

Proof of Delivery: Mechanized Mailbox Types

Edgard Schiebelbein¹[0009-0002-9476-5727],
Annette Bieniusa¹[0000-0002-1654-6118], and Simon Fowler²[0000-0001-5143-5475]

¹ RPTU University Kaiserslautern-Landau, Kaiserslautern, Germany
{edgard.schiebelbein, annette.bieniusa}@cs.rptu.de

² University of Glasgow, Glasgow, UK
simon.fowler@glasgow.ac.uk

Abstract. Programming languages based on the actor model such as Erlang and Elixir rely on message passing for communication. A core concept of this communication model are mailboxes: all actors are allowed to send messages to a mailbox, but only its owner may retrieve them. A major challenge in such systems is detecting and eliminating protocol violations and deadlocks. Mailbox types are a novel type system that enables static reasoning about the contents of mailboxes, but the metatheory of mailbox type systems is complex and requires extensive reasoning about subtyping. This paper establishes a machine-checked foundation for mailbox types. As a basis for this, we formalize Pat, the first programming language to use mailbox types, in Rocq. With its help, we identified and corrected several oversights in the original definitions. Furthermore, we provide mechanized proofs of several properties of the semantics of mailbox types and the substitution lemma for Pat.

Keywords: Mailbox types · Behavioural type systems · Formal verification · Actor languages

1 Introduction

Programming concurrent and distributed systems is notoriously difficult. Communication-centric programming languages such as Erlang or Elixir, and frameworks like MPI, are based around *message passing* concurrency, where processes send and receive messages to coordinate computation, rather than relying on shared memory. These actor languages have seen much industrial interest due to their suitability for distribution and compatibility with failure recovery idioms such as supervision hierarchies [2]. However, these languages are also susceptible to issues like communication mismatches and deadlocks which pose significant problems at runtime, and are also difficult to detect and fix.

Mailbox types [6,9] are a novel behavioural type system with a focus on mailboxes, as used in actor languages like Erlang and Elixir. Mailboxes are incoming message queues: multiple processes can *write* to a mailbox, but only the process that owns a mailbox can read from it.

As an example of mailbox typing, consider the following Erlang implementation of a *future variable* [9], which is a write-once placeholder variable.

```

1 empty_future() →
2   receive
3     {put, X} → full_future(X)
4   end.
5 full_future(X) →
6   receive
7     {get, Pid} → Pid ! {reply, X},
8     full_future(X);
9     {put, _} →
10    erlang:error("Multiple writes")
11  end.
12 client() →
13   Future = spawn(future, empty_future, []),
14   Future ! {put, 5},
15   Future ! {get, self()},
16   receive
17     {reply, Result} →
18     io:fwrite("~w~n", [Result])
19  end.

```

The process `empty_future` waits to receive a `put` message containing the value to hold. This is a *selective receive*: other messages may be present in the mailbox but they will not be retrieved. After the `put` message has been received, the process then calls `full_future` with the provided value `X`. In `full_future` the process waits to receive `get` messages containing the process ID of the requesting actor, and responds by sending value `X` to the sender. Since a future variable should only be set once, the process throws an error if more `put` messages are received. Even in small examples such as this, protocol violations such as unexpected messages, forgotten replies and self-deadlocks may occur. All of these problems can be addressed and statically detected using mailbox typing [6,9].

A mailbox type consists of a *capability* and a *pattern*. The capability indicates whether the type describes messages that must be sent (!) or describe messages to be received from a mailbox (?). A pattern is a commutative regular expression that describes the contents of a mailbox. For the above example, we can define the following mailbox types:

$$\text{EmptyFuture} \triangleq ?(\text{Put}[\text{Int}] \odot \text{Get}[\text{!Reply}[\text{Int}]]^*) \quad \text{FullFuture} \triangleq ?\text{Get}[\text{!Reply}[\text{Int}]]^*$$

Type `EmptyFuture` describes a mailbox that may contain a single `Put` message with payload of type `Int`, and a potentially arbitrary amount of `Get` messages. Each `Get` message contains a mailbox reference that expects a `Reply` message to be sent. Type `FullFuture` is more restrictive in that it describes a mailbox only containing an arbitrary amount of `Reply` messages. By using these types, the type system can statically require that exactly one `Put` message is sent, and can ensure that every `Get` request receives a `Reply` message.

Of course, a type system can only guarantee correctness if its own specification is correct. It is therefore crucial to rigorously prove the properties a type system gives. Verifying that these proofs are indeed correct can be challenging in itself. Errors or oversights can remain undetected, resulting in false promises about type checked programs. For example, a mechanization of multiparty session typing [13] uncovered errors nearly a decade later [22].

In this paper, we report on our mechanization of the semantics of mailbox types, and of the *Pat* [9] programming language that incorporates mailbox types. Although *Pat* includes extensive pen-and-paper proofs of correctness, the proofs have not been mechanized. In particular, we have mechanized numerous properties about mailbox types themselves, and have mechanized the static semantics

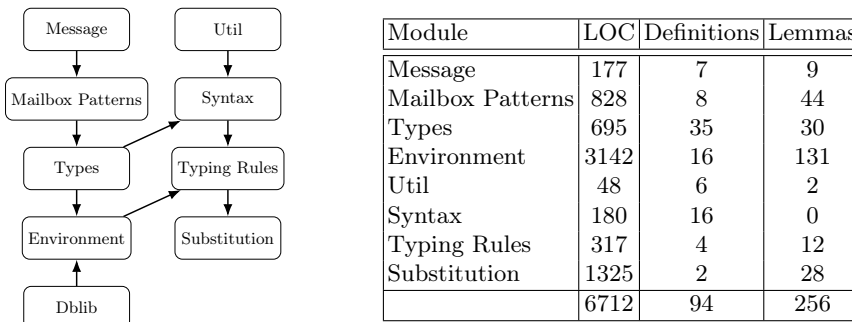



Fig. 1. Dependency graph of the modules in the Rocq project. Transitive dependencies are omitted.

of Pat. During our mechanization, we discovered several oversights in the original definitions and proofs which have since been fixed.

We use the Rocq [20] proof assistant for our mechanization. Figure 1 summarizes the structure of our Rocq project and shows the lines of code, number of definitions and lemmas for each file. Due to the necessity of *type combination* and *environment subtyping*, the `Environment` module contains, by far, the largest number of lemmas.

Throughout this paper, definitions and lemmas that are accompanied with the -symbol link to the corresponding code in the online documentation. The project can be found at <https://github.com/edgardSchi/mailbox-types-rocq>.

Methodology. Our mechanization includes both the syntax and type system of Pat, and prove several properties of the language such as the substitution lemma. We limit the scope of this paper to Pat’s static semantics, which alone requires a significant mechanization effort; a full account of Pat’s operational semantics and preservation theorem would require significant additional work and is left as a future direction. We concentrate on properties that are needed for proving preservation, but also do not require a description of Pat’s operational semantics.

We use Rocq due to its maturity and active development: several successful mechanizations have been conducted in Rocq, such as *CompCert* [15], the first fully formally verified real-world compiler, and formalizations of session types [22] and process calculi [1], modelling communication between processes.

Contributions. The main contributions of this paper are as follows:

- The first mechanization of the syntax and semantics of mailbox types (§2).
- A formalization of the static semantics of Pat, including definitions that are more amenable to mechanization (§3).
- Proofs of metatheoretical properties of the static semantics, including the nontrivial substitution lemma (§3), requiring over 200 auxiliary lemmas.
- We identify missing cases in the original definitions and places where the original proofs could be improved, that have since been fixed (§2 and §3).

Mailbox configurations (☞)	$m ::= \langle \rangle \mid \langle \mathbf{m} \rangle \mid m_1 \uplus m_2$	Pattern power (☞)
Mailbox patterns (☞)	$E, F ::= \emptyset \mid \mathbb{1} \mid \langle \mathbf{m} \rangle \mid E \odot F$ $\mid E \oplus F \mid E^*$	$E^0 := \mathbb{1}$ $E^{n+1} := E \odot E^n$

Fig. 2. The syntax of mailbox configurations and patterns.

2 Mailbox Types

We begin by describing mailbox configurations and mailbox patterns, before formalising mailbox types along with the necessary notion of type combination. Due to the potentially infinite representation of mailbox semantics, we require novel definitions that are more amenable to mechanization.

2.1 Mailboxes and Mailbox Patterns

Like Pat but unlike in the original definition of mailbox types [6], we do not include message payloads as part of a mailbox type, but rely only on *message tags* \mathbf{m} like **Put** and **Get** and associate types with messages when defining the typing rules for Pat (§3.5).

Mailbox Configurations. A mailbox represents an unordered collection of messages, where a message can be included multiple times. Figure 2 shows the definition of a *mailbox configuration* m . It is either the empty mailbox $\langle \rangle$, a singleton mailbox $\langle \mathbf{m} \rangle$ containing only message \mathbf{m} , or the union of two mailboxes $m_1 \uplus m_2$. In Rocq, we model mailbox configurations as lists, where \uplus is implemented as append. To reason about the equality of mailboxes, i.e. mailboxes with the same content, we need to take into account that order does not matter. Mailbox configurations m_1 and m_2 are equal, written $m_1 \approx m_2$, if they are permutations of each other (☞). We then have $\langle \mathbf{m} \rangle \uplus \langle \mathbf{m} \rangle \uplus \langle \mathbf{n} \rangle \approx \langle \mathbf{m} \rangle \uplus \langle \mathbf{n} \rangle \uplus \langle \mathbf{m} \rangle$, but $\langle \mathbf{n} \rangle \uplus \langle \mathbf{n} \rangle \not\approx \langle \mathbf{n} \rangle$.

