1 = c!add \quad \rho_2 = c!sub \quad \rho_3 = \rho_1.a?add \quad \rho_4 = \rho_2.a?add \\
\Sigma_1 = [\langle \epsilon \parallel T \rangle \leq \langle \epsilon \parallel T' \rangle \mapsto \epsilon] \quad \Sigma_2 = \Sigma_1 \left[\langle \epsilon \parallel c! \left\{ \begin{array}{l} add.a?add.T \\ sub.a?add.T \end{array} \right\} \rangle \leq \langle \epsilon \parallel T' \rangle \mapsto \epsilon \right]
\end{array}$$

Figure A.15. Derivation trees to verify the subtyping of the Ring protocol

```

1 struct Previous {
2   visits: usize,
3   snapshots: Option<Pair<Snapshot>>,
4 }

```

In the theory, termination is guaranteed by allowing recursions to be unrolled only n times. Here, our ‘ n ’ is the value of `visits`, which limits how many times the same combination of states can be visited. Since M and M' each contain a finite number of states and their cross product is also finite, this will achieve termination just as in the theory (provided that n is also finite). Otherwise, this is identical to the theory—our history matrix corresponds to the map of assumptions Σ and the `Previous` struct represents a single mapping (we use snapshots in place of prefixes).

Each state in an `Fsm` is given a unique `StateIndex` that identifies it. Our `Visitor` is executed using its recursive `visit` method, which takes a mutable reference to the `Visitor` and a `StateIndex` for each `Fsm`.

```

1 impl Visitor {
2   fn visit(&mut self, states: Pair<StateIndex>) ->
3     bool {
4     [...]
5   }
6 }

```

This `visit` method performs our asynchronous subtyping algorithm as follows.

1. We look up the current combination states in our history to ensure visits is positive, as in `[μL]` and `[μR]`. If it is not then our bound has been exhausted and we return with false.
2. We attempt to reduce the pair of prefixes, as in `[SUB]`. This reduction process follows precisely the same rules as in the theory, lazily removing transitions where appropriate.
3. If the current combination of states has been visited before, we attempt to use our assumptions map to return true, as in `[ASM]`. The method we use to check the actions sets, explained below, differs slightly from the theory.
4. If both FSMs are in a *terminal* state and the prefixes are empty then we return true, as in `[END]`.
5. If both FSMs are in a *non-terminal* state then we
 - take a snapshot of the current prefixes;
 - update the history matrix for the current combination of states, setting `visits` to `visits - 1` and `snapshots` to the snapshots we just took;
 - for each pair of transitions we can take from the current combination of states we
 - add each transition in the pair to its corresponding prefix;
 - recurse using the `visit` method, setting the `states` argument to the pair of end states corresponding

