Speak Now

Safe Actor Programming with Multiparty Session Types

DRAFT: November 2024

SIMON FOWLER, University of Glasgow, United Kingdom
RAYMOND HU, Queen Mary University of London, UK, United Kingdom

Actor languages such as Erlang and Elixir are widely used for implementing scalable and reliable distributed applications, but the informally-specified nature of actor communication patterns leaves systems vulnerable to costly errors such as communication mismatches and deadlocks. *Multiparty session types* (MPSTs) rule out communication errors early in the development process, but until now, the nature of actor communication has made it difficult for actor languages to benefit from session types.

This paper introduces Maty, the first actor language design supporting both *static* multiparty session typing and the full power of actors taking part in *multiple sessions*. Our main insight is to enforce session typing through a flow-sensitive type-and-effect system, combined with an event-driven programming style and first-class message handlers. Using MPSTs allows us to guarantee communication safety: a process will never send or receive an unexpected message, nor will it ever get stuck waiting for a message that will never arrive.

We extend Maty to support both Erlang-style cascading failure handling and the ability to proactively switch between sessions. We implement Maty in Scala using an API generation approach, and evaluate our implementation on a series of microbenchmarks, a factory scenario, and a chat server.

1 INTRODUCTION

The infrastructure underpinning our daily lives is powered by distributed software. Unfortunately, writing distributed software is difficult: developers must reason about a host of issues such as deadlocks, failures, and adherence to complex communication protocols. Actor languages such as Erlang and Elixir, and frameworks such as Akka, are popular tools for writing scalable and resilient distributed applications. Processes in actor languages communicate using message passing rather than shared memory, and are inspired by the actor model of computation [2, 20]. Because actor communication is asynchronous and every message is stored locally to the actor that will process it, actor languages support idioms such as *supervision hierarchies* that allow a failed process to be restarted if it crashes. Erlang in particular has been used to implement real-time systems such as telephone switches [3] and powers the servers of WhatsApp, which has billions of users worldwide.

In spite of these advantages, the communication-centric nature of actor languages is not a silver bullet: it is still possible—*easy*, even—to introduce subtle bugs that can lead to errors that are difficult to detect, debug, and fix: for example, waiting for a message that will never arrive, sending a message that cannot be handled, or sending an incorrect payload.

Multiparty session types (MPSTs) allow a developer to check, at compile time, that they have correctly followed all communication protocols, in turn avoiding costly bugs manifesting themselves late in the development process. However, MPSTs were originally designed for communication channels, which poses challenges when applying them to actor languages. Existing approaches either check session typing dynamically, catching errors only at run time, or sacrifice expressiveness by emulating binary session types or limiting each actor to a single session. Furthermore, no existing work on statically-typed actors has shown how session typing can be used alongside the *let-it-crash* philosophy and supervision hierarchies that make actor languages so popular in practice.

This paper presents Maty, the first actor-based programming language fully supporting statically-checked multiparty session types and failure handling, allowing developers to benefit from both the error prevention mechanism of session types and the scalability and fault tolerance of actor languages.

Authors' addresses: Simon Fowler, University of Glasgow, United Kingdom; Raymond Hu, Queen Mary University of London, UK, United Kingdom.

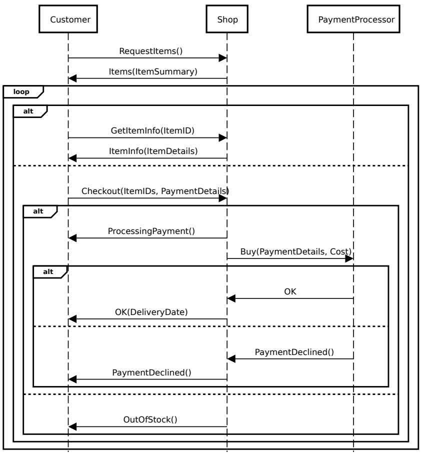

(a) Sequence diagram for Shop example

```
handle_cast({checkout, ItemIDs, PaymentDetails,
             ReplyTo}, Items) ->
  AllItemsInStock = check_stock(ItemIDs, Items),
  if
    AllItemsInStock ->
      ReplyTo ! processing_payment,
      case gen_server:call(payment_processor,
          { buy, PaymentDetails,
            cost(ItemIDs, Items) }) of
        ok ->
          NewItems =
            decrease_stock(ItemIDs, Items),
          ReplyTo ! { ok, tomorrow },
          { noreply, NewItems };
        payment_declined ->
          ReplyTo ! payment_declined,
          { noreply, Items }
      end;
    true ->
      ReplyTo ! out_of_stock,
      {noreply, Items}
  end;
```

(b) Handling the `checkout` message

Fig. 1. Protocol and implementation for Shop example

1.1 Motivating Example

Consider the following scenario, depicted in Figure 1a:

- A **Shop** can serve many **Customer**s at once.
- The **Customer** begins by requesting a list of items from the **Shop**, which sends back a list of pairs of an item's identifier and name.
- The **Customer** can then repeatedly either request full details (including description and cost) of an item, or proceed to checkout.
- To check out, the **Customer** sends their payment details and a list of item IDs to the **Shop**.
- If any items are out of stock, then the **Shop** notifies the customer who can then try again. Otherwise, the **Shop** notifies the **Customer** that it is processing the payment, and forwards the payment details and total cost to the **Payment Processor**.
- The **Payment Processor** responds to the **Shop** with whether the payment was successful.
- The **Shop** relays the result to the **Customer**, with a delivery date if the purchase was successful.
- Separately, **Staff** can also log in using a different system and adjust the stock.

Erlang applications often make use of the Erlang/OTP framework [8] which includes pre-defined *behaviours*. In Erlang we would implement the **Shop** using the `gen_server` behaviour, which encapsulates the common use case of a server that can react to asynchronous messages and synchronous calls from multiple clients while maintaining some state.

As an example, the **Shop** handles the `checkout` message via the `handle_cast` callback for asynchronous messages (Figure 1b) by firstly checking if all items are in stock. If so, then it notifies the customer with a `processing_payment` message and makes a call to the `payment_processor`. The payment processor replies either with `ok` (indicating that the payment was successful), in which case the stock is decreased; or `payment_declined`. In both cases the result is relayed to the customer.

