

# Ferrite: A Judgmental Embedding of Session Types in Rust

Ruo Fei Chen ✉ 

Independent Researcher

Stephanie Balzer ✉

Carnegie Mellon University

Bernardo Toninho ✉ 

Universidade Nova de Lisboa and NOVA LINCS

---

## Abstract

*Session types* have proved viable in expressing and verifying the protocols of message-passing systems. While message passing is a dominant concurrency paradigm in practice, real world software is written without session types. A limitation of existing session type libraries in mainstream languages is their restriction to linear session types, precluding application scenarios that demand sharing and thus aliasing of channel references.

This paper introduces Ferrite, a shallow embedding of session types in Rust that supports both *linear* and *shared* sessions. The formal foundation of Ferrite constitutes the shared session type calculus SILL<sub>s</sub>, which Ferrite encodes via a novel *judgmental embedding* technique. The fulcrum of the embedding is the notion of a typing judgment that allows reasoning about shared and linear resources to type a session. Typing rules are then encoded as functions over judgments, with a valid typing derivation manifesting as a well-typed Rust program. This Rust program generated by Ferrite serves as a *certificate*, ensuring that the application will proceed according to the protocol defined by the session type. The paper details the features and implementation of Ferrite and includes a case study on implementing Servo’s canvas component in Ferrite.

**2012 ACM Subject Classification** Theory of computation → Linear logic; Theory of computation → Type theory; Software and its engineering → Domain specific languages; Software and its engineering → Concurrent programming languages

**Keywords and phrases** Session Types, Rust, DSL

**Digital Object Identifier** 10.4230/LIPIcs.ECOOP.2022.22

**Funding** *Stephanie Balzer*: National Science Foundation Award No. CCF-1718267

*Bernardo Toninho*: FCT/MCTES grant NOVALINCS/BASE UIDB/04516/2020

## 1 Introduction

Message-passing is a dominant concurrency paradigm, adopted by mainstream languages such as Erlang, Scala, Go, and Rust, putting the slogan “*Do not communicate by sharing memory; instead, share memory by communicating*” [11] into practice. In this setting, messages are exchanged along channels, which can be shared by several senders and receivers. Type systems in such languages typically allow channels to be typed, specifying and constraining the types of messages they may carry (e.g. integers, strings, sums, references, etc.).

An aspect inherent to message-passing concurrency that is not captured in mainstream type systems, however, is the idea of a *protocol*. Protocols dictate the sequencing and types of messages to be exchanged. To express and enforce such protocols, *session types* [13, 14, 15] were introduced. Session typing disciplines assign types to channel endpoints according to their intended usage protocols in terms of sequencing of input/output actions (e.g. “send an integer and, afterwards, receive a string”) and branching/selection actions (e.g. “receive either a buy message and process the payment; or a cancellation message and abort the



© R. F. Chen, S. Balzer, and B. Toninho;

licensed under Creative Commons License CC-BY 4.0

36th European Conference on Object-Oriented Programming (ECOOP 2022).

Editors: Karim Ali and Jan Vitek; Article No. 22; pp. 22:1–22:47



Leibniz International Proceedings in Informatics

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

transaction”), ensuring the action sequence is followed correctly and thus, adherence to the protocol. Thanks to their correspondence to *linear* logic [4, 49, 48, 47, 28, 5] session types enjoy a strong logical foundation and ensure, in addition to protocol adherence (*session fidelity*), the existence of a communication partner (*progress*). Session types have also been extended with safe *sharing* [1, 2, 3] to accommodate multi-client scenarios that are rejected by exclusively linear session types.

Despite these theoretical advances, session types have not (yet) been adopted at scale. While various session type embeddings exist in mainstream languages such as Java [17, 16], Scala [43], Haskell [42, 38, 22, 29], OCaml [36, 21], and Rust [23, 27, 6, 7], all of these embeddings lack support for multi-client scenarios that mandate controlled aliasing in addition to linearity. This paper introduces *Ferrite*, a shallow embedding of session types in Rust. In contrast to prior work, Ferrite supports *both* linear and shared session types, with protocol adherence guaranteed statically by the Rust compiler.

Ferrite’s underlying theory is based on the calculus  $SILL_{\mathcal{S}}$  introduced in [1], which develops the logical foundation of shared session types. As a matter of fact, Ferrite encodes  $SILL_{\mathcal{S}}$  typing derivations as Rust functions, through a technique we dub *judgmental embedding*. Through our judgmental embedding, a type-checked Ferrite program yields a Rust program that corresponds to a  $SILL_{\mathcal{S}}$  typing derivation and thus the *proof* of protocol adherence.

In order to faithfully encode  $SILL_{\mathcal{S}}$  typing in Rust, this paper further makes several technical contributions to emulate advanced typing features, such as higher-kinded types, by a skillful combination of traits (type classes) and associated types (type families). For example, Ferrite supports recursive (session) types in this way, which are limited to recursive structs of a fixed size in plain Rust. A combination of type-level natural numbers with ideas from profunctor optics [37] are also used to support named channels and labeled choices. We adopt the idea of *lenses* [9] for selecting and updating individual channels in an arbitrary-length linear context. Similarly, we use *prisms* for selecting a branch out of arbitrary-length choices. Whereas `session-ocaml` [36] has previously explored the use of n-ary choice through extensible variants in OCaml, we are the first to connect n-ary choice to prisms and non-native implementation of extensible variants. Remarkably, the Ferrite codebase remains entirely in the safe fragment of Rust, with no direct use of unsafe features.

Given its support of both linear and shared session types, Ferrite is capable of expressing any session typed program in Rust. We substantiate this claim by providing an implementation of Servo’s production canvas component with the communication layer entirely within Ferrite. We report on our findings, including benchmarks in Section 7.

In summary, this paper makes the following contributions:

- design and implementation of *Ferrite*, an embedded domain-specific language (EDSL) for writing session-typed programs in Rust;
- support of both *linear* and *shared* sessions, guaranteed to be observed by type checking;
- a novel *judgmental embedding* of custom typing rules in a host language with the resulting program carrying the proof of successful type checking;
- an encoding of *arbitrary-length choice* in terms of prisms and extensible variants in Rust;
- an *empirical evaluation* based on a full implementation of Servo’s canvas component in Ferrite.

*Outline.* Section 2 gives a brief account of session types and sharing, as found in the  $SILL_{\mathcal{S}}$  calculus [1]. Section 3 tours through the key ideas underlying Ferrite, which are refined in subsequent sections. Section 4 introduces the technical aspects of Ferrite’s type system, focusing on the judgmental embedding and enforcement of linearity. Section 5 explains how

■ **Table 1** Overview of session types and terms in  $\text{SILL}_S$  together with their operational meaning. Subscripts L and S denote linear and shared sessions, resp., where  $m, n \in \{\text{L}, \text{S}\}$ .

Session type		Process term		Description
current	cont	current	cont	
$c_L : \oplus \{\overline{l : A_L}\}$	$c_L : A_{L_h}$	$c_L.l_h; P$	$P$	provider sends label $l_h$ along $c_L$
$c_L : \& \{\overline{l : A_L}\}$	$c_L : A_{L_h}$	case $c_L$ of $\overline{l} \Rightarrow \overline{Q}$	$Q_h$	client receives label $l_h$ along $c_L$
$c_L : A_m \otimes B_L$	$c_L : B_L$	$c_L.l_h; Q$	$Q$	client sends label $l_h$ along $c_L$
$c_L : A_m \multimap B_L$	$c_L : B_L$	send $c_L d_m; P$	$P$	provider sends channel $d_m : A_m$ along $c_L$
		$y_m \leftarrow \text{recv } c_L; Q_{y_m}$	$Q_{d_m}$	client receives channel $d_m : A_m$ along $c_L$
		send $c_L d_m; Q$	$Q$	client sends channel $d_m : A_m$ along $c_L$
$c_L : \mathbf{1}$	-	close $c_L$	-	provider sends “end” along $c_L$
		wait $c_L; Q$	$Q$	provider receives “end” along $c_L$
$c_L : \downarrow_L^S A_S$	$c_S : A_S$	$c_S \leftarrow \text{detach } c_L; P_{x_S}$	$P_{c_S}$	provider sends “detach $c_S$ ” along $c_L$
		$x_S \leftarrow \text{release } c_L; Q_{x_S}$	$Q_{c_S}$	client receives “detach $c_S$ ” along $c_L$
$c_S : \uparrow_L^S A_L$	$c_L : A_L$	$c_L \leftarrow \text{acquire } c_S; Q_{x_L}$	$Q_{c_L}$	client sends “acquire $c_L$ ” along $c_S$
		$x_L \leftarrow \text{accept } c_S; P_{x_L}$	$P_{c_L}$	provider receives “acquire $c_L$ ” along $c_S$
$c_m : A_m$	$c_m : A_m$	$z_n \leftarrow X \leftarrow \overline{d_m}; P_{z_n}$	$P_{z_n}$	spawn (“cut”) $X$ along $z_n : B_n$ with $\overline{d_m} : \overline{D_m}$
$c_m : A_m$	-	fwd $c_m d_m$	-	forward to channel $d_m : A_m$ and terminate

Ferrite addresses Rust’s limited support of recursive data types to allow for arbitrary recursive and shared session types. Section 6 describes the implementation of n-ary choice using prisms and extensible variants. Section 7 provides an evaluation of Ferrite via a re-implementation of the Servo canvas component. Section 8 reports on related and future work.

An anonymized version of Ferrite’s source code with examples is provided as an artifact. All typing rules and their encoding as well as further materials of interest to an inquisitive reader are provided in the appendix.

## 2 Background

This section gives a brief tour of linear and shared session types. The presentation is based on the intuitionistic session-typed process calculus  $\text{SILL}_S$  [1], which Ferrite builds upon. We consider the protocol governing the interaction between a queue and its client:

$$\text{queue } A = \&\{\text{enq} : A \multimap \text{queue } A, \text{deq} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \otimes \text{queue } A\}\}$$

Table 1 provides an overview of the types used in the example. Since  $\text{SILL}_S$  is based on a Curry-Howard correspondence between intuitionistic linear logic and the session-typed  $\pi$ -calculus [4, 5] it uses linear logic connectives ( $\oplus$ ,  $\&$ ,  $\otimes$ ,  $\multimap$ ,  $\mathbf{1}$ ) as session types. The remaining connectives concern shared sessions, a feature we remark on shortly. A crucial—and probably unusual—characteristic of session-typed processes is that a process *changes* its typing along with the messages it exchanges. As a result, a process’ typing always reflects the current protocol state. Table 1 lists state transitions inflicted by a message exchange in the first and second column and corresponding process terms in the third and fourth column. The fifth column provides the operational meaning of a type.

Consulting Table 1, we gather that the above polymorphic session type  $\text{queue } A$  imposes the following recursive protocol: A client may either send the label  $\text{enq}$  or  $\text{deq}$  to the queue,

depending on whether the client wishes to enqueue or dequeue an element of type  $A$ , resp. In the former case, the client sends the element to be enqueued, after which the queue recurs. In the latter case, the queue indicates to the client whether it is empty (**none**) or not (**some**), and proceeds by either terminating or sending the dequeued element and recurring, resp.

A linear typing discipline is beneficial because it immediately guarantees session fidelity—even in the presence of perpetual protocol change—by ensuring that a channel connects exactly two processes. Unfortunately, linearity also rules out various practical programming scenarios that demand sharing and thus aliasing of channel references. For example, the above linear session type `queue`  $A$  is limited to a *single* client. To support safe sharing of stateful channel references while upholding session fidelity,  $\text{SILL}_5$  extends linear session types with shared session types ( $\downarrow_L^S A_S, \uparrow_L^S A_L$ ). These two connectives mediate between shared and linear sessions by requiring that clients of shared sessions interact in *mutual exclusion* from each other. Concretely, a type  $\uparrow_L^S A_L$  mandates a client to *acquire* the process offering the shared session. If the request is successful, the client receives a linear channel to the acquired process along which it must proceed as detailed by the session type  $A_L$ . A type  $\downarrow_L^S A_S$ , on the other hand, mandates a client to *release* the linear process, relinquishing ownership of the linear channel and only being left with a shared channel alias to the now shared process at type  $A_S$ .

Using these connectives, we can turn the above linear queue into a shared one, bracketing enqueue and dequeue operations within acquire-release:

$$\text{squeue } A_S = \uparrow_L^S \& \{ \text{enq} : A_S \multimap \downarrow_L^S \text{squeue } A_S, \text{deq} : \oplus \{ \text{none} : \downarrow_L^S \text{squeue } A_S, \text{some} : A_S \otimes \downarrow_L^S \text{squeue } A_S \} \}$$

In contrast to the linear queue, the above version recurs in the **none** branch and thus keeps the queue alive to serve the next client. For convenience,  $\text{SILL}_5$  allows the connectives  $\otimes$  and  $\multimap$  to be used to transport both linear and shared channels along a linear carrier channel.

To provide a flavor of session-typed programming in  $\text{SILL}_5$ , we briefly comment on the below processes *empty* and *elem*, which implement the shared queue session type as a sequence of *elem* processes, ended by an *empty* process. A process implementation consists of its signature (first two lines) and body (after  $=$ ). The first line indicates the typing of channel variables used by the process (left of  $\vdash$ ) and the type of the providing channel variable (right of  $\vdash$ ). The second line binds the channel variables. In  $\text{SILL}_5$ ,  $\leftarrow$  generally denotes variable bindings. We leave it to the reader to convince themselves, consulting Table 1, that the code in the body of the two processes executes the protocol defined by session type `squeue`  $A_S$ .

$$\begin{array}{ll} \cdot \vdash \text{empty} :: q : \text{squeue } A_S & x : A_S, t : \text{squeue } A_S \vdash \text{elem} :: q : \text{squeue } A_S \\ q \leftarrow \text{empty} \leftarrow \cdot = & q \leftarrow \text{elem} \leftarrow x, t = \\ q' \leftarrow \text{accept } q ; & q' \leftarrow \text{accept } q ; \\ \text{case } q' \text{ of} & \text{case } q' \text{ of} \\ | \text{enq} \rightarrow x \leftarrow \text{recv } q' ; & | \text{enq} \rightarrow y \leftarrow \text{recv } q' ; \\ & \quad q \leftarrow \text{detach } q' ; \\ & \quad e \leftarrow \text{empty} ; q \leftarrow \text{elem} \leftarrow x, e \\ | \text{deq} \rightarrow q'.\text{none} ; & \quad t' \leftarrow \text{acquire } t ; \\ & \quad t'.\text{enq} ; \text{send } t' y ; \\ & \quad t \leftarrow \text{release } t' ; q \leftarrow \text{detach } q' ; \\ & \quad q \leftarrow \text{elem} \leftarrow x, t \\ & | \text{deq} \rightarrow q'.\text{some} ; \text{send } q' x ; \\ & \quad q \leftarrow \text{detach } q' ; \text{fwd } q t \end{array}$$

Imposing acquire-release not only as a programming methodology but also as a *typing discipline* has the advantage of recovering session fidelity for shared sessions. To this end, shared session types in  $\text{SILL}_5$  must be *strictly equi-synchronizing* [1, 3], imposing the

invariant that an acquired session is released to the type at which previously acquired. For example, the shared session type `squeue`  $A_S$  is strictly equi-synchronizing whereas the type `invalid`  $= \uparrow_L^s \& \{ \text{left} : \downarrow_L^s \uparrow_L^s \oplus \{ \text{yes} : \downarrow_L^s \text{invalid}, \text{no} : \mathbf{1} \}, \text{right} : \downarrow_L^s \text{invalid} \}$  is not.

It is instructive to review the typing rules for acquire-release:

$$\begin{array}{c}
 \text{(T-}\uparrow_L^s\text{)} \\
 \frac{\Psi, x_S : \uparrow_L^s A_L; \Delta, y_L : A_L \vdash Q_{y_L} :: (z_L : C_L)}{\Psi, x_S : \uparrow_L^s A_L; \Delta \vdash y_L \leftarrow \text{acquire } x_S; Q_{y_L} :: (z_L : C_L)}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{(T-}\uparrow_L^s\text{R)} \\
 \frac{\Psi; \cdot \vdash P_{y_L} :: (y_L : A_L)}{\Psi \vdash y_L \leftarrow \text{accept } x_S; P_{y_L} :: (x_S : \uparrow_L^s A_L)}
 \end{array}$$
  

$$\begin{array}{c}
 \text{(T-}\downarrow_L^s\text{)} \\
 \frac{\Psi, x_S : A_S; \Delta \vdash Q_{x_S} :: (z_L : C_L)}{\Psi; \Delta, y_L : \downarrow_L^s A_S \vdash x_S \leftarrow \text{release } y_L; Q_{x_S} :: (z_L : C_L)}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{(T-}\downarrow_L^s\text{R)} \\
 \frac{\Psi \vdash P_{x_S} :: (x_S : A_S)}{\Psi; \cdot \vdash x_S \leftarrow \text{detach } y_L; P_{x_S} :: (y_L : \downarrow_L^s A_S)}
 \end{array}$$

Due to its foundation in intuitionistic linear logic,  $\text{SILL}_S$ ' typing rules are phrased using a *sequent calculus*, leading to *left* and *right* rules for each connective. Left rules describe the interaction from the point of view of the client, right rules from the point of view of the provider. The typing judgments  $\Psi; \Delta \vdash P :: (x_L : A_L)$  and  $\Psi \vdash P :: (x_S : A_S)$  read as "process  $P$  offers a session of type  $A$  along channel  $x$  using sessions offered along channels in  $\Psi$  (and  $\Delta$ )." The typing contexts  $\Psi$  and  $\Delta$  provide the typing of shared and linear channels, resp. Whereas  $\Psi$  is a structural context,  $\Delta$  is a linear context, forbidding channels to be dropped (weakened) or duplicated (contracted). In contrast to linear processes, shared processes must not use any linear channels, a requirement crucial for type safety. The notions of acquire and release are naturally formulated from the point of view of a client, so these terms appear in the left rules. The right rules use the terms *accept* and *detach* with the meaning that an accept accepts an acquire and a detach initiates a release. The rules are read bottom-up, where the premise denotes the next action to be taken after the message exchange.

### 3 Key Ideas

This section introduces the key ideas underlying Ferrite. Subsequent sections provide further details.

#### 3.1 $\text{SILL}_R$ – A stepping stone from $\text{SILL}_S$ to Ferrite

In Section 2, we reviewed  $\text{SILL}_S$  and its typing judgment. Our goal with Ferrite is to faithfully and compositionally encode  $\text{SILL}_S$  typing derivations in Rust. However, when viewed under the lens of a general purpose programming language, most readers will find  $\text{SILL}_S$  a prohibitively austere formalism, lacking most facilities needed to write realistic programs (e.g. basic data types, pattern matching, etc.) and provided by a convenient and usable programming language like Rust. From an ergonomics standpoint alone it would be unreasonably prohibitive for our embedding to forbid the use of Rust features such as functions, traits and enumerations, only for the sake of precisely mirroring  $\text{SILL}_S$ . Moreover, to realize such an embedding we must be able to account for both  $\text{SILL}_S$ ' linear session discipline (i.e. the *linear* context  $\Delta$ ) and shared session discipline (i.e. the *structural* context  $\Psi$ ) within Rust's usage discipline. Since Rust's typing discipline is essentially *affine*, its treatment of variable usage is neither linear nor purely structural, and so both shared and linear channels must be treated explicitly in the encoding.