$$\begin{array}{c}
\frac{\text{act}(s?d0.s!a0.\text{end}) \supseteq \text{act}(\text{end})}{s?d0.s!a0; \Sigma_4 \vdash \langle \epsilon, T, 0 \rangle \leq \langle \epsilon, T', 0 \rangle} \text{[ASM]} \\
\frac{s?d0.s!a0; \Sigma_4 \vdash \langle \epsilon, T, 0 \rangle \leq \langle \epsilon, T', 0 \rangle}{s?d0.s!a0; \Sigma_4 \vdash \langle s!a1, T, 0 \rangle \leq \langle s!a1, T', 0 \rangle} \text{[SUB]} \quad \dots \text{[OUT-OUT]} \\
\frac{s?d0.s!a0; \Sigma_4 \vdash \langle \epsilon, s!a1.T, 0 \rangle \leq \langle \epsilon, s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 0 \rangle}{s?d0.s!a0; \Sigma_4 \vdash \langle s?d1, s!a1.T, 0 \rangle \leq \langle s?d1, s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 0 \rangle} \text{[SUB]} \\
\frac{s?d0.s!a0; \Sigma_4 \vdash \langle s?d1, s!a1.T, 0 \rangle \leq \langle s?d1, s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 0 \rangle}{s?d0.s!a0; \Sigma_4 \vdash \langle \epsilon, s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\}, 0 \rangle \leq \langle \epsilon, s?d1.s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 0 \rangle} \text{[IN-IN]} \\
\frac{s?d0.s!a0; \Sigma_4 \vdash \langle \epsilon, s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\}, 0 \rangle \leq \langle \epsilon, s?d1.s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 0 \rangle}{s?d0.s!a0; \Sigma_3 \vdash \langle \epsilon, s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\}, 0 \rangle \leq \langle \epsilon, T_1, 1 \rangle} \text{[}\mu\text{R]} \\
\frac{s?d0.s!a0; \Sigma_3 \vdash \langle \epsilon, s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\}, 0 \rangle \leq \langle \epsilon, T_1, 1 \rangle}{s?d0.s!a0; \Sigma_2 \vdash \langle \epsilon, T, 1 \rangle \leq \langle \epsilon, T_1, 1 \rangle} \text{[}\mu\text{L]} \\
\frac{s?d0.s!a0; \Sigma_2 \vdash \langle \epsilon, T, 1 \rangle \leq \langle \epsilon, T_1, 1 \rangle}{s?d0.s!a0; \Sigma_2 \vdash \langle s!a0, T, 1 \rangle \leq \langle s!a0, T_1, 1 \rangle} \text{[SUB]} \quad \dots \text{[OUT-OUT]} \\
\frac{s?d0; \Sigma_2 \vdash \langle \epsilon, s!a0.T, 1 \rangle \leq \langle \epsilon, s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 1 \rangle}{s?d0; \Sigma_2 \vdash \langle s?d0, s!a0.T, 1 \rangle \leq \langle s?d0, s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 1 \rangle} \text{[SUB]} \\
\frac{s?d0; \Sigma_2 \vdash \langle s?d0, s!a0.T, 1 \rangle \leq \langle s?d0, s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 1 \rangle}{\epsilon; \Sigma_2 \vdash \langle \epsilon, s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\}, 1 \rangle \leq \langle \epsilon, s?d0.s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 1 \rangle} \text{[IN-IN]} \\
\frac{\epsilon; \Sigma_2 \vdash \langle \epsilon, s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\}, 1 \rangle \leq \langle \epsilon, s?d0.s! \left\{ \begin{array}{l} a0.T_1 \\ a1.T' \end{array} \right\}, 1 \rangle}{\epsilon; \Sigma_1 \vdash \langle \epsilon, s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\}, 1 \rangle \leq \langle \epsilon, T', 2 \rangle} \text{[}\mu\text{R]} \\
\frac{\epsilon; \Sigma_1 \vdash \langle \epsilon, s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\}, 1 \rangle \leq \langle \epsilon, T', 2 \rangle}{\epsilon; \emptyset \vdash \langle \epsilon, T, 2 \rangle \leq \langle \epsilon, T', 2 \rangle} \text{[}\mu\text{L]} \\
T_1 = \mu x. s?d1.s! \left\{ \begin{array}{l} a0.x \\ a1.T' \end{array} \right\} \\
\Sigma_1 = [\langle \epsilon \parallel T \rangle \leq \langle \epsilon \parallel T' \rangle \mapsto \epsilon] \quad \Sigma_2 = \Sigma_1 \left[\langle \epsilon \parallel s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\} \rangle \leq \langle \epsilon \parallel T' \rangle \mapsto \epsilon \right] \\
\Sigma_3 = \Sigma_2 [\langle \epsilon \parallel T \rangle \leq \langle \epsilon \parallel T_1 \rangle \mapsto s?d0.s!a0] \quad \Sigma_4 = \Sigma_3 \left[\langle \epsilon \parallel s? \left\{ \begin{array}{l} d0.s!a0.T \\ d1.s!a1.T \end{array} \right\} \rangle \leq \langle \epsilon \parallel T_1 \rangle \mapsto s?d0.s!a0 \right]
\end{array}$$

Figure A.16. Derivation trees to verify the subtyping of the Alternating-Bit protocol

- to our transitions; and after the recursive call returns
- revert the changes made to the prefixes by using the snapshot we took previously;
 - restore the current history matrix entry to its original value; and
 - return a value depending on the results of the recursive calls and whether the current combination of states performs input or output actions, as described by the quantifiers in $\{[\text{IN}, \text{OUT}]\text{-}\{\text{IN}, \text{OUT}\}\}$.
6. Otherwise, one of the FSMs has reached a terminal state but the other has not. In this case, there is no way to progress and we return false.