Mailbox Patterns. The crucial building block of mailbox types are *mailbox patterns*. They can be thought of as a commutative regular expression that describes the content of a mailbox. Figure 2 shows the definition of mailbox patterns. The pattern \emptyset denotes an *unreliable* mailbox, i.e. one that received an unexpected message and which cannot be used for further communication. Pat’s type system ensures that it is not possible to create such a mailbox using well-typed terms. Pattern $\mathbb{1}$ represents an *empty* mailbox, containing zero messages. The pattern $\langle \mathbf{m} \rangle$ denotes a mailbox containing a *single message* \mathbf{m} . Mailbox *composition* is described by pattern $E \odot F$, combining both patterns E and F . As an example, pattern $\langle \mathbf{m} \rangle \odot \langle \mathbf{n} \rangle$ would describe a mailbox containing both \mathbf{m} and \mathbf{n} . On the other hand, $E \oplus F$ denotes *choice* between patterns. In particular, it describes that the mailbox is either described by pattern E or by pattern F . Lastly, pattern E^* describes repeated composition of pattern E . For example, $\langle \mathbf{m} \rangle^*$ denotes a mailbox containing zero or more \mathbf{m} messages.

$$\frac{}{\langle \rangle \in \mathbb{1}} \quad \frac{}{\langle \mathbf{m} \rangle \in \langle \langle \mathbf{m} \rangle \rangle} \quad \frac{m \in E}{m \in E \oplus F} \quad \frac{m \in F}{m \in E \oplus F} \quad \frac{\exists n. m \in E^n}{m \in E^*} \quad \frac{m \approx m_1 \uplus m_2 \quad m_1 \in E \quad m_2 \in F}{m \in E \odot F}$$

Fig. 3. The semantics of mailbox patterns (↗).

Pattern Semantics. Since patterns describe the contents of an unordered mailbox, multiple patterns can describe the same mailbox. For example, pattern $\langle \langle \mathbf{m} \rangle \odot \langle \langle \mathbf{n} \rangle \rangle$ describes the same mailbox as $\langle \langle \mathbf{n} \rangle \odot \langle \langle \mathbf{m} \rangle \rangle$. Thus, checking equality of mailbox patterns requires a semantic definition. The semantics of mailbox types are usually given by a denotational semantics that map patterns to sets of multisets of atoms [6]. Since we are working with Rocq, which is based on constructive logic, we require a method of constructing these sets; a particular challenge is representing the E^* pattern whose denotation is an infinite union of multisets and not easily representable in Rocq. Instead, we use a relational approach to connect mailbox configurations to patterns.

Instead of building possibly infinite sets of possible configurations, we relate a mailbox configuration to a pattern, written $m \in E$. A configuration is included in the set induced by a pattern if the configuration can be constructed using the pattern. For example, we have $\langle \mathbf{m} \rangle \uplus \langle \mathbf{n} \rangle \in \langle \langle \mathbf{m} \rangle \odot \langle \langle \mathbf{n} \rangle \rangle$ and $\langle \mathbf{n} \rangle \uplus \langle \mathbf{m} \rangle \in \langle \langle \mathbf{m} \rangle \odot \langle \langle \mathbf{n} \rangle \rangle$, but also $\langle \mathbf{m} \rangle \uplus \langle \mathbf{n} \rangle \in \langle \langle \mathbf{m} \rangle \odot \langle \langle \mathbf{n} \rangle \rangle \odot \mathbb{1}$.

The inference rules for this relation are given in Figure 3. The first two rules describe that the empty mailbox configuration $\langle \rangle$ is the pattern $\mathbb{1}$ and a configuration containing a single message \mathbf{m} is part of pattern $\langle \langle \mathbf{m} \rangle \rangle$. For the choice operator $E \oplus F$, there are two rules: one for when m is part of E and one for when m is part of F . The rule for composition is the most complex one. If m is part of a composition of patterns E and F , then m is required to be split into two sub-configurations, where one is part of E and one is part of F . Lastly, the rule for E^* uses the *power* operation (defined in Figure 2) to state that any mailbox configuration that can be built using a pattern E^n (for some n) is included in E^* .

One important observation to make is that we do not include a rule for pattern \emptyset . According to the definition of de'Liguoro and Padovani [6], \emptyset should map to the empty set, containing no configurations. Thus, since we are relating configurations and patterns, there must be no relation between any configuration and \emptyset .

We define *pattern inclusion* $E \sqsubseteq F$ as the proposition $\forall m. m \in E \Rightarrow m \in F$ (↗). Taking inclusion in both directions yields *pattern equivalence* $E \cong F$ (↗).

To verify this formalization has the required properties described by de'Liguoro and Padovani [6], we prove the following properties:

Lemma 1. *The relation \sqsubseteq is a precongruence, i.e. the following properties hold:*

- *Reflexive* (↗): $E \sqsubseteq E$
- *Transitive* (↗): if $E \sqsubseteq F$ and $F \sqsubseteq G$, then $E \sqsubseteq G$

$$\begin{array}{c}
\text{Pattern residuals } (\overline{\mathfrak{M}}) \quad \boxed{E \setminus \mathfrak{m} \triangleq F} \\
\\
\frac{}{0 \setminus \mathfrak{m} \triangleq 0} \quad \frac{}{1 \setminus \mathfrak{m} \triangleq 0} \quad \frac{}{\langle\langle \mathfrak{m} \rangle\rangle \setminus \mathfrak{m} \triangleq 1} \quad \frac{\mathfrak{m} \neq \mathfrak{n}}{\langle\langle \mathfrak{m} \rangle\rangle \setminus \mathfrak{n} \triangleq 0} \quad \frac{E \setminus \mathfrak{m} \triangleq E'}{E^* \setminus \mathfrak{m} \triangleq E' \odot E^*} \\
\\
\frac{E \setminus \mathfrak{m} \triangleq E' \quad F \setminus \mathfrak{m} \triangleq F'}{E \oplus F \setminus \mathfrak{m} \triangleq E' \oplus F'} \quad \frac{E \setminus \mathfrak{m} \triangleq E' \quad F \setminus \mathfrak{m} \triangleq F'}{E \odot F \setminus \mathfrak{m} \triangleq (E' \odot F) \oplus (E \odot F')}
\end{array}$$

Fig. 4. The definition of pattern residuals.

- *Compatible wrt. \odot* ($\overline{\mathfrak{M}}$): if $E_1 \sqsubseteq F_1$ and $E_2 \sqsubseteq F_2$, then $(E_1 \odot E_2) \sqsubseteq (F_1 \odot F_2)$
- *Compatible wrt. \oplus* ($\overline{\mathfrak{M}}$): if $E_1 \sqsubseteq F_1$ and $E_2 \sqsubseteq F_2$, then $(E_1 \oplus E_2) \sqsubseteq (F_1 \oplus F_2)$

Lemma 2. *The relation \cong is an equivalence relation, i.e. it is*

- *Reflexive* ($\overline{\mathfrak{M}}$): $E \cong E$
- *Transitive* ($\overline{\mathfrak{M}}$): if $E \cong F$ and $F \cong G$, then $E \cong G$
- *Symmetric* ($\overline{\mathfrak{M}}$): if $E \cong F$, then $F \cong E$

Additionally, we show that we indeed work with a commutative Kleene algebra $(M, \oplus, \odot, *, 0, 1)$, where M is the set of messages ($\overline{\mathfrak{M}}$).

Pattern Residuals. Mailboxes are dynamic objects, not static. Their content changes dynamically while receiving and consuming messages. For example, when consuming message \mathfrak{m} from a mailbox described by pattern $\langle\langle \mathfrak{m} \rangle\rangle$ the resulting mailbox should be empty, described by pattern 1 . To model this behavior, de'Liguoro and Padovani introduce the concept of *pattern residuals* [6] (Figure 4). In Pat [10], pattern residuals correspond exactly to the Brzozowski derivative of a commutative regular expression [3]. A pattern residual describes the resulting pattern when a message is consumed, i.e. removed from the mailbox.

An important property of residuals is *balancing*. For a configuration described by pattern $\langle\langle \mathfrak{m} \rangle\rangle \odot E$, which is included in a pattern F , E should be included in the residual $F \setminus \mathfrak{m} \triangleq F'$. This property is needed for reasoning about the type of a mailbox after a message has been received.

Lemma 3 (Balancing ($\overline{\mathfrak{M}}$)). *If $\langle\langle \mathfrak{m} \rangle\rangle \odot F \sqsubseteq E$ where $F \not\sqsubseteq 0$ and $E \setminus \mathfrak{m} \triangleq E'$, then $F \sqsubseteq E'$.*

Mechanization. Mailbox configurations mostly rely on the definitions of list permutations already provided by Rocq's standard library, but the formalization of mailbox patterns requires more effort. Describing the semantics of mailbox patterns as a relation between a configuration and a pattern, instead of possibly infinite multisets, yields a mechanization that is adequate to work with. The proofs about pattern equivalence and inclusion are typically inductions over

$$\begin{array}{ll}
\text{Mailbox types } (\bar{\mathfrak{A}}) & J, K ::= ?E \mid !E \\
\text{Base types } (\bar{\mathfrak{A}}) & C ::= \mathbf{1} \mid \text{Bool} \\
\text{Types } (\bar{\mathfrak{A}}) & T, U ::= C \mid J
\end{array}$$

Fig. 5. The syntax of types.

a pattern, but proving balancing (Lemma 3) requires several technical lemmas that rely on proven properties about permutations of mailbox configurations.