■ **Table 2** Overview of  $\text{SILL}_R$  types and terms and their encoding in Ferrite. Note that  $\text{SILL}_R$  uses  $\tau \triangleleft A_L$  and  $\tau \triangleright A_L$  for shared channel output and input, resp., and  $\epsilon$  for termination.

Type Ferrite	$\text{SILL}_R$	Terms ( $\text{SILL}_R$ ) provider	client
InternalChoice<Row>	$\oplus\{\overline{l_i : A_{L_i}}\}$	offer $l_i; K$	case $a \{\overline{l_i : K_i}\}$
ExternalChoice<Row>	$\&\{\overline{l_i : A_{L_i}}\}$	offer_choice $\{\overline{l_i : K_i}\}$	choose $a l_i; K$
SendChannel<A, B>	$A_L \otimes B_L$	send_channel_from $a; K$	$a \leftarrow$ receive_channel_from $f a; K$
ReceiveChannel<A, B>	$A_L \multimap B_L$	$a \leftarrow$ receive_channel; $K$	send_channel_to $f a; K$
SendValue<T, A>	$\tau \triangleleft A_L$	send_value $x; K$	$x \leftarrow$ receive_value_from $a x; K$
ReceiveValue<T, A>	$\tau \triangleright A_L$	$x \leftarrow$ receive_value; $K$	send_value_to $a x; K$
End	$\epsilon$	terminate	wait $a; K$
SharedToLinear<A>	$\downarrow_L^S A_S$	detach_shared_session; $K_s$	release_shared_session $a; K_l$
LinearToShared<A>	$\uparrow_L^S A_L$	accept_shared_session; $K_l$	$a \leftarrow$ acquire_shared_session $s; K_l$

The two points above naturally lead us to the language  $\text{SILL}_R$  as a formal stepping stone between  $\text{SILL}_S$  and our embedding, Ferrite.  $\text{SILL}_R$  is, in its essence, a pragmatic extension of  $\text{SILL}_S$  with Rust (type and term) constructs, allowing us to intersperse Rust code with the communication primitives of  $\text{SILL}_S$ . In  $\text{SILL}_R$  we use the judgment

$$\Gamma; \Delta \vdash \text{expr} :: A,$$

denoting that expression  $\text{expr}$  has session type  $A$ , using the sessions tracked by  $\Gamma$  and  $\Delta$ . This judgment differs from that of  $\text{SILL}_S$  in its context region  $\Gamma$  and term  $\text{expr}$ , with the latter permitting arbitrary Rust expressions in addition to  $\text{SILL}_S$  communication primitives. Whereas  $\text{SILL}_S$ 's structural context  $\Psi$  exclusively tracks shared channels,  $\text{SILL}_R$ 's  $\Gamma$  tracks *both* shared sessions (subject to weakening and contraction) and plain Rust (affine) variables. A shared channel type in both  $\text{SILL}_R$  and  $\text{SILL}_S$  is always of the form  $\uparrow_L^S A$ , so there is no confusion among the affine and shared contents of  $\Gamma$ . As we discuss in Section 5.2, the distinction between a plain Rust variable, which is treated as affine, and a shared channel, which is treated structurally, is modelled in Ferrite by making shared channels implement Rust's `Clone` trait.

Table 2 provides an overview of  $\text{SILL}_R$  types and terms and their Ferrite encoding.  $\text{SILL}_R$  types stand in direct correspondence with  $\text{SILL}_S$  types (see Table 1), apart from shared channel output and input. The  $\text{SILL}_S$  types for sending and receiving shared channels ( $A_S \otimes A_L$  and  $A_S \multimap A_L$ ) correspond to  $\text{SILL}_R$  types for sending and receiving values ( $T \triangleleft A$  and  $T \triangleright A$ , resp.), which support *both* Rust values and shared channels. Their typing rules are:

$$\begin{array}{c}
(T \triangleleft_R) \\
\frac{\Gamma; \Delta \vdash K :: A}{\Gamma, x : \tau; \Delta \vdash \text{send\_value } x; K :: \tau \triangleleft A}
\end{array}
\qquad
\begin{array}{c}
(T \triangleleft_L) \\
\frac{\Gamma, x : \tau; \Delta, a : A \vdash K :: B}{\Gamma; \Delta, a : \tau \triangleright A \vdash x \leftarrow \text{receive\_value\_from } a; K :: B}
\end{array}$$

Rule  $T \triangleleft_R$  indicates that the value bound to variable  $x$  of type  $\tau$  will be sent, after which the continuation  $K$  will execute, offering type  $A$ . Dually, rule  $T \triangleleft_L$  states that using such a provider bound to  $a$  will bind  $x$  of type  $\tau$  in continuation  $K$ , which must now use the channel bound to  $a$  according to  $A$ .

### 3.2 Judgmental Embedding

Having introduced the  $\text{SILL}_R$  typing judgment and illustrated some of its typing rules, we can now clarify the idea behind our notion of *judgmental embedding*, which enables the Rust

■ **Table 3** Judgmental embedding of  $\text{SILL}_R$  in Ferrite.

$\text{SILL}_R$	Ferrite	Description
$\Gamma; \cdot \vdash A$	<code>Session&lt;A&gt;</code>	Typing judgment for top-level session (i.e. closed program).
$\Gamma; \Delta \vdash A$	<code>PartialSession&lt;C, A&gt;</code>	Typing judgment for partial session.
$\Delta$	<code>C: Context</code>	Linear context; explicitly encoded.
$\Gamma$	-	Shared / Affine context; delegated to Rust.
$A$	<code>A: Protocol</code>	Session type.

compiler to typecheck  $\text{SILL}_R$  programs by encoding typing derivations as Rust programs. The basic idea underlying this encoding can be schematically described as follows:

$$\frac{\Gamma; \Delta_2 \vdash \text{cont} :: A_2}{\Gamma; \Delta_1 \vdash \text{expr}; \text{cont} :: A_1} \quad \begin{array}{l} \text{fn expr}\langle \dots \rangle \\ \quad (\text{cont}: \text{PartialSession}\langle C_2, A_2 \rangle) \\ \quad \rightarrow \text{PartialSession}\langle C_1, A_1 \rangle \end{array}$$

On the left we show a  $\text{SILL}_R$  typing rule and on the right its encoding in Ferrite. Ferrite encodes a  $\text{SILL}_R$  *typing judgment*  $\Gamma; \Delta \vdash \text{expr} :: A$  as a value of Rust *type* `PartialSession<C, A>`, where `C` encodes the linear context  $\Delta$  and `A` the session type  $A$ , standing for any of the Ferrite types of Table 2. Ferrite then encodes a  $\text{SILL}_R$  *typing rule* for an expression  $\text{expr}$  as a Rust *function* `expr` that accepts a `PartialSession<C2, A2>` and returns a `PartialSession<C1, A1>`, where  $\text{expr}$  stands for any of the  $\text{SILL}_R$  terms of Table 2. The encoding makes use of *continuation passing style* (arising from the sequent calculus-based formulation of  $\text{SILL}_R$ ), with the return type being the conclusion of the rule and the argument type being its premise. Table 3 summarizes the judgmental embedding; Section 4.1 provides further details. Whereas Ferrite explicitly performs a type-level encoding of the linear context  $\Delta$ , the representation of the shared and affine context region  $\Gamma$  is achieved through Rust’s normal binding structure, with the obligation that shared channels implement Rust’s `Clone` trait to permit contraction. To type a closed program, Ferrite defines the type `Session<A>`, which stands for a  $\text{SILL}_R$  judgment with an empty linear context.

Adopting a judgmental embedding technique for implementing a DSL delivers the benefits of proof-carrying code: the `PartialSession<C1, A1>` returned from a well-typed Ferrite `expr` is the typing derivation of the corresponding  $\text{SILL}_R$  term. In case the  $\text{SILL}_R$  term is a  $\text{SILL}_S$  term, its typing derivation certifies protocol adherence by virtue of the type safety proof of  $\text{SILL}_S$  [1]. In case the  $\text{SILL}_R$  term includes Rust code, its typing derivation certifies protocol adherence modulo the possibility of a panic raised by the Rust code.

### 3.3 Recursive and Shared Session Types in Ferrite

Rust’s support for recursive types is limited to recursive struct definitions of a known size. To circumvent this restriction and support arbitrary recursive session types, Ferrite introduces a type-level fixed-point combinator `Rec<F>` to obtain the fixed point of a type function `F`. Since Rust lacks higher-kinded types such as `Type → Type`, we use *defunctionalization* [40, 51] by accepting any Rust type `F` implementing the trait `RecApp` with a given associated type `F::Applied`, as shown below. Section 5.1 provides further details.

```
trait RecApp<X> { type Applied; }
struct Rec<F: RecApp<Rec<F>>> { unfold: Box<F::Applied> }
```

Recursive types are also vital for encoding shared session types. In line with [3], we restrict shared session types to be recursive, making sure that a shared component is continuously available. To guarantee type preservation, recursive session types must be *strictly equi-synchronizing* [1, 3], requiring an acquired session to be released to the same type at which



it was previously acquired. Ferrite enforces this invariant by defining a specialized trait `SharedRecApp` which omits an implementation for `End`:

```
trait SharedRecApp<X> { type Applied; } trait SharedProtocol { ... }
struct SharedToLinear<F> { ... } struct SharedChannel<S: SharedProtocol> { ... }
struct LinearToShared<F: SharedRecApp<SharedToLinear<LinearToShared<F>>>> { ... }
```

Ferrite achieves safe communication for shared sessions by imposing an acquire-release discipline [1] on shared sessions, establishing a critical section for the linear portion of the process enclosed within the acquire and release. `SharedChannel` denotes the shared process running in the background, and clients with a reference to it can *acquire* an exclusive linear channel to communicate with it. As long as the linear channel exists, the shared process is locked and cannot be acquired by any other client. With the strictly equi-synchronizing constraint in place, the now linear process must eventually be released (`SharedToLinear`) back to the same shared session type at which it was previously acquired, giving turn to another client waiting to acquire. Section 5.2 provides further details on the encoding.

### 3.4 N-ary Choice and Linear Context

Ferrite implements n-ary *choices* and linear typing *contexts* as extensible *sums* and *products* of session types, resp. Ferrite uses heterogeneous lists [24] to annotate a list of session types of arbitrary length. The notation  $\text{HList!}[A_0, A_1, \dots, A_{N-1}]$  denotes a heterogeneous list of  $N$  session types, with  $A_i$  being the session type at the  $i$ -th position of the list. The `HList!` macro acts as syntactic sugar for the heterogeneous list, which in its raw form is encoded as  $(A_0, (A_1, (\dots, (A_{N-1}, ())))))$ . Ferrite uses the Rust tuple constructor `(,)` for `HCons`, and unit `()` for `HNil`. The heterogeneous list itself can be directly used to represent an n-ary product. Using an associated type, the list can moreover be transformed into an n-ary sum.

One disadvantage of using heterogeneous lists is that its elements have to be addressed by position rather than a programmer-chosen label. To recover labels for accessing list elements, we use optics [37]. More precisely, Ferrite uses *lenses* [9] to access a channel in a linear context and *prisms* to select a branch of a choice. We further combine the optics abstraction with *de Bruijn levels* and implement lenses and prisms using type level natural numbers. Given an inductive trait definition of natural numbers as zero (`z`) and successor (`s<N>`), a natural number  $\mathbb{N}$  implements the lens to access the  $N$ -th element in the linear context, and the prism to access the  $N$ -th branch in a choice. Schematically, the lens encoding can be captured as follows:

$$\frac{\Gamma; \Delta, l_n : B_2 \vdash K :: A_2}{\Gamma; \Delta, l_n : B_1 \vdash \text{expr } l_n; K :: A_1} \quad \begin{array}{l} \text{fn expr}\langle \dots \rangle \\ ( \ 1 : \mathbb{N}, \text{cont} : \text{PartialSession}\langle C1, A2 \rangle ) \\ \rightarrow \text{PartialSession}\langle C2, A1 \rangle \\ \text{where } \mathbb{N} : \text{ContextLens}\langle C1, B1, B2, \text{Target}=C2 \rangle \end{array}$$

The index  $\mathbb{N}$  amounts to the type of the variable `1` that the programmer chooses as a name for a channel in the linear context. Ferrite handles the mapping, supporting random access to programmer-named channels. Section 4.2 provides further details, including the support of higher-order channels. Similarly, prisms allow choice selection in constructs such as `offer_case` to be encoded as follows:

$$\frac{\Gamma; \Delta \vdash K :: A_n}{\Gamma; \Delta \vdash \text{offer\_case } l_n; K :: \oplus\{\dots, l_n : A_n, \dots\}} \quad \begin{array}{l} \text{fn offer\_case}\langle \mathbb{N}, \text{Row}, C, A \rangle \\ ( \ 1 : \mathbb{N}, \text{cont} : \text{PartialSession}\langle C, A \rangle ) \\ \rightarrow \text{PartialSession}\langle C, \text{InternalChoice}\langle \text{Row} \rangle \rangle \\ \text{where } \mathbb{N} : \text{Prism}\langle \text{Row}, \text{Elem}=A \rangle, \dots \end{array}$$

Ferrite maps a choice label to a constant having the singleton value of a natural number  $\mathbb{N}$ , which implements the prism to access the  $N$ -th branch of a choice. In addition to prisms, Ferrite implements a version of *extensible variants* [30] to support polymorphic operations



on arbitrary sums of session types representing choices. Finally, the `define_choice!` macro is used as a helper to export type aliases as programmer-friendly identifiers. Details are reported in Section 6.

## 4 Ferrite – A Judgmental Embedding of $SILL_R$

Having introduced some of the key concepts to the implementation of Ferrite, we now cover in detail the implementation of Ferrite’s core constructs, building up the knowledge required for Section 5 and Section 6. Ferrite, like any other DSL, has to tackle the various technical challenges encountered when embedding a DSL in a host language. In doing so, we take inspiration from the range of embedding techniques developed for Haskell and adjust them to the Rust setting. The lack of higher-kinded types, limited support of recursive types, and presence of weakening, in particular, make the development far from trivial. A more conceptual contribution of this work is thus to demonstrate how existing Rust features can be combined to emulate many of the missing features that are beneficial to DSL embeddings and how to encode custom typing rules in Rust or any similarly expressive language. The techniques described in this and subsequent sections also serve as a reference for embedding other DSLs in a host language like Rust.

### 4.1 Encoding Typing Rules via Judgmental Embedding

A distinguishing characteristic of Ferrite is its *propositions as types* approach, yielding a direct correspondence between  $SILL_R$  notions and their Ferrite encoding. This correspondence was introduced in Section 3.2 (see Table 3) and we now discuss it in more detail. To this end, let’s consider the typing of value input. We remind the reader of Table 2 in Section 3, which provides a mapping between  $SILL_R$  and Ferrite session types. Interested readers can find a corresponding mapping on the term level in Table 5 in the supplement.

$$\frac{\Gamma, a : \tau; \Delta \vdash K :: A}{\Gamma; \Delta \vdash a \leftarrow \text{receive\_value}; K :: \tau \triangleright A} \text{ (T}\triangleright_R\text{)}$$

The  $SILL_R$  right rule  $T\triangleright_R$  types the expression  $a \leftarrow \text{receive\_value}; K$  as the session type  $\tau \triangleright A$  and the continuation  $K$  as the session type  $A$ , where  $a$  is now in scope with type  $\tau$ . Following the schema hinted in Section 3.2, Ferrite encodes this rule as the function `receive_value`, parameterized by a value type  $T$  ( $\tau$ ), a linear context  $C$  ( $\Delta$ ), and an offered session type  $A$ .

```
fn receive_value<T, C:Context, A:Protocol>(cont:impl FnOnce(T) -> PartialSession<C, A>)
    -> PartialSession<C, ReceiveValue<T, A>>
```

The function yields a value of type `PartialSession<C, ReceiveValue<T, A>>`, i.e. the conclusion of the rule, given an (affine) closure of type  $T \rightarrow \text{PartialSession}\langle C, A \rangle$ , encoding the premise of the rule. Notably, Ferrite uses plain Rust binding (through function types) to encode the contents of  $\Gamma$ , as illustrated for the received value above. The use of a closure reveals the continuation-passing-style of the encoding, where the received value of type  $T$  is passed to the continuation closure. The affine closure implements the `FnOnce` trait, ensuring that it can only be called once.

The type `PartialSession` is a core construct of Ferrite that enables the judgmental embedding of  $SILL_R$ . A Rust value of type `PartialSession<C, A>` represents a Ferrite program that guarantees linear usage of session type channels in the linear context  $C$  and offers the linear session type  $A$ , corresponding to the  $SILL_R$  typing judgment  $\Gamma; \Delta \vdash \text{expr} :: A$ . The type

## 22:10 Ferrite: A Judgmental Embedding of Session Types in Rust

parameters  $C$  and  $A$  are constrained to implement the traits `Context` and `Protocol` – two other Ferrite constructs representing a linear context and linear session type, resp.:

```
trait Context { ... }    trait Protocol { ... }
struct PartialSession<C: Context, A: Protocol> { ... }
```

For each  $SILL_R$  session type, Ferrite defines a corresponding Rust struct that implements the trait `Protocol`, yielding the listing shown in Table 2. Implementations for  $\epsilon$  (`End`) and  $\tau \triangleright A$  (`ReceiveValue<T, A>`) are shown below. When a session type is nested within another session type, such as in the case of `ReceiveValue<T, A>`, the constraint to implement `Protocol` is propagated to the inner session type, requiring  $A$  to also implement `Protocol`:

```
struct End { ... }      struct ReceiveValue<T, A> { ... }
impl Protocol for End { ... }
impl<A: Protocol> Protocol for ReceiveValue<T, A> { ... }
```

Thus, while Ferrite delegates the handling of the shared/structural context  $\Gamma$  to Rust, the encoding of the linear context  $\Delta$  is explicit. Being affine, the Rust type system permits weakening, a structural property rejected by linear logic. Ferrite encodes a linear context as a heterogeneous (type-level) list [24] of the form `HList![A0, A1, ..., AN-1]`, with all its type elements  $A_i$  implementing `Protocol`. Internally, the `HList` macro desugars the type-level list into a nested tuple `(A0, (A1, (... (AN-1, ())))`). The unit type `()` is used as the empty list (`HNil`) and the tuple constructor `(,)` is used as the `HCons` constructor. The implementation for `Context` is defined inductively as follows:

```
impl Context for () { ... }    impl<A: Protocol, C: Context> Context for (A, C) { ... }
```

To represent a closed program, i.e. a program without free channel variables, we define a type alias `Session<A>` for `PartialSession<C, A>`, with  $C$  restricted to the empty context:

```
type Session<A> = PartialSession<(), A>;
```

A complete session type program in Ferrite is thus of type `Session<A>` and amounts to the  $SILL_R$  typing derivation proving that the program adheres to the defined protocol. Below we show a “hello world”-style program in Ferrite:

```
let hello_provider = receive_value(|name| {
  println!("Hello, {}", name); terminate() });
```

The Ferrite program `hello_provider` has an inferred Rust type `Session<ReceiveValue<String, End>>`. It offers the type `ReceiveValue<String, End>` by first receiving a string value using `receive_value`, binding it to `name` in the continuation closure. Upon receiving the name string, It prints out the name with a “Hello” greeting, and terminates using `terminate()`.