By performing a depth-first search we can make changes to the history and prefixes fields of our visitor and revert them later, using a snapshot for each prefix. This method improves the efficiency of our algorithm by avoiding copying memory. If we instead used a breadth-first search, for instance, we would need to store a separate visitor for each

frontier of our search. This would require an expensive copy of the history and prefixes.

Checking actions. In the theory, our [ASM] rule compares two sets of actions to ensure that it is safe to apply an assumption. Specifically, it checks that the actions of the supposed supertype’s prefix (π') are a subset of the actions performed by the subtype since the assumption was made (ρ'). In our algorithm, we can actually perform a far cheaper but equivalent check thanks to our use of lazy removal. We need only to confirm that

$$\begin{aligned}
&\text{transitions}[\text{start}..] == \\
&\text{transitions}[\dots \text{snapshot.size}][\text{snapshot.start}..] \tag{2}
\end{aligned}$$

for each prefix/snapshot combination. The syntax $x[i..]$ evaluates to x with the first i elements removed and $y[. . . j]$ evaluates to the first j elements of y . Surprisingly, this check is identical in effect to the one performed in the theory due to two observations.

1. Comparing the full list of transitions (which include labels and sorts) rather than only their actions is sound since the reduction rules do not allow sends or receives to or from the same participant to be reordered.

We can easily prove this by contradiction. Suppose $p?l(S) \in \pi'$ and $p?l'(S') \in \rho'$ and we can apply [ASM]. Clearly, $p?l(S)$ has not been reduced by [REF-IN], otherwise, it would not still be in π' . Therefore, at some point since the assumption was added to Σ , [REF- \mathcal{A}] must have been used to move $p?l(S)$ before $p?l'(S')$. This is a contradiction because $\mathcal{A}^{(p)}$ cannot contain $p?l'(S')$ by definition so [REF- \mathcal{A}] cannot have been applied. A similar argument can be made for the output case.

2. The version in the theory is intuitively checking whether there is an action that ‘hangs on’ to the far left of π' for multiple iterations of a recursive type without ever being reduced. If this is the case, then the action will not be matched by any of the actions in ρ' (otherwise it would have been reduced) so $\rho' \not\preceq \pi'$.

In our implementation, if an action hangs on to the supertype’s prefix then it will never be lazily removed. This means that the size of the prefix will grow on each iteration of the FSM since `start` is never advanced. Since

```
transitions[start..].len() !=
transitions[..snapshot.size][snapshot.start..].len()
```

Eq. (2) is trivially false. Note that the full check in Eq. (2) must be performed, rather than only comparing the lengths, to ensure that the prefixes do actually match those of the assumption, as in [ASM].

Fail-early reductions. Our practical implementation performs the same reduction rules on prefixes as described in the theory. However, we add a practical optimisation to, in some cases, determine that a particular path cannot succeed before even reaching the bound.

For example, consider the pair $\langle p?l(S).\pi \parallel q!l'(S').p?l(S).\pi' \rangle$. Regardless of what π and π' are set to, this pair cannot be reduced as it will require using the [REF- \mathcal{A}] but $q!l'(S')$ cannot be contained in $\mathcal{A}^{(p)}$. Therefore, if at some point we reach a pair of prefixes which looks like $\langle p?l(S).p?l(S) \parallel q!l'(S').p?l(S) \rangle$, we can immediately return `false` as there is no way that it can ever be reduced by adding more terms.