When mechanizing the pattern residual relation, we discovered that in the definition of Pat [9] the rule for $*$ was missing. This issue was consequently fixed in the journal version [8].

2.2 Types

A mailbox type (Figure 5) consists of a capability and a pattern. A capability is either *input* $?$ or *output* $!$. Together, capability and pattern describe what messages are expected to be received from ($?$) or stored in ($!$) the mailbox. To build an intuition for mailbox types, consider the following examples: A mailbox with mailbox type $?\langle\mathfrak{m}\rangle$ gives a process the right to receive a message \mathfrak{m} from that mailbox, while a mailbox with type $!\langle\mathfrak{m}\rangle$ requires the process to send a message \mathfrak{m} to that mailbox. For a mailbox with mailbox type $!\langle\langle\mathfrak{m}\rangle \oplus \langle\mathfrak{n}\rangle$, a process may choose whether to send message \mathfrak{m} or \mathfrak{n} , whereas a mailbox with type $?\langle\langle\mathfrak{m}\rangle \oplus \langle\mathfrak{n}\rangle$ allows a process to receive either message \mathfrak{m} or \mathfrak{n} from that mailbox. The type $?\langle\langle\mathfrak{m}\rangle \odot \langle\mathfrak{n}\rangle$ is associated to a mailbox that allows a process to receive both \mathfrak{m} and \mathfrak{n} , while a mailbox with mailbox type $!\langle\langle\mathfrak{m}\rangle \odot \langle\mathfrak{n}\rangle$ requires a process to send both \mathfrak{m} and \mathfrak{n} to the mailbox, in whichever order. Finally, a mailbox with type $!\langle\langle\mathfrak{m}\rangle^*\rangle$ allows a process to send an arbitrary number of \mathfrak{m} messages to the mailbox, while $?\langle\langle\mathfrak{m}\rangle^*\rangle$ describes a mailbox from which a process is able to receive an arbitrary number of \mathfrak{m} messages.

In addition to mailbox types, we also include *base types*. While the unit type $\mathbf{1}$ is required by the type system, additional base types can be added. For demonstration purposes we also include a Boolean type `Bool`. A type T is then either some base type or a mailbox type.

Type Combination. To model the dynamic nature of mailboxes receiving and sending messages, mailbox types should be able to be combined to create a new mailbox type. Mailbox types need to ensure that sends and receives are matched, i.e. a send is required to have a corresponding receive. For example, given mailbox types $!\langle\mathfrak{m}\rangle$ and $?\langle\langle\mathfrak{m}\rangle \odot \langle\mathfrak{n}\rangle$, the resulting type from the type combination should be $?\langle\mathfrak{n}\rangle$.

In the original definitions of type combinations [6,9], types are identified up to commutativity and associativity. In practice, this means that types such as $!\langle\langle\mathfrak{m}\rangle \odot \langle\mathfrak{n}\rangle$ and $!\langle\langle\mathfrak{n}\rangle \odot \langle\mathfrak{m}\rangle$ are considered to be equivalent. While this is true from a semantic point of view, syntactically these are different types and assuming equality up to a certain property creates problems when mechanizing

$$\begin{array}{c}
\text{Pattern permutation } (\text{⌘}) \quad \boxed{E \sim F} \quad \text{Type permutation } (\text{⌘}) \quad \boxed{T \sim U} \\
\\
\frac{}{(E \odot F) \sim (F \odot E)} \quad \frac{F \sim F'}{E \odot F \sim E \odot F'} \quad \frac{}{C \sim C} \quad \frac{E \sim F}{!E \sim !F} \\
\\
\frac{}{E \sim E} \quad \frac{F \sim E}{E \sim F} \quad \frac{E \sim F' \quad F' \sim F}{E \sim F} \quad \frac{E \sim F}{?E \sim ?F} \\
\\
\frac{}{E \odot (F \odot F') \sim (E \odot F) \odot F'} \\
\\
\text{Type combinations } (\text{⌘}) \quad \boxed{T \boxplus U = T'} \\
\\
\frac{}{C \boxplus C = C} \quad \frac{}{!E \boxplus !F = !(E \odot F)} \quad \frac{}{!E \boxplus ?(E \odot F) = ?F} \\
\\
\frac{}{?(E \odot F) \boxplus !E = ?F} \quad \frac{T'_1 \boxplus T'_2 = T'_3 \quad T_1 \sim T'_1 \quad T_2 \sim T'_2 \quad T_3 \sim T'_3}{T_1 \boxplus T_2 = T_3}
\end{array}$$

Fig. 6. The definitions of pattern and type permutations, as well as type combinations.

such definitions. We require a way to express equivalence up to commutativity and associativity in the context of type combinations. To this end, we define a relation between mailbox types, holding if they are permutations of each other.

The relation $E \sim F$ (Figure 6) is expressing that two patterns can be rewritten into each other using commutativity and associativity of the \odot operator. Note how these rules do not include the \oplus operator. This is because type combination only considers type composition, and it is not required to include pattern choice in the definition of \sim to achieve the desired properties. Note the difference between the relations $E \sim F$ and $E \cong F$. The former is expressing that E and F can be rewritten into each other by swapping around patterns that appear between the composition operator. The latter relates two semantically equal patterns. For example, consider pattern $\langle\langle m \rangle\rangle \odot \langle\langle n \rangle\rangle$. Both $\langle\langle m \rangle\rangle \odot \mathbb{1} \sim \mathbb{1} \odot \langle\langle m \rangle\rangle$ and $\langle\langle m \rangle\rangle \odot \mathbb{1} \cong \mathbb{1} \odot \langle\langle m \rangle\rangle$ hold, while $\langle\langle m \rangle\rangle \odot \mathbb{1} \cong \langle\langle m \rangle\rangle$, but $\langle\langle m \rangle\rangle \odot \mathbb{1} \not\sim \langle\langle m \rangle\rangle$.

The relation \sim is extended to mailbox types by relating the underlying patterns. Base types are always considered to be permutations of each other.

Type combinations model the concurrent interaction between processes acting on the same mailbox. For example, assume processes P_1 and P_2 that both interact with mailbox m . Process P_1 stores some message m , while P_2 stores some message n . Then, mailbox m can be described by the type combination $!\langle\langle m \rangle\rangle \boxplus !\langle\langle n \rangle\rangle = !(\langle\langle m \rangle\rangle \odot \langle\langle n \rangle\rangle)$. If instead, we have type $?\langle\langle m \rangle\rangle \odot \langle\langle n \rangle\rangle$, then we get the type combination $!\langle\langle m \rangle\rangle \boxplus ?(\langle\langle m \rangle\rangle \odot \langle\langle n \rangle\rangle) = ?\langle\langle n \rangle\rangle$. Figure 7 shows an example with three mailbox types. There are different ways of arranging the combination, but in the end, the result always balances out to an empty mailbox.

Figure 6 shows the definition of the *type combination relation*. Note how there is no rule for type combination of two types with input capabilities. This




Fig. 7. Example of possible type combinations for three mailbox types. Each top node represents the type of the same mailbox in three different processes that run in parallel. Edges indicate the type combination relation.

would be allowing parallel reads of the same mailbox, which is not safe [6], and thus not permitted. The last rule in the definition of type combinations allows us to substitute types for their permutations. Note how the previous definition of pattern equality $E \cong F$ is specifically *not* used in the definition of type combination. The combination operation is a *syntactic* operation, i.e. we can only combine types that are *syntactically equal* to the forms specified by the type combination operation (up to commutativity and associativity). To combine other mailbox types, for example $!\langle m \rangle$ and $?(\langle m \rangle \oplus \langle m \rangle)$ we must firstly use *subtyping* to rewrite the latter mailbox type to $?(\langle m \rangle \odot 1)$ before using the combination operator. We will discuss subtyping further in Section 3.

Mechanization. Our mechanization of mailbox types differs from the original formalizations in that we treat partial operations as relations, and include the notion of pattern and type permutations. This addition allows us to prove desired properties, such as commutativity of type combinations. Naturally, proofs about type combinations become more complex since we have to cover additional cases.

3 The Pat Programming Language

Pat [9] is the first programming language to support mailbox types. In contrast to languages like Erlang or Elixir, where mailboxes are implicit, *Pat* supports first-class explicit mailboxes. We can express the Future example in *Pat* as follows:

```

1  def emptyFuture(self : EmptyFuture) : 1 {
2    guard self : Put ⊙ Get* {
3      receive Put[x] from self ↦
4        fullFuture(self, x)
5    }
6  }
7  def fullFuture(self : FullFuture, value : Int) : 1 {
8    guard self : Get* {
9      free ↦ ()
10     receive Get[user] from self ↦
11       user!Reply[value];
12     fullFuture(self, value)
13   }
14 }

15 def client() : 1 {
16   let future = new in
17   spawn emptyFuture(future);
18   let self = new in
19   future!Put[5];
20   future!Get[self];
21   guard self : Reply {
22     receive Reply[result]
23   from self ↦
24     free self;
25     print(intToString(result))
26   }
27 }

```

To support dynamic mailbox creation, new mailboxes are created using the **new** keyword and bound to names using **let**-bindings. *Pat* programs can then

send to mailboxes (e.g., $future!Put[5]$ sends a **Put** message with payload 5 to mailbox $future$) and receive from mailboxes using the **guard** keyword. A guard is always accompanied by a mailbox pattern, representing the mailbox’s contents, and contains *guard clauses* including **receive** which receives a message from the mailbox and binds the payload and new mailbox name; and **free** which is invoked when a mailbox will not receive any more messages and can be deallocated.