## 4.2 Manipulating the Linear Context

### Context Lenses

The use of a type-level list to encode the linear context has the advantage of allowing contexts of arbitrary length. However, the list imposes an order on the context’s elements, disallowing exchange. To allow exchange, we make use of the concept of *lenses* [9] to define a `ContextLens` trait, which is implemented using type-level natural numbers.

```
#[derive(Copy)] struct Z;      #[derive(Copy)] struct S<N>(PhantomData<N>);
trait ContextLens<C: Context, A1: Protocol, A2: Protocol> { type Target: Context; ... }
```

The `ContextLens` trait defines the read and update operations on a linear context, such that given a *source* context  $C = HList![\dots, A_N, \dots]$ , the source element of interest,  $A_N$

at position  $N$ , can be updated to the target element  $B$  to form the *target* context  $\text{Target} = \text{HList!}[\dots, B, \dots]$ , with the remaining elements unchanged. We use natural numbers to inductively implement `ContextLens` at each position in the linear context, such that it satisfies all constraints of the form:

$$N: \text{ContextLens}\langle \text{HList!}[\dots, A_N, \dots], A_N, B, \text{Target}=\text{HList!}[\dots, B, \dots] \rangle$$

The implementation of natural numbers as context lenses is done by first considering the base case, with  $Z$  used to access the first element of any non-empty linear context:

```
impl<A1: Protocol, A2: Protocol, C: Context> ContextLens<(A1, C), A1, A2>
  for Z { type Target = (A2, C); ... }
impl<A1: Protocol, A2: Protocol, B: Protocol, C: Context, N: ContextLens<C, A1, A2>>
  ContextLens <(B, C), A1, A2> for S<N> { type Target = (B, N::Target); ... }
```

In the inductive case, for any natural number  $N$  implementing the context lens for a context  $\text{HList!}[A_0, \dots, A_N, \dots]$ , its successor  $S\langle Z \rangle$  implements the context lens for  $\text{HList!}[A_{-1}, A_0, \dots, A_N, \dots]$ , with a new element  $A_{-1}$  appended to the head of the linear context. Using context lenses, we can encode the  $\text{SILL}_R$  left rule  $\text{T}_{\triangleright_L}$  shown below, which types sending an ambient value  $x$  to a channel  $a$  in the linear context that expects to receive a value.

$$\frac{\Gamma; \Delta, a: A \vdash K :: B}{\Gamma, x: \tau; \Delta, a: \tau \triangleright A \vdash \text{send\_value\_to } a \ x; K :: B} (\text{T}_{\triangleright_L})$$

In Ferrite,  $\text{T}_{\triangleright_L}$  is implemented as the function `send_value_to`, which uses a context lens  $N$  to send a value of type  $T$  to the  $N$ -th channel in the linear context  $C_1$ . This requires the  $N$ -th channel to have type `ReceiveValue<T,A>`. A continuation `cont` is then given with the linear context  $C_2$ , which has the  $N$ -th channel updated to type  $A$ .

```
fn send_value_to<N, T, C1: Context, C2: Context, A: Protocol, B: Protocol>
  ( n: N, x: T, cont: PartialSession<C2, B> ) -> PartialSession <C1, B>
where N: ContextLens<C1, ReceiveValue<T, A>, A, Target=C2>
```

## Channel Removal

The above definition of a context lens is suited for *updating* channel types in a context. However, we have not addressed how channels can be *removed* or *added* to the linear context. These operations are required to implement session termination and higher-order channel constructs such as  $\otimes$  and  $\multimap$ . To support channel removal, we introduce a special `Empty` element to denote the *absence* of a channel at a given position in the linear context:

```
struct Empty; trait Slot { ... }
impl Slot for Empty { ... } impl<A: Protocol> Slot for A { ... }
```

To allow `Empty` to be present in a linear context, we introduce a new `Slot` trait and make both `Empty` and `Protocol` implement `Slot`. The original definition of `Context` is then updated to allow types that implement `Slot` instead of `Protocol`.

$$\frac{\Gamma; \Delta \vdash K :: A}{\Gamma; \Delta, a: \epsilon \vdash \text{wait } a; K :: A} (\text{T}_{1_L}) \quad \frac{}{\Gamma; \cdot \vdash \text{terminate}; :: \epsilon} (\text{T}_{1_R})$$

Using `Empty`, it is straightforward to implement  $\text{SILL}_R$ 's session termination. Rule  $\text{T}_{1_L}$  is encoded via a context lens that replaces a channel of session type `End` with the `Empty` slot. The function `wait` shown below does not really remove a slot from a linear context, but merely replaces the slot with `Empty`. The use of `Empty` is necessary, because we want to preserve the position of channels in a linear context in order for the context lens for a channel to work across continuations.

## 22:12 Ferrite: A Judgmental Embedding of Session Types in Rust

```
fn wait<C1: Context, C2: Context, A: Protocol, N>
  ( n: N, cont: PartialSession<C2, A> ) -> PartialSession<C1, A>
where N: ContextLens<C1, End, Empty, Target=C2>
```

With `Empty` introduced, an empty linear context may now contain any number of `Empty` slots, such as `HList![Empty, Empty]`. We introduce a new `EmptyContext` trait to abstract over the different forms of empty linear contexts and provide an inductive definition as its implementation:

```
trait EmptyContext: Context { ... } impl EmptyContext for () { ... }
impl<C: EmptyContext> EmptyContext for (Empty, C) { ... }
```

Given the empty list `()` as the base case, the inductive case `(Empty, C)` is an empty linear context, if `C` is also an empty linear context. Using the definition of an empty context, the  $SILL_R$  right rule  $T1_R$  can then be easily encoded as the function `terminate`, which works generically for all contexts that implement `EmptyContext` as shown below:

```
fn terminate<C: EmptyContext>() -> PartialSession<C, End>
```

### Channel Addition

The Ferrite function `wait` removes a channel from the linear context by replacing it with `Empty`. Dually, the function `receive_channel`, adds a new channel to the linear context. The  $SILL_R$  rule  $T\multimap_R$  for channel input is shown below. It binds the received channel of session type  $A$  to the channel variable  $a$  and adds it to the linear context  $\Delta$  of the continuation.

$$\frac{\Gamma; \Delta, a : A \vdash K :: B}{\Gamma; \Delta \vdash a \leftarrow \text{receive\_channel}; K :: A \multimap B} (T\multimap_R)$$

To encode  $T\multimap_R$ , an append operation on contexts is defined via the `AppendContext` trait:

```
trait AppendContext<C: Context>: Context { type Appended: Context; ... }
impl<C: Context> AppendContext<C> for () { type Appended = C; ... }
impl<A: Slot, C1: Context, C2: Context, C3: Context> AppendContext<C2>
  for (A, C1) where C1: AppendContext<C2, Appended=C3> { type Appended = (A, C3); ... }
```

The `AppendContext` trait is parameterized by a linear context `C` and an associated type `Appended`. If a linear context `C1` implements the trait `AppendContext<C2>`, it means that context `C2` can be appended to `C1`, with `C3 = C1::Appended` being the result of the append operation. The implementation of `AppendContext` is defined inductively, with the empty list `()` implementing the base case and the cons cell `(A, C)` implementing the inductive case.

Using `AppendContext`, a channel `B` can be appended to the end of a linear context `c`, if `c` implements `AppendContext<HList![B]>`. The new linear context after the append operation is given in the associated type `C::Appended`. We then observe that the position of channel `B` in `C::Appended` is the same as the length of the original linear context `c`. In other words, the context lens for channel `B` in `C::Appended` can be generated by obtaining the length of `c`. In Ferrite, the length operation is implemented by adding an associated type `Length` to the `Context` trait. The implementation of `Context` for `()` and `(A, C)` is updated correspondingly.

```
trait Context { type Length; ... } impl Context for () { type Length = Z; ... }
impl<A: Slot, C: Context> Context for (A, C) { type Length = S<C::Length>; ... }
```

The  $SILL_R$  right rule  $T\multimap_R$  is then encoded as follows:

```
fn receive_channel<A: Protocol, B: Protocol, C1: Context, C2: Context>(
  cont: impl FnOnce(C1::Length) -> PartialSession<C2, B>) ->
  PartialSession<C1, ReceiveChannel<A, B>> where C1: AppendContext<(A, ()), Appended=C2>
```

The function `receive_channel` is parameterized by a linear context `C1` implementing `AppendContext` to append the session type `A` to `C1`. The continuation argument `cont` is a closure that is given a context lens `C::Length`, and returns a `PartialSession` with `C2=C1::Appended` as its linear context. The function returns a `PartialSession` with linear context `C1`, offering session type `ReceiveChannel<A, B>`.

It is worth noting that in the type signature of `receive_channel`, the type `C1::Length` is not shown to have any `ContextLens` implementation. However when `C1::Length` is instantiated into the concrete types `Z`, `S<Z>`, etc in the continuation body, Rust will use the appropriate implementations of `ContextLens` so that they can be used to access the appended channel in the linear context.

The use of `receive_channel` is illustrated with the `hello_client` example below:

```
let hello_client = receive_channel(|a| {
  send_value_to(a, "Alice".to_string(), wait(a, terminate())) });
```

The `hello_client` program is inferred to have the Rust type `Session<ReceiveChannel<ReceiveValue<String, End>, End>>`. It is written to communicate with the `hello_provider` program defined earlier in Section 4.1. The interaction is achieved by having `hello_client` offering the session type `ReceiveChannel<ReceiveValue<String, End>, End>`. In its body, `hello_client` uses `receive_channel` to receive channel `a` of type `ReceiveValue<String, End>` from `hello_provider`. The continuation closure is given an argument `a:Z`, denoting the context lens generated by `receive_channel` for accessing the received channel in the linear context. The context lens `a:Z` is then used for sending a string value, after which we `wait` for `hello_provider` to terminate. We note that the type `Z` of channel `a` (i.e. the channel position in the context) is automatically inferred by Rust and not exposed to the user.

### 4.3 Communication

At this point we have defined the necessary constructs to build and typecheck both `hello_provider` and `hello_client`, but the two are separate Ferrite programs that are yet to be linked with each other and executed.

$$\frac{\Gamma; \Delta_1 \vdash K_1 :: A \quad \Gamma; \Delta_2, a : A \vdash K_2 :: B}{\Gamma; \Delta_1, \Delta_2 \vdash a \leftarrow \text{cut } K_1; K_2 :: B} \text{ (T-CUT)} \quad \frac{}{\Gamma; a : A \vdash \text{forward } a :: A} \text{ (T-FWD)}$$

In `SILLR`, rule `T-CUT` allows two session-typed programs to run in parallel, with the channel offered by `K1` added to the linear context of program `K2`. Together with the forward rule `T-FWD`, we can use `cut` twice to run both `hello_provider` and `hello_client` in parallel, and have a third program that sends the channel offered by `hello_provider` to `hello_client`. The program `hello_main` would have the following pseudo code in `SILLR`:

```
hello_main : ε = f ← cut hello_client;
              a ← cut hello_provider;
              send_channel_to f a;
              forward f
```

To implement `cut` in Ferrite, we need a way to split a linear context `C = Δ1, Δ2` into two sub-contexts `c1 = Δ1` and `c2 = Δ2` so that they can be passed to the respective continuations. Moreover, since Ferrite programs use context lenses to access channels, the ordering of channels inside `c1` and `c2` must be preserved. We can preserve the ordering by replacing the corresponding slots with `Empty` during the splitting. Ferrite defines the `SplitContext` trait to implement the splitting as follows:

```
enum L {}      enum R {}
trait SplitContext<C: Context> { type Left: Context; type Right: Context; ... }
```

## 22:14 Ferrite: A Judgmental Embedding of Session Types in Rust

We first define two (uninhabited) marker types `L` and `R`. We then use type-level lists consisting of elements `L` and `R` to implement the `SplitContext` trait for a given linear context `C`. The `SplitContext` implementation contains the associated types `Left` and `Right`, representing the contexts `C1` and `C2` after splitting. As an example, the type `HList![L, R, L]` would implement `SplitContext<HList![A1, A2, A3]>` for any slot `A1`, `A2` and `A3`, with the associated type `Left` being `HList![A1, Empty, A3]` and `Right` being `HList![Empty, A2, Empty]`. We omit the implementation details of `SplitContext` for brevity. Using `SplitContext`, the function `cut` can be implemented as follows:

```
fn cut<XS, C: Context, C1: Context, C2: Context, C3: Context, A: Protocol, B: Protocol>
  ( cont1: PartialSession<C1, A>,
    cont2: impl FnOnce(C2::Length) -> PartialSession<C3, B> ) -> PartialSession<C, B>
where XS: SplitContext<C, Left=C1, Right=C2>, C2: AppendContext<HList![A], Appended=C3>
```

The function `cut` works by using the heterogeneous list `XS` that implements `SplitContext` to split a linear context `C` into `C1` and `C2`. To pass on the channel `A` that is offered by `cont1` to `cont2`, `cut` uses a similar technique to `receive_channel` to append the channel `A` to the end of `C2`, resulting in `C3`. Using `cut`, we can write `hello_main` in Ferrite as follows:

```
let hello_main: Session<End> =
  cut::<HList![]>(hello_client, |f| { cut::<HList![R]>(hello_provider, |a| {
    send_channel_to(f, a, forward(f)) }) });
```

Due to ambiguous instances for `SplitContext`, the type parameter `XS` has to be annotated explicitly for Rust to know in which context a channel should be placed. In the first use of `cut`, the context is empty, so we call `cut` with the empty list `HList![]`. We pass `hello_client` as the first continuation to run in parallel, and name the channel offered by `hello_client` as `f`. In the second use of `cut`, the linear context would be `HList![ReceiveValue<String, End>]`, with one channel `f`. We then have `cut` move `f` to the right side using `HList![R]`. On the left continuation, we have `hello_provider` run in parallel, and name the offered channel as `a`. In the right continuation, we use `send_channel_to` to send channel `a` to `f`. Finally, we forward the continuation of `f`, which now has type `End`.

Although `cut` provides the primitive way for Ferrite programs to communicate, its use can be cumbersome and requires a lot of boilerplate. For simplicity, Ferrite provides a specialized construct `apply_channel` that abstracts over the common pattern usage of `cut` described earlier. `apply_channel` takes a client program `f` offering session type `ReceiveChannel<A, B>` and a provider program `a` offering session type `A`, and sends `a` to `f` using `cut`. The use of `apply_channel` is akin to regular function application, making it more intuitive for programmers to use:

```
fn apply_channel<A: Protocol, B: Protocol>(
  f: Session<ReceiveChannel<A, B>>, a: Session<A>) -> Session<B>
```

### 4.4 Executing Ferrite Programs

To actually *execute* a Ferrite program, the program must offer some specific session types. In the simplest case, Ferrite provides the function `run_session` for running a top-level Ferrite program offering `End`, with an empty linear context:

```
async fn run_session(session: Session<End>) { ... }
```

Function `run_session` executes the session *asynchronously* using Rust's `async/await` infrastructure. Internally, the struct `PartialSession<C, A>` implements the dynamic semantics of the Ferrite program, which is only accessible by public functions such as `run_session`. Ferrite currently uses the `tokio` [46] runtime for asynchronous execution, as well as the one shot

channels from `tokio::sync::oneshot` to implement the low-level communication of Ferrite channels.

Since `run_session` accepts an argument of type `Session<End>`, this means that programmers must first use `cut` or `apply_channel` to fully link partial Ferrite programs with free channel variables, or Ferrite programs that offer session types other than `End` before they can be executed. This restriction ensures that all linear channels created by a Ferrite program are consumed. For example, the programs `hello_provider` and `hello_client` cannot be executed individually, but the linked program resulting from composing `hello_provider` with `hello_client` can be executed:

```
async fn main() { run_session(apply_channel(hello_client, hello_provider)).await; }
```

We omit the implementation details of the dynamics of Ferrite, which use low-level primitives such as Rust channels while carefully ensuring that the requirements and invariants of session types are satisfied. Interested readers can find more details in Appendix B.

## 5 Recursive and Shared Session Types

Many real world applications, such as web services and instant messaging, implement protocols that are recursive in nature. As a result, it is essential for Ferrite to support recursive session types. In this section, we first report on Rust’s limited support for recursive types and how Ferrite addresses this limitation. We then discuss how Ferrite encodes *shared* session types, which are recursive.

### 5.1 Recursive Session Types

Consider a simple example of a counter session type, which sends an infinite stream of integer values, incrementing each by one. To write a Ferrite program that offers such a session type, we may attempt to define the counter session type as follows:

```
type Counter = SendValue<u64, Counter>;
```

If we try to use the type definition above, the compiler will emit the error "cycle detected when processing `Counter`". The problem with the above definition is that it is a directly self-referential type alias, which is not supported in Rust. Rust imposes various restrictions on the legal forms of recursive types to ensure that the memory layout of data is known at compile-time.

#### Type-Level Fixed Points

To address this limitation, we implement type-level fixed points using *defunctionalization* [40, 51]. This is done by introducing a `RecApp` trait that is implemented by defunctionalized types that can be “applied” with a type parameter:

```
trait RecApp<X> { type Applied; }    type AppRec<F, X> = <F as RecApp<X>>::Applied;
struct Rec<F: RecApp<Rec<F>>> { unfold: Box<AppRec<F, Rec<F>>> }
```

The `RecApp` trait is parameterized by a type `X`, which serves as the type argument to be applied to. This makes it possible for a Rust type `F` that implements `RecApp` to act as if it has the higher-kinded type `Type → Type`, and be “applied” to type `X`. We define a type alias `AppRec<F, X>` to refer to the associated type `Applied` resulting from “applying” `F` to `X` via `RecApp`. Using `RecApp`, we can now define a type-level recursor `Rec` as a struct parameterized by a type `F` that implements `RecApp<Rec<F>>`. The body of `Rec` contains a boxed value `Box<AppRec<F, RecApp<Rec<F>>>>` to make it have a fixed size in Rust.



## 22:16 Ferrite: A Judgmental Embedding of Session Types in Rust

Ferrite implements `RecApp` for all `Protocol` types, with the type `Z` used to denote the recursion point. With that, the example `Counter` type would be defined as follows:

```
type Counter = Rec<SendValue<u64, Z>>;
```

The type `Rec<SendValue<T, Z>>` is unfolded into `SendValue<T, Rec<SendValue<T, Z>>>`. This is achieved by having the following generic implementations of `RecApp` for `SendValue` and `Z`:

```
impl<X> RecApp<X> for Z { type Applied = X; }
impl<X, T, A: RecApp<X>> RecApp<X> for SendValue<T, A> {
    type Applied = SendValue<T, AppRec<A, X>; }

```

Inside `RecApp`, `Z` simply replaces itself with the type argument `X`. `SendValue<T, A>` delegates the type application of `X` to `A`, provided that the session type `A` also implements `RecApp` for `X`.

The session type `Counter` is iso-recursive, as the rolled type `Rec<SendValue<u64, Z>>` and the folded type `SendValue<u64, Rec<SendValue<u64, Z>>` are considered distinct types in Rust. As a result, Ferrite provides the constructs `fix_session` and `unfix_session` for converting between the rolled and unfolded versions of a recursive session type.

### Nested Recursive Session Types

The use of `RecApp` is akin to emulating the higher-kinded type (HKT) `Type → Type` in Rust. As of this writing, HKTs are only available in the nightly (unstable) version of Rust through *generic associated types*. However even with support for HKTs, our defunctionalization-based approach via `RecApp` allows us to generalize to *nested* recursive types.