C Benchmarking results

C.1 Session-Based Rust Implementations

Results for the stream benchmark.

n	Throughput ($n/\mu s$)				
	SESH	MULTICRUSTY	FERRITE	RUMPSTEAK	RUMPSTEAK (opt.)
10	0.019389	0.011678	0.011386	0.202587	0.215583
20	0.028142	0.014325	0.012994	0.336988	0.356978
30	0.034193	0.015160	0.013463	0.427489	0.437795
40	0.036566	0.016072	0.013671	0.488886	0.517468
50	0.040315	0.016577	0.014126	0.545378	0.583366

Results for the double buffering benchmark.

n	Throughput ($n/\mu s$)				
	SESH	MULTICRUSTY	FERRITE	RUMPSTEAK	RUMPSTEAK (opt.)
5000	6.929567	5.675414	7.617643	27.704354	32.340989
10000	13.138401	11.254181	14.649028	44.154722	50.126532
15000	18.739983	16.187341	20.429845	56.813002	67.884430
20000	24.103215	20.481378	25.506427	67.595301	82.039366
25000	28.609966	25.050058	29.629025	75.848611	96.010424

Results for the FFT benchmark.

n	Throughput ($n/\mu s$)				
	SESH	MULTICRUSTY	FERRITE	RUSTFFT	RUMPSTEAK
1000	0.551154	0.810134	1.458279	9.320778	5.038554
2000	1.050958	1.515538	2.513855	9.313359	7.206404
3000	1.510567	2.163629	3.496405	9.333569	8.421026
4000	1.935263	2.783617	4.198723	9.336939	9.262763
5000	2.303627	3.261020	4.811375	9.323199	9.316716

C.2 Verifying Asynchronous Message Reordering

Results for the stream benchmark.

n	Running time (s)		
	SOUNDBINARY	k -MC	RUMPSTEAK
0	0.003476	0.005504	0.001872
10	0.008556	0.019316	0.001899
20	0.020673	0.057417	0.001848
30	0.041673	0.142145	0.001906
40	0.076425	0.276446	0.001874
50	0.127865	0.496929	0.002080
60	0.198541	0.805577	0.002083
70	0.292471	1.233327	0.002064
80	0.422571	1.780778	0.002178
90	0.583863	2.475443	0.002190
100	0.767426	3.349204	0.002249

Results for the nested choice benchmark.

n	Running time (s)		
	SOUNDBINARY	k -MC	RUMPSTEAK
1	0.002295	0.006554	0.000702
2	0.004504	0.014901	0.000755
3	0.016347	0.072423	0.001745
4	0.224858	1.515528	0.007656
5	4.692525	41.688068	0.157548

Results for the ring benchmark.

n	Running time (s)	
	k -MC	RUMPSTEAK
2	0.004007	0.000675
4	0.007239	0.000731
6	0.011806	0.000701
8	0.018822	0.000835
10	0.024842	0.000757
12	0.049232	0.000777
14	0.102257	0.000744
16	0.191078	0.000813
18	0.340262	0.000817
20	0.570656	0.000766
22	0.913412	0.000911
24	1.391075	0.000737
26	2.042452	0.000752
28	2.918943	0.000732
30	4.099072	0.000769

Results for k -buffering benchmark.

n	Running time (s)	
	k -MC	RUMPSTEAK
0	0.004825	0.000630
5	0.007668	0.000747
10	0.013613	0.000705
15	0.018770	0.000667
20	0.031376	0.000825
25	0.054910	0.000718
30	0.080879	0.000760
35	0.122315	0.000853
40	0.170533	0.000802
45	0.236354	0.000792
50	0.305749	0.000916
55	0.406071	0.000882
60	0.506069	0.000959
65	0.639521	0.001028
70	0.773931	0.001057
75	0.954399	0.001045
80	1.127240	0.001125
85	1.359600	0.001120
90	1.571745	0.001164
95	1.869339	0.001156
100	2.111687	0.001234