A mechanization of the operational semantics of Pat is outside of the scope of this paper, but we give an outline of the semantics here for context. Pat’s semantics makes use of a *frame stack* representation that simplifies specifying inductive invariants for the type preservation proof. There are two reduction relations: a deterministic β -reduction relation on functional terms (like **let**-bindings and function calls) that is entirely standard, and a nondeterministic reduction relation on a language of *configurations* that closely resemble processes in the π -calculus. The communication and concurrency constructs reduce as follows:

- Evaluating the **new** construct creates a fresh runtime mailbox name and returns it to the calling thread.
- Evaluating a message send $a!m[v]$ returns the unit value to the calling thread and creates an *in-flight* message $a \leftarrow m[v]$.
- Evaluating **spawn** t returns the unit value to the calling thread and evaluates t in a new process.
- The behaviour of a **guard** $a : E \{ \vec{g} \}$ expression depends on the messages that have been sent to mailbox a along with the guards contained in \vec{g} :
 - If there are no messages sent to a ; no references to a in other threads; and \vec{g} contains a **free** $\mapsto t$ guard, then mailbox name a will be deleted and the thread will run computation t .
 - If there is a message $a \leftarrow m[v]$ and \vec{g} contains a clause **receive** $m \mapsto t$, then the thread will run computation t with bindings for the received value and updated mailbox name.

A major problem in Pat is *mailbox name aliasing*, where multiple static names could be given to the same underlying mailbox name. This could be used to break type soundness. The Mailbox Calculus addresses this issue using a dependency graph, which can additionally rule out cyclic dependencies, but this approach does not scale to programming languages (see [9] for a detailed discussion). Instead, Pat uses *quasi-linear typing* [14,7], where types are equipped with *usage annotations*. If a type is defined as *returnable*, it is allowed to appear in the return type of an expression and be returned from an evaluation frame. Otherwise, if a type is defined as *second-class*, it can be used several times, but is never allowed to escape its scope. A returnable occurrence of a mailbox name must be its last lexical usage. Quasilinearity annotations, in combination with syntactic restrictions, eliminates aliasing issues and self-deadlocks.

The authors of Pat present both a *declarative* and an *algorithmic* version of their type system. The algorithmic version (which is needed because the declarative system makes use of environment subtyping and non-deterministic context splits) makes use of *bidirectional typing* and *constraint solving*, and is at the core of Pat’s implementation. This paper focuses on the declarative type system.

		Type operations (⊠, ⊡, ⊢, ⊣)
Usage annotations (⊠)	$\eta ::= \circ \mid \bullet$	$[C] = C$ $[C] = C$
Usage-annotated types (⊠)	$A, B ::= C \mid J^\eta$	$[J] = J^\circ$ $[J] = J^\bullet$
		$[J^\eta] = J^\circ$ $[J^\eta] = J^\bullet$
Usage-annotated type combinations (⊠)		
	$A \triangleright B = A'$	$\eta_1 \triangleright \eta_2 = \eta_3$
	$\frac{}{C \triangleright C = C}$	$\frac{\eta_1 \triangleright \eta_2 = \eta \quad J_1 \boxplus J_2 = K}{J_1^{\eta_1} \triangleright J_2^{\eta_2} = K^\eta}$
	$\circ \triangleright \circ = \circ$	$\circ \triangleright \bullet = \bullet$
Usage subtyping (⊠)		
	$\eta_1 \leq \eta_2$	$A \leq B$
	$\eta \leq \eta$	$\bullet \leq \circ$
	$\frac{}{C \leq C}$	$\frac{E \sqsubseteq F \quad \eta_1 \leq \eta_2}{?E^{\eta_1} \leq ?F^{\eta_2}}$
		$\frac{F \sqsubseteq E \quad \eta_1 \leq \eta_2}{!E^{\eta_1} \leq !F^{\eta_2}}$

Fig. 8. Definitions of usage-annotated types.

3.1 Quasi-linear Mailbox Types

Pat uses quasi-linear typing [14,7] to solve the problem of name aliasing and cyclic dependency between mailboxes. To this end, every mailbox type is equipped with a *usage annotation*: either *second class* (\circ) or *returnable* (\bullet). Values of second class types can be used several times within the scope they are defined, however, they cannot leave that scope. On the other hand, values of returnable types can be let-bound and used as the subject of a guard, but must appear as the last lexical occurrence of the name. Base types are not equipped with usage, as they do not suffer from the aliasing problems.

Figure 8 shows definitions of usage annotations, usage-annotated types and relations on them. We also define operators $[\cdot]$ and $[\cdot]$ to turn types into second class and returnable respectively. Additionally, we define a mailbox type to be *unrestricted* (written $\text{un}(A)$) (⊠) if it is a base type, or $A = !\mathbb{1}^\circ$.

Usage-annotated Type Combination. We extend type combination to *usage-annotated type combination* by incorporating usage combination. For usage annotations, we have to obey the rule that a \circ variable has to occur before a \bullet variable. Otherwise, we would violate the property that any \bullet variable appears as the final one in a process. Combining two \bullet types is not allowed as there could otherwise be multiple instances of a returnable mailbox name per thread.

Subtyping. Subtyping is needed in mailbox typing to rewrite mailbox types so that they can be combined by the type combination operator. For mailbox types, subtyping is *covariant* for types with receive capability and *contravariant* for types with send capability [6]. Subtyping also allows us to treat returnable types as second-class types.

Definition names (☞)	f
Function definitions (☞)	$d ::= \mathbf{def} f(A) : B \{t\}$
Booleans	$b ::= \mathbf{true} \mid \mathbf{false}$
Values (☞)	$v ::= b \mid () \mid \mathbf{Var} n$
Terms (☞)	$t ::= v \mid \mathbf{let} t_1 \mathbf{in} t_2 \mid f(v) \mid \mathbf{spawn} t$ $\quad \mid \mathbf{new} \mid v_1!m[v_2] \mid \mathbf{guard} v : E \{\vec{g}\}$
Guards (☞)	$g ::= \mathbf{fail} \mid \mathbf{free} \mapsto t \mid \mathbf{receive} m \mapsto t$
Environments (☞)	$\Gamma ::= \emptyset \mid \perp, \Gamma \mid A, \Gamma$

Fig. 9. The syntax of Pat

3.2 Syntax

The Pat syntax we are working with is close to the original definitions by Fowler et al. [9], making use of de Bruijn indices [5] to represent variables. We use de Bruijn indices as they are well-studied for proof assistants, and in particular Rocq [1,19]. We make use of the Rocq library *dblib* [18] that automatically handles de Bruijn indices and operations on them such as shifting.

To ease the burden of proving properties about messages, we restrict messages to only carry a single payload, and consider only functions that take a single argument. We expect that extending our proofs to include arbitrary numbers of payloads and function arguments will not significantly affect the proofs.

Figure 9 shows our syntax of Pat. A value v is either a Boolean value, the unit value $()$, or a variable $\mathbf{Var} n$ represented as a de Bruijn index, with natural number n . Including variables into the definition of values is fairly nonstandard. The inspiration for this approach is the *fine-grain call-by-value* evaluation strategy [16], that makes a syntactic difference between values and computations. Values, including variables that can also represent mailbox names, are immutable and cannot be reduced, while other terms produce values and can be reduced further.

A term is either a value, a let-expression or a function application $f(v)$ of function f on value v . Additionally, there are terms specifically for interacting with mailboxes and processes. The term $\mathbf{spawn} t$ can be used to create a new process running term t . Dynamic mailbox name creation is done with the **new** keyword, generating a fresh mailbox name. Term $v_1!m[v_2]$ sends message m with payload value v_2 to the mailbox with name v_1 . A guard term $\mathbf{guard} v : E \{\vec{g}\}$ asserts that the mailbox associated with name V has pattern E , and invokes some guard from the list of guards \vec{g} . Guards perform actions on a mailbox: the guard **fail** is triggered when a mailbox receives an unexpected message, and should be evaluated by a well-typed program. In contrast, guard **free** $\mapsto t$ is triggered when a mailbox is empty and not referenced elsewhere in the system; the corresponding mailbox is freed and we proceed with continuation t . Finally, the guard **receive** $m \mapsto t$ is triggered when a mailbox contains the message m . The contents of the message are then available in the continuation t .

A Pat program $\mathcal{P} = (\mathcal{S}, \mathcal{D}, t_0)$ (☞) is a tuple consisting of a signature \mathcal{S} for mapping messages to their payload type, a mapping \mathcal{D} from definition names to

function definitions, and an initial term t_0 . For simplicity, given a program \mathcal{P} , we write $\mathcal{S}_{\mathcal{P}}$ for the signature, and $\mathcal{D}_{\mathcal{P}}$ for the function definitions.

Mechanization. The mechanization of the syntax of Pat is close to its pen-and-paper counterpart, with the exception of using de Bruijn indices. Defining the syntax poses no problem after settling on the representation of variables. While the choice of using de Bruijn indices is standard [1,22,19], it comes with the downside of extensive reasoning about indices and lifting. While *dblib* alleviates some of this burden, a substantial amount of manual work is still required.

Definition names are treated similar to message tags, being represented as an arbitrary type equipped decidable equality.