Even though the scenario is quite small, there is a lot of room for error in the implementation: for example, forgetting the `ReplyTo ! out_of_stock` line would result in the customer waiting indefinitely, whereas a payload or arity mismatch on any of the messages would result in a runtime error.

```
(Shop interacting with customer)
Customer → Shop : requestItems() .
Shop → Customer : items([(ItemID × ItemName)]) .
μ loop.
  Customer → Shop : {
    getItemInfo(ItemID) .
      Shop → Customer : itemInfo(Description) . loop,
    checkout(([ItemID] × PaymentDetails)) .
      Shop → Customer : {
        paymentProcessing() .
          Shop → PaymentProcessor :
            buy(PaymentDetails × Price) .
          PaymentProcessor → Shop : {
            ok() .
              Shop → Customer : ok(DeliveryDate) . loop,
            paymentDeclined() .
              Shop → Customer :
                paymentDeclined() . loop
          },
        outOfStock() .
          Shop → Customer : outOfStock() . loop
      }
  }

(Shop interacting with staff)
μ loop.
  Staff → Shop : {
    addItem((Name × Description × Price × Stock)) . loop,
    removeItem(ItemID) . loop
  }
```

(a) Global types

```
(Shop interacting with customer)
ShopTy ≜
  Customer & requestItems() .
  Customer ⊕ items([(ItemID × ItemName)]) .
  μ loop.
    Customer &{
      getItemInfo(ItemID) .
        Customer ⊕ itemInfo(Description) . loop
      checkout(([ItemID] × PaymentDetails)) .
        Customer ⊕{
          paymentProcessing() .
            PaymentProcessor ⊕
              buy((PaymentDetails × Price)) .
            PaymentProcessor &{
              ok() . Customer ⊕ ok(DeliveryDate) . loop
              paymentDeclined() .
                Customer ⊕ paymentDeclined() . loop
            },
          outOfStock() .
            Customer ⊕ outOfStock() . loop
        }
    }

(Shop interacting with staff)
ShopTy2 ≜
  μ loop.
    Staff &{
      addItem((Name × Description × Price × Stock)) . loop,
      removeItem(ItemID) . loop
    }
```

(b) Local types

Fig. 2. Global and local types for shop scenario

1.2 Multiparty session types

Structured interactions, like those in our motivating example, can be captured by *multiparty session types* (MPSTs) [21]. A *global type* is a sequence of interactions between participants. Figure 2a shows the global types for our scenario: the first global type shows the interactions between the customer, shop, and payment processor, whereas the second shows the (simpler) interactions between the staff and the shop. For example, Shop ⟶ Customer : items([(ItemID × ItemName)]) . G denotes the Shop role sending an items message to the Customer role before proceeding as G. Interactions may potentially have different branches (e.g., getItemInfo and checkout).

A global type can be projected to *local types* that describe the protocol from the perspective of each participant. Figure 2b shows the corresponding local types for the Shop; selection (output) actions are denoted with ⊕, whereas offering (receive) actions are denoted with &.

MPSTs are a convenient and successful approach that allow us to statically check conformance to communication protocols. However, there are significant challenges to applying MPSTs to actor-style programming: session types focus on communication channels, whereas actors use many-sender single-receiver communication where an actor receives from an implicit mailbox (see [15] for a detailed comparison). For example: we can type a channel endpoint with (binary) session type !Int.?Int.End (send then receive an integer), but since we can only *send* to an actor reference (resp. *receive* from a mailbox), it is non-obvious how to apply session types directly.

Multiparty session actors. Neykova and Yoshida [37] introduced a programming methodology for actor programming with multiparty session types, which was later applied to Erlang [12]. The idea (Figure 3) is that each actor can be involved in *multiple sessions*, with incoming and outgoing messages passing through FSM-based monitors. However, the dynamic approach detects violations late, and monitoring incurs performance overheads. Furthermore, these works have not been formalised nor distilled into a language design, so there is no formal account of their metatheoretical properties.

There are also significant gaps between existing work on languages with statically-checked MPSTs and the actor paradigm: key to the session actor paradigm is the ability to take part in *multiple sessions*. However, most existing systems offer few guarantees when a process is involved with more than one session: without sophisticated type systems [9] or more restrictive communication topologies that enforce separation between sessions [26] we cannot rule out deadlocks.

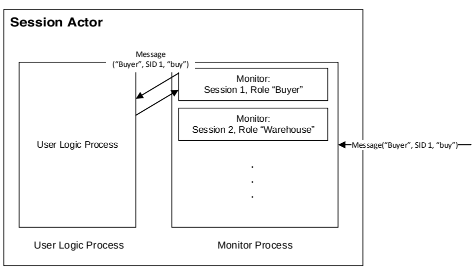

Fig. 3. Multiparty session actors as introduced by Neykova and Yoshida [37] (image from [12])

The majority of session-typed languages and frameworks do not allow a process to listen for messages on multiple sessions, in spite of this being a common pattern in concurrent programming. The notable exceptions are systems that combine session types with event-driven programming (e.g., [47, 49]), but these require a *full* inversion of control, leading to code that can be difficult to follow. These approaches are also formalised as process calculi, leaving a significant gap between the formalism and the concrete programming language design.

Key ideas. The main contribution of this work is the first *statically-typed* language design, Maty, that fully supports actor programming with multiparty session types (in contrast to more limited designs [16, 19] that only allow actors to take part in a single session). Our key insight is to combine a *flow-sensitive* effect type system [31] (similar to work on parameterised monads [4]) with *event-driven programming* and *first-class message handlers*. Actors register to participate in a session through an *access point* [17]. After session establishment, each actor can perform computations and send messages in direct style, and suspend with a message handler when it is waiting for a message from another participant.

1.3 Maty by example

Maty is a functional programming language with support for lightweight processes that communicate using session-typed message passing. In this section we will show how our Erlang shop actor can be written in Maty; the other components can be written similarly.

The entry point of our program creates two *access points* [17]: one for customer sessions, and one for staff sessions. An access point can be thought of as a "matchmaking service": different participants can register their intention to take part in a session, and the session is established once all participants are available. We spawn customer and payment processor actors (details omitted), and also a shop actor. Each access point is created by specifying the set of roles involved with the session along with their local types; ShopTy is above but we omit the other local types.

```
main ≜
 let custAP = newAP  Shop : ShopTy,
                     Customer : CustTy,
                     PaymentProcessor : PPTy
 in
 let staffAP = newAP  Shop : ShopTy2,
                      Staff : StaffTy
 in
 spawn shop(custAP, staffAP) initialState;
 spawn staff(staffAP) ();
 spawn customer(custAP) ()
```

```
registerForever(ap, role, callback) ≜
  rec install(_).
    register ap role (install ());
    callback ()

shop(custAP, staffAP) ≜
  register custAP Shop
    (registerForever(custAP, Shop, λ_. suspend itemReqHandler) ());
  register staffAP Shop
    (registerForever(staffAP, Shop, λ_. suspend staffReqHandler) ())
```

```
itemReqHandler ≜
    handler Customer {
      requestItems() ↦
        let items = get in
        Customer ! itemSummary(summary(items));
        suspend custReqHandler
    }

custReqHandler ≜
    handler Customer {
      getItemInfo(itemID) ↦
        let items = get in
        Customer ! itemInfo(lookupItem(itemID, items));
        suspend custReqHandler
      checkout((itemIDs, details)) ↦
        let items = get in
        if inStock(itemIDs, items) then
          Customer ! paymentProcessing();
          let total = cost(itemIDs, items) in
          set decreaseStock(itemIDs, items);
          PaymentProcessor ! buy((total, details));
          suspend paymentResponseHandler(itemIDs)
        else
          Customer ! outOfStock();
          suspend custReqHandler
    }

paymentResponseHandler(itemIDs) ≜
    handler PaymentProcessor {
      ok() ↦
          Customer ! ok(deliveryDate(itemIDs));
          suspend custReqHandler
      paymentDeclined() ↦
          Customer ! paymentDeclined();
          let items = get in
          set increaseStock(itemIDs, items);
          suspend custReqHandler
    }

staffReqHandler ≜
    handler Staff {
      addItem((name, description, price, stock)) ↦
        let items = get in
        set add(name, description, price, stock, items)
        suspend stockHandler
      removeItem(itemID) ↦
        let items = get in
        set remove(itemID, items);
        suspend stockHandler
    }
```

Fig. 4. Implementation of Shop message handlers in Maty

The shop definition takes the two access points and then proceeds to *register* to take part both in a session to interact with customers, and also to interact with staff. In general, by evaluating **register** V p M an actor registers with access point V to take part in the session as role p, storing computation M to be invoked when the session is established. The registerForever meta-level definition ensures that the actor re-registers whenever a session is established, meaning that the shop can accept an unlimited number of clients. After each session has been established, the session type for the shop states that it needs to receive a message from a client, so the shop suspends with *message handlers* itemReqHandler and staffReqHandler respectively. Suspending places the actor in an idle state and installs the handler to be invoked when a message arrives.

Figure 4 shows the implementation of the shop's message handlers. A message handler expecting to receive from role p has the form **handler** p $\{\ell_i(x_i) \mapsto M_i\}_{i \in I}$; each $\ell_i(x_i)$ denotes a message tag ℓ_i (for example buy or removeItem) and a variable x_i to bind the message's payload in the computation M_i. For simplicity, we use meta-level recursion as a shorthand for a (mutually)-recursive definition; we assume the usual encoding using anonymous recursive functions. The structure of the program closely mirrors that of the corresponding `gen_server` code. The main differences are:

- Communication takes place using *role names* as opposed to process IDs. This avoids the need to pass PIDs as part of messages.
- Unlike the `gen_server` code, where code for all messages must be specified in the `handle_call` and `handle_cast` functions, our structured programming model means that each handler only needs to consider messages that are relevant at that point of the session.
- The code makes use of an effectful treatment of state through the **get** and **set** constructs.

Importantly, each actor can be in *multiple sessions at once*, so can handle requests from many clients, with *all communication checked statically*.

1.4 Contributions

This paper introduces Maty: the first statically-typed actor language design with multiparty session types where each actor can be involved in multiple sessions and can gracefully handle failure. Concretely, we make three specific contributions:

(1) We introduce Maty, the first actor language design with full support for multiparty session types (§2). We show that Maty enjoys a strong metatheory including type preservation, progress, and global progress; in practice this means that Maty programs are free of communication mismatches and deadlocks (§3).
(2) We describe three extensions to Maty: state; the ability to proactively switch to another session; and support for process supervision and cascading failure as in Erlang (§4).
(3) We detail our implementation of Maty using an API generation approach in Scala (§5), and demonstrate our implementation on a real-world case study from the factory domain as well as a chat server application.

Section 6 discusses related work, and Section 7 concludes. **We will submit our implementation and examples as an artifact.**

2 MATY: A CORE ACTOR LANGUAGE WITH MULTIPARTY SESSION TYPES

In this section we introduce Maty, giving its syntax, typing rules, and semantics.

2.1 Syntax and Typing Rules

Figure 5 shows the static semantics of Maty. We omit state (as in our Shop example) from our core calculus as it is orthogonal, considering it as an extension. We let p, q range over roles, and x, y, z, f range over variables. We stratify the calculus into values V, W and computations M, N in the style of *fine-grain call-by-value* [29], with different typing judgements for each. Unlike many session type systems, we do not need linear types when typing values or computations as session typing is enforced by effect typing; our approach is inspired by that of Harvey et al. [19].

Session types. Although global types are convenient for describing protocols, we instead follow Scalas and Yoshida [44] and base our formalism around local types (*projection* of global types onto roles is standard [21, 42]; the local types resulting from a projecting a global type satisfy the properties that we will see in §3 [44]). *Selection* session types $\mathsf{p} \oplus \{\ell_i(A_i) \,.\, S_i\}_{i \in I}$ indicate that a process can choose to send a message with label ℓ_j and payload type A_j to role p, and continue as session type S_j (assuming $j \in I$). *Branching* session types $\mathsf{p}\,\&\{\ell_i(A_i) \,.\, S_i\}_{i \in I}$ indicate that a process must *receive* a message. We let $S^!$ range over selection (or *output*) session types, and let $S^?$ range over branching (or *input*) session types. Session type $\mu X.S$ indicates a recursive session type that binds variable X in S; we take an equi-recursive view of session types and identify each recursive session type with its unfolding. Finally, end denotes a session type that has finished.

Types. Base types C are standard. Since our type system enforces session typing by pre- and postconditions (c.f. parameterised monads [4]), a function type $A \xrightarrow{S,T} B$ states that the function takes an argument of type A where the current session type is S, and produces a result of type B with resulting session type T. An access point has type $\mathsf{AP}((\mathsf{p}_i : S_i)_i)$, mapping each role to a local type. Finally, a message handler has type $\mathsf{Handler}(S^?)$ where $S^?$ is an *input* session type.

Values. The value typing judgement has the form $\Gamma \vdash_\varphi V : A$ (we will return to behavioural properties φ in §3, and omit φ from the rules to avoid clutter). Typing rules for variables and constants are standard (we assume constants include at least the unit value () of type $\mathbf{1}$), and typing rules for anonymous functions and anonymous recursive functions are adapted to include

Syntax of types and type environments

$$
\begin{array}{llcl}
\text{Output session types} & S^! & ::= & \mathsf{p} \oplus \{\ell_i(A_i).S_i\}_i \\
\text{Input session types} & S^? & ::= & \mathsf{p} \,\&\{\ell_i(A_i).S_i\}_i \\
\text{Session types} & S, T & ::= & S^! \mid S^? \mid \mu X.S \\
 & & \mid & X \mid \mathsf{end}
\end{array}
\qquad
\begin{array}{llcl}
\text{Types} & A, B & ::= & C \mid A \xrightarrow{S,T} B \mid \mathsf{AP}((\mathsf{p}_i : S_i)_i) \\
 & & \mid & \mathsf{Handler}(S^?) \\
\text{Base types} & C & ::= & \mathbf{1} \mid \mathsf{Bool} \mid \mathsf{Int} \mid \cdots \\
\text{Type envs.} & \Gamma & ::= & \cdot \mid \Gamma, x : A
\end{array}
$$

Value typing $\boxed{\Gamma \vdash_\varphi V : A}$

$$
\text{TV-Var}\ \frac{x : A \in \Gamma}{\Gamma \vdash x : A}
\qquad
\text{TV-Const}\ \frac{c \text{ has base type } C}{\Gamma \vdash c : C}
\qquad
\text{TV-Lam}\ \frac{\Gamma, x : A \mid S \triangleright M : B \triangleleft T}{\Gamma \vdash \lambda x.M : A \xrightarrow{S,T} B}
$$

$$
\text{TV-Rec}\ \frac{\Gamma, x : A, f : A \xrightarrow{S,T} B \mid S \triangleright M : A \xrightarrow{S,T} B \triangleleft T}{\Gamma \vdash \mathbf{rec}\, f(x).M : A \xrightarrow{S,T} B}
$$

$$
\text{TV-Handler}\ \frac{(\Gamma, x : A_i \mid S_i \triangleright M_i : \mathbf{1} \triangleleft \mathsf{end})_i}{\Gamma \vdash \mathbf{handler}\ \mathsf{p}\ \{\ell_i(x_i) \mapsto M_i\}_i : \mathsf{Handler}(\mathsf{p}\,\&\{\ell_i(A_i).S_i\}_i)}
$$

Computation typing $\boxed{\Gamma \mid S \triangleright_\varphi M : A \triangleleft T}$

$$
\text{T-Let}\ \frac{\Gamma \mid S_1 \triangleright M : A \triangleleft S_2 \qquad \Gamma, x : A \mid S_2 \triangleright N : B \triangleleft S_3}{\Gamma \mid S_1 \triangleright \mathbf{let}\ x \Leftarrow M\ \mathbf{in}\ N : B \triangleleft S_3}
\qquad
\text{T-Return}\ \frac{\Gamma \vdash V : A}{\Gamma \mid S \triangleright \mathbf{return}\ V : A \triangleleft S}
$$

$$
\text{T-App}\ \frac{\Gamma \vdash V : A \xrightarrow{S,T} B \qquad \Gamma \vdash W : A}{\Gamma \mid S \triangleright V\ W : B \triangleleft T}
\qquad
\text{T-Spawn}\ \frac{\Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma \mid S \triangleright \mathbf{spawn}\ M : \mathbf{1} \triangleleft S}
$$

$$
\text{T-If}\ \frac{\Gamma \vdash V : \mathsf{Bool} \qquad \Gamma \mid S_1 \triangleright M : A \triangleleft S_2 \qquad \Gamma \mid S_1 \triangleright N : A \triangleleft S_2}{\Gamma \mid S_1 \triangleright \mathbf{if}\ V\ \mathbf{then}\ M\ \mathbf{else}\ N : A \triangleleft S_2}
$$

$$
\text{T-Send}\ \frac{j \in I \qquad \Gamma \vdash V : A_j}{\Gamma \mid \mathsf{p} \oplus \{\ell_i(A_i).S_i\}_{i \in I} \triangleright \mathsf{p}\,!\,\ell_j(V) : \mathbf{1} \triangleleft S_j}
\qquad
\text{T-Suspend}\ \frac{\Gamma \vdash V : \mathsf{Handler}(S^?)}{\Gamma \mid S^? \triangleright \mathbf{suspend}\ V : A \triangleleft S'}
$$

$$
\text{T-NewAP}\ \frac{\varphi \text{ is a safety property} \qquad \varphi((\mathsf{p}_i : T_i)_{i \in I})}{\Gamma \mid S \triangleright \mathbf{newAP}_{(\mathsf{p}_i : T_i)_{i \in I}} : \mathsf{AP}((\mathsf{p}_i : T_i)_{i \in I}) \triangleleft S}
$$

$$
\text{T-Register}\ \frac{j \in I \qquad \Gamma \vdash V : \mathsf{AP}((\mathsf{p}_i : T_i)_{i \in I}) \qquad \Gamma \mid T_j \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma \mid S \triangleright \mathbf{register}\ V\ \mathsf{p}_j\ M : \mathbf{1} \triangleleft S}
$$

Fig. 5. Maty Static Semantics

session pre- and postconditions. A *message handler* **handler** p $\{\ell_i(x_i) \mapsto M_i\}_i$ specifies an actor's behaviour when a message is received from role p; each clause states that when a message with label ℓ_i is received, the payload is bound to x_i in M_i. Rule TV-Handler states that the handler is typable with type $\mathsf{Handler}(\mathsf{p}\,\&\{\ell_i(A_i)\,.\,S_i\}_i)$ if each continuation M_i is typable with session precondition S_i where the environment is extended with x_i of type A_i, and all branches have the postcondition end.

Computations. The computation typing judgement has the form $\Gamma \mid S \triangleright_\varphi M : A \triangleleft T$, read as "under type environment Γ, and session precondition S, term M has type A and postcondition T". Again, φ refers to a behavioural property and will be discussed in §3.

A let-binding **let** $x \Leftarrow M$ **in** N evaluates M and binds its result to x in N, with the session postcondition from typing M used as the precondition when typing N (T-Let); note that this is the only evaluation context in the system. The **return** V expression is a trivial computation returning value V and has type A if V also has type A (T-Return). A function application $V\ W$ is typable by T-App provided that the precondition in the function type matches the current precondition, and advances the postcondition to that of the function type. Rule T-If types a conditional if its condition is of type Bool and both continuations have the same return type and postcondition.

The **spawn** M term spawns a new actor that evaluates term M; rule T-Spawn types **spawn** M with the unit type if the spawned thread M must has return type $\mathbf{1}$ and pre- and postconditions end (since the spawned computation is not yet in a session and so cannot communicate). Rule T-Send types a send computation $\mathsf{p} \,!\, \ell(V)$ if ℓ is contained within the selection session precondition, and if V has the corresponding type; the postcondition is the session continuation for the specified branch. There is *no **receive** construct*, since receiving messages is handled by the event loop. Instead, when an actor wishes to receive a message, it must *suspend* itself and install a message handler using **suspend** V. The T-Suspend rule states that **suspend** V is typable if the handler is compatible with the current session type precondition; since the computation does not return, it can be given an arbitrary return type and postcondition.

Sessions are initiated using *access points*: we create an access point for a session with roles and types $(\mathsf{p}_i : S_i)_i$ using $\mathbf{newAP}_{(\mathsf{p}_i:S_i)_i}$, which must annotated with the set of roles and local types to be involved in the session (T-NewAP). The rule ensures that the session types satisfy a *safety property*; we will describe this further in §3, but at a high level, if a set of session types is safe then the types are guaranteed never to cause a runtime type error due to a communication mismatch.

An actor can *register* to take part in a session as role p on access point V using **register** V p M; term M is a callback to be invoked once the session is established. Rule T-Register ensures that the access point must contain a session type T associated with role p, and since the initiation callback will be evaluated when the session is established, M must be typable under session type T. Since neither **newAP** nor **register** perform any communication, the session types are unaltered.

2.2 Operational semantics

Figure 6 introduces runtime syntax, structural congruence, and configuration reduction rules.

Runtime syntax. To model the concurrent behaviour of Maty processes, we require additional runtime syntax. Runtime names are identifiers for runtime entities: actor names a identify actors; session names s identify established sessions; access points p identify access points; and *initialisation tokens* ι associate registration entries in an access point with registered initialisation continuations.

We model communication and concurrency through a language of *configurations* (reminiscent of π-calculus processes). A *name restriction* $(\nu\alpha)\mathcal{C}$ binds runtime name α in configuration $\mathcal{C}$, and the right-associative parallel composition $\mathcal{C} \parallel \mathcal{D}$ denotes configurations $\mathcal{C}$ and $\mathcal{D}$ running in parallel.

An actor is represented as a 4-tuple $\langle a, \mathcal{T}, \sigma, \rho \rangle$, where $\mathcal{T}$ is a thread that can either be **idle**; a term M that is not involved in a session; or $(M)^{s[\mathsf{p}]}$ denoting that the actor is evaluating term M playing role p in session s. We say that an actor is *active* if its thread is M or $(M)^{s[\mathsf{p}]}$ (for some s, p, and M), and *idle* otherwise. A handler state σ maps endpoints to handlers, which are invoked when an incoming message is received and the actor is idle. Finally ρ is an initialisation state that maps initialisation tokens to callbacks to be invoked whenever a session is established. Our reduction rules make use of indexing notation as syntactic sugar for parallel composition: for example, $\langle a_i, \mathcal{T}_i, \sigma_i, \rho_i \rangle_{i \in 1..n}$ is syntactic sugar for the configuration $\langle a_1, \mathcal{T}_1, \sigma_1, \rho_1 \rangle \parallel \cdots \parallel \langle a_n, \mathcal{T}_n, \sigma_n, \rho_n \rangle$.

An access point $p(\chi)$ has name p and state χ, where the state maps roles to lists of initialisation tokens for actors that have registered to take part in the session. Finally, each session s is associated with a queue $s \triangleright \delta$, where δ is a list of entries $(\mathsf{p}, \mathsf{q}, \ell(V))$ denoting a message $\ell(V)$ sent from p to q.

Structural congruence and term reduction. Structural congruence is the smallest congruence relation defined by the axioms in Figure 6. As with the π-calculus, parallel composition is associative and commutative, and we have the usual scope extrusion rule; we write $\mathsf{fn}(\mathcal{C})$ to refer to the set of free names in a configuration $\mathcal{C}$. We also include a structural congruence rule on queues that allows us to reorder unrelated messages; notably this rule maintains message ordering between pairs of participants. Consequently, the session-level queue representation is isomorphic to a set of queues between each pair of roles. Term reduction $M \longrightarrow_{\mathsf{M}} N$ is standard β-reduction (omitted).

Runtime syntax

Actor names	a, b		
Session names	s		
AP names	p		
Init. tokens	ι		
Runtime names	α	::=	$a \mid s \mid p \mid \iota$
Values	V	::=	$\cdots \mid p$
Type env.	Γ	::=	$\cdots$
		$\mid$	$\Gamma, p : \mathsf{AP}((\mathsf{p}_i : S_i)_i)$
Reduction labels	l	::=	$s \mid \tau$

Configurations	$\mathcal{C}, \mathcal{D}$	::=	$(\nu\alpha)\mathcal{C} \mid \mathcal{C} \parallel \mathcal{D}$
		$\mid$	$\langle a, \mathcal{T}, \sigma, \rho \rangle \mid p(\chi) \mid s \triangleright \delta$
Message queues	δ	::=	$\epsilon \mid (\mathsf{p}, \mathsf{q}, \ell(V)) \cdot \delta$
Stored handlers	σ	::=	$\epsilon \mid \sigma, s[\mathsf{p}] \mapsto V$
Initialisation states	ρ	::=	$\epsilon \mid \rho, \iota \mapsto M$
Thread states	$\mathcal{T}$	::=	$\mathbf{idle} \mid (M)^{s[\mathsf{p}]} \mid M$
Access point states	χ	::=	$(\mathsf{p}_i \mapsto \widetilde{\iota_i})_i$
Evaluation contexts	$\mathcal{E}$	::=	$[\,] \mid \mathbf{let}\ x \Leftarrow \mathcal{E}\ \mathbf{in}\ M$
Thread contexts	$\mathcal{M}$	::=	$\mathcal{E} \mid (\mathcal{E})^{s[\mathsf{p}]}$
Top-level contexts	$\mathcal{Q}$	::=	$[\,] \mid ([\,])^{s[\mathsf{p}]}$

Structural congruence (configurations) $\boxed{\mathcal{C} \equiv \mathcal{D}}$

$$\mathcal{C} \parallel \mathcal{D} \equiv \mathcal{D} \parallel \mathcal{C} \qquad \mathcal{C} \parallel (\mathcal{D} \parallel \mathcal{D}') \equiv (\mathcal{C} \parallel \mathcal{D}) \parallel \mathcal{D}' \qquad \frac{\alpha \notin \mathsf{fn}(\mathcal{C})}{\mathcal{C} \parallel (\nu\alpha)\mathcal{D} \equiv (\nu\alpha)(\mathcal{C} \parallel \mathcal{D})} \qquad \frac{}{(\nu s)(s \triangleright \epsilon) \parallel \mathcal{C} \equiv \mathcal{C}}$$

$$\frac{\mathsf{p}_1 \neq \mathsf{p}_2 \vee \mathsf{q}_1 \neq \mathsf{q}_2}{s \triangleright \sigma_1 \cdot (\mathsf{p}_1, \mathsf{q}_1, \ell_1(V_1)) \cdot (\mathsf{p}_2, \mathsf{q}_2, \ell_2(V_2)) \cdot \sigma_2 \equiv s \triangleright \sigma_1 \cdot (\mathsf{p}_2, \mathsf{q}_2, \ell_2(V_2)) \cdot (\mathsf{p}_1, \mathsf{q}_1, \ell_1(V_1)) \cdot \sigma_2}$$

Configuration reduction $\boxed{\mathcal{C} \xrightarrow{l} \mathcal{D}}$

E-Send
$$\frac{}{\langle a, (\mathcal{E}[\mathsf{q}\,!\,\ell(V)])^{s[\mathsf{p}]}, \sigma, \rho \rangle \parallel s \triangleright \delta \xrightarrow{s} \langle a, (\mathcal{E}[\mathbf{return}\ ()])^{s[\mathsf{p}]}, \sigma, \rho \rangle \parallel s \triangleright \delta \cdot (\mathsf{p}, \mathsf{q}, \ell(V))}$$

E-React
$$\frac{(\ell(x) \mapsto M) \in \overrightarrow{H}}{\langle a, \mathbf{idle}, \sigma[s[\mathsf{p}] \mapsto \mathbf{handler}\ \mathsf{q}\ \{\overrightarrow{H}\}], \rho \rangle \parallel s \triangleright (\mathsf{q}, \mathsf{p}, \ell(V)) \cdot \delta \xrightarrow{s} \langle a, (M\{V/x\})^{s[\mathsf{p}]}, \sigma, \rho \rangle \parallel s \triangleright \delta}$$

E-Suspend
$$\frac{}{\langle a, (\mathcal{E}[\mathbf{suspend}\ V])^{s[\mathsf{p}]}, \sigma, \rho \rangle \xrightarrow{\tau} \langle a, \mathbf{idle}, \sigma[s[\mathsf{p}] \mapsto V], \rho \rangle}$$

E-Spawn
$$\frac{}{\langle a, \mathcal{M}[\mathbf{spawn}\ M], \sigma, \rho \rangle \xrightarrow{\tau} (\nu b)(\langle a, \mathcal{M}[\mathbf{return}\ ()], \sigma, \rho \rangle \parallel \langle b, M, \epsilon, \epsilon \rangle)}$$

E-Reset
$$\frac{}{\langle a, \mathcal{Q}[\mathbf{return}\ ()], \sigma, \rho \rangle \xrightarrow{\tau} \langle a, \mathbf{idle}, \sigma, \rho \rangle}$$

E-NewAP
$$\frac{p \text{ fresh}}{\langle a, \mathcal{M}[\mathbf{newAP}_{(\mathsf{p}_i : S_i)_{i \in I}}], \sigma, \rho \rangle \xrightarrow{\tau} (\nu p)(\langle a, \mathcal{M}[\mathbf{return}\ p], \sigma, \rho \rangle \parallel p((\mathsf{p}_i \mapsto \epsilon)_{i \in I}))}$$

E-Register
$$\frac{\iota \text{ fresh}}{\langle a, \mathcal{M}[\mathbf{register}\ p\ \mathsf{p}\ M], \sigma, \rho \rangle \parallel p(\chi[\mathsf{p} \mapsto \widetilde{\iota'}]) \xrightarrow{\tau} (\nu\iota)(\langle a, \mathcal{M}[\mathbf{return}\ ()], \sigma, \rho[\iota \mapsto M] \rangle \parallel p(\chi[\mathsf{p} \mapsto \widetilde{\iota'} \cup \{\iota\}]))}$$

E-Init
$$\frac{s \text{ fresh}}{(\nu\iota_{\mathsf{p}_i})_{i \in 1..n}(p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}} \cup \{\iota_{\mathsf{p}_i}\})_{i \in 1..n}) \parallel \langle a_i, \mathbf{idle}, \sigma_i, \rho_i[\iota_{\mathsf{p}_i} \mapsto M_i] \rangle_{i \in 1..n}) \xrightarrow{\tau} (\nu s)(p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}})_{i \in 1..n}) \parallel s \triangleright \epsilon \parallel \langle a_i, (M_i)^{s[\mathsf{p}_i]}, \sigma_i, \rho_i \rangle_{i \in 1..n})}$$

E-Lift
$$\frac{M \longrightarrow_{\mathsf{M}} N}{\langle a, \mathcal{M}[M], \sigma, \rho \rangle \xrightarrow{\tau} \langle a, \mathcal{M}[N], \sigma, \rho \rangle}$$

E-Nu
$$\frac{\mathcal{C} \xrightarrow{l} \mathcal{D}}{(\nu\alpha)\mathcal{C} \xrightarrow{l - \alpha} (\nu\alpha)\mathcal{D}}$$

E-Par
$$\frac{\mathcal{C} \xrightarrow{l} \mathcal{C}'}{\mathcal{C} \parallel \mathcal{D} \xrightarrow{l} \mathcal{C}' \parallel \mathcal{D}}$$

E-Struct
$$\frac{\mathcal{C} \equiv \mathcal{C}' \qquad \mathcal{C}' \xrightarrow{l} \mathcal{D}' \qquad \mathcal{D}' \equiv \mathcal{D}}{\mathcal{C} \xrightarrow{l} \mathcal{D}}$$

where $l - \alpha = \tau$ if $l = \alpha$, and l otherwise

Fig. 6. Operational semantics

Communication and concurrency. It is convenient for our metatheory to annotate each communication reduction with the name of the session in which the communication occurs, although we sometimes omit the label where it is not relevant. Rule E-Send describes a process playing role p in session s sending a message $\ell(V)$ to role q: the message is appended to the session queue and the operation reduces to **return** (). The E-React rule captures the event-driven nature of the system: if an actor is idle, has a stored handler $\ell(x) \mapsto M$ for $s[\mathsf{p}]$, and there exists a matching message in the session queue, then the message is dequeued and the message handler is activated. If an actor

Runtime types, environments, and labels

Polarised initialisation tokens	$\iota^{\pm}$	::=	$\iota^{+} \mid \iota^{-}$
Queue types	Q	::=	$\epsilon \mid (\mathtt{p}, \mathtt{q}, \ell(A)) \cdot Q$
Runtime type environments	Δ	::=	$\cdot \mid \Delta, a \mid \Delta, p \mid \Delta, \iota^{\pm} : S \mid \Delta, s[\mathtt{p}] : S \mid \Delta, s : Q$
Labels	γ	::=	$s : \mathtt{p} \uparrow \mathtt{q}{::}\ell \mid s : \mathtt{p} \downarrow \mathtt{q}{::}\ell \mid \mathsf{end}(s, \mathtt{p})$

Structural congruence (queue types) $\boxed{Q \equiv Q'}$

$$\frac{\mathtt{p}_1 \neq \mathtt{p}_2 \vee \mathtt{q}_1 \neq \mathtt{q}_2}{Q_1 \cdot (\mathtt{p}_1, \mathtt{q}_1, \ell_1(A_1)) \cdot (\mathtt{p}_2, \mathtt{q}_2, \ell_2(A_2)) \cdot Q_2 \equiv Q_1 \cdot (\mathtt{p}_2, \mathtt{q}_2, \ell_2(A_2)) \cdot (\mathtt{p}_1, \mathtt{q}_1, \ell_1(A_1)) \cdot Q_2}$$

Runtime type environment reduction $\boxed{\Delta \xrightarrow{\gamma} \Delta'}$

Lbl-Send	$\Delta, s[\mathtt{p}] : \mathtt{q} \oplus \{\ell_i(A_i).S_i\}_{i \in I}, s : Q$	$\xrightarrow{s:\mathtt{p}\uparrow\mathtt{q}::\ell_j}$	$\Delta, s[\mathtt{p}] : S_j, s : Q \cdot (\mathtt{p}, \mathtt{q}, \ell_j(A_j))$ (if $j \in I$)
Lbl-Recv	$\Delta, s[\mathtt{p}] : \mathtt{q} \& \{\ell_i(A_i).S_i\}_{i \in I}, s : (\mathtt{q}, \mathtt{p}, \ell_j(A_j)) \cdot Q$	$\xrightarrow{s:\mathtt{q}\downarrow\mathtt{p}::\ell_j}$	$\Delta, s[\mathtt{p}] : S_j, s : Q$ (if $j \in I$)
Lbl-End	$\Delta, s[\mathtt{p}] : \mathsf{end}$	$\xrightarrow{\mathsf{end}(s,\mathtt{p})}$	Δ
Lbl-Rec	$\Delta, s[\mathtt{p}] : \mu X.S$	$\xrightarrow{\gamma}$	Δ' (if $\Delta, s[\mathtt{p}] : S\{\mu X.S/X\} \xrightarrow{\gamma} \Delta'$)

Fig. 7. Labelled transition system on runtime type environments

is currently evaluating a computation in the context of a session $s[\mathtt{p}]$, rule E-Suspend evaluates **suspend** V by installing handler V for $s[\mathtt{p}]$ and returning the actor to the **idle** state.

Rule E-Spawn spawns a fresh actor with empty handler and initialisation state, and E-Reset returns an actor to the **idle** state once it has finished evaluating.

Session initialisation. Rule E-NewAp creates fresh access point name p and an access point with empty mappings for each role. Rule E-Register evaluates **register** p $\mathtt{p}$ M to register an actor to play role $\mathtt{p}$ with access point p: the rule creates an initialisation token ι, storing a mapping from ι to the callback M in the local initialisation environment, and appending ι to the participant set for $\mathtt{p}$ in p. Finally, E-Init establishes a session when idle participants are registered for all roles: in this case, the initialisation tokens are discarded; a session name restriction and empty queue is created; and each initialisation callback is invoked in the context of the newly-created session. The remaining rules are administrative.

3 METATHEORY

Following Scalas and Yoshida [44] we begin by showing a type semantics for sets of local types. Using this semantics we can define behavioural properties on types (such as *safety*, which ensures that communicated messages are always compatible; and *progress*, which ensures communication is deadlock-free). By making our runtime typing rules *parametric* in the particular behavioural property used, we can customise the property to show that behavioural properties on types give rise to corresponding guarantees about the behaviour of configurations.

Relations. We write $\mathcal{R}^?$, $\mathcal{R}^+$, and $\mathcal{R}^*$ for the reflexive, transitive, and reflexive-transitive closures of a relation $\mathcal{R}$ respectively. We write $\mathcal{R}_1\mathcal{R}_2$ for the composition of relations $\mathcal{R}_1$ and $\mathcal{R}_2$.

Runtime types and environments. Runtime environments are used to type configurations and to define behavioural properties on sets of local types. Unlike type environments Γ, runtime type environments Δ are *linear* to ensure safe use of session channel endpoints, and also to ensure that there is precisely one instance of each actor and access point. Runtime type environments can contain access point names p; *polarised* initialisation tokens $\iota^{\pm} : S$ (since each initialisation token is used twice: once in the access point and one inside an actor's initialisation environment); session channel endpoints $s[\mathtt{p}] : S$; and finally session queue types $s : Q$. Queue types mirror the structure of queue entries and are a triple $(\mathsf{p}, \mathsf{q}, \ell(A))$. We include structural congruence on queue types to match structural congruence on queues, and extend this to runtime environments.

Labelled transition system on environments. Figure 7 shows the LTS on runtime type environments. The Lbl-Send reduction gives the behaviour of an output session type interacting with a queue: supposing we send a message with some label ℓ_j from p to q, we advance the session type for p to the continuation S_j and add the message to the end of the queue. The Lbl-Recv rule handles receiving and works similarly, instead *consuming* the message from the queue. Rule Lbl-End allows us to discard a session endpoint from the environment if it does not support any further communication, and Lbl-Rec allows reduction of recursive session types by considering their unrolling. We write $\Delta \Longrightarrow \Delta'$ if $\Delta \equiv\xrightarrow{\gamma}\equiv \Delta'$ for some synchronisation label γ.

Safety property. *Safety* is the minimum property we require for type preservation: it ensures that communication does not introduce type errors. Intuitively a safety property ensures that a message received from a queue is of the expected type, thereby ruling out communication mismatches; safety properties must also hold under unfoldings of recursive session types and safety must be preserved by environment reduction.

Definition 3.1 (Safety property). φ is a *safety property* of runtime type environments Δ if:

(1) $\varphi(\Delta, s[\mathsf{p}] : \mathsf{q}\,\&\{\ell_i(A_i).S_i\}_{i\in I}, s : Q)$ with $Q \equiv (\mathsf{q}, \mathsf{p}, \ell_j(B_j)) \cdot Q'$ implies $j \in I$ and $B_j = A_j$;
(2) $\varphi(\Delta, s[\mathsf{p}] : \mu X.S)$ implies $\varphi(\Delta, s[\mathsf{p}] : S\{\mu X.S/X\})$; and
(3) $\varphi(\Delta)$ and $\Delta \Longrightarrow \Delta'$ implies $\varphi(\Delta')$.

A runtime environment is *safe*, written $\mathsf{safe}(\Delta)$, if $\varphi(\Delta)$ for a safety property φ.

We henceforth assume that all other properties are safety properties.

3.1 Runtime typing

In order to reason about the metatheory we must firstly define an extrinsic [41] type system for configurations; note that this is used only for reasoning and need not be implemented as part of a typechecker. Figure 8 shows the runtime typing rules for the system.

Runtime typing rules. The runtime typing judgement $\Gamma; \Delta \vdash_\varphi C$ can be read, "under type environment Γ and runtime type environment Δ, where the session types used in each session must satisfy behavioural property φ, configuration C is well typed". We omit φ from the rules to avoid clutter, and write $\Gamma; \Delta \vdash C$ when we wish to consider the largest safety property.

We have three rules for name restrictions: read bottom-up, T-APName adds p to both the type and runtime environments, and rule T-InitName adds tokens of both polarities to the runtime type environment. Rule T-SessionName is key to the generalised multiparty session typing approach introduced by Scalas and Yoshida [44]: the type environment Δ' consists of a set of session channel endpoints $\{s[\mathsf{p}_i]\}_i$ with session types S_{p_i}, along with a session queue $s : Q$. Environment Δ' must satisfy φ, where φ is at least a safety property.

Rule T-Par types the two parallel subconfigurations under disjoint runtime environments. Rule T-AP types an access point: it requires that the access point reference is included in Γ and through the auxiliary judgement $\{(\mathsf{p}_i : S_i)_i\}\ \Delta \vdash \chi$ ensures that each initialisation token in the access point state has a compatible type. We also require that the collection of roles that make up the access point satisfy a safety property in order to ensure that any established session is safe.

Rule T-Actor types an actor $\langle a, \mathcal{T}, \sigma, \rho\rangle$ and makes use of three auxiliary judgements. The thread state typing judgement $\Gamma; \Delta \vdash \mathcal{T}$ states that the **idle** state is always well typed (TT-Idle); a thread $(M)^{s[\mathsf{p}]}$ is well typed if given a singleton runtime environment $s[\mathsf{p}] : S$, term M is typable with session precondition S, return type $\mathbf{1}$, and post condition end; and a non-session thread M is typable

Runtime typing rules $\boxed{\Gamma;\, \Delta \vdash_{\varphi} C}$

$$\textsc{T-APName}\ \dfrac{\Gamma, p : \mathsf{AP}((\mathsf{p}_i : S_i)_i);\, \Delta, p \vdash C}{\Gamma;\, \Delta \vdash (\nu p)C} \qquad \textsc{T-InitName}\ \dfrac{\Gamma;\, \Delta, \iota^+ : S, \iota^- : S \vdash C}{\Gamma;\, \Delta \vdash (\nu \iota)C}$$

$$\textsc{T-SessionName}\ \dfrac{\Delta' = \{s[\mathsf{p}_i] : S_{\mathsf{p}_i}\}_i, s : Q \qquad \varphi(\Delta') \qquad s \notin \Delta \qquad \Gamma;\, \Delta, \Delta' \vdash C \qquad \varphi \text{ is a safety property}}{\Gamma;\, \Delta \vdash (\nu s)C}$$

$$\textsc{T-ActorName}\ \dfrac{\Gamma;\, \Delta, a \vdash C}{\Gamma;\, \Delta \vdash (\nu a)C} \qquad \textsc{T-Par}\ \dfrac{\Gamma;\, \Delta_1 \vdash C \qquad \Gamma;\, \Delta_2 \vdash \mathcal{D}}{\Gamma;\, \Delta_1, \Delta_2 \vdash C \parallel \mathcal{D}}$$

$$\textsc{T-AP}\ \dfrac{p : \mathsf{AP}((\mathsf{p}_i : S_i)_i) \in \Gamma \qquad \{(\mathsf{p}_i : S_i)_i\}\ \Delta \vdash \chi \qquad \varphi((\mathsf{p}_i : S_i)_i) \qquad \varphi \text{ is a safety property}}{\Gamma;\, \Delta, p \vdash p(\chi)}$$

$$\textsc{T-Actor}\ \dfrac{\Gamma; \Delta_1 \vdash \mathcal{T} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma;\, \Delta_1, \Delta_2, \Delta_3, a \vdash \langle a, \mathcal{T}, \sigma, \rho \rangle} \qquad \textsc{T-EmptyQueue}\ \dfrac{}{\Gamma;\, s : \epsilon \vdash s \triangleright \epsilon}$$

$$\textsc{T-ConsQueue}\ \dfrac{\Gamma \vdash V : A \qquad \Gamma;\, s : Q \vdash s \triangleright \sigma}{\Gamma;\, s : ((\mathsf{p}, \mathsf{q}, \ell(A)) \cdot Q) \vdash s \triangleright (\mathsf{p}, \mathsf{q}, \ell(V)) \cdot \sigma}$$

Access point state typing $\boxed{\{(\mathsf{p}_i : S_i)_i\}\ \Delta \vdash \chi}$

$$\textsc{TA-Empty}\ \dfrac{}{\{(\mathsf{p}_i : S_i)_i\}\ \cdot \vdash S} \qquad \textsc{TA-Entry}\ \dfrac{j \in I \qquad \{(\mathsf{p}_i : S_i)_{i \in I}\}\ \Delta \vdash \chi}{\{(\mathsf{p}_i : S_i)_{i \in I}\}\ \Delta, \widetilde{\iota^- : S_j} \vdash \chi[\mathsf{p}_j \mapsto \widetilde{\iota}]}$$

Thread state typing $\boxed{\Gamma; \Delta \vdash \mathcal{T}}$

$$\textsc{TT-Idle}\ \dfrac{}{\Gamma; \cdot \vdash \mathbf{idle}} \qquad \textsc{TT-Sess}\ \dfrac{\Gamma \mid S \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : S \vdash (M)^{s[\mathsf{p}]}} \qquad \textsc{TT-NoSess}\ \dfrac{\Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; \cdot \vdash M}$$

Handler state typing $\boxed{\Gamma; \Delta \vdash \sigma}$

$$\textsc{TH-Empty}\ \dfrac{}{\Gamma; \cdot \vdash \epsilon} \qquad \textsc{TH-Handler}\ \dfrac{\Gamma \vdash V : \mathsf{Handler}(S^?) \qquad \Gamma; \Delta \vdash \sigma}{\Gamma; \Delta, s[\mathsf{p}] : S^? \vdash \sigma[s[\mathsf{p}] \mapsto V]}$$

Initialisation state typing $\boxed{\Gamma; \Delta \vdash \rho}$

$$\textsc{TI-Empty}\ \dfrac{}{\Gamma; \cdot \vdash \epsilon} \qquad \textsc{TI-Callback}\ \dfrac{\Gamma \mid S \triangleright M : \mathbf{1} \triangleleft \mathsf{end} \qquad \Gamma; \Delta \vdash \rho}{\Gamma; \Delta, \iota^+ : S \vdash \rho[\iota \mapsto M]}$$

Fig. 8. Runtime typing

if it has session pre- and postconditions end and return type **1**. In turn this ensures that all session actions are used, or the thread suspends. The handler typing judgement ensures that the stored handlers match the types in the runtime environments, and the initialisation state typing judgement ensures that all initialisation callbacks match the session type of the initialisation token.

Finally, T-EmptyQueue and T-ConsQueue ensure that queued messages match the queue type.

3.2 Properties

With runtime typing defined, we can begin to describe the properties enjoyed by Maty.

3.2.1 Preservation. Tying is preserved by reduction; consequently we know that communication actions must match those specified by the session type. Full proofs can be found in Appendix C.

Theorem 3.2 (Preservation). *Typability is preserved by structural congruence and reduction.*

(≡) *If* $\Gamma;\, \Delta \vdash C$ *and* $C \equiv \mathcal{D}$ *then there exists some* $\Delta' \equiv \Delta$ *such that* $\Gamma;\, \Delta' \vdash \mathcal{D}$.

(⟶) *If* $\Gamma;\, \Delta \vdash C$ *with* $\mathit{safe}(\Delta)$ *and* $C \longrightarrow \mathcal{D}$, *then there exists some* Δ' *such that* $\Delta \Longrightarrow^? \Delta'$ *and* $\Gamma;\, \Delta' \vdash \mathcal{D}$.

3.2.2 Progress. The next major property we can show is *progress*, which states that if a configuration is typable using a property that guarantees session type reduction, then the configuration can either make progress, or is in a position where no actors are involved in a session and no further sessions can be established. We start by classifying a *canonical form* for configurations.

Definition 3.3 (Canonical form). A configuration C is in *canonical form* if it can be written:

$$(\nu \tilde{\iota})(\nu p_{i \in 1..l})(\nu s_{j \in 1..m})(\nu a_{k \in 1..n})(p_i(\chi_i)_{i \in 1..l} \parallel (s_j \triangleright \delta_j)_{j \in 1..m} \parallel \langle a_k, \mathcal{T}_k, \sigma_k, \rho_k \rangle_{k \in 1..n})$$

Every well typed configuration can be written in canonical form; the result follows from the structural congruence rules and Theorem 3.2.

Proposition 3.4. *If* $\Gamma; \Delta \vdash C$ *then there exists a* $\mathcal{D} \equiv C$ *where* $\mathcal{D}$ *is in canonical form.*

Next, we define *progress* on runtime environments, which is a safety property on types that ensures all sent messages are eventually received.

Definition 3.5 (Progress). A runtime environment Δ *satisfies progress*, written $\text{prog}(\Delta)$, if $\Delta \Longrightarrow^{*} \Delta' \not\Longrightarrow$ implies that $\Delta' = s : \epsilon$.

If we require the session *types* in every session to satisfy progress, the property transfers to *configurations*: a non-reducing closed configuration cannot be blocked on any session communication and so cannot contain any sessions.

Theorem 3.6 (Progress). *If* $\cdot; \cdot \vdash_{prog} C$*, then either there exists some* $\mathcal{D}$ *such that* $C \longrightarrow \mathcal{D}$*, or* C *is structurally congruent to the following canonical form:*

$$(\nu\tilde{\iota})(\nu p_{i\in 1..m})(\nu a_{j\in 1..n})(p_1(\chi_1)_{i\in 1..m} \parallel \langle a_j, \textbf{\textit{idle}}, \epsilon, \rho_j\rangle_{j\in 1..n})$$

3.2.3 Global Progress. The progress theorem shows that session typing can rule out deadlocks. In the absence of general recursion, the system in fact enjoys *global progress*: every session will be able to reduce after a finite number of steps. The restriction on general recursion aligns with the expectation that message handlers should not run indefinitely and block the event loop. Nevertheless, finite recursive behaviour can be achieved using for example structural recursion [32] or a natural number recursor as in System T (c.f. [18]). Let $\Gamma \vdash^{\mathsf{f}} V : A$, $\Gamma \mid S \triangleright^{\mathsf{f}} M : A \triangleleft T$, and $\Gamma; \Delta \vdash^{\mathsf{f}} C$ be type judgements for finite values, terms, and configurations respectively, where terms cannot contain recursive functions. Given a configuration typing derivation it is sometimes useful to annotate session name restrictions with their associated runtime environments, i.e., $(\nu s : \Delta)C$. The key *session progress* theorem shows that for every session, any reduction in its associated session typing environment can be (eventually) reflected by a session reduction in the configuration.

Definition 3.7 (Active environment / session). We say that a runtime type environment Δ is *active*, written $\text{active}(\Delta)$, if it contains at least one entry of the form $s[\mathtt