References

- [1] [n.d.]. Introduction to Protocol Engineering. <http://cs.uccs.edu/~cs522/pe/pe.htm>
- [2] [n.d.]. *vScr*. <https://github.com/nuscr/nuscr>
- [3] [n.d.]. *Rumpsteak*. <https://github.com/zakcutner/rumpsteak>
- [4] Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-Malo Deniérou, Simon J. Gay, Nils Geshbert, Elena Giachino, Raymond Hu, Einar Broch Johnsen, Francisco Martins, Viviana Mascardi, Fabrizio Montesi, Rumyana Neykova, Nicholas Ng, Luca Padovani, Vasco T. Vasconcelos, and Nobuko Yoshida. 2016. Behavioral Types in Programming Languages. *Foundations and Trends in Programming Languages* 3, 2-3 (2016), 95–230. <https://doi.org/10.1561/25000000031>
- [5] Stephanie Balzer and Frank Pfenning. 2017. Manifest Sharing with Session Types. *Proceedings of the ACM on Programming Languages* 1, ICFP (2017), 1–29. <https://doi.org/10.1145/3110281>
- [6] Ahmed Bouajjani, Constantín Enea, Kailiang Ji, and Shaz Qadeer. 2018. On the Completeness of Verifying Message Passing Programs Under Bounded Asynchrony. In *Computer Aided Verification (LNCS, Vol. 10982)*. Springer, 372–391. https://doi.org/10.1007/978-3-319-96142-2_23
- [7] Mario Bravetti, Marco Carbone, Julien Lange, Nobuko Yoshida, and Gianluigi Zavattaro. 2021. A Sound Algorithm for Asynchronous Session Subtyping and its Implementation. *Logical Methods in Computer Science* 17 (2021), Issue 1. [https://doi.org/10.23638/LMCS-17\(1:20\)2021](https://doi.org/10.23638/LMCS-17(1:20)2021) (repository is found at <https://github.com/julien-lange/asynchronous-subtyping>).
- [8] Mario Bravetti, Marco Carbone, and Gianluigi Zavattaro. 2017. Undecidability of Asynchronous Session Subtyping. *Information and Computation* 256, C (2017), 300–320. <https://doi.org/10.1016/j.ic.2017.07.010>
- [9] Mario Bravetti, Marco Carbone, and Gianluigi Zavattaro. 2018. On the Boundary Between Decidability and Undecidability of Asynchronous Session Subtyping. *Theoretical Computer Science* 722 (2018), 19–51. <https://doi.org/10.1016/j.tcs.2018.02.010>
- [10] David Castro, Raymond Hu, Sung-Shik Jongmans, Nicholas Ng, and Nobuko Yoshida. 2019. Distributed Programming Using Role-Parametric Session Types in Go: Statically-Typed Endpoint APIs