3.3 Environments

Pat’s type system makes use of *environments* for keeping track of types of variables while type checking terms. The implementation of environments depends on how variables are represented. For traditional strings as variables, it is common to use a map, mapping variables to types [17]. In the case of de Bruijn indices, an environment is typically represented as a list of types where a variable indicates the position of its type in the list [17,23]. For Pat, the representation of an environment as a simple list of types is unfortunately not enough, due to its parallel fragment and use of quasilinearity. To illustrate the problem, consider the type derivation of a mailbox calculus term on the left-hand side:

$$\frac{A_1, \emptyset \vdash \text{Var } 0 \quad A_2, \emptyset \vdash \text{Var } 1}{A_2, A_1, \emptyset \vdash \text{Var } 0 \parallel \text{Var } 1} \quad \frac{\perp, A_1, \emptyset \vdash \text{Var } 0 \quad A_2, \perp, \emptyset \vdash \text{Var } 1}{A_2, A_1, \emptyset \vdash \text{Var } 0 \parallel \text{Var } 1}$$

We assume variable $\text{Var } 0$ has type A_1 and variable $\text{Var } 1$ has type A_2 . The term expresses that there are two processes running in parallel, each with only a single variable. Additionally, the current environment consists of A_2, A_1, \emptyset . By the typing rules of the mailbox calculus (read bottom-up), given by de’Liguoro and Padovani [6], when typing parallel processes, the environment needs to be split into two. Assuming variable $\text{Var } 0$ has indeed type A_1 , the left subtree works without problems. Looking up the type at position $\text{Var } 0$ in environment A_1, \emptyset yields A_1 . However, the right subtree poses a problem. There is no type at position $\text{Var } 1$ in environment A_2, \emptyset , as it only contains a single element.

One way of solving this issue is the inclusion of entries in the environment representing the absence of a type at a certain position. The example on the right-hand side adds an additional \perp -entry, representing an empty position. With this approach, whenever an environment is split, the size of the two new environments is equal to the size of the original one. The *dblib* library we use for managing de Bruijn indices also includes an implementation of environments based on this exact approach. It provides the needed definitions and operations, such as lookup and insertion, together with proofs about properties on environments.

Fowler et al. [9] describe several relations and operations on environments, defined in Figure 10. Our definitions are essentially equal to the original ones,

$$\begin{array}{c}
\textbf{Environment subtyping} \quad (\text{⊆}) \quad \boxed{\Gamma_1 \leq \Gamma_2} \\
\frac{}{\Gamma \leq \Gamma} \quad \frac{\Gamma_1 \leq \Gamma_2 \quad \Gamma_2 \leq \Gamma_3}{\Gamma_1 \leq \Gamma_3} \quad \frac{\Gamma_1 \leq \Gamma_2}{\perp, \Gamma_1 \leq \perp, \Gamma_2} \quad \frac{\Gamma_1 \leq \Gamma_2 \quad \text{un}(A)}{A, \Gamma_1 \leq \perp, \Gamma_2} \quad \frac{\Gamma_1 \leq \Gamma_2 \quad A_1 \leq A_2}{A_1, \Gamma_1 \leq A_2, \Gamma_2} \\
\textbf{Environment combination} \quad (\text{⊳}) \quad \boxed{\Gamma_1 \triangleright \Gamma_2 = \Gamma_3} \\
\frac{}{\emptyset \triangleright \emptyset = \emptyset} \quad \frac{\Gamma_1 \triangleright \Gamma_2 = \Gamma_3}{\perp, \Gamma_1 \triangleright \perp, \Gamma_2 = \perp, \Gamma_3} \quad \frac{\Gamma_1 \triangleright \Gamma_2 = \Gamma_3}{A, \Gamma_1 \triangleright \perp, \Gamma_2 = A, \Gamma_3} \quad \frac{\Gamma_1 \triangleright \Gamma_2 = \Gamma_3}{\perp, \Gamma_1 \triangleright A, \Gamma_2 = A, \Gamma_3} \\
\frac{A_1 \triangleright A_2 = A_3 \quad \Gamma_1 \triangleright \Gamma_2 = \Gamma_3}{A_1, \Gamma_1 \triangleright A_2, \Gamma_2 = A_3, \Gamma_3} \\
\textbf{Disjoint environment combination} \quad (\text{⊕}) \quad \boxed{\Gamma_1 + \Gamma_2 = \Gamma_3} \\
\frac{}{\emptyset + \emptyset = \emptyset} \quad \frac{\Gamma_1 + \Gamma_2 = \Gamma_3}{\perp, \Gamma_1 + \perp, \Gamma_2 = \perp, \Gamma_3} \quad \frac{\Gamma_1 + \Gamma_2 = \Gamma_3}{A, \Gamma_1 + \perp, \Gamma_2 = A, \Gamma_3} \quad \frac{\Gamma_1 + \Gamma_2 = \Gamma_3}{\perp, \Gamma_1 + A, \Gamma_2 = A, \Gamma_3} \\
\frac{\Gamma_1 + \Gamma_2 = \Gamma_3}{C, \Gamma_1 + C, \Gamma_2 = C, \Gamma_3}
\end{array}$$

Fig. 10. The definitions of relations on environments.

with modifications to incorporate \perp -entries and some corrections that are further discussed at the end of this section. In each relation we add a rule for the case where the first entry in both environments is \perp . Similar to partial operations on types, we turn partial operations on environments into relations.

Pat requires subtyping on *environments* rather than just types in order to rewrite types so that they can be used by the type combination operator. Sequential environment combination $\Gamma_1 \triangleright \Gamma_2 = \Gamma$ makes use of type combination and is used when type checking let-expressions. This is the main rule used to ensure that subsequent uses of a mailbox variable “balance out”. Additionally, we use *disjoint environment combination* to reason about expressions where subterms may not share mailbox-typed variables (e.g. a variable may not be used as the target of a message send and also be contained in its payload). Environments combined using $+$ are only allowed to share variables of base types and be disjoint otherwise.

As our definition of environments includes \perp entries, it is possible for an environment to be empty without being \emptyset , necessitating a new notion of emptiness of environments $\text{Empty}(\Gamma)$ (⊆) that holds when Γ only contains \perp -entries.

Mechanization. Formalizing environments is, by far, the most complex part of this mechanization. The library *dblib* reduces some of the complexity around environments by proving operations and properties on them. However, reasoning about environment combinations and environment subtyping still is far from trivial. Independent of the representation of environments that were considered,

Pattern normal form (P)

 $E \models F$

$$\frac{}{E \models \mathbb{0}} \quad \frac{}{E \models \mathbb{1}} \quad \frac{E \setminus \mathbf{m} \triangleq E' \quad F \cong E'}{E \models \langle \mathbf{m} \rangle \odot F} \quad \frac{E \models F_1 \quad E \models F_2}{E \models F_1 \oplus F_2}$$

Fig. 11. Pattern normal form.

environment splitting and subtyping require dozens of technical lemmas. While the proofs of these properties follow the same general approach, they are lengthy and tedious to perform. For example, for the relation $\Gamma_1 \triangleright \Gamma_2 = \Gamma_3$ we have proved that there exists an environment that is the environment combination of both Γ_1 and Γ_2 with the entry at the same position replaced by \perp in both environments (P). Similar lemmas had to be proven for environment subtyping and disjoint combination. While mechanizing subtyping for environments according to the definition of Pat [9], we noticed that the relation was missing an inference rule, making environment subtyping not transitive. By explicitly including the transitivity rule we were able to prove desired properties of environment subtyping. This issue was subsequently fixed in a revised version [8].

3.4 Pattern Normal Form

When reasoning about the typing of guards, both the Mailbox Calculus and Pat require the type of a mailbox to be in *pattern normal form* (PNF), which is a particular structure of pattern that reflects the subpatterns expected by each guard clause. Specifically, a pattern E is in PNF if it is a sum of subpatterns of the form $\mathbb{0}$ (for a **fail** guard), $\mathbb{1}$ (for a **free** guard) and $\langle \mathbf{m} \rangle \odot F$ (for a **receive** guard), where F is equivalent to the residual of E with respect to \mathbf{m} . This allows a correct type to be given to a re-bound mailbox name after receiving.

Figure 11 shows the formal definitions for PNF; we write $\models E$ for $E \models E$. The following example derivation shows that pattern $(\langle \mathbf{m} \rangle \odot \mathbb{1}) \oplus \mathbb{1}$ is in PNF:

$$\frac{\frac{\langle \mathbf{m} \rangle \setminus \mathbf{m} \triangleq \mathbb{1} \quad \mathbb{1} \setminus \mathbf{m} \triangleq \mathbb{0}}{\langle \mathbf{m} \rangle \odot \mathbb{1} \oplus \mathbb{1} \setminus \mathbf{m} \triangleq (\mathbb{1} \odot \mathbb{1}) \oplus (\langle \mathbf{m} \rangle \odot \mathbb{0})} \quad \mathbb{1} \cong (\mathbb{1} \odot \mathbb{1}) \oplus (\langle \mathbf{m} \rangle \odot \mathbb{0})}{\langle \mathbf{m} \rangle \odot \mathbb{1} \oplus \mathbb{1} \models (\langle \mathbf{m} \rangle \odot \mathbb{1})} \quad \frac{}{\langle \mathbf{m} \rangle \odot \mathbb{1} \oplus \mathbb{1} \models \mathbb{1}}}{\langle \mathbf{m} \rangle \odot \mathbb{1} \oplus \mathbb{1} \models (\langle \mathbf{m} \rangle \odot \mathbb{1}) \oplus \mathbb{1}}$$

3.5 Typing Rules

We now have all the necessary definitions for specifying the typing rules of Pat. These rules are similar to the original ones by Fowler et al. [9], with a few modifications to fit our definitions of variables, environments and terms.

Figure 12 shows the typing rules for Pat, parameterised by a program \mathcal{P} ; we omit \mathcal{P} in the rules for readability. We highlight some of the typing rules. Whenever an empty environment is required, we use predicate $\text{Empty}(T)$ instead of checking for the empty environment \emptyset . The insertion operation $[x \mapsto A]T$ ensures that T contains type A at position x .