To account for a recursive type with multiple recursion points, we introduce a *recursion context* `R` as a type-level list of elements (c.f. the linear context of Section 4.2). The type-level natural numbers `Z`, `S<Z>`, etc. are now used as de Bruijn indices to unfold to the elements in the recursion context. The type-level fixed point combinator `Rec` is redefined as `RecX`, containing the recursion context:

```
struct RecX<R, F: RecApp<(RecX<R, F>, R)>> { unfix: Box<AppRec<F, (RecX<R, F>, R)>> }
type Rec<F> = RecX<(), F>;
impl<R, F: RecApp<(RecX<R, F>, R)>> RecApp<R> for RecX<(), F> {
    type Applied = RecX<R, F>; }

```

A recursive session type is defined starting with an empty recursion context. Since nested recursive session types allow a `RecX` to be embedded inside another `RecX`, we have `RecX` also implement `RecApp`, provided it has an empty recursion context. When unfolded from another recursion context `R`, `RecX` simply saves `R` as its own recursion context and does not unfold further in `F`. The inner type `F` is only unfolded once with the full recursion context after all surrounding `RecX` types are unfolded.

The recursive marker `Z` is modified to unfold to the first element of the recursion context. We then implement `S<N>` to unfold to the  $(N+1)$ -th position in the recursion context:

```
impl<A, R> RecApp<(A, R)> for Z { type Applied = A; }
impl<A, R, N: RecApp<R>> RecApp<(A, R)> for S<N> { type Applied = N::Applied; }

```

## 5.2 Shared Session Types

In the previous section we explored a recursive session type `Counter`, which is defined using `Rec` and `Z`. Since `Counter` is defined as a linear session type, it cannot be shared among multiple clients. Shared communication, however, is essential to implement many practical applications. For instance, we may want to implement a simple counter web-service, to send a unique count for each request. To support such shared communication, we introduce *shared session types* in Ferrite, enabling *safe* shared communication in the presence multiple clients.

## Shared Session Types in Ferrite

As introduced in Section 2, the  $SILL_S$  (and  $SILL_R$ ) notion of shared session types is recursive in nature, as a shared session type must offer the same linear critical section to all clients that acquire a shared resource. For instance, a shared version of the `Counter` type in  $SILL_R$  is:

$$\text{SharedCounter} = \uparrow_L^S \text{Int} \triangleleft \downarrow_L^S \text{SharedCounter}$$

The linear portion of `SharedCounter` in between  $\uparrow_L^S$  (acquire) and  $\downarrow_L^S$  (release) amounts to a critical section. When a `SharedCounter` is *acquired*, it offers a linear session type  $\text{Int} \triangleleft \downarrow_L^S \text{SharedCounter}$ , willing to send an integer value, after which it must be *released* to become available again as a `SharedCounter` to the next client.

The recursive aspect of shared session types in  $SILL_R$  means that we can reuse the implementation technique that we use for recursive session types. The type `SharedCounter` can be defined in Ferrite as follows:

```
type SharedCounter = LinearToShared<SendValue<u64, Release>>;
```

Compared to linear recursive session types, the main difference is that instead of using `Rec`, a shared session type is defined using the `LinearToShared` construct. This corresponds to  $\uparrow_L^S$  in  $SILL_R$ , with the inner type `SendValue<u64, Release>` corresponding to the linear portion of the shared session type. At the point of recursion, the type `Release` is used in place of  $\downarrow_L^S \text{SharedCounter}$ . As a result, the type `LinearToShared<SendValue<u64, Release>>` is unfolded into `SendValue<u64, SharedToLinear<LinearToShared<SendValue<u64, Release>>>>` after being acquired. Type unfolding is implemented as follows:

```
trait SharedRecApp<X> { type Applied; }    trait SharedProtocol { ... }
struct SharedToLinear<F> { ... }         struct LinearToShared<F> { ... }
impl<F> Protocol for SharedToLinear<LinearToShared<F>>
  where F: SharedRecApp<SharedToLinear<LinearToShared<F>>> { ... }
impl<F> SharedProtocol for LinearToShared<F>
  where F: SharedRecApp<SharedToLinear<LinearToShared<F>>> { ... }
```

The struct `LinearToShared` is parameterized by a linear session type `F` that implements the trait `SharedRecApp<SharedToLinear<LinearToShared<F>>>`. It uses the `SharedRecApp` trait instead of the `RecApp` trait to ensure that the session type is *strictly equi-synchronizing* [3], requiring an acquired session to be released to the same type at which it was previously acquired. Ferrite enforces this requirement by omitting an implementation of `SharedRecApp` for `End`, ruling out invalid shared session types such as `LinearToShared<SendValue<u64, End>>`. We note that the type argument to `F`'s `SharedRecApp` is another struct `SharedToLinear`, which corresponds to  $\downarrow_L^S$  in  $SILL_R$ . A `SharedProtocol` trait is also defined to identify shared session types, i.e. `LinearToShared`.

Once a shared process is started, a shared channel is created to allow multiple clients to access the shared process through the use of shared channel:

```
struct SharedChannel<S: SharedProtocol> { ... }
impl<S> Clone for SharedChannel<S> { ... };
```

The code above shows the definition of the `SharedChannel` struct. Unlike linear channels, shared channels follow structural typing, i.e. they can be weakened or contracted. This means that we can delegate the handling of shared channels to Rust, given that `SharedChannel` implements Rust's `Clone` trait to allow contraction. Whereas  $SILL_S$  provides explicit constructs for sending and receiving shared channels, Ferrite's shared channels can be sent as regular Rust values using `Send/ReceiveValue`.

On the client side, a `SharedChannel` serves as an endpoint for interacting with a shared process running in parallel. To start the execution of such a shared process, a corresponding

## 22:18 Ferrite: A Judgmental Embedding of Session Types in Rust

Ferrite program has to be defined and executed. Similar to `PartialSession`, we define `SharedSession` as shown below to represent such a shared Ferrite program.

```
struct SharedSession<S: SharedProtocol> { ... }
fn run_shared_session<S: SharedProtocol>(session: SharedSession<S>) -> SharedChannel<S>
```

Just as `PartialSession` encodes linear Ferrite programs without executing them, `SharedSession` encodes shared Ferrite programs without executing them. Since `SharedSession` does not implement the `Clone` trait, the shared Ferrite program is itself affine and cannot be shared. To enable sharing, the shared Ferrite program must first be executed with `run_shared_session`. The function `run_shared_session` takes a shared Ferrite program of type `SharedSession<S>` and starts it in the background as a shared process. Then, in parallel, the shared channel of type `SharedChannel<S>` is returned to the caller, which can then be sent to multiple clients for access to the shared process.

The details of each shared Ferrite construct are described in Appendix A.2.11. Below we demonstrate how a program with a shared session can be defined and used by multiple clients:

```
type SharedCounter = LinearToShared<SendValue<u64, Release>>;
fn counter_producer(current_count: u64) -> SharedSession<SharedCounter> {
  accept_shared_session(async move {
    send_value(current_count, detach_shared_session(
      counter_producer(current_count + 1))) }) }

fn counter_client(counter: SharedChannel<SharedCounter>) -> Session<End> {
  acquire_shared_session(counter, move | chan | {
    receive_value_from(chan, move | count | {
      println!("received count: {}", count);
      release_shared_session(chan, terminate()) }) }) }
```

The recursive function `counter_producer` creates a `SharedSession` program that, when executed, offers a shared channel of session type `SharedCounter`. On the provider side, a shared session is defined using the `accept_shared_session` construct, with a continuation given as an `async` thunk that is executed when a client acquires the shared session and enters the linear critical section (of type `SendValue<u64, SharedToLinear<SharedCounter>>`). Inside the closure, the producer uses `send_value` to send the current count to the client and then uses `detach_shared_session` to exit the linear critical section. The construct `detach_shared_session` offers the linear session type `SharedToLinear<SharedCounter>` and expects a continuation that offers the shared session type `SharedCounter` to serve the next client. We generate the continuation by recursively calling the `counter_producer` function.

The `counter_client` function takes a shared channel of session type `SharedCounter` and returns a session type program that acquires the shared channel and prints the received count value to the terminal. A linear Ferrite program can acquire a shared session using the `acquire_shared_session` construct, which accepts a `SharedChannel` object and adds the acquired linear channel to the linear context. In this case, the continuation closure is given the context lens `Z`, which provides access to the linear channel of session type `SendValue<u64, SharedToLinear<SharedCounter>>` in the first slot of the linear context. It then uses `receive_value_from` to receive the value sent by the shared provider and then prints the value. On the client side, the linear session of type `SharedToLinear<SharedCounter>` must be released using the `release_shared_session` construct. After releasing the shared session, other clients will then be able to acquire the shared session.

```
async fn main () {
  let counter1: SharedChannel<SharedCounter> = run_shared_session(counter_producer(0));
  let counter2 = counter1.clone();
  let child1 = task::spawn(async move { run_session(counter_client(counter1)).await; });
```

```
let child2 = task::spawn(async move { run_session(counter_client(counter2)).await; });
join!(child1, child2).await; }
```

To illustrate a use of `SharedCounter`, we have a `main` function that initializes a shared producer with an initial value of 0 and then runs the shared provider using the `run_shared_session` construct. The returned `SharedChannel` is then cloned, making the shared counter accessible via aliases `counter1` and `counter2`. It then uses `task::spawn` to spawn two async tasks that run `counter_client` twice. A key observation is that multiple Ferrite programs that are executed independently can access *the same* shared producer through a reference to the shared channel.

## 6 N-ary Choice

Session types support *internal* and *external* choice, leaving the choice among several options to the provider or the client, resp. (see Table 2). When restricted to binary choice, the implementation is relatively straightforward, as shown below by the two right rules for internal choice in  $\text{SILL}_R$ . The `offer_left` and `offer_right` constructs allow a provider to offer an internal choice  $A \oplus B$  by offering either  $A$  or  $B$ , resp.

$$\frac{\Gamma; \Delta \vdash K :: A}{\Gamma; \Delta \vdash \text{offer\_left}; K :: A \oplus B} (\text{T}\oplus\text{2L}_R) \quad \frac{\Gamma; \Delta \vdash K :: B}{\Gamma; \Delta \vdash \text{offer\_right}; K :: A \oplus B} (\text{T}\oplus\text{2R}_R)$$

It is straightforward to implement the two versions of the right rules by writing the two respective functions `offer_left` and `offer_right`:

```
fn offer_left<C: Context, A: Protocol, B: Protocol>
  ( cont: PartialSession<C, A> ) -> PartialSession<C, InternalChoice2<A, B>>
fn offer_right < C: Context, A: Protocol, B: Protocol >
  ( cont: PartialSession<C, B> ) -> PartialSession<C, InternalChoice2<A, B>>
```

However, this approach does not scale if we want to generalize choice beyond two options. To support N-ary choice, the functions would have to be explicitly reimplemented N times. Instead, we implement a single `offer_case` function which allows selection from n-ary branches.

### 6.1 Prisms

In Section 4.2, we explored heterogeneous list to encode the linear context, i.e. *products* of session types of arbitrary lengths. We then implemented context *lenses* to access and update individual channels in the linear context. Observing that n-ary choices can be encoded as *sums* of session types, we now use *prisms* to implement the selection of an arbitrary-length branch. Instead of having a binary choice type `InternalChoice2<A, B>`, we can define an n-ary choice type `InternalChoice<HList![...]>`, with `InternalChoice<HList![A, B]>` being the special case of a binary choice. To select a branch out of the heterogeneous list, we define the `Prism` trait as follows:

```
trait Prism<Row> { type Elem; ... }
impl<A, R> Prism<(A, R)> for Z { type Elem = A; ... };
impl<N, A, R> Prism<(A, R)> for S<N> where N: Prism<R> { type Elem = N::Elem; ... }
```

The `Prism` trait is parameterized over a row type `Row=HList![...]`, with the associated type `Elem` being the element type that has been selected from the list by the prism. We then inductively implement `Prism` using type-level natural numbers, with the number `N` used for selecting the N-th element of the heterogeneous list. The definition of `Prism` is similar to `ContextLens`, with the main difference being that we only need `Prism` to support extraction and injections operations on the sum types that are derived from the heterogeneous list. Using `Prism`, a generalized `offer_case` function is implemented as follows:

## 22:20 Ferrite: A Judgmental Embedding of Session Types in Rust

```
fn offer_case<C: Context, A: Protocol, Row, N: Prism<Row, Elem=A>>
  (n: N, cont: PartialSession<C, A>) -> PartialSession<C, InternalChoice<Row>>
```

The function accepts a natural number  $N$  as the first parameter, which acts as the *prism* for selecting a session type  $A_N$  out of the row type  $\text{Row}=\text{HList!}[\dots, A_N, \dots]$ . Through the associated type  $A=N::\text{Elem}$ , `offer_case` forces the programmer to provide a continuation that offers the chosen session type  $A$ .

### 6.2 Binary Branching

While `offer_case` is a step in the right direction, it only allows the selection of a specific choice, but not the provision of *all* possible choices. The latter, however, is necessary to encode the  $\text{SILL}_R$  left rule of internal choice and right rule of external choice. To illustrate the problem, let's consider the right rule of a binary external choice,  $\text{T\&2}_R$ :

$$\frac{\Gamma; \Delta \vdash K_l :: A \quad \Gamma; \Delta \vdash K_r :: B}{\Gamma; \Delta \vdash \text{offer\_choice\_2 } K_l K_r :: A\&B} (\text{T\&2}_R)$$

The `offer_choice_2` construct has two possible continuations  $K_l$  and  $K_r$ , with only one of them being executed, depending on the selection by the client. In a naive implementation, we can define the construct to accept two continuations as follows:

```
fn offer_choice_2<C: Context, A: Protocol, B: Protocol>
  ( cont_left: PartialSession<C, A>, cont_right: PartialSession<C, B> )
  -> PartialSession<C, ExternalChoice2<A, B>>
```

While the above implementation works in most languages, it is not adequate in Rust. Since Rust's type system is *affine*, variables can only be captured by one of the continuation closures, but not both. As far as the compiler is aware, both closures can potentially be called, and we cannot state that one of the branches is guaranteed to never run.

In order for `offer_choice_2` to work in Rust's affine typing, it has to accept only one continuation closure and have it return either `PartialSession<C, A>` or `PartialSession<C, B>`, depending on the client's selection. It is not as straightforward to express such behavior as a valid type in a language like Rust. If Rust supported dependent types, `offer_choice_2` could be implemented along the following lines:

```
fn offer_choice_2<C: Context, A: Protocol, B: Protocol>
  ( cont: impl FnOnce(first: bool) ->
    { if first { PartialSession<C, A> } else { PartialSession<C, B> } } )
  -> PartialSession<C, ExternalChoice2<A, B>>
```

That is, the return type of the `cont` closure depends on the whether the *value* of the `first` argument is true or false. However, since Rust does not support dependent types, we emulate a dependent sum in a non-dependent language, using a CPS transformation:

```
fn offer_choice_2<C: Context, A: Protocol, B: Protocol>
  ( cont: impl FnOnce(InjectSum2<C, A, B>) -> ContSum2<C, A, B> )
  -> PartialSession<C, ExternalChoice2<A, B>>
```

The function `offer_choice_2` accepts a continuation function `cont` that is given a value of type `InjectSum2<C, A, B>` and returns a value of type `ContSum2<C, A, B>`. We will now look at the definitions of `ContSum2` and `InjectSum2`. First, we observe that the different return types for the two branches can be unified with a type `ContSum2`:

```
struct ContSum2<C: Context, A: Protocol, B: Protocol> { ... }
async fn run_cont_sum<C: Context, A: Protocol, B: Protocol>(cont: ContSum2<C, A, B>)
```

The type `ContSum2` contains the necessary data for executing either a `PartialSession<C, A>` or a `PartialSession<C, B>`, together with the runtime data for the linear context `c`. For brevity, the implementation details of `ContSum2` are omitted, with the private function `run_cont_sum` provided as an abstraction for Ferrite to execute the continuation.

We then define `InjectSum2` as a sum of boxed closures that would construct a `ContSum2` from either a `PartialSession<C, A>` or a `PartialSession<C, B>`:

```
enum InjectSum2<C, A, B> {
  InjectLeft(Box<dyn FnOnce(PartialSession<C, A>) -> ContSum2<C, A, B>>),
  InjectRight(Box<dyn FnOnce(PartialSession<C, B>) -> ContSum2<C, A, B>>) }
```

When the `cont` passed to `offer_choice_2` is given a value of type `InjectSum2<C, A, B>`, it has to branch on it and match on whether the `InjectLeft` or `InjectRight` constructors are used. Since the return type of `cont` is `ContSum2<C, A, B>` and the constructor for `ContSum2` is private, there is no other way for `cont` to construct the return value other than to call either `InjectLeft` or `InjectRight` with the appropriate continuation.

The use of `InjectSum2` prevents the programmer from providing the wrong branch in the continuation by keeping the constructor private. However a private constructor alone cannot prevent two uses of `InjectSum2` to be *deliberately* interchanged, causing a protocol violation. To fully ensure that there is no way for the user to provide a `ContSum2` from elsewhere, we instead use a technique from `GhostCell` [52] that uses *higher-ranked trait bounds* (HTRB) to mark a phantom invariant lifetime on both `InjectSum2` and `ContSum2`:

```
fn offer_choice_2<C: Context, A: Protocol, B: Protocol>
  ( cont: for <'r> impl FnOnce(InjectSum2<'r, C, A, B>) -> ContSum2<'r, C, A, B> )
  -> PartialSession<C, ExternalChoice2<A, B>>
```

The use of HRTB ensures that each call of `offer_choice_2` would generate a unique lifetime `'r` for the continuation. Using that, Ferrite can ensure that a value of type `InjectSum2<'r1, C, A, B>` cannot be used to construct the return value of type `ContSum2<'r2, C, A, B>`, if the lifetimes `<'r1>` and `<'r2>` are different. An example use of `offer_choice_2` is as follows:

```
let choice_provider: Session<ExternalChoice2<SendValue<u64, End>, End>>
  = offer_choice_2(| b | { match b {
    InjectLeft(ret) => ret(send_value(42, terminate())),
    InjectRight(ret) => ret(terminate()) } });
```

The example code above requires some boilerplate code to call the session injector `ret` to wrap around the continuation expression. To free the programmer from writing such boilerplate, Ferrite also provides a macro `offer_choice` that translates into the underlying pattern matching syntax, which is explained in the next section.