- for Dynamically-Instantiated Communication Structures. *Proceedings of the ACM on Programming Languages* 3, POPL (2019), 1–30. <https://doi.org/10.1145/3290342>
- [11] David Castro-Perez and Nobuko Yoshida. 2020. CAMP: Cost-Aware Multiparty Session Protocols. *Proceedings of the ACM on Programming Languages* 4, OOPSLA (2020), 1–30. <https://doi.org/10.1145/3428223>
- [12] Ruofei Chen and Stephanie Balzer. 2021. Ferrite: A Judgmental Embedding of Session Types in Rust. arXiv:2009.13619 (repository is found at <https://github.com/ferrite-rs/ferrite>).
- [13] Tzu-Chun Chen, Mariangiola Dezani-Ciancaglini, Alceste Scalas, and Nobuko Yoshida. 2017. On the Preciseness of Subtyping in Session Types. *Logical Methods in Computer Science* 13 (2017), Issue 2. [https://doi.org/10.23638/LMCS-13\(2:12\)2017](https://doi.org/10.23638/LMCS-13(2:12)2017)
- [14] Jaemin Choi, David F. Richards, and Laxmikant V. Kale. 2020. Achieving Computation-Communication Overlap with Overdecomposition on GPU Systems. In *2020 IEEE/ACM Fifth International Workshop on Extreme Scale Programming Models and Middleware (ESPM2)*. IEEE, 1–10. <https://doi.org/10.1109/ESPM251964.2020.00006>
- [15] Alex Crichton. [n.d.]. *Futures*. <https://github.com/rust-lang/futures-rs>
- [16] Zak Cutner, Nobuko Yoshida, and Martin Vassor. 2021. *Artifact: Deadlock-Free Asynchronous Message Reordering in Rust with Multiparty Session Types*. <https://doi.org/10.5281/zenodo.5786034>
- [17] Zak Cutner, Nobuko Yoshida, and Martin Vassor. 2021. Deadlock-Free Asynchronous Message Reordering in Rust with Multiparty Session Types. (repository is found at <https://arxiv.org/abs/2112.12693>).
- [18] Romain Demangeon, Kohei Honda, Raymond Hu, Rumyana Neykova, and Nobuko Yoshida. 2015. Practical interruptible conversations: Distributed dynamic verification with multiparty session types and Python. *FMSD* (2015), 1–29.
- [19] Ryan Donovan. 2020. *Why the Developers Who Use Rust Love It so Much*. <https://stackoverflow.blog/2020/06/05/why-the-developers-who-use-rust-love-it-so-much/>
- [20] José Duarte and António Ravara. 2021. Retrofitting Typestates into Rust. In *Proceedings of the 24th Brazilian Symposium on Context-Oriented Programming and Advanced Modularity (SBLP)*. ACM. <https://github.com/rusttype/typestate-rs/blob/main/paper/sblp21.pdf> (to appear).
- [21] Stefan Edelkamp and Richard E. Korf. 1998. The Branching Factor of Regular Search Spaces. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI)*. AAAI, 299–304. <https://www.aaai.org/Papers/AAAI/1998/AAAI98-042.pdf>
- [22] Simon Gay and Malcolm Hole. 2005. Subtyping for Session Types in the Pi Calculus. *Acta Informatica* 42, 2–3 (2005), 191–225. <https://doi.org/10.1007/s00236-005-0177-z>
- [23] Simon Gay and António Ravara. 2017. Behavioural Types: from Theory to Tools. 1–412. <https://doi.org/10.13052/rp-9788793519817>
- [24] Simon Gay and António Ravara. 2017. *Behavioural Types: from Theory to Tools*. River Publisher, 1–412 pages. <https://doi.org/10.13052/rp-9788793519817>
- [25] Silvia Ghilezan, Jovanka Pantović, Ivan Prokić, Alceste Scalas, and Nobuko Yoshida. 2021. Precise Subtyping for Asynchronous Multiparty Sessions. *Proceedings of the ACM on Programming Languages* 5, POPL (2021), 1–28. <https://doi.org/10.1145/3434297>
- [26] Paul Harvey, Simon Fowler, Ornella Dardha, and Simon J. Gay. 2021. Multiparty Session Types for Safe Runtime Adaptation in an Actor Language. In *35th European Conference on Object-Oriented Programming (LIPICs, Vol. 194)*. Schloss Dagstuhl, 10:1–10:30. <https://doi.org/10.4230/LIPICs.ECOOP.2021.10>
- [27] Brook Heisler. [n.d.]. *Criterion.rs*. <https://github.com/bheisler/criterion.rs>
- [28] Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Disciplines for Structured Communication-based Programming. In *Programming Languages and Systems (LNCS, Vol. 1381)*. Springer, 122–138. <https://doi.org/10.1007/bfb0053567>