Typing rules for programs and functions (⊢_ℙ) (⊢_ℙ) $\vdash_{\mathcal{P}}$ $\vdash_{\mathcal{P}} \mathbf{def} f(A) : B \{t\}$

$$\frac{\mathcal{P} = (\mathcal{S}, \mathcal{D}, t_0) \quad \forall f. \vdash_{\mathcal{P}} \mathcal{D}(f) \quad \emptyset \vdash t_0 : \mathbf{1}}{\vdash_{\mathcal{P}}}$$

$$\frac{A, \emptyset \vdash t : B}{\vdash_{\mathcal{P}} \mathbf{def} f(A) : B \{t\}}$$

Typing rules for terms (⊢_ℙ) $\Gamma \vdash_{\mathcal{P}} t : A$

VAR $\text{Empty}(\Gamma)$ $\Gamma \vdash_{\mathcal{P}} \mathbf{Var} x : A$

BOOL $\text{Empty}(\Gamma)$ $\Gamma \vdash_{\mathcal{P}} b : \mathbf{Bool}$

UNIT $\text{Empty}(\Gamma)$ $\Gamma \vdash_{\mathcal{P}} () : \mathbf{1}$

APP $\mathcal{D}_{\mathcal{P}}(f) = \mathbf{def} f(A) : B \{t\}$ $\Gamma \vdash_{\mathcal{P}} v : A$ $\Gamma \vdash_{\mathcal{P}} f(v) : B$

LET $\Gamma_1 \triangleright \Gamma_2 = \Gamma \quad \Gamma_1 \vdash t_1 : [A]$ $\Gamma \vdash_{\mathcal{P}} \mathbf{let} t_1 \mathbf{in} t_2 : B$

$[0 \mapsto [A]]\Gamma_2 \vdash t_2 : B$

NEW $\text{Empty}(\Gamma)$ $\Gamma \vdash_{\mathcal{P}} \mathbf{new} : ?\mathbf{1}^{\bullet}$

SEND $\mathcal{S}_{\mathcal{P}}(\mathbf{m}) = T \quad \Gamma_1 \vdash v_1 : !\langle\mathbf{m}\rangle^{\circ}$ $\Gamma \vdash_{\mathcal{P}} v_1 !\mathbf{m}[v_2] : \mathbf{1}$

$\Gamma_2 \vdash v_2 : [T] \quad \Gamma_1 + \Gamma_2 = \Gamma$

GUARD $E \sqsubseteq F \quad \models F \quad \Gamma_1 \vdash v : ?F^{\bullet}$ $\Gamma \vdash_{\mathcal{P}} \mathbf{guard} v : E \{\vec{g}\} : A$

$\Gamma_2 \vdash \vec{g} : A :: F \quad \Gamma_1 + \Gamma_2 = \Gamma$

SUB $\Gamma \leq \Gamma' \quad A \leq B \quad \Gamma' \vdash t : A$ $\Gamma \vdash_{\mathcal{P}} t : B$

SPAWN $[\Gamma] = \Gamma' \quad \Gamma' \vdash t : \mathbf{1}$ $\Gamma \vdash_{\mathcal{P}} \mathbf{spawn} t : \mathbf{1}$

Typing rules for guard sequences (⊢_ℙ) $\Gamma \vdash_{\mathcal{P}} \vec{g} : A :: E$

$\frac{\Gamma \vdash g : A :: E}{\Gamma \vdash \{g\} : A :: E}$ **SINGLE** $\frac{\Gamma \vdash g : A :: E \quad \Gamma \vdash \vec{g}_s : A :: E'}{\Gamma \vdash \{g, \vec{g}_s\} : A :: E \oplus E'}$ **SEQ**

Typing rules for guards (⊢_ℙ) $\Gamma \vdash_{\mathcal{P}} g : A :: E$

FAIL $\Gamma \vdash_{\mathcal{P}} \mathbf{fail} : A :: \emptyset$

FREE $\Gamma \vdash_{\mathcal{P}} \mathbf{free} \mapsto t : A :: \mathbf{1}$

RECEIVE $\mathcal{S}_{\mathcal{P}}(\mathbf{m}) = T \quad \mathbf{base}(\mathbf{T}) \vee \mathbf{base}(\Gamma)$ $\Gamma \vdash_{\mathcal{P}} \mathbf{receive} \mathbf{m} \mapsto t : B :: \langle\mathbf{m}\rangle \odot E$

$[0 \mapsto [T]]([0 \mapsto ?E^{\bullet}]\Gamma) \vdash t : B$

Fig. 12. Pat typing rules

The SEND rule types a send expression $v_1 !\mathbf{m}[v_2]$, where message \mathbf{m} with payload v_2 is sent to mailbox v_1 . The mailbox associated with v_1 has to be of type $!\langle\mathbf{m}\rangle^{\circ}$, i.e. it must be capable of sending the message \mathbf{m} . Payload v_2 is required to be of the type given by the message's signature with second-class usage. Both derivations for v_1 and v_2 must occur under separate environments Γ_1 and Γ_2 , i.e. they only share base types and their combination yields Γ .

Rule RECEIVE types the receive guard $\mathbf{receive} \mathbf{m} \mapsto t$, retrieving message \mathbf{m} from the mailbox, with type B and pattern $\langle\mathbf{m}\rangle \odot E$. Either Γ must only consist of base types or the payload's type T is a base type. This is a syntactic restriction, ensuring that no existing mailbox name is shadowed by a message's

payload. The continuation t must be typeable under Γ extended by $?E^\bullet$ and $[T]$. Type $?E^\bullet$ represents the mailbox after \mathfrak{m} is consumed. Because of the usage of de Bruijn indices, the order in which the types are inserted into the environment is important. Since inserting a type at position 0 corresponds to simply adding it to the front of the environment, the environment $[0 \mapsto [T]]([0 \mapsto ?E^\bullet]\Gamma)$ is equal to $[T], ?E^\bullet, \Gamma$. Thus, a variable 0 would map to type $[T]$ and a variable 1 would map to type $?E^\bullet$.

The typing relation for a sequence of guards enables us to typecheck each guard individually, while also ensuring the correct pattern structure of the overall expression. In the original definition of typing for guard sequences, Fowler et al. [9] provide only a single rule of the following form:

$$\frac{(\Gamma \vdash g_i : A :: E_i)_i}{\Gamma \vdash \vec{g} : A :: E_1 \oplus \dots \oplus E_n}$$

By splitting this single rule into two distinct ones (SINGLE and SEQ), we provide a recursive definition for typing of sequences. This avoids using lists of well-typed relations as a premise in the rule, and in turn, Rocq is able to automatically come up with appropriate inductive schemes.

An example of a type derivation in Pat can be found in Appendix A.

Mechanization. The mechanization of Pat’s typing rules closely follows the pen-and-paper definitions. However, proving properties about well-typed terms requires technical lemmas, such as inversion lemmas or canonical forms of terms. For example, we had to prove that every well-typed Boolean value under environment Γ implies that Γ is a subenvironment of an empty environment ($\bar{\mathfrak{A}}$). These types of proofs would usually be omitted in classical pen-and-paper proofs, but in our case, even these kinds of proofs required technical lemmas such as properties about environment subtyping.

3.6 Substitution

The call-by-value operational semantics of Pat ensures that substitution only needs to replace a variable with a value. However, since Pat enforces a syntactic separation of values and computations (following *fine-grain call-by-value* [16]), variables are classed as values, which makes the proof of substitution nontrivial.

The Rocq library *dblib* [18], which we use for de Bruijn indices, provides a definition of substitution. To use it, a user defines substitution as a *traverse* function ($\bar{\mathfrak{A}}$) over terms. It applies a function f at every variable and provides f with the number of binders that have been visited. To use the substitution operation, several lemmas about the traverse function need to be proven.

Lemma 4 (Substitution ($\bar{\mathfrak{A}}$)). *If $[x \mapsto A]\Gamma_1 \vdash t : B$ and $\Gamma_2 \vdash v : A'$, with $A \leq A'$ and $\Gamma_1 + \Gamma_2 = \Gamma$, then $\Gamma \vdash \{v/x\}t : B$.*

Mechanization. Originally, the proof of this lemma was assumed to be standard [6,9]. The mechanization revealed the opposite to be the case, and the proof is by far the most complex and largest in the mechanization. One of the most complex parts of the proof is the LET case, which requires extensive reasoning about variable lifting and environment combination. Additionally, due to quasi-linear typing, environment combinations and environment subtyping, one has to consider more cases than initially expected. Although *dblib* alleviates some of the proof burden, manual work is still required. Since variables are considered values in Pat, reasoning about substituting one variable for another is required. While *dblib* provides tactics for automatically proving some of these properties, it often would not be able to do so. We are unsure as to why this is the case.

4 Related Work

Since little research exists on mailbox types, we discuss related work with a focus on mechanizations of the π -calculus and session types [12]. These are tailored to languages that make use of *communication channels* rather than mailboxes.

Ambal et al. [1] present a Rocq formalization of a higher-order π -calculus, where messages are able to carry processes. Their work focuses on the representation of binders in name restrictions of processes. The authors provide four versions of their mechanization, each with a different representation of variables. They conclude that their de Bruijn formalization is the most concise, but requires complex manipulation of de Bruijn indices and intricate proofs about renaming lemmas. This fact can also be observed in our work.

Goto et al. [11] provide a mechanization of polymorphic session types for the π -calculus in Rocq. Their mechanization uses the locally-nameless [4] approach for variables. Similar to our work, the authors first provide their definitions for processes and types independently. They are able to prove that their type system guarantees preservation and safety properties.