### 6.3 N-ary Branching

To generalize `offer_choice_2` to n-ary choices, Ferrite has its own version of polymorphic variants implemented in Rust. Our implementation specifically targets Rust, and is based on similar work by [30] and [31]. The base variant types are as follows:

```
enum Bottom {}          enum Sum<A, B> { Inl(A), Inr(B) }
trait TypeApp<A> { type Applied; }      trait SumApp<F> { type Applied; }
type App<F, A> = <F as TypeApp<A>>::Applied;
type AppSum<Row, F> = <Row as SumApp<F>>::Applied;
impl<F> SumApp<F> for () { type Applied = Bottom; }
impl<A, F: TypeApp<A>, R: SumApp<F>> SumApp<F> for (A, R) {
  type Applied = Sum<F::Applied, R::Applied>; }
```

Similar to `RecApp` described in Section 5.1, `TypeApp` is used to represent a Rust type emulating the kind `Type → Type` for non-recursive usage. Furthermore, the `SumApp` trait is



## 22:22 Ferrite: A Judgmental Embedding of Session Types in Rust

used to represent a Rust type emulating the kind  $(\text{Type} \rightarrow \text{Type}) \rightarrow \text{Type}$ . The type alias `App<F, A>` is used to extract the associated type `Applied` when `F` is applied to `A` via `TypeApp`. The type alias `AppSum<Row, F>` is used to extract the associated type `Applied` when a row type `Row` is applied to a type constructor `F`, which implements `TypeApp<A>` for all `A`.<sup>1</sup>

Using `SumApp`, we map an heterogeneous list to nested sums such that `AppSum<HList![A0, A1, ...], F> = Sum![App<F, A0>, App<F, A1>, ...]`, with the macro `Sum!` used to expand the macro arguments into nested sums, i.e. `Sum![A0, A1, ...] = Sum<A0, Sum<A1, ..., Bottom>>`. We then define the n-ary versions of `InjectSum2` and `ContSum2` as follows:

```
struct InjectSessionF<'r, Row, C> {}      struct InjectSession<'r, Row, C, A> { ... }
struct ContSum<'r, Row, C: Context> { ... }
impl<'r, Row, C: Context> TypeApp<A> for InjectSessionF<C> {
    type Applied = InjectSession<C, A>; }
impl<'r, Row, C: Context, A: Protocol> FnOnce(PartialSession<C, A>)
    -> ContSum<'r, Row, C> for InjectSession<'r, Row, C, A> { ... }
```

The type `InjectSessionF<'r, Row, C>` serves as a marker type for `TypeApp`, such that when applied to a type `A`, we get the struct `InjectSession<'r, Row, C, A>`. Conceptually, the struct implements the trait `FnOnce(PartialSession<C, A>) -> ContSum<'r, Row, C>`, so that we can apply a `PartialSession<C, A>` to it and get back a `ContSum<'r, Row, C>`.<sup>2</sup> The composed type `AppSum<Row, InjectSessionF<'r, Row, C>>` represents a row of `InjectSession`, with `Row` being a heterogeneous list in the form `HList![A0, A1, ..., AN-1]`. For example, the type `AppSum<HList![A, B], InjectSessionF<'r, Row, C>>` evaluates to `Sum![InjectSession<'r, Row, C, A>, InjectSession<'r, Row, C, B>]`, which is isomorphic to the type `InjectSum2<C, A, B>` that we defined for the binary case. Using the row constructs, we can define n-ary version `offer_choice` as follows:

```
fn offer_choice<C: Context, Row>(cont1 : impl for <'r>
    FnOnce(AppSum<Row, InjectSessionF<'r, Row, C>>) -> ContSum<'r, Row, C>
) -> PartialSession<C, ExternalChoice<Row>>
where Row: SumApp<InjectSessionF<'r, Row, C>>, ...
```

With the n-ary version of `offer_choice` available, we can re-implement binary choice as a specialized version. To do that, we only need a few type aliases and struct definitions to make the syntax more pleasing:

```
enum EitherSum<A, B> { Left(A), Right(B) };      type Either<A, B> = HList![A, B];
const LeftLabel: Z = Z::new();                  const RightLabel: S<Z> = <S<Z>>::new();
impl<A, B> std::convert::From<Sum![A, B]> for EitherSum<A, B> { ... }
```

We first define an `EitherSum` enum, and a `std::convert::From` instance that converts an unlabeled nested sum `Sum![A, B]` into the labeled sum `EitherSum<A, B>`. The conversion allows users to use a flat list of labeled match arms during branching, and give meaningful labels `Left` and `Right` to each branch. We also define `Either<A, B>` as a type alias to the row type `HList![A, B]`, to give a meaningful name to the choice protocol. Finally we define the constants `LeftLabel` and `RightLabel` to refer to the prisms `Z` and `S<Z>`, resp. Ferrite also provides a helper macro `define_choice!` to help users define custom choice protocols that look similar to the above. This is used in conjunction with macros such as `offer_choice!`, which

---

<sup>1</sup> For brevity, we omit some details that the types `App` and `AppSum` are actually implemented as structs that are *isomorphic* to `<F as TypeApp<A>>::Applied` and `<Row as SumApp<F>>::Applied`, resp. The main difference is that the actual structs *hide* the propagation of the trait bound requirements of `TypeApp` and `SumApp` from their callers, resulting in much cleaner code. This does not affect the understanding of the core concepts introduced in this section.

<sup>2</sup> Technically, Rust does not allow custom implementation of `FnOnce`, so Ferrite defines a custom trait with the same behavior.



```

1 enum CanvasMsg { Canvas2d(Canvas2dMsg, CanvasId), Close(CanvasId), ... }
2 enum Canvas2dMsg { LineTo(Point2D), GetTransform(Sender<Transform2D>),
3   IsPointInPath(f64, f64, FillRule, IpcSender<bool>), ... }
4 enum ConstellationCanvasMsg { Create { id_sender: Sender<CanvasId>, size: Size2D } }
5 struct CanvasPaintThread { canvases: HashMap<CanvasId, CanvasData>, ... }
6 impl CanvasPaintThread { ...
7   fn start() -> (Sender<ConstellationCanvasMsg>, Sender<CanvasMsg>) {
8     let (msg_sender, msg_receiver) = channel();
9     let (create_sender, create_receiver) = channel();
10    thread::spawn(move || { loop { select! {
11      recv(canvas_msg_receiver) -> { ...
12        CanvasMsg::Canvas2d(message, canvas_id) => { ...
13          Canvas2dMsg::LineTo(point) => self.canvas(canvas_id).move_to(point),
14          Canvas2dMsg::GetTransform(sender) =>
15            sender.send(self.canvas(canvas_id).get_transform().unwrap(), ... }
16          CanvasMsg::Close(canvas_id) => canvas_paint_thread.canvases.remove(&canvas_id) }
17      recv(create_receiver) -> { ...
18        ConstellationCanvasMsg::Create { id_sender, size } => {
19          let canvas_id = ...; self.canvases.insert(canvas_id, CanvasData::new(size, ...));
20          id_sender.send(canvas_id); } } } });
21    (create_sender, msg_sender) }
22  fn canvas(&mut self, canvas_id: CanvasId) -> &mut CanvasData {
23    self.canvases.get_mut(&canvas_id).expect("Bogus canvas id") } }

```

■ **Figure 1** Message-passing concurrency in Servo’s canvas component (simplified for illustration purposes).

cleans up the boilerplate required to enable different match branches to return different types. Using the macros, users can define the same `Either` protocol and write an external choice provider as follows:

```

define_choice!{ Either<A, B>; Left: A, Right: B }
// Inferred type: Session<ExternalChoice<Either<SendValue<u64, End>, End>>>
let provider = offer_choice!{
  Left => send_value(42, terminate()), Right => terminate() };

```

For convenience, Ferrite exports the choice definition for `Either` for the anonymous declaration of binary choice in session types.

## 7 Evaluation

The Ferrite library is more than just a research prototype. It is designed for practical use in real world applications. To evaluate the design and implementation of Ferrite, we re-implemented the communication layer of the canvas component of Servo [32] entirely in Ferrite. Servo is an under development browser engine that uses message-passing for heavy task parallelization. Canvas provides 2D graphic rendering, allowing clients to create new canvases and perform operations on a canvas such as moving the cursor and drawing shapes.

The canvas component is a good target for evaluation as it is sufficiently complex and also very demanding in terms of performance. Canvas is commonly used for animations in web applications. For an animation to look smooth, a canvas must render at least 24 frames per second, with potentially thousands of operations to be executed per frame.

The changes we made are fairly minimal, consisting of roughly 750 lines of additions and 620 lines of deletions, out of roughly 300,000 lines of Rust code in Servo. The sources of our implementation are provided as an artifact. To differentiate the two versions of code snippets, we use [blue](#) for the original code, and [green](#) for the code using Ferrite.

### 7.1 Servo Canvas Component

Figure 1 provides a sketch of the main communication paths in Servo’s canvas component [33]. The canvas component is implemented by the `CanvasPaintThread`, whose function `start` contains the main communication loop running in a separate thread (lines 10–20). This

loop processes client requests received along `canvas_msg_receiver` and `create_receiver`, which are the receiving endpoints of the channels created prior to spawning the loop (lines 8–9). The channels are typed with the enumerations `ConstellationCanvasMsg` and `CanvasMsg`, defining messages for creating and terminating the canvas component and for executing operations on an individual canvas, resp. When a client sends a message that expects a response from the recipient, such as `GetTransform` and `IsPointInPath` (lines 2–3), it sends a channel along with the message to be used by the recipient to send back the result. Canvases are identified by an id, which is generated upon canvas creation (line 19) and stored in the thread’s `canvases` hash map (line 5). If a client requests an invalid id, for example after prior termination and removal of the canvas (line 16), the failed assertion `expect("Bogus canvas id")` (line 23) will result in a `panic!`, causing the canvas component to crash and subsequent calls to fail.

The code in Figure 1 uses a clever combination of enumerations to type channels and ownership to rule out races on the data sent along channels. Nonetheless, Rust’s type system is not expressive enough to enforce the intended *protocol* of message exchange and existence of a communication partner. The latter is a consequence of Rust’s type system being *affine*, which permits “dropping of a resource”. The dropping or premature closure of a channel, however, can result in a proliferation of `panic!` and thus cause an entire application to crash. In fact, while refactoring Servo to use Ferrite, we were able to uncover a protocol violation in Servo, caused by one of the nested match arms of the provider doing an early return before sending back any result to the client.

## 7.2 Canvas Protocol in Ferrite

In the original canvas component, the provider `CanvasPaintThread` accepts messages of type `CanvasMsg`, made up of a combination of smaller sub-message types such as `Canvas2dMsg`. We note that the majority of the sub-message types have the following trivial form:

```
enum CanvasMsg { Canvas2d(Canvas2dMsg, CanvasId), Close(CanvasId), ... }
enum Canvas2dMsg { BeginPath, ClosePath, Fill(FillOrStrokeStyle), ... }
```

The trivial sub-message types such as `BeginPath`, `Fill`, and `LineTo` do not require a response from the provider, so the client can simply fire them and proceed. Although we can offer all sub-message types as separate branches in an external choice, it is more efficient to keep trivial sub-messages in a single enum. In our implementation, we define `CanvasMessage` to have similar sub-messages as `Canvas2dMsg`, with non-trivial messages such as `IsPointInPath` moved to separate branches.

```
enum CanvasMessage { BeginPath, ClosePath, Fill(FillOrStrokeStyle), ... }
define_choice! { CanvasOps; Message: ReceiveValue<CanvasMessage, Release>, ... }
type Canvas = LinearToShared<ExternalChoice<CanvasOps>>;
```

We use the `define_choice!` macro described in Section 6 to define an n-ary choice `CanvasOps`. The first branch of `CanvasOps` is labelled `Message`, and the only action is for the provider to receive a `CanvasMessage`. The choices are offered as an external choice, and the session type `CanvasProtocol` is defined as a shared protocol that offers the choices in the critical section.

The original design of the `CanvasPaintThread` would be sufficient if the only messages being sent were trivial messages. However, `Canvas2dMsg` also contains non-trivial sub-messages, such as `GetImageData` and `IsPointInPath`, demanding a response from the provider:

```
enum Canvas2dMsg { ..., GetImageData(Rect<u64>, Size2D<u64>, IpcBytesSender),
  IsPointInPath(f64, f64, FillRule, IpcSender<bool>), ... }
```

To obtain the result from the original canvas, clients must create a new inter-process communication (IPC) channel and bundle the channel’s sender endpoint with the message. In our implementation, we define separate branches in `CanvasOps` to handle non-trivial cases:

■ **Table 4** MotionMark Benchmark scores in fps (higher is better)

Benchmark Name	Servo	Servo/Ferrite	Firefox	Chrome
Arcs	12.21 ± 6.75%	11.83 ± 11.49%	52.61 ± 32.88%	46.00 ± 9.00%
Paths	43.76 ± 10.66%	40.98 ± 18.94%	55.59 ± 28.80%	59.50 ± 14.90%
Lines	7.48 ± 7.06%	11.47 ± 12.74%	14.35 ± 6.65%	32.43 ± 6.48%
Bouncing clipped rects	18.43 ± 7.06%	18.23 ± 11.00%	34.82 ± 7.76%	58.07 ± 19.85%
Bouncing gradient circles	8.02 ± 7.74%	7.72 ± 12.63%	58.79 ± 21.03%	59.77 ± 10.07%
Bouncing PNG images	7.97 ± 5.91%	6.31 ± 10.26%	24.61 ± 6.35%	59.94 ± 13.04%
Stroke shapes	10.60 ± 3.95%	10.35 ± 10.96%	51.21 ± 11.25%	59.38 ± 16.87%
Put/get image data	60.01 ± 3.81%	32.08 ± 10.83%	59.66 ± 20.16%	60.00 ± 5.00%

```
define_choice! { CanvasOps; Message: ReceiveValue<CanvasMessage, Release>,
  GetImageData: ReceiveValue<(Rect<u64>, Size2D<u64>), SendValue<ByteBuf, Release>>,
  IsPointInPath: ReceiveValue<(f64, f64, FillRule), SendValue<bool, Release>>, ... }
```

The original `GetImageData` accepts an `IpcBytesSender`, which sends raw bytes back to the client. In Ferrite, we translate the use of `IpcBytesSender` to the type `SendValue<ByteBuf, Z>`, which sends the raw bytes wrapped in a `ByteBuf` type. We discuss possible performance penalties of this approach in Section 7.3.

Aside from the `Canvas` protocol, we also redesign the use of `ConstellationCanvasMsg` into its own shared protocol, `ConstellationCanvas`:

```
type ConstellationCanvas = LinearToShared<ReceiveValue<Size2D,
  SendValue<SharedChannel<Canvas>, Release>>>;
```

To create a new canvas, a client first acquires the shared channel of type `SharedChannel<ConstellationCanvas>`. Afterwards, the client sends the `Size2D` parameter to specify the canvas size. The constellation canvas provider then spawns a new canvas shared process through `run_shared_session` and sends back the shared channel of type `SharedChannel<Canvas>` as a value. Finally, the session is released, allowing other clients to acquire the shared provider.

### 7.3 Performance Evaluation

To evaluate the performance of the canvas component, we use the MotionMark benchmark suite [50]. MotionMark is a web benchmark that focuses on graphics performance of web browsers. It contains benchmarks for various web components, including canvas, CSS, and SVG. As MotionMark does not yet support Servo, we modified the benchmark code to make it work in the absence of features that are not implemented in Servo.

We provide the modified benchmark source code along with instructions for running it as an artifact. Appendix D is also provided to highlight some of the implementation challenges in porting Servo to use Ferrite, in particular on the latency incurred by inter-process communication, and our workaround to compensate the complication.

For the purpose of this evaluation, we focused on benchmarks that target the canvas component and skipped benchmarks that fail in Servo due to missing features. We ran each benchmark in a fixed 1600x800 resolution for 30 seconds, on a Core i7 Linux desktop machine. We ran the benchmarks against the original Servo, modified Servo with Ferrite canvas (Servo/Ferrite), Firefox, and Chrome. Our performance scores are measured in the fixed mode version of MotionMark, which measures frames per second (fps) performance of executing the same set of canvas operations per frame.

The benchmark results are shown in Table 4, with the performance scores in fps (higher fps is better). It is worth noting that a benchmark can achieve at most 60 fps. Our goal in this benchmark is to keep the scores of Servo/Ferrite close to those of Servo, *not* to achieve better performance than the original. This is shown to be the case in most of the benchmarks.

The only benchmark with a large difference between Servo and Servo/Ferrite is *Put/get image data*, with Ferrite performing 2x worse. This is because in Servo/Ferrite, we use `ByteBuf` to transfer the images as raw bytes within the same shared channel. In contrast, Servo uses a specialized structure `IpBytesSender` for transferring of raw bytes in parallel to other messages. As a result, the communication in Servo/Ferrite is congested during the transfer of the image data, while the original Servo can process new messages in parallel to the image data being transmitted.

We also observe that there are significant performance differences in the scores between Servo and those in Firefox and Chrome, indicating that there exist performance bottlenecks in Servo unrelated to communication protocols.

## 8 Related and Future Work

Session type embeddings exist for various languages, including Haskell [38, 22, 29, 34], OCaml [36, 21], Java [17, 16], and Scala [43]. Functional languages like ML, OCaml, and Haskell, in particular, are ideal host languages for creating EDSLs thanks to their advanced features (e.g. type classes, type families, higher-rank and higher-kinded types and GADTs). [38] first demonstrated the feasibility of embedding session types in Haskell, with refinements done in later works [22, 29, 34]. Similar embeddings have also been contributed in the context of OCaml by `FuSe` [36] and `session-ocaml` [21].

Aside from Ferrite, there are other implementations of session types in Rust, including `session_types` [23], `sesh` [27], and `rumpsteak` [6, 7]. `session_types` were the first implementation to make use of affinity to provide a session type library in Rust. `sesh` emphasizes this aspect by embedding the affine session type system Exceptional GV [10] in Rust. Both `session_types` and `sesh` adopt a classical perspective, requiring the endpoints of a channel to be typed with dual types. `rumpsteak` develops an embedding of multiparty session types by generating Rust types derived from multiparty session types defined in Scribble [53].

Due to their reliance on Rust's affine type system, neither `session_types` nor `sesh` prevents a channel endpoint from being dropped prematurely, relegating the handling of such errors to the runtime. `rumpsteak` uses some type-level techniques similar to Ferrite to enforce a channel's linear usage in the continuation passed to the `try_session` function. This ensures that a linear channel in `rumpsteak` is always fully consumed, if it is ever consumed. However, prior to the call to `try_session`, the linear channel exist as an affine value, which may be dropped by the Rust program without being consumed at all, thereby causing deadlock. In comparison, Ferrite enforces linearity at all level, including safe linking of multiple linear processes using `cut`.

In terms of concurrency, `session_types`, `sesh`, and `rumpsteak` all require the programmer to manually manage concurrency, either by spawning threads or async tasks. This introduces potential failure when the code fails follow the requirement to spawn all processes. On the other hand, the simplicity of such a model allows relatively few threads or async tasks to be spawned, thereby allowing the underlying runtime to execute the processes more efficiently. In comparison, Ferrite offers fully managed concurrency, without the programmer having to worry about how to spawn the processes and execute them in parallel.

In terms of performance, the downside of Ferrite's concurrency approach is that it aggressively spawns new async tasks in each use of `cut`. Although async tasks in Rust are much more lightweight than OS threads, there is still a significant overhead in spawning and managing many async tasks, especially in micro-benchmarks. As a result, Ferrite tends to perform slower than alternative Rust implementations in settings where only a fixed small

number of processes need to be spawned. Nevertheless, it is worth noting that the async ecosystem in Rust is still relatively immature, with many potential improvements to be made. In practice, the overhead of the async runtime may also be negligible when compared to the core application logic. In such cases, Ferrite would also allow applications to scale more easily by allowing many more processes to be spawned and managed concurrently without requiring additional effort from the programmer.

In terms of DSL design, Ferrite is more closely related to the embeddings in OCaml and Haskell, as it fully enforces a linear treatment of session type channels and thus *statically* rules out any panics arising from dropping a channel prematurely. Ferrite also differs from other libraries in that it adopts intuitionistic typing [4], allowing the typing of a channel rather than its two endpoints. On the use of profunctor optics, our work is the first to connect n-ary choice to prisms, while prior work by `session-ocaml` [22] has only established the connection between lenses, the dual of prisms, and linear contexts. `FuSe` [36] and `session-ocaml` [21] have previously explored the use of n-ary (generalized) choice through extensible variants available only in OCaml. Our work demonstrates that it is possible to encode extensible variants, and thus n-ary choice, as type-level constructs using features available in Rust.