- [29] Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008. Multiparty Asynchronous Session Types. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. ACM, 273–284. <https://doi.org/10.1145/1328438.1328472>
- [30] Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2016. Multiparty Asynchronous Session Types. *Journal of the ACM* 63, 1 (2016), 1–67. <https://doi.org/10.1145/2827695>
- [31] Raymond Hu and Nobuko Yoshida. 2016. Hybrid Session Verification through Endpoint API Generation. In *Fundamental Approaches to Software Engineering (LNCS, Vol. 9633)*. Springer, 401–418. https://doi.org/10.1007/978-3-662-49665-7_24
- [32] Raymond Hu and Nobuko Yoshida. 2017. Explicit Connection Actions in Multiparty Session Types. In *Fundamental Approaches to Software Engineering (LNCS, Vol. 10202)*. Springer, 116–133. https://doi.org/10.1007/978-3-662-54494-5_7
- [33] Hai Huang, Padmanabhan Pillai, and Kang G. Shin. 2002. Improving Wait-Free Algorithms for Interprocess Communication in Embedded Real-Time Systems. In *2002 USENIX Annual Technical Conference*. USENIX Association. [https://www.usenix.org/legacy/events/usenix02/Huang.html](https://www.usenix.org/legacy/events/usenix02/ Huang.html)
- [34] Thomas Bracht Laumann Jespersen, Philip Munksgaard, and Ken Friis Larsen. 2015. Session Types for Rust. In *Proceedings of the 11th ACM SIGPLAN Workshop on Generic Programming (WGP)*. ACM, 13–22. <https://doi.org/10.1145/2808098.2808100>
- [35] Jonas Kastberg Hinrichsen. 2021. *Sessions and Separation*. Ph.D. Dissertation. IT University of Copenhagen, Copenhagen. <http://itu.dk/people/jkas/papers/thesis.pdf>
- [36] Ralf Jung. 2020. Understanding and evolving the Rust programming language. <https://doi.org/10.22028/D291-31946>
- [37] Ki-Hwan Kim and Q-Han Park. 2012. Overlapping Computation and Communication of Three-Dimensional FDTD on a GPU Cluster. *Computer Physics Communications* 183, 11 (2012), 2364–2369. <https://doi.org/10.1016/j.cpc.2012.06.003>
- [38] Wen Kokke. 2019. Rusty Variation: Deadlock-free Sessions with Failure in Rust. *Electronic Proceedings in Theoretical Computer Science* 304 (2019), 48–60. <https://doi.org/10.4204/eptcs.304.4> (repository is found at <https://github.com/wenkokke/sesh>).
- [39] Richard E. Korf. 1985. Depth-First Iterative-Deepening: An Optimal Admissible Tree Search. *Artificial Intelligence* 27, 1 (1985), 97–109. [https://doi.org/10.1016/0004-3702\(85\)90084-0](https://doi.org/10.1016/0004-3702(85)90084-0)
- [40] Dimitrios Kouzapas, Ornella Dardha, Roly Perera, and Simon J. Gay. 2018. Typechecking Protocols with Mungo and StMungo: A Session Type Toolchain for Java. *Science of Computer Programming* 155 (2018), 52–75. <https://doi.org/10.1016/j.scico.2017.10.006>
- [41] Nicolas Lagailardie, Rumyana Neykova, and Nobuko Yoshida. 2020. Implementing Multiparty Session Types in Rust. In *Coordination Models and Languages (LNCS, Vol. 12134)*. Springer, 127–136. https://doi.org/10.1007/978-3-030-50029-0_8 (repository is found at https://github.com/NicolasLagailardie/mpst_rust_github).
- [42] Julien Lange and Nobuko Yoshida. 2017. On the Undecidability of Asynchronous Session Subtyping. In *Foundations of Software Science and Computation Structures (LNCS, Vol. 10203)*. Springer, 441–457. https://doi.org/10.1007/978-3-662-54458-7_26
- [43] Julien Lange and Nobuko Yoshida. 2019. Verifying Asynchronous Interactions via Communicating Session Automata. In *Computer Aided Verification (LNCS, Vol. 11561)*. Springer, 97–117. https://doi.org/10.1007/978-3-030-25540-4_6 (repository is found at <https://bitbucket.org/julien-lange/kmc-cav19>).
- [44] Message Passing Interface Forum. 2021. *MPI: A Message-Passing Interface Standard*. <https://www.mpi-forum.org/docs/mpi-4.0/mpi40-report.pdf> (version 4.0).
- [45] Anson Miu, Francisco Ferreira, Nobuko Yoshida, and Fangyi Zhou. 2021. Communication-Safe Web Programming in TypeScript with Routed Multiparty Session Types. In *Proceedings of the 30th ACM*