Thiemann [21] presents a mechanization of session types in Agda. In contrast to other mechanizations, his approach makes use of *intrinsically* typed terms, meaning that only well-typed terms can be constructed. As is common with intrinsically typed mechanizations, variables are represented by de Bruijn indices [23]. The preservation and progress properties are proven by construction. Intrinsic typing can lead to reduced proof sizes and fewer lemmas [23]. Whether such a representation would also work for Pat is worth exploring.

Tirole et al. [22] provide a mechanization of multiparty asynchronous session types in Rocq. They rely on the original definitions by Honda et al. [13]. During their mechanization they discovered flaws with the original definitions, such that type preservation does not hold in that system. They propose a new type system for which preservation and other properties hold, formally proving this claim. To manage the complexity of variables, they use the *Autosubst2* library [19], which generates Rocq code based on the de Bruijn representation. In contrast to our formalization, the authors choose to represent environments as a list of

pairs, with name and type. They claim that this representation, in particular, simplified context splitting for parallel composition.

Zalakain and Dardha [24] present a mechanization of a type system for the π -calculus in Agda. Their approach facilitates *leftover typing*, where the typing judgment includes a second leftover environment. The leftover environment contains types not used in the typing of the term, leaving them available for parallel processes. This technique eliminates the need for environment splits but would likely be difficult to use due to the need for type combination.

5 Conclusion and Future Work

In this paper, we have described the first mechanization of mailbox types and their semantics, as well as the syntax and semantics of the Pat programming language, in the Rocq Prover. We show that mechanizing Pat is nontrivial: in particular, environment combinations and subtyping pose a big challenge due to the numerous technical lemmas needed. Our work successfully provides a significant number of machine-checked proofs for various properties about mailbox types and Pat’s static semantics. During the mechanization process, we also identified and fixed multiple oversights in the original definitions and proofs.

The next step for future work is the mechanization of the operational semantics, which would allow us to formulate machine-checked proofs of properties such as preservation and progress. Initial work has shown that a primary difficulty in formalizing Pat’s operational semantics is the constant manipulation of de Bruijn indices. While tedious, previous work [24] was able to successfully formalize a similar semantics with this representation of variables. Alternatively, we could make heavier use of Rocq’s dependent types. In some mechanizations of the π -calculus [1], the type system keeps track of available channel endpoints for each process, which could potentially reduce the proofs regarding de Bruijn indices.

Acknowledgements. We thank the anonymous reviewers for their helpful comments. Fowler was supported by EPSRC grant EP/T014628/1 (STARDUST).

References

1. Ambal, G., Lenglet, S., Schmitt, A.: $\text{HO}\pi$ in Coq. *Journal of Automated Reasoning* **65**(1), 75–124 (Jan 2021). <https://doi.org/10.1007/s10817-020-09553-0>
2. Armstrong, J.: Making reliable distributed systems in the presence of software errors. Ph.D. thesis, Royal Institute of Technology, Stockholm, Sweden (2003), <https://nbn-resolving.org/urn:nbn:se:kth:diva-3658>
3. Brzozowski, J.A.: Derivatives of Regular Expressions. *J. ACM* **11**(4), 481–494 (Oct 1964). <https://doi.org/10.1145/321239.321249>
4. Charguéraud, A.: The Locally Nameless Representation. *Journal of Automated Reasoning* **49**(3), 363–408 (Oct 2012). <https://doi.org/10.1007/s10817-011-9225-2>
5. de Bruijn, N.G.: Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae (Proceedings)* **75**(5), 381–392 (Jan 1972). [https://doi.org/10.1016/1385-7258\(72\)90034-0](https://doi.org/10.1016/1385-7258(72)90034-0)
6. de'Liguoro, U., Padovani, L.: Mailbox Types for Unordered Interactions. *LIPICs, Volume 109, ECOOP 2018* **109**, 15:1–15:28 (2018). <https://doi.org/10.4230/LIPICs.ECOOP.2018.15>
7. Ennals, R., Sharp, R., Mycroft, A.: Linear Types for Packet Processing. In: Schmidt, D. (ed.) *Programming Languages and Systems*. pp. 204–218. Springer, Berlin, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24725-8_15
8. Fowler, S., Attard, D.P., Marshall, D., Gay, S.J., Trinder, P.: Special Delivery: Programming with Mailbox Types (Extended Version) (Sep 2025). <https://doi.org/10.48550/arXiv.2306.12935>
9. Fowler, S., Attard, D.P., Sowul, F., Gay, S.J., Trinder, P.: Special Delivery: Programming with Mailbox Types. *Proceedings of the ACM on Programming Languages* **7**(ICFP), 78–107 (Aug 2023). <https://doi.org/10.1145/3607832>
10. Fowler, S., Attard, D.P., Sowul, F., Gay, S.J., Trinder, P.: Special Delivery: Programming with Mailbox Types (Extended Version) (Jun 2023). <https://doi.org/10.48550/arXiv.2306.12935>
11. Goto, M., Jagadeesan, R., Jeffrey, A., Pitcher, C., Riely, J.: An extensible approach to session polymorphism. *Mathematical Structures in Computer Science* **26**(3), 465–509 (Mar 2016). <https://doi.org/10.1017/S0960129514000231>
12. Honda, K.: Types for dyadic interaction. In: Best, E. (ed.) *CONCUR'93*. pp. 509–523. Springer, Berlin, Heidelberg (1993). https://doi.org/10.1007/3-540-57208-2_35
13. Honda, K., Yoshida, N., Carbone, M.: Multiparty Asynchronous Session Types. *J. ACM* **63**(1), 9:1–9:67 (Mar 2016). <https://doi.org/10.1145/2827695>
14. Kobayashi, N.: Quasi-linear types. In: *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. pp. 29–42. POPL '99, Association for Computing Machinery, New York, NY, USA (Jan 1999). <https://doi.org/10.1145/292540.292546>
15. Leroy, X.: Formal verification of a realistic compiler. *Communications of the ACM* **52**(7), 107–115 (Jul 2009). <https://doi.org/10.1145/1538788.1538814>
16. Levy, P., Power, J., Thielecke, H.: Modelling environments in call-by-value programming languages. *Information and Computation* **185**(2), 182–210 (Sep 2003). [https://doi.org/10.1016/S0890-5401\(03\)00088-9](https://doi.org/10.1016/S0890-5401(03)00088-9)
17. Pierce, B.C.: *Types and Programming Languages*. MIT Press, Cambridge, Mass (2002)

18. Pottier, F., Orr, K.: Dbliib. <https://github.com/rocq-community/dbliib> (2021), commit: 25469872c0ba99b046f7e5b8608205eeea5ac077
19. Stark, K.: Mechanising Syntax with Binders in Coq. Ph.D. thesis, Universität des Saarlandes, Saarbrücken, Germany (2019)
20. Team, T.R.D.: The Rocq Prover. Zenodo (Apr 2025). <https://doi.org/10.5281/zenodo.15149629>
21. Thiemann, P.: Intrinsically-Typed Mechanized Semantics for Session Types. In: Proceedings of the 21st International Symposium on Principles and Practice of Declarative Programming. pp. 1–15. PPDP '19, Association for Computing Machinery, New York, NY, USA (Oct 2019). <https://doi.org/10.1145/3354166.3354184>
22. Tireore, D., Bengtson, J., Carbone, M.: Multiparty Asynchronous Session Types: A Mechanised Proof of Subject Reduction. In: Aldrich, J., Silva, A. (eds.) 39th European Conference on Object-Oriented Programming (ECOOP 2025). Leibniz International Proceedings in Informatics (LIPIcs), vol. 333, pp. 31:1–31:30. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2025). <https://doi.org/10.4230/LIPIcs.ECOOP.2025.31>
23. Wadler, P., Kokke, W., Siek, J.G.: Programming Language Foundations in Agda (Aug 2022), <https://plfa.inf.ed.ac.uk/22.08/>
24. Zalakain, U., Dardha, O.: π with Leftovers: A Mechanisation in Agda. In: Formal Techniques for Distributed Objects, Components, and Systems: 41st IFIP WG 6.1 International Conference, FORTE 2021, Held as Part of the 16th International Federated Conference on Distributed Computing Techniques, DisCoTec 2021, Valletta, Malta, June 14–18, 2021, Proceedings. pp. 157–174. Springer-Verlag, Berlin, Heidelberg (Jun 2021). https://doi.org/10.1007/978-3-030-78089-0_9

A Type Derivation Example

We show an example of how type checking works in Pat. To this end, we model a small system with two processes where one process sends a **Ping** message to the other process. As a response, a **Pong** message is sent back. The messages have signature $\mathcal{S} = \{\text{Ping} \mapsto !\langle\langle\text{Pong}\rangle\rangle, \text{Pong} \mapsto \mathbf{1}\}$, i.e. the payload of **Ping** contains a reference to a mailbox expecting a **Pong** message, and the payload of **Pong** is of the unit type.