A major difference in terms of implementation is that Ferrite uses a continuation-passing style, whereas Haskell and OCaml embeddings commonly use (indexed) monads and notation style. This technical difference amounts to a key conceptual one: a direct correspondence between the Rust programs generated from Ferrite constructs and the  $SILL_R$  typing derivation. As a result, the generated Rust code can be viewed as carrying the proof of protocol adherence.

The embeddings of `ESJ` [16] and `lchannels` [43] also adopt a continuation-passing style, but do not faithfully embed typing derivations (i.e. they do not statically enforce linearity). These approaches follow an encoding of session types using linear types [8] first proposed by Kobayashi [25] in the setting of  $\pi$ -calculus. Type systems for message-passing in  $\pi$ -calculus have a long history, dating back to the work of Kobayashi and Igarashi [18, 19, 20]. These systems often focus on (but are not limited to) deadlock-freedom and lock-freedom [26] by enforcing a partial order on matching communication. This approach has been studied for the linear  $\pi$ -calculus [35] and in the presence of interrupts [44] or unbounded process networks [12]. While session types are generally less powerful than the approaches of Kobayashi et al., they provide a useful compromise between expressiveness and simplicity, being more amenable to embeddings in general-purpose language constructs and type systems.

In terms of expressiveness, Ferrite contributes over all prior session-based works in its support for shared session types [1], allowing it to express real-world protocols, as demonstrated by our implementation of Servo’s canvas component. Shared session types reclaim the expressiveness of the untyped asynchronous  $\pi$ -calculus in session-typed languages [2], at the cost of deadlock-freedom. Recent extensions of classical linear logic session types contribute another approach to softening the rigidity of linear session types to support multiple client sessions and nondeterminism [39] and memory cells and nondeterministic updates [41], resp.

Our technique of a judgmental embedding opens up new possibilities for embedding type systems other than session types in Rust. Although we have demonstrated that the judgmental embedding is sufficiently powerful to encode a type system like session types, the embedding is currently *shallow*, with the implementation hardcoded to use the channels and async run-time from `tokio`. Rust comes with unique features such as affine types and lifetimes that makes it especially suited for implementing concurrency primitives, as evidenced by the wealth of channel and async run-time implementations available. As discussed in Section 7, one of our future goals is to explore the possibility of making Ferrite a *deep* embedding of

session types in Rust, so that users can choose from multiple low-level implementations. Although deep embeddings have extensively been explored for languages like Haskell [45, 29], it remains an open question to find suitable approaches that work well in Rust.

---

## References

- 1 Stephanie Balzer and Frank Pfenning. Manifest sharing with session types. *Proceedings of the ACM on Programming Languages (PACMPL)*, 1(ICFP):37:1–37:29, 2017.
- 2 Stephanie Balzer, Frank Pfenning, and Bernardo Toninho. A universal session type for untyped asynchronous communication. In *29th International Conference on Concurrency Theory (CONCUR)*, LIPIcs, pages 30:1–30:18. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.
- 3 Stephanie Balzer, Bernardo Toninho, and Frank Pfenning. Manifest deadlock-freedom for shared session types. In *28th European Symposium on Programming (ESOP)*, volume 11423 of *Lecture Notes in Computer Science*, pages 611–639. Springer, 2019.
- 4 Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In *21st International Conference on Concurrency Theory (CONCUR)*, pages 222–236. Springer, 2010.
- 5 Luís Caires, Frank Pfenning, and Bernardo Toninho. Linear logic propositions as session types. *Mathematical Structures in Computer Science*, 26(3):367–423, 2016.
- 6 Zak Cutner and Nobuko Yoshida. Safe session-based asynchronous coordination in rust. In Ferruccio Damiani and Ornela Dardha, editors, *Coordination Models and Languages - 23rd IFIP WG 6.1 International Conference, COORDINATION 2021, Held as Part of the 16th International Federated Conference on Distributed Computing Techniques, DisCoTec 2021, Valletta, Malta, June 14-18, 2021, Proceedings*, volume 12717 of *Lecture Notes in Computer Science*, pages 80–89. Springer, 2021. doi:10.1007/978-3-030-78142-2\\_5.
- 7 Zak Cutner, Nobuko Yoshida, and Martin Vassor. Deadlock-free asynchronous message reordering in rust with multiparty session types. *CoRR*, abs/2112.12693, 2021. URL: <https://arxiv.org/abs/2112.12693>, arXiv:2112.12693.
- 8 Ornela Dardha, Elena Giachino, and Davide Sangiorgi. Session types revisited. In *Principles and Practice of Declarative Programming (PPDP)*, pages 139–150, 2012.
- 9 J. Nathan Foster, Michael B. Greenwald, Jonathan T. Moore, Benjamin C. Pierce, and Alan Schmitt. Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem. *ACM Trans. Program. Lang. Syst.*, 29(3):17, 2007. doi:10.1145/1232420.1232424.
- 10 Simon Fowler, Sam Lindley, J. Garrett Morris, and Sára Decova. Exceptional asynchronous session types: Session types without tiers. *Proceedings of the ACM on Programming Languages*, 3(POPL):28:1–28:29, 2019. doi:10.1145/3290341.
- 11 Andrew Gerrand. The go blog: Share memory by communicating, 2010. URL: <https://blog.golang.org/share-memory-by-communicating>.
- 12 Elena Giachino, Naoki Kobayashi, and Cosimo Laneve. Deadlock analysis of unbounded process networks. In *CONCUR 2014 - Concurrency Theory - 25th International Conference, CONCUR 2014, Rome, Italy, September 2-5, 2014. Proceedings*, pages 63–77, 2014.
- 13 Kohei Honda. Types for dyadic interaction. In *4th International Conference on Concurrency Theory (CONCUR)*, pages 509–523. Springer, 1993.
- 14 Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In *7th European Symposium on Programming (ESOP)*, pages 122–138. Springer, 1998.
- 15 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In *35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 273–284. ACM, 2008.
- 16 Raymond Hu, Dimitrios Kouzapas, Olivier Pernet, Nobuko Yoshida, and Kohei Honda. Type-safe eventful sessions in Java. In *24th European Conference on Object-Oriented Programming*



- (*ECOOP*), volume 6183 of *Lecture Notes in Computer Science*, pages 329–353. Springer, 2010. doi:10.1007/978-3-642-14107-2\_16.
- 17 Raymond Hu, Nobuko Yoshida, and Kohei Honda. Session-based distributed programming in Java. In *22nd European Conference on Object-Oriented Programming (ECOOP)*, volume 5142 of *Lecture Notes in Computer Science*, pages 516–541. Springer, 2008. doi:10.1007/978-3-540-70592-5\_22.
  - 18 Atsushi Igarashi and Naoki Kobayashi. Type-based analysis of communication for concurrent programming languages. In *Static Analysis, 4th International Symposium, SAS '97*, pages 187–201, 1997.
  - 19 Atsushi Igarashi and Naoki Kobayashi. A generic type system for the pi-calculus. In *Conference Record of POPL 2001: The 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 128–141, 2001.
  - 20 Atsushi Igarashi and Naoki Kobayashi. A generic type system for the pi-calculus. *Theor. Comput. Sci.*, 311(1-3):121–163, 2004.
  - 21 Keigo Imai, Nobuko Yoshida, and Shoji Yuen. Session-ocaml: a session-based library with polarities and lenses. *Science of Computer Programming*, 172:135–159, 2019. doi:10.1016/j.scico.2018.08.005.
  - 22 Keigo Imai, Shoji Yuen, and Kiyoshi Agusa. Session type inference in haskell. In *3rd Workshop on Programming Language Approaches to Concurrency and Communication-cEntric Software (PLACES) 2010, Paphos, Cyprus, 21st March 2011*, volume 69 of *EPTCS*, pages 74–91, 2010. doi:10.4204/EPTCS.69.6.
  - 23 Thomas Bracht Laumann Jespersen, Philip Munksgaard, and Ken Friis Larsen. Session types for Rust. In *11th ACM SIGPLAN Workshop on Generic Programming (WGP)*, 2015. doi:10.1145/2808098.2808100.
  - 24 Oleg Kiselyov, Ralf Lämmel, and Kean Schupke. Strongly typed heterogeneous collections. In Henrik Nilsson, editor, *Proceedings of the ACM SIGPLAN Workshop on Haskell, Haskell 2004, Snowbird, UT, USA, September 22-22, 2004*, pages 96–107. ACM, 2004. doi:10.1145/1017472.1017488.
  - 25 Naoki Kobayashi. Type systems for concurrent programs. In Bernhard K. Aichernig and T. S. E. Maibaum, editors, *Formal Methods at the Crossroads. From Panacea to Foundational Support, 10th Anniversary Colloquium of UNU/IIST, the International Institute for Software Technology of The United Nations University, Lisbon, Portugal, March 18-20, 2002, Revised Papers*, volume 2757 of *Lecture Notes in Computer Science*, pages 439–453. Springer, 2002. doi:10.1007/978-3-540-40007-3\_26.
  - 26 Naoki Kobayashi and Davide Sangiorgi. A hybrid type system for lock-freedom of mobile processes. *ACM Trans. Program. Lang. Syst.*, 32(5):16:1–16:49, 2010. doi:10.1145/1745312.1745313.
  - 27 Wen Kokke. Rusty variation: Deadlock-free sessions with failure in rust. In *12th Interaction and Concurrency Experience, ICE 2019*, pages 48–60, 2019.
  - 28 Sam Lindley and J. Garrett Morris. A semantics for propositions as sessions. In *24th European Symposium on Programming (ESOP)*, volume 9032 of *Lecture Notes in Computer Science*, pages 560–584, 2015. doi:10.1007/978-3-662-46669-8\_23.
  - 29 Sam Lindley and J. Garrett Morris. Embedding session types in Haskell. In *9th International Symposium on Haskell*, pages 133–145. ACM, 2016. doi:10.1145/2976002.2976018.
  - 30 J. Garrett Morris. Variations on variants. In Ben Lippmeier, editor, *Proceedings of the 8th ACM SIGPLAN Symposium on Haskell, Haskell 2015, Vancouver, BC, Canada, September 3-4, 2015*, pages 71–81. ACM, 2015. doi:10.1145/2804302.2804320.
  - 31 J. Garrett Morris and James McKinna. Abstracting extensible data types: or, rows by any other name. *PACMPL*, 3(POPL):12:1–12:28, 2019. doi:10.1145/3290325.
  - 32 Mozilla. Servo, the Parallel Browser Engine Project. <https://servo.org/>, 2012.



- 33 Mozilla. Servo source code – canvas paint thread, 2021. URL: [https://github.com/servo/servo/blob/d13a9355b8e66323e666dde7e82ced7762827d93/components/canvas/canvas\\_paint\\_thread.rs](https://github.com/servo/servo/blob/d13a9355b8e66323e666dde7e82ced7762827d93/components/canvas/canvas_paint_thread.rs).
- 34 Dominic A. Orchard and Nobuko Yoshida. Effects as sessions, sessions as effects. In *43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 568–581. ACM, 2016. doi:10.1145/2837614.2837634.
- 35 Luca Padovani. Deadlock and lock freedom in the linear  $\pi$ -calculus. In *Computer Science Logic – Logic in Computer Science (CSL-LICS)*, pages 72:1–72:10, 2014.
- 36 Luca Padovani. A simple library implementation of binary sessions. *J. Funct. Program.*, 27:e4, 2017. doi:10.1017/S0956796816000289.
- 37 Matthew Pickering, Jeremy Gibbons, and Nicolas Wu. Profunctor optics: Modular data accessors. *Programming Journal*, 1(2):7, 2017. doi:10.22152/programming-journal.org/2017/1/7.
- 38 Riccardo Pucella and Jesse A. Tov. Haskell session types with (almost) no class. In *1st ACM SIGPLAN Symposium on Haskell*, pages 25–36. ACM, 2008. doi:10.1145/1411286.1411290.
- 39 Zesen Qian, G. A. Kavvos, and Lars Birkedal. Client-server sessions in linear logic. *CoRR*, abs/2010.13926, 2020. URL: <https://arxiv.org/abs/2010.13926>, arXiv:2010.13926.
- 40 John C. Reynolds. Definitional interpreters for higher-order programming languages. In *ACM Annual Conference*, volume 2, pages 717–740. ACM, 1972. doi:10.1145/800194.805852.
- 41 Pedro Rocha and Luís Caires. Propositions-as-types and shared state. *Proc. ACM Program. Lang.*, 5(ICFP):1–30, 2021.
- 42 Matthew Sackman and Susan Eisenbach. Session types in haskell: Updating message passing for the 21st century. Technical report, Imperial College, 2008. URL: <http://hdl.handle.net/10044/1/5918>.
- 43 Alceste Scalas and Nobuko Yoshida. Lightweight session programming in Scala. In *30th European Conference on Object-Oriented Programming (ECOOP)*, volume 56 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:28. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2016.
- 44 Kohei Suenaga and Naoki Kobayashi. Type-based analysis of deadlock for a concurrent calculus with interrupts. In Rocco De Nicola, editor, *Programming Languages and Systems, 16th European Symposium on Programming, ESOP 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007, Braga, Portugal, March 24 - April 1, 2007, Proceedings*, volume 4421 of *Lecture Notes in Computer Science*, pages 490–504. Springer, 2007. doi:10.1007/978-3-540-71316-6\_33.
- 45 Josef Svenningsson and Emil Axelsson. Combining deep and shallow embedding for EDSL. In Hans-Wolfgang Loidl and Ricardo Peña, editors, *Trends in Functional Programming - 13th International Symposium, TFP 2012, St. Andrews, UK, June 12-14, 2012, Revised Selected Papers*, volume 7829 of *Lecture Notes in Computer Science*, pages 21–36. Springer, 2012. doi:10.1007/978-3-642-40447-4\_2.
- 46 Tokio. Tokio Homepage. <https://tokio.rs/>, 2021.
- 47 Bernardo Toninho. *A Logical Foundation for Session-based Concurrent Computation*. PhD thesis, Carnegie Mellon University and New University of Lisbon, 2015.
- 48 Bernardo Toninho, Luís Caires, and Frank Pfenning. Higher-order processes, functions, and sessions: a monadic integration. In *22nd European Symposium on Programming (ESOP)*, pages 350–369. Springer, 2013. doi:[https://doi.org/10.1007/978-3-642-37036-6\\_20](https://doi.org/10.1007/978-3-642-37036-6_20).
- 49 Philip Wadler. Propositions as sessions. In *17th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 273–286. ACM, 2012.
- 50 WebKit. MotionMark Homepage. <https://browserbench.org/MotionMark/>, 2021.
- 51 Jeremy Yallop and Leo White. Lightweight higher-kinded polymorphism. In *Functional and Logic Programming - 12th International Symposium, FLOPS 2014, Kanazawa, Japan, June 4-6, 2014. Proceedings*, pages 119–135, 2014. doi:10.1007/978-3-319-07151-0\_8.

- 52 Joshua Yanovski, Hoang-Hai Dang, Ralf Jung, and Derek Dreyer. Ghostcell: separating permissions from data in rust. *Proc. ACM Program. Lang.*, 5(ICFP):1–30, 2021. doi:10.1145/3473597.
- 53 Nobuko Yoshida, Raymond Hu, Romyana Neykova, and Nicholas Ng. The scribble protocol language. In Martín Abadi and Alberto Lluch-Lafuente, editors, *Trustworthy Global Computing - 8th International Symposium, TGC 2013, Buenos Aires, Argentina, August 30-31, 2013, Revised Selected Papers*, volume 8358 of *Lecture Notes in Computer Science*, pages 22–41. Springer, 2013. doi:10.1007/978-3-319-05119-2\_3.

■ **Table 5** Overview of session terms in SILL<sub>R</sub> and Ferrite.

SILL <sub>R</sub>	Ferrite	Term	Description
$\epsilon$	End	terminate; wait $a$ ; $K$	Terminate session. Wait for channel $a$ to close.
$\tau \triangleright A$	ReceiveValue<T, A>	$x \leftarrow \text{receive\_value}; K$	Receive value $x$ of type $\tau$ .
$\tau \triangleleft A$	SendValue<T, A>	send_value_to $a$ $x$ ; $K$	Send value $x$ of type $\tau$ to $a$ .
		send_value $x$ ; $K$	Send value of type $\tau$ .
		$x \leftarrow \text{receive\_value\_from } a$ $x$ ; $K$	Receive value of type $\tau$ from channel $a$ .
$A \multimap B$	ReceiveChannel<A, B>	$a \leftarrow \text{receive\_channel}; K$	Receive channel $a$ of session type $A$ .
		send_channel_to $f$ $a$ ; $K$	Send channel $a$ to channel $f$ of session type $A \multimap B$ .
$A \otimes B$	SendChannel<A, B>	send_channel_from $a$ ; $K$	Send channel $a$ of session type $A$ .
		$a \leftarrow \text{receive\_channel\_from } f$ $a$ ; $K$	Receive channel $a$ from channel $f$ of session type $A \otimes B$ .
$\uparrow_l^s A$	LinearToShared<A>	accept_shared_session; $K_l$	Accept an acquire, then continue as linear session $K_l$ .
		$a \leftarrow \text{acquire\_shared\_session } s$ ; $K_l$	Acquire shared channel $s$ as linear channel $a$ .
$\downarrow_l^s S$	SharedToLinear<S>	detach_shared_session; $K_s$	Detach linear session and continue as shared session $K_s$ .
$A \& B$	ExternalChoice<Either<A, B>>	release_shared_session $a$ ; $K_l$	Release acquired linear session.
		offer_choice_2 $K_l$ $K_r$	Offer either continuation $K_l$ or $K_r$ based on client's choice.
		choose_left $a$ ; $K$	Choose the left branch offered by channel $a$
		choose_right $a$ ; $K$	Choose the right branch offered by channel $a$
$A \oplus B$	InternalChoice<Either<A, B>>	offer_left; $K$	Offer the left branch
		offer_right; $K$	Offer the right branch
		case_2 $a$ $K_l$ $K_r$	Branch to either $K_l$ or $K_r$ based on choice offered by channel $a$ .
$\&\{\overline{l_i : A_i}\}$	ExternalChoice<Row>	offer_choice $\{\overline{l_i : K_i}\}$	Offer continuation $K_i$ when the client selects $l_i$ .
		choose $a$ $l_i$ ; $K$	Choose the $l_i$ branch offered by channel $a$
$\oplus\{\overline{l_i : A_i}\}$	InternalChoice<Row>	offer $l_i$ ; $K$	Offer the $l_i$ branch
		case $a$ $\{\overline{l_i : K_i}\}$	Branch to continuation $K_i$ when channel $a$ offers $l_i$ .
-	Rec<F>	fix_session(cont)	Fold session type $F::\text{Applied}$ offered by cont.
		unfix_session( $a$ , cont)	Unfold channel $a$ to session type $F::\text{Applied}$ in cont.

## A Typing Rules

### A.1 Typing Rules for SILL<sub>R</sub>

Following is a list of inference rules in SILL<sub>R</sub>.

#### Communication

$$\frac{\Gamma; \Delta_1 \vdash a :: A \quad \Gamma; \Delta_2, a' : A \vdash b :: B}{\Gamma; \Delta_1, \Delta_2 \vdash a' \leftarrow \text{cut } a; b :: B} \text{ (T-CUT)} \quad \frac{\Gamma; \cdot \vdash a :: A \quad \Gamma; \Delta, a' : A \vdash b :: B}{\Gamma; \Delta \vdash a' \leftarrow \text{include } a; b :: B} \text{ (T-INCL)}$$

$$\frac{\Gamma; \cdot \vdash f :: A \multimap B \quad \Gamma; \cdot \vdash a :: A}{\Gamma; \cdot \vdash \text{apply\_channel } f a :: B} \text{(T-APP)} \quad \frac{}{\Gamma; a : A \vdash \text{forward } a :: A} \text{(T-FWD)}$$

**Termination**

$$\frac{}{\Gamma; \cdot \vdash \text{terminate}; :: \epsilon} \text{(T1R)} \quad \frac{\Gamma; \Delta \vdash K :: A}{\Gamma; \Delta, a : \epsilon \vdash \text{wait } a; K :: A} \text{(T1L)}$$

**Receive Value**

$$\frac{\Gamma, x : \tau; \Delta \vdash K :: A}{\Gamma; \Delta \vdash x \leftarrow \text{receive\_value}; K :: \tau \triangleright A} \text{(T}\triangleright\text{R)} \quad \frac{\Gamma; \Delta, a : A \vdash K :: B}{\Gamma, x : \tau; \Delta, a : \tau \triangleright A \vdash \text{send\_value\_to } a x; K :: B} \text{(T}\triangleright\text{L)}$$

**Send Value**

$$\frac{\Gamma; \Delta \vdash K :: A}{\Gamma, x : \tau; \Delta \vdash \text{send\_value } x; K :: \tau \triangleleft A} \text{(T}\triangleleft\text{R)} \quad \frac{\Gamma, a : \tau; \Delta, a : A \vdash K :: B}{\Gamma; \Delta, a : \tau \triangleright A \vdash x \leftarrow \text{receive\_value\_from } a; K :: B} \text{(T}\triangleleft\text{L)}$$

**Receive Channel**

$$\frac{\Gamma; \Delta, a : A \vdash K :: B}{\Gamma; \Delta \vdash a \leftarrow \text{receive\_channel}; K :: A \multimap B} \text{(T}\multimap\text{R)} \quad \frac{\Gamma; \Delta, f : A_2 \vdash K :: B}{\Gamma; \Delta, f : A_1 \multimap A_2, a : A_1 \vdash \text{send\_channel\_to } f a; K :: B} \text{(T}\multimap\text{L)}$$

**Send Channel**

$$\frac{\Gamma; \Delta \vdash K :: B}{\Gamma; \Delta, a : A \vdash \text{send\_channel\_from } a; K :: A \otimes B} \text{(T}\otimes\text{R)} \quad \frac{\Gamma; \Delta, f : A_2, a : A_1 \vdash K :: B}{\Gamma; \Delta, f : A_1 \otimes A_2 \vdash a \leftarrow \text{receive\_channel\_from } f; K :: B} \text{(T}\otimes\text{L)}$$

**Shared Session Types**

$$\frac{\Gamma; \cdot \vdash K :: A}{\Gamma; \cdot \vdash \text{accept\_shared\_session}; K :: \uparrow_l^s A} \text{(T}\uparrow_l^s\text{R)} \quad \frac{\Gamma; \Delta, a : A \vdash K :: B}{\Gamma, s : \uparrow_l^s A; \Delta \vdash a \leftarrow \text{acquire\_shared\_session } s; K :: B} \text{(T}\uparrow_l^s\text{L)}$$

$$\frac{\Gamma; \cdot \vdash K :: S}{\Gamma; \cdot \vdash \text{detach\_shared\_session}; K :: \downarrow_l^s S} \text{(T}\downarrow_l^s\text{R)} \quad \frac{\Gamma, s : S; \Delta \vdash K :: B}{\Gamma; \Delta, a : \downarrow_l^s S \vdash s \leftarrow \text{release\_shared\_session } a; K :: B} \text{(T}\downarrow_l^s\text{L)}$$

**External Choice (Binary)**

$$\frac{\Gamma; \Delta \vdash K_l :: A \quad \Gamma; \Delta \vdash K_r :: B}{\Gamma; \Delta \vdash \text{offer\_choice } K_l K_r :: A \& B} \text{(T}\&2\text{R)} \quad \frac{\Gamma; \Delta, a : A_1 \vdash K :: B}{\Gamma; \Delta, a : A_1 \& A_2 \vdash \text{choose\_left } a; K :: B} \text{(T}\&2\text{L)}$$

$$\frac{\Gamma; \Delta, a : A_2 \vdash K :: B}{\Gamma; \Delta, a : A_2 \& A_2 \vdash \text{choose\_right } a; K :: B} \text{(T}\&2\text{L)}$$

**Internal Choice (Binary)**

$$\frac{\Gamma; \Delta \vdash K :: A}{\Gamma; \Delta \vdash \text{offer\_left}; K :: A \oplus 2B} \text{(T}\oplus2\text{R)} \quad \frac{\Gamma; \Delta, a : A_1 \vdash K_l :: B \quad \Gamma; \Delta, a : A_2 \vdash K_r :: B}{\Gamma; \Delta, a : A_1 \oplus A_2 \vdash \text{case } a K_l K_r :: B} \text{(T}\oplus2\text{L)}$$

$$\frac{\Gamma; \Delta \vdash K :: B}{\Gamma; \Delta \vdash \text{offer\_right}; K :: A \oplus B} \text{(T}\oplus2\text{R)}$$

**External Choice**

$$\frac{\Gamma; \Delta \vdash \overline{K_i} :: \overline{A_i}}{\Gamma; \Delta \vdash \text{offer\_choice } \{\overline{l_i} : \overline{K_i}\} :: \& \{\overline{l_i} : \overline{A_i}\}} \text{(T}\&\text{R)} \quad \frac{\Gamma; \Delta, a : A_i \vdash K :: B}{\Gamma; \Delta, a : \& \{\overline{l_i} : \overline{A_i}\} \vdash \text{choose } a l_i; K :: B} \text{(T}\&\text{L)}$$

**Internal Choice**

$$\frac{\Gamma; \Delta \vdash K :: A}{\Gamma; \Delta \vdash \text{offer } l_i; K :: \oplus \{\overline{l_i} : \overline{A_i}\}} \text{(T}\oplus\text{R)} \quad \frac{\Gamma; \Delta, a : A_i \vdash K_i :: B}{\Gamma; \Delta, a : \oplus \{\overline{l_i} : \overline{A_i}\} \vdash \text{case } a \{\overline{l_i} : \overline{K_i}\} :: B} \text{(T}\oplus2\text{L)}$$

## A.2 Typing Constructs in Ferrite

Following is a list of function signatures of the term constructors provided in Ferrite.

### A.2.1 Forward

```
fn forward<N, C, A>(_ : N) -> PartialSession<C, A>
where
  A : Protocol,
  C : Context,
  N :: Target: EmptyContext,
  N : ContextLens<C, A, Empty>
```

### A.2.2 Termination

```
pub struct End;
impl Protocol for End { ... }

fn terminate<C>() -> PartialSession<C, End>
where
  C : EmptyContext

fn wait<N, C, A>(
  _ : N,
  cont : PartialSession<N::Target, A>
) -> PartialSession<C, A>
where
  C : Context,
  A : Protocol,
  N : ContextLens<C, End, Empty>
```

### A.2.3 Communication

```
fn cut<X, C, C1, C2, A, B>(
  cont1 : PartialSession<C1, A>,
  cont2 : impl FnOnce(C2::Length) -> PartialSession<C2::Appended, B>
) -> PartialSession<C, B>
where
  A : Protocol,
  B : Protocol,
  C : Context,
  C1 : Context,
  C2 : Context,
  X : SplitContext<C, Left = C1, Right = C2>,
  C2 : AppendContext<(A, ())>

fn include_session<C, A, B>(
  session : Session<A>,
  cont : impl FnOnce(C::Length) -> PartialSession<C::Appended, B>
) -> PartialSession<C, B>
where
  A : Protocol,
  B : Protocol,
  C : Context,
  C : AppendContext<(A, ())>

fn apply_channel<A, B>(
  f : Session<ReceiveChannel<A, B>>,
  a : Session<A>,
```

```

) -> Session<B>
where
  A : Protocol,
  B : Protocol

```

## A.2.4 Receive Value

```

struct ReceiveValue<T, A> { ... }
impl<T, A> Protocol for ReceiveValue<T, A>
where
  T: Send + 'static,
  A: Protocol
{ ... }

fn receive_value<T, C, A>(
  cont : impl FnOnce(T) -> PartialSession<C, A> + Send + 'static
) -> PartialSession<C, ReceiveValue<T, A>>
where
  T : Send + 'static,
  A : Protocol,
  C : Context

fn send_value_to<N, C, A, B, T>(
  _ : N,
  val : T,
  cont : PartialSession<N::Target, A>
) -> PartialSession<C, A>
where
  A : Protocol,
  B : Protocol,
  C : Context,
  T : Send + 'static,
  N : ContextLens<C, ReceiveValue<T, B>, B>

```

## A.2.5 Send Value

```

struct SendValue<T, A> { ... }
impl<T, A> Protocol for SendValue<T, A>
where
  T: Send + 'static,
  A: Protocol
{ ... }

fn send_value<T, C, A>(
  val : T,
  cont : PartialSession<C, A>
) -> PartialSession<C, SendValue<T, A>>
where
  T : Send + 'static,
  A : Protocol,
  C : Context

fn receive_value_from<N, C, T, A, B>(
  _ : N,
  cont : impl FnOnce(T) -> PartialSession<N::Target, B> + Send + 'static
) -> PartialSession<C, B>
where
  A : Protocol,
  B : Protocol,
  C : Context,
  T : Send + 'static,
  N : ContextLens<C, SendValue<T, A>, A>

```

## A.2.6 Receive Channel

```

pub struct ReceiveChannel<A, B> { ... }
impl<A: Protocol, B: Protocol> Protocol for ReceiveChannel<A, B> { ... }

fn receive_channel<C, A, B>(
  cont : impl FnOnce(C::Length) -> PartialSession<C::Appended, B>
) -> PartialSession<C, ReceiveChannel<A, B>>
where
  A : Protocol,
  B : Protocol,
  C : Context,
  C : AppendContext<(A, ())>

fn send_channel_to<N1, N2, C, A1, A2, B>(
  _ : N1,
  _ : N2,
  cont : PartialSession<N1::Target, B>
) -> PartialSession<C, B>
where
  C : Context,
  A1 : Protocol,
  A2 : Protocol,
  B : Protocol,
  N2 : ContextLens<C, A1, Empty>,
  N1 : ContextLens<N2::Target, ReceiveChannel<A1, A2>, A2>

```

## A.2.7 Send Channel

```

struct SendChannel<A, B> { ... }
impl<A: Protocol, B: Protocol> Protocol for SendChannel<A, B>

fn send_channel_from<C, A, B, N>(
  _ : N,
  cont : PartialSession<N::Target, B>
) -> PartialSession<C, SendChannel<A, B>>
where
  A : Protocol,
  B : Protocol,
  C : Context,
  N : ContextLens<C, A, Empty>

fn receive_channel_from<C1, C2, A1, A2, B, N>(
  _ : N,
  cont_builder : impl FnOnce(C2::Length) -> PartialSession<C2::Appended, B>
) -> PartialSession<C1, B>
where
  A1 : Protocol,
  A2 : Protocol,
  B : Protocol,
  C1 : Context,
  C2 : AppendContext<(A1, ())>,
  N : ContextLens<C1, SendChannel<A1, A2>, A2, Target = C2>

```

## A.2.8 External Choice

```

struct ExternalChoice<Row> { ... }
impl<Row> Protocol for ExternalChoice<Row>
where
  Row: ToRow + Send + 'static
{ ... }

```



```

fn offer_choice<C, Row1, Row2>(
  cont1: impl for<'r> FnOnce(
    AppSum<'r, Row2, InjectSessionF<'r, Row1, C>>,
    ) -> ContSum<'r, Row1, C>
    + Send
    + 'static
) -> PartialSession<C, ExternalChoice<Row1>>
where
  C: Context,
  Row1: Send + 'static,
  Row2: Send + 'static,
  Row1: ToRow<Row = Row2>,
  Row2: RowCon,
  Row2: SumFunctor

fn choose<N, M, C1, C2, A, B, Row1, Row2>(
  _: N,
  _: M,
  cont: PartialSession<C2, A>
) -> PartialSession<C1, A>
where
  C1: Context,
  C2: Context,
  A: Protocol,
  B: Protocol,
  Row2: RowCon,
  Row1: Send + 'static,
  Row2: Send + 'static,
  Row1: ToRow<Row = Row2>,
  N: ContextLens<C1, ExternalChoice<Row1>, B, Target = C2>,
  M: Prism<Row2, Elem = B>

```

## A.2.9 Internal Choice

```

struct InternalChoice<Row> { ... }
impl<Row> Protocol for InternalChoice<Row>
where
  Row: ToRow + Send + 'static
{ ... }

fn offer_case<N, C, A, Row1, Row2>(
  _: N,
  cont: PartialSession<C, A>
) -> PartialSession<C, InternalChoice<Row1>>
where
  C: Context,
  A: Protocol,
  Row1: Send + 'static,
  Row2: Send + 'static,
  Row2: RowCon,
  Row1: ToRow<Row = Row2>,
  N: Prism<Row2, Elem = A>

fn case<N, C1, C2, B, Row1, Row2>(
  _: N,
  cont1: impl for<'r> FnOnce(
    AppSum<'r, Row2, ContF<'r, N, C2, B>>,
    ) -> ChoiceRet<'r, N, C2, B>
    + Send
    + 'static
) -> PartialSession<C1, B>
where
  B: Protocol,
  C1: Context,

```

```

C2: Context,
Row1: Send + 'static,
Row2: Send + 'static,
Row1: ToRow<Row = Row2>,
N: ContextLens<C1, InternalChoice<Row1>, Empty, Target = C2>

```

### A.2.10 Recursive Session Types

```

fn fix_session<R, F, A, C>(
  cont: PartialSession<C, A>
) -> PartialSession<C, RecX<R, F>>
where
  C: Context,
  R: Context,
  F: Protocol,
  A: Protocol,
  F: RecApp<(RecX<R, F>, R), Applied = A>

fn unfix_session<N, C1, C2, A, B, R, F>(
  _n: N,
  cont: PartialSession<C2, B>
) -> PartialSession<C1, B>
where
  B: Protocol,
  C1: Context,
  C2: Context,
  F: Protocol,
  R: Context,
  F: RecApp<(RecX<R, F>, R), Applied = A>,
  A: Protocol,
  N: ContextLens<C1, RecX<R, F>, A, Target = C2>

```

### A.2.11 Shared Session Types

```

struct LinearToShared<F> { ... }
struct SharedToLinear<F> { ... }
struct Lock<F> { ... }

impl<F> SharedProtocol for LinearToShared<F>
where
  F: Protocol,
  F: SharedRecApp<SharedToLinear<LinearToShared<F>>>,
  F::Applied: Protocol
{ ... }

impl<F> Protocol for SharedToLinear<LinearToShared<F>>
where
  F: SharedRecApp<SharedToLinear<LinearToShared<F>>> + Send + 'static
{ ... }

impl<F> Protocol for Lock<F>
where
  F: Protocol,
  F: SharedRecApp<SharedToLinear<LinearToShared<F>>>,
  F::Applied: Protocol
{ ... }

```

A detail we omitted in the main text is that we introduced a special linear session type called `Lock`, internal to the library. The `Lock` type holds the underlying shared Rust channel that connects to the corresponding endpoint held by `SharedChannel`. This allows multiple uses

of `accept_shared_session` and `detach_shared_session` to all access the same underlying Rust channel without having to rely on global state.

An additional role of the linear session type `Lock` is that it also enforces the equi-synchronizing constraint of shared session type, by requiring all use of `accept_shared_session` to always be accompanied by `detach_shared_session` with the same shared session type. This provides the same functionality of enforcing the equi-synchronizing constraint as specified in Section 3.3 in [1].

```
fn accept_shared_session<F>(
  cont: impl Future<Output = PartialSession<(Lock<F>, ()), F::Applied>>
  + Send
  + 'static
) -> SharedSession<LinearToShared<F>>
where
  F: Protocol,
  F: SharedRecApp<SharedToLinear<LinearToShared<F>>>,
  F::Applied: Protocol
```

The `accept_shared_session` construct is parameterized over a shared session type `LinearToShared<F>`. The type `F` is required to implement `SharedRecApp<SharedToLinear<LinearToShared<F>>>`, which unfolds the shared session type by applying the type `SharedToLinear<LinearToShared<F>>` to `F`. The continuation is an `async` block with `PartialSession` result that offers the linear session type `F::Applied`. It also has an internal session type `Lock<F>`, which is described next. The construct returns a shared session type program of type `SharedSession<LinearToShared<F>>`. This needs to be passed to `run_shared_session` to execute the program and get back a shared channel of type `SharedChannel<LinearToShared<F>>`.

```
fn detach_shared_session<F, C>(
  cont: SharedSession<LinearToShared<F>>
) -> PartialSession<(Lock<F>, C), SharedToLinear<LinearToShared<F>>>
where
  F: Protocol,
  F: SharedRecApp<SharedToLinear<LinearToShared<F>>>,
  F::Applied: Protocol,
  C: EmptyContext
```

The `detach_shared_session` construct is parameterized by a linear session type `LinearToShared<F>` and an empty linear context `C`. The type `F` is required to implement `SharedRecApp<SharedToLinear<LinearToShared<F>>>` to unfold `F` recursively. This is required for `LinearToShared<F>` to satisfy the `SharedProtocol` constraint. The construct accepts a `SharedSession` continuation with the offered shared session type `LinearToShared<F>`. Note that this is the only continuation that is *not* a `PartialSession`. It is also *not* a `SharedChannel`, as this is a shared Ferrite program that is yet to be executed. The construct returns a `PartialSession` that offers the linear session type `SharedToLinear<F>`. It also has a linear context with `Lock<F>` being the first linear channel, and the tail `C` being an empty linear context of arbitrary length.

```
fn acquire_shared_session<C, F, A>(
  shared: SharedChannel<LinearToShared<F>>,
  cont1: impl FnOnce(C::Length) -> PartialSession<C::Appended, A> + Send + 'static
) -> PartialSession<C, A>
where
  C: Context,
  F: Protocol,
  A: Protocol,
  F::Applied: Protocol,
  F: SharedRecApp<SharedToLinear<LinearToShared<F>>>,
  C: AppendContext<F::Applied, ()>
```

The `acquire_shared_session` construct is parameterized over a shared session type `LinearToShared<F>`, a linear context `C`, and an offered session type `A`. The type `F` is required

## 22:40 Ferrite: A Judgmental Embedding of Session Types in Rust

to implement `SharedRecApp<SharedToLinear<LinearToShared<F>>>`, which unfolds the shared session type by applying the type `SharedToLinear<LinearToShared<F>>` to `F`. The unfolded session type `F::Applied` is a linear session type implementing `Protocol`, and it is appended to the end of `C` using `AppendContext`, with `C::Appended` being the result.

The first argument to `acquire_shared_session` is a cloneable `SharedChannel` of (shared) session type `LinearToShared<F>`. The second argument is the continuation closure. It is given the context lens `C::Length`, which implements the context lens to access the linear channel `F::Applied` in `C::Appended`. The continuation closure returns a `PartialSession` with `C::Appended` being the linear context, and `A` being the offered session type. The construct returns a `PartialSession` that works with the original linear context `C`, and offers the session type `A`.

```
fn release_shared_session<N, C1, C2, A, B>(
  _n: N,
  cont: PartialSession<C2, B>,
) -> PartialSession<C1, B>
where
  A: Protocol,
  B: Protocol,
  C1: Context,
  C2: Context,
  A: SharedRecApp<SharedToLinear<LinearToShared<A>>>,
  N: ContextLens<C1, SharedToLinear<LinearToShared<A>>>, Empty, Target = C2>
```

The `release_shared_session` construct is parameterized over a linear session type `SharedToLinear<A>`, a linear context `C`, a context lens `N` for accessing `SharedToLinear<LinearToShared<A>>` from `C`, and an offered session type `B`. The continuation is a `PartialSession` with `N::Target` being the linear context `C` with `SharedToLinear<LinearToShared<A>>` removed, and offers the session type `B`. The construct returns a `PartialSession` with the original linear context `C`, and offers the session type `B`.

## B Dynamics

Section 4 introduced the type system of Ferrite, based on the constructs `End`, `ReceiveValue`, and `ReceiveChannel`. This section revisits those constructs and fills in the missing implementations to make the constructs executable, amounting to the *dynamic semantics* of Ferrite.

### B.1 One-shot Channels

Internally, Ferrite uses `tokio`'s `oneshot` [46] channels as the primitive building block for session-typed channels. A one-shot channel with a payload type `P` consists of a pair of sender and receiver, of type `Sender<P>` and `Receiver<P>`, resp., denoting the two endpoints of the channel. The channel is one-shot in the sense that at most one value of type `P` can be sent across the channel. However since the one-shot channel is affine, it is also possible to have no value being sent over the channel.

The one-shot channel can be used directly by Rust programmers to emulate simple session types. As an example, the session type `ReceiveValue<i32, End>` can be implemented using one-shot channels as follows:

```
use tokio::{task, try_join};
use tokio::sync::oneshot::{channel, Sender, Receiver};

async fn receive_int_provider(value_receiver: Receiver<i32, Sender<()>>) {
  let (value, end_sender) = value_receiver.await.unwrap()
  println!("provider received value: {}", value);
}
```

```

    end_sender.send().unwrap();
}
async fn receive_int_client(value_sender: Sender<i32, Receiver<>>>) {
    let (end_sender, end_receiver) = channel();
    value_sender.send((42, end_sender));
    end_receiver.await.unwrap();
}
async fn main() {
    let (value_sender, value_receiver) = channel();
    let child1 = spawn(async move {
        receive_int_provider(value_receiver).await;
    });
    let child2 = spawn(async move {
        receive_int_client(value_sender).await;
    });
}

```

The code above defines the `receive_int_provider` and `receive_int_client` functions to execute the provider and client processes corresponding to the session type `ReceiveValue<i32, End>`, resp. On the provider side, it needs to first receive an `i32` value and then send back an end signal to the client when it is terminating. This corresponds to the one-shot channel type `Receiver<i32, Sender<>>>`, with the `Sender<>` used to send a unit `()` as termination signal. On the receiver side, the polarity of the one-shot channel is switched and become `Sender<i32, Receiver<>>>`. This indicates that the client first sends an `i32` value, together with a `Receiver<>` for the provider to send back the termination signal.

## B.2 Protocol Definitions

The above example demonstrates that even a simple session type like `ReceiveValue<i32, End>` requires non-trivial effort to be implemented manually using one-shot channels. To automate this in Ferrite, we need to derive the one-shot channel types `Receiver<i32, Sender<>>>` and `Sender<i32, Receiver<>>>` from the session type `ReceiveValue<i32, End>`. This is achieved by defining some associated types and methods in the `Protocol` trait:

```

trait Protocol {
    type ProviderEndpoint;
    type ClientEndpoint;

    fn create_endpoints() -> (Self::ProviderEndpoint, Self::ClientEndpoint);
}

```

The associated types `ProviderEndpoint` and `ClientEndpoint` are used to define the one-shot channel types for the provider end and consumer end, resp. The trait method `create_endpoints` is used to create a channel pair which connects both the provider and client endpoints. Following the previous example, the implementation should derive the type `<ReceiveValue<i32, End>>::ProviderEndpoint` to be `Receiver<i32, Sender<>>>`, and `<ReceiveValue<i32, End>>::ClientEndpoint` to be `Sender<i32, Receiver<>>>`. This is implemented by first implementing `Protocol` for `End`:

```

impl Protocol for End
{
    type ProviderEndpoint = Sender<>>;
    type ClientEndpoint = Receiver<>>;

    fn create_endpoints() -> (Self::ProviderEndpoint, Self::ClientEndpoint)
    {
        {
            channel()
        }
    }
}

```

## 22:42 Ferrite: A Judgmental Embedding of Session Types in Rust

In the implementation of the `End` protocol, the provider end is the party that needs to send the termination signal `()` to the client end. Hence its `ProviderEndpoint` type is `Sender<()>`, and vice versa for the client end. The implementation of the `create_endpoints` method is to simply call `channel()` to create the one-shot channel pair.

To implement `Protocol` for a session type `ReceiveValue<T, A>`, we would need to make use of the `Protocol` implementation for the continuation session type `A`:

```
impl<T, A> Protocol for ReceiveValue<T, A>
{
    type ProviderEndpoint = Receiver<(T, A::ProviderEndpoint)>;
    type ClientEndpoint = Sender<(T, A::ProviderEndpoint)>;

    fn create_endpoints() -> (Self::ProviderEndpoint, Self::ClientEndpoint)
    {
        let (sender, receiver) = channel();
        (receiver, sender)
    }
}
```

The provider end is given a receiver for the value `T`, together with its continuation endpoint for `A`. Given that the continuation for the provider also needs the provider endpoint, and it has to be extracted from the receiver, the provider would need to receive `A::ProviderEndpoint` alongside with the value `T`. Hence the associated type `<ReceiveValue<T, A>>::ProviderEndpoint` becomes `Receiver<(T, A::ProviderEndpoint)>`. On the client side, the value `T` needs to be sent alongside with `A::ProviderEndpoint`, hence the associated type `<ReceiveValue<T, A>>::ClientEndpoint` is `Sender<(T, A::ProviderEndpoint)>`. Notice that both the `ProviderEndpoint` and `ClientEndpoint` associated types for `ReceiveValue<T, A>` contains `A::ProviderEndpoint`, but not `A::ClientEndpoint`.

In the implementation of `create_endpoints` for `ReceiveValue`, the ordering of the sender and receiver pair returned from calling `channel()` is flipped. This is because `create_endpoints` always return the provider endpoint first followed by the client endpoint. And since the provider endpoint is a receiver in this case, it needs to be returned in the first position.

With the `Protocol` definitions of both `End` and `ReceiveValue`, we can follow that the associated types and channel creation for `ReceiveValue<i32, End>` matches the channel types and behavior of the example at the beginning of this section.

### B.3 Linear Context

The linear context of a Ferrite program comprises the client endpoints for the session types. Conceptually, Ferrite needs to derive from a session type list `HList![A0, A1, ...]` into a list of client endpoint list `HList![A0::ClientEndpoint, A1::ClientEndpoint, ...]`. However a linear context may also contain the special `Empty` element, which do not implement `Protocol`. To allow the transformation of the linear context, we need to first add an associated type to the `Slot` trait as follows:

```
trait Slot {
    type Endpoint: Send;
}
impl<A: Protocol> Slot for A {
    type Endpoint = A::ClientEndpoint;
}
impl Slot for Empty {
    type Endpoint = ();
}
```

We define the associated type `Endpoint` in `Slot` such that if a type `A` implements `Protocol`, then `A::Endpoint` is simply `A::ClientEndpoint`. We also define the special case for `Empty`, which

the `Endpoint` associated type is `()` to represent the absence of a client endpoint. With that, we can extend the `Context` trait to include the `Endpoints` associated type:

```
trait Context {
  type Endpoints;
}
impl Context for () {
  type Endpoints = ();
}
impl<A: Slot, C: Context> Context for (A, C) {
  type Endpoints = (A::Endpoint, C::Endpoints);
}
```

For the base case of an empty list `()` (`HList![]`), the result `Endpoints` is also an empty list. For the inductive case, if the tail `C` of a linear context `(A, C)` implements `Context`, and the head `A` implements `Slot`, then the associated type `(A, C)::Endpoints` is `(A::Endpoint, C::Endpoints)`.

## B.4 Session Dynamics

Ferrite generates session type programs by composing `PartialSession` objects generated by constructs such as `receive_value`. To enable execution of the Ferrite program, the `PartialSession` struct contains an internal `executor` field that is defined as follows:

```
struct PartialSession<C: Context, A: Protocol> {
  executor: Box<
    dyn FnOnce(
      C::Endpoints,
      A::ProviderEndpoint,
    ) -> Pin<Box<dyn Future<Output = ()> + Send>>
    + Send,
  >
}
```

The `executor` field contains an `FnOnce` closure that accepts two arguments – the endpoints for the linear context `C::Endpoints`, and the provider endpoint for the offered session type `A::ProviderEndpoint`. When called, the closure executes asynchronously by returning a *future* with the type `Pin<Box<dyn Future<Output = ()> + Send>>`. The boilerplate signature is required, as Rust has not stabilized the syntactic sugar for `async` closures. Conceptually, the closure signature is equivalent to the `async fn` signature `async fn(C::Endpoints, A::ProviderEndpoint)`.

Ferrite keeps the `executor` field private within the library to prevent end users from constructing new `PartialSession` values or running the `executor` closure. This is because the creation and execution of `PartialSession` may be unsafe. We demonstrate two simple examples of unsafe (i.e. non-linear) usage of `PartialSession`.

Below shows an example Ferrite program `p1` of type `Session<SendValue<String, End>>` is constructed, but in the `executor` closure both the client endpoints and the provider endpoint are ignored. As a result, `p1` violates the linearity constraint of session types and never sends any string value or signal for termination.

```
let p1: Session<SendValue<String, End>>
  = PartialSession { executor: Box::pin(async |_ctx, _provider_end| { }) };
```

Below shows an example client, which calls a Ferrite program `p2` of type `ReceiveValue<String, End>` by directly running its `executor`. The client creates an endpoint pair but drops the client endpoint. It then executes `p2` with the provider endpoint. However because the client endpoint is dropped, `p2` fails to receive any value, and the program results in a deadlock.

```
let p2: Session<ReceiveValue<String, End>> = ...;
let (provider_end, _client_end) = <ReceiveValue<String, End>>::create_endpoints();
(p2.executor)((), provider_end).await;
```



From the examples above we can see that direct access to the `executor` field is unsafe. The `PartialSession` is used with care within Ferrite to ensure that linearity is enforced in the implementation. Externally, the `run_session` function is provided for executing Ferrite programs of type `Session<End>`, as only such programs can be executed safely without additional safe guard.

## C Rust as a Host Language

In this section, we address some common questions arise from the choice of using Rust as a host language for Ferrite.

### C.1 Benefits of Affine Type System

The affine type system in Rust helps Ferrite to better verify the correctness of its underlying implementation. Internally, Ferrite uses one-shot Rust channels to implement the communication. The affine property in Rust helps us guarantee that our underlying implementation cannot accidentally send two payloads through the one-shot channels.

Ferrite user programs also benefit from the affine type system in Rust. Ferrite constructs accept continuation closures with the `FnOnce` trait bound, to guarantee that the continuation cannot be called more than once. As a result, Rust values can be moved inside the continuation closures and work more efficiently without requiring copies to be made. Similarly, the `send/receive` value constructs works with the affine type system in Rust, so values such as byte arrays can be sent efficiently in Ferrite without requiring copying.

In comparison, while previous works in Haskell and OCaml are able to enforce the linear usage in session type programs, the structural semantics of these languages may impose challenge on the compiler from being able to optimize the use of linear resources inside the program. In particular, the indexed monad that encapsulates the session type program is itself copyable. As a result, continuations cannot guarantee that the variables they capture cannot be used more than once.

### C.2 Support for Lifetime

At the moment, Ferrite requires the continuations to have `'static` lifetime. This is due to the underlying async implementations requiring spawned async tasks to have `'static` lifetime. We plan to overcome this limitation in the future by finding ways to spawn async tasks with a scoped lifetime. Once that limitation is overcome, it will also be possible to access mutable references inside scoped Ferrite programs.

### C.3 Type Errors

Type error messages in Ferrite are expressed in terms of the structs and traits of Ferrite. As a result it is not difficult for users to read and understand the error messages, provided they are familiar with the basic terminology used by Ferrite.

Consider the example `hello_client` from section 4

```
let hello_client: Session<
  ReceiveChannel<ReceiveValue<String, End>, End>>
  = receive_channel(| a | {
    send_value_to(a, "Alice".to_string(),
      wait(a, terminate())
    ) });
```

If we were to forget to wait for channel `a` and terminate immediately, the following error is generated:

```
let hello_client: Session<
  ReceiveChannel<ReceiveValue<String, End>, End>>
= receive_channel(| a | {
  send_value_to(a, "Alice".to_string(),
    // the trait `EmptyContext` is not implemented for `(End, ())`
    terminate()
  } });
```

This indicates that the linear context `(End, ())` is not empty, and as a result the `terminate` construct cannot be used.

If we try to wait for `a` to terminate before sending a value to `a`, we get a different error:

```
let hello_client: Session<
  ReceiveChannel<ReceiveValue<String, End>, End>>
= receive_channel(| a | {
  // the trait `ContextLens<(ReceiveValue<String, End>, ()), End, Empty>`
  // is not implemented for `Z`
  wait(a, terminate())
});
```

The error message indicates an invalid use of a context lens to update a channel of the wrong session type in the linear context. Recall from section 3.2 that the constraint `Z: ContextLens<(ReceiveValue<String, End>, ()), End, Empty>` would require the first channel (`Z`) in the linear context `((ReceiveValue<String, End>, ()))` to be of session type `End`, but here the session type of the first channel in the linear context is `ReceiveValue<String, End>`.

Error messages such as the above are commonly generated by non-linear use of channels or a mismatch in session types. While they require some understanding of the concepts such as linear context and context lenses, the error messages are not too difficult to decipher.

## C.4 Hole Driven Development

Aside from designing readable error messages, we recommend a *hole-driven* approach of writing Ferrite programs to minimize the chance of users encountering complex type errors. In this approach, the user would implement a Ferrite program in small steps, with the continuation filled with `todo!()` as a placeholder. We demonstrate this by showing how a new user would implement the `hello_provider` program in section 4:

```
let hello_provider: Session<SendValue<String, End>> = todo!();
```

The `todo!()` macro allows us to put a placeholder in unfinished Rust code so that we can try and compile the code and see if there is any type error. By writing our code step by step and filling the blank with `todo!()`, we can narrow down the potential places where our code is incorrect. At this stage, we should be able to compile our program with no error. This shows that the protocol that we have defined, `SendValue<String, End>`, is a valid session type. If we have gotten a compile error otherwise, it could have been caused by us trying to write an invalid protocol like `SendValue<String, String>`.

We can try to compile our code again, and Rust will accept the code we have written. However the use of `todo!()` does not tell us how we should continue our program. In Rust, we could use the unit type `()` to deliberately cause a compile error:

```
let hello_provider: Session<SendValue<String, End>> =
  send_value("Hello World!".to_string(), ());
```

Now if we compile our code, we would get a compile error from Rust:

## 22:46 Ferrite: A Judgmental Embedding of Session Types in Rust

```
error[E0308]: mismatched types
  |
  |   send_value("Hello World!".to_string(), ());
  |                                             ^^ expected struct PartialSession, found ()
  |
= note: expected struct PartialSession<(), End>
       found unit type ()
```

With this compile error, we can know that we are supposed to fill in the hole with Rust expression that has the type `PartialSession<(), End>`. Sometimes we may also intuitively think of a type that should be in a hole. In such case, we can also use the `todo!() as T` pattern to verify if our intuition is correct. So we can for example write:

```
let hello_provider: Session<SendValue<String, End>> =
  send_value("Hello World!".to_string(), todo!() as Session<End>);
```

And our code will compile successfully. If we were to annotate it with an invalid type, such as `todo!() as Session<ReceiveValue<String, End>>` again, Rust will also return a compile error. Now that we know the continuation needs to have the type `Session<End>`, we can then fill in the blank with `terminate()` and complete our program.

### **D** Challenges in Using Ferrite on Servo

We report on some of the challenges that we faced when implementing the Servo canvas component in Ferrite in Section 7, and how the challenges are addressed.

#### D.1 Interprocess Communication

As a browser rendering engine, Servo puts much emphasis on security, using sandboxing to ensure that malicious web applications cannot easily compromise a user's computer. A main design outcome of this emphasis is that the provider and client are executed in separate OS processes. Regular Rust channels cannot be used for communication between different processes, because the underlying implementation requires a common address space. As a result, Servo uses the `ipc_channel` crate to create inter-process communication (IPC) channels for communication between the provider and client of its components. The IPC channels in Servo create a local file socket and serialize the Rust messages to send them over the socket as raw bytes. This requires the channel payload types to implement the `Serialize` and `Deserialize` traits for them to be usable in the IPC channels. IPC channels are themselves serializable, so it is possible to send an IPC channel over another IPC channel.

Since Ferrite internally makes use of `tokio` channels for communication, this presents challenges since they cannot be serialized and sent through Servo's IPC channels. For the purpose of the evaluation, we implemented our own serialization of `SharedChannel`. Our serialization involves creating a bidirectional pair of opaque (untyped) IPC channels, and forwards all communication from the regular Rust channels to the IPC channels. This approach works, albeit inefficiently, as there needs to be two background tasks in the provider and client processes to perform the actual serialization and forwarding. We benchmarked the performance of our implementation, revealing a decrease of about a factor of ten. We have not spent much effort on fine-tuning our serialization implementation because the primary purpose of this study is to show that the message protocols underlying Servo's canvas component can be made explicit and verified in Ferrite.

## D.2 Latency in Acquire-Release

Servo's canvas component has very high performance demands, requiring the sending of thousands of messages in a few milliseconds. In our initial implementation, we found the Ferrite implementation to be lacking in performance, despite not saturating the CPU usage. A closer inspection revealed that the bottleneck was in the latency caused by the acquire-release cycle introduced in the implementation of shared session types. In Ferrite, the client of a shared channel needs to first send an acquire to the shared provider and then wait for the acknowledgment before it can start communicating through the acquired linear channel. This round trip latency becomes significant if the communication frequency is high. Consider two canvas messages being sent right after each other. In the original design, the second message can be sent immediately after the first message has been sent. In the Ferrite implementation, on the other hand, the two messages are sent in two separate acquire-release cycles, interspersing additional acquire and release messages and possibly delays because of blocking acquires.

The latency is aggravated by the use of IPC channels. Since IPC channels are mapped to file sockets, efficient parallel communications must be multiplexed among a small number of channels. For the case of Ferrite shared channels, the multiplexing currently is done by queuing and forwarding the requests in serial, which can be inefficient. As a workaround, we batch messages on the client side, such that trivial messages like `LineTo` are stored in a local `Vec<CanvasMessage>` buffer before being sent to the provider in a new `Messages` branch in `CanvasOps`. The buffered messages are sent in batch every ten milliseconds, or when a non-trivial protocol such as `GetImageData` is called. With batching, we have gained enough performance to render complex canvases smoothly.