- SIGPLAN International Conference on Compiler Construction (CC). ACM, 94–106. <https://doi.org/10.1145/3446804.3446854>
- [46] Romyana Neykova, Raymond Hu, Nobuko Yoshida, and Fahd Abdelljallal. 2018. A Session Type Provider: Compile-time API Generation of Distributed Protocols with Refinements in F#. In *Proceedings of the 27th International Conference on Compiler Construction (CC)*. ACM, 128–138. <https://doi.org/10.1145/3178372.3179495>
- [47] Romyana Neykova and Nobuko Yoshida. 2017. Let It Recover: Multi-party Protocol-Induced Recovery. In *26th International Conference on Compiler Construction*. ACM, 98–108.
- [48] Romyana Neykova, Nobuko Yoshida, and Raymond Hu. 2013. SPY: Local Verification of Global Protocols. In *Runtime Verification (LNCS, Vol. 8174)*. Springer, 358–363. https://doi.org/10.1007/978-3-642-40787-1_25
- [49] Nicholas Ng, Jose Gabriel de Figueiredo Coutinho, and Nobuko Yoshida. 2015. Protocols by Default. In *Compiler Construction (LNCS, Vol. 9031)*. Springer, 212–232. https://doi.org/10.1007/978-3-662-46663-6_11
- [50] J. Palsberg and C.B. Jay. 1998. The Essence of the Visitor Pattern. In *Proceedings of the 37th Annual Computer Software and Applications Conference*. IEEE Computer Society, 9–15. <https://doi.org/10.1109/CMPASAC.1998.716629>
- [51] David Peter. [n.d.]. *Hyperfine*. <https://github.com/sharkdp/hyperfine>
- [52] Boqin Qin. 2020. rust-lock-bug-detector: Statically Detect Double-Lock & Conflicting-Lock Bugs on MIR. <https://github.com/BurtonQin/rust-lock-bug-detector>
- [53] Boqin Qin, Yilun Chen, Zeming Yu, Linhai Song, and Yiying Zhang. 2020. Understanding Memory and Thread Safety Practices and Issues in Real-World Rust Programs. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. ACM, 763–779. <https://doi.org/10.1145/3385412.3386036>
- [54] Alceste Scalas, Ornela Dardha, Raymond Hu, and Nobuko Yoshida. 2017. A Linear Decomposition of Multiparty Sessions for Safe Distributed Programming. In *31st European Conference on Object-Oriented Programming (LIPICs, Vol. 74)*. Schloss Dagstuhl, 24:1–24:31. <https://doi.org/10.4230/LIPICs.ECOOP.2017.24>
- [55] Scribble Authors. [n.d.]. *Scribble: Describing Multi Party Protocols*. <http://www.scribble.org/>
- [56] Marc Sergent, Mario Dagrada, Patrick Carribault, Julien Jaeger, Marc Pérache, and Guillaume Papauré. 2018. Efficient Communication-Computation Overlap with MPI+OpenMP Runtimes Collaboration. In *Parallel Processing (LNCS, Vol. 11014)*. Springer, 560–572. https://doi.org/10.1007/978-3-319-96983-1_40
- [57] Kaku Takeuchi, Kohei Honda, and Makoto Kubo. 1994. An Interaction-based Language and its Typing System. In *Parallel Architectures and Languages Europe (LNCS, Vol. 817)*. Springer, 398–413. https://doi.org/10.1007/3-540-58184-7_118
- [58] The Rust Project Developers. [n.d.]. *Procedural Macros*. <https://doc.rust-lang.org/reference/procedural-macros.html>
- [59] The Rust Project Developers. [n.d.]. *The MIR (Mid-level IR)*. <https://rustc-dev-guide.rust-lang.org/mir/index.html>
- [60] Tokio Contributors. [n.d.]. Tokio. <https://github.com/tokio-rs/tokio>
- [61] Allen Welkie and Elliott Mahler. [n.d.]. *RustFFT*. <https://github.com/ejmahler/RustFFT>
- [62] Nobuko Yoshida and Lorenzo Gheri. 2020. A Very Gentle Introduction to Multiparty Session Types. In *Distributed Computing and Internet Technology (LNCS, Vol. 11969)*. Springer, 73–93. https://doi.org/10.1007/978-3-030-36987-3_5
- [63] Nobuko Yoshida, Raymond Hu, Romyana Neykova, and Nicholas Ng. 2013. The Scribble Protocol Language. In *Trustworthy Global Computing (LNCS, Vol. 8358)*. Springer, 22–41. https://doi.org/10.1007/978-3-319-05119-2_3
- [64] Fangyi Zhou, Francisco Ferreira, Raymond Hu, Romyana Neykova, and Nobuko Yoshida. 2020. Statically Verified Refinements for Multiparty Protocols. *Proceedings of the ACM on Programming Languages* 4 (2020),