For better readability we use explicit names as variables in this example and further syntactic sugar. We write $t_1 ; t_2$ for **let** $x = t_1$ **in** t_2 , where x has type $\mathbf{1}$ and does not occur in t_2 . For directly freeing a mailbox V , we write **free** V as a shorthand for **guard** $V : \mathbf{1}$ {**free** $\mapsto ()$ }.

```

1  def client() :  $\mathbf{1}$  {
2    let server = new in
3    spawn serverListen(server);
4    let self = new in
5    server!Ping[self];
6    guard self :  $\langle\langle\text{Pong}\rangle\rangle$  {
7      receive Ping[reply] from self  $\mapsto$ 
8      free self; ()
9    }
10 }

11 def serverListen(self :  $? \langle\langle\text{Ping}\rangle\rangle \oplus \mathbf{1}^\bullet$ ) :  $\mathbf{1}$  {
12   guard self :  $\langle\langle\text{Ping}\rangle\rangle \oplus \mathbf{1}$  {
13     receive Ping[reply] from self  $\mapsto$ 
14     reply!Pong[self];
15     free self; ()
16   }
17 }
18 }
```

The program consists of two functions: `client` and `serverListen`. Function `client` is the initial starting point, where a new mailbox name is bound to variable `server`. Then, a new process with `serverListen` is created, providing `server` as argument. The function continues by again creating a new mailbox name, binding it to `self`, and sending a **Ping** message to `server` with `self` as payload. Then, in the guard-expression, a receive guard is triggered when `self` contains a **Pong** message. The mailbox is freed and the unit value is returned.

Function `serverListen` requires an argument of type $? \langle\langle\text{Ping}\rangle\rangle \oplus \mathbf{1}^\bullet$, indicating a mailbox containing a **Ping** message or being empty. The function's body contains only a single guard-expression consisting of two guards. When a **Ping** message with some payload `reply` is received, a **Pong** message is sent to `reply`. After that, `self` is freed and the unit value is returned. The other guard is triggered when `self` does not contain any messages; in this case, the mailbox is freed.

Using the syntax we defined with de Bruijn indices and without syntactic sugar, both functions look as follows:

```

1  def client() :  $\mathbf{1}$  {
2    let new in
3    let spawn serverListen(Var 0) in
4    let new in
5    let (Var 2)!Ping[Var 0] in
6    guard Var 1 :  $\langle\langle\text{Pong}\rangle\rangle$  {
7      receive Ping  $\mapsto$ 
8      guard Var 1 :  $\mathbf{1}$  {free  $\mapsto ()$ }
9    }
10 }

11 def serverListen(self :  $? \text{Ping} \oplus \mathbf{1}^\bullet$ ) :  $\mathbf{1}$  {
12   guard Var 0 :  $\langle\langle\text{Ping}\rangle\rangle \oplus \mathbf{1}$  {
13     receive Ping  $\mapsto$ 
14     let (Var 0)!Pong[()] in
15     guard Var 2 :  $\mathbf{1}$  {free  $\mapsto ()$ }
16   }
17 }
18 }
```

Next, we present a type derivation of the `client` definition. Due to the amount of rules needed for this derivation, we focus on the subtrees with the essential

The derivation starts by applying the LET rule, where **new** from the left subterm of the let-expression is trivially typed using NEW and the empty environment. For the right subterm, another LET rule is used, creating one subtree for typing the send-term and one subtree for typing the guard. The guard is typed using another derivation \mathbf{D}_2 . To type the sending of message **Ping** successfully, we require two mailbox types for variables **Var 2** and **Var 0**. Since **Var 2** is the receiving mailbox, it must be able to receive a **Ping** message, i.e. have type $!\langle\langle\mathbf{Ping}\rangle\rangle^\circ$. Because the current environment contains the similar type $!\langle\langle\mathbf{Ping}\rangle\rangle \oplus \mathbf{1}^\bullet$, we are able to type **Var 2** with the help of subtyping, since $!\langle\langle\mathbf{Ping}\rangle\rangle \oplus \mathbf{1}^\bullet \leq !\langle\langle\mathbf{Ping}\rangle\rangle^\circ$. For **Var 0**, the signature of **Ping** tells us that **Var 0** has to have type $!\langle\langle\mathbf{Pong}\rangle\rangle^\circ$. Thus, the environment combination from the previous LET rule is $!\langle\langle\mathbf{Pong}\rangle\rangle^\circ, \mathbf{1}, !\langle\langle\mathbf{Ping}\rangle\rangle \oplus \mathbf{1}^\bullet, \perp \triangleright ?\langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}^\bullet, \perp^3 = ?\mathbf{1}^\bullet, \mathbf{1}, !\langle\langle\mathbf{Ping}\rangle\rangle \oplus \mathbf{1}^\bullet, \perp$. This introduces $\langle\langle\mathbf{Pong}\rangle\rangle$ into the environment, letting us correctly type **Var 0** with an intermediate SUB step to get rid of the unnecessary $\mathbf{1}$. Tree \mathbf{D}_2 describes the typing derivation of the guard:

$$\begin{array}{c}
\frac{}{\perp^7 \vdash () : \mathbf{1}} \text{UNIT} \\
\frac{}{\perp^7 \vdash \mathbf{free} \mapsto () : \mathbf{1} :: \mathbf{1}} \text{FREE} \\
\frac{}{\perp^7 \vdash \mathbf{free} \mapsto () : \mathbf{1} :: \mathbf{1}} \text{SINGLE} \\
\frac{}{\perp^7 \vdash \mathbf{free} \mapsto () : \mathbf{1} :: \mathbf{1}} \text{GUARD} \\
\frac{}{\perp, ?\mathbf{1}^\bullet, \perp^5 \vdash \text{Var } 1 : ?\mathbf{1}^\bullet} \text{VAR} \\
\frac{}{\mathbf{1}, ?\mathbf{1}^\bullet, \perp^5 \vdash \text{Var } 1 : ?\mathbf{1}^\bullet} \text{SUB} \\
\frac{}{\mathbf{1}, ?\mathbf{1}^\bullet, \perp^5 \vdash \text{guard Var } 1 : \mathbf{1} \{ \text{free} \mapsto () : \mathbf{1} :: ?\langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}^\bullet \}} \text{RECEIVE} \\
\frac{}{\perp^5 \vdash \text{guard Var } 1 : \mathbf{1} \{ \text{free} \mapsto () : \mathbf{1} :: ?\langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}^\bullet \}} \text{SINGLE} \\
\frac{}{\perp^5 \vdash \text{guard Var } 1 : \mathbf{1} \{ \text{free} \mapsto () : \mathbf{1} :: ?\langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}^\bullet \}} \text{GUARD} \\
\frac{}{\perp, ?\langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}^\bullet, \perp^3 \vdash \text{Var } 1 : ?\langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}^\bullet} \text{VAR} \\
\frac{}{\perp, ?\langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}^\bullet, \perp^3 \vdash \text{guard Var } 1 : \langle\langle\mathbf{Pong}\rangle\rangle \{ \text{receive Ping} \mapsto \text{guard Var } 1 : \mathbf{1} \{ \text{free} \mapsto () \} : \mathbf{1} \}} \text{SUB} \\
\frac{}{\mathbf{1}, ?\langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}^\bullet, \perp^3 \vdash \text{guard Var } 1 : \langle\langle\mathbf{Pong}\rangle\rangle \{ \text{receive Ping} \mapsto \text{guard Var } 1 : \mathbf{1} \{ \text{free} \mapsto () \} : \mathbf{1} \}} \text{SUB}
\end{array}$$

Before typing the guard, subtyping is used to remove the newly introduced, but unused $\mathbf{1}$ from the environment. Then, the GUARD rule can be applied, where the guarded mailbox name **Var 1** can be typed using the current environment. GUARD requires some pattern F in pattern normal form and a superpattern of $\langle\langle\mathbf{Pong}\rangle\rangle$. Thus, we choose $F = \langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}$, since $\langle\langle\mathbf{Pong}\rangle\rangle \sqsubseteq \langle\langle\mathbf{Pong}\rangle\rangle \odot \mathbf{1}$, which is in PNF:

$$\frac{\frac{\langle\langle \text{Pong} \rangle\rangle \setminus \text{Pong} \triangleq \mathbf{1} \quad \mathbf{1} \setminus \text{Pong} \triangleq \mathbf{0}}{\langle\langle \text{Pong} \rangle\rangle \odot \mathbf{1} \setminus \text{Pong} \triangleq \mathbf{1}} \quad \frac{\mathbf{1} \odot \mathbf{1} \oplus \langle\langle \text{Pong} \rangle\rangle \odot \mathbf{0} \cong \mathbf{1} \oplus \mathbf{0} \cong \mathbf{1}}{\langle\langle \text{Pong} \rangle\rangle \odot \mathbf{1} \models \langle\langle \text{Pong} \rangle\rangle \odot \mathbf{1}}$$

For typing the actual guard, we start with `SINGLE`, as the guard sequence only consists of a single **receive** guard. The `RECEIVE` rule adds two new types to the environment when typing the continuation term after receive message **Pong**: type $?1^\bullet$ and $\mathbf{1}$. The unit type $\mathbf{1}$ comes from the payload type associated with **Pong**, while $?1^\bullet$ is derived from the pattern $? \langle\langle \text{Pong} \rangle\rangle \odot 1^\bullet$ and describes the mailbox after receiving a **Pong** message. The last guard freeing the mailbox is typed trivially using the appropriate rules, with an additional subtyping step to remove the previously added $\mathbf{1}$ from the environment.

This example illustrates how challenging it can be to find the correct typing derivation. The main difficulty lies in combining environments. While it is relatively straightforward to construct disjoint environment combinations, the general environment combinations required by the `LET` rule are much harder to determine. One must choose environments that allow each subtree to be typed successfully while ensuring that their combination reconstructs the original environment. Another complication is that the typing rules are not deterministic. In this example, the `SUB` rule could be applied at several different points in the derivation tree. For example, in \mathbf{D}_2 the first rule used is `SUB` to get rid of $\mathbf{1}$ in the environment. An alternative is to first use `GUARD` and after that `SUB` in the left subtree when typing `Var 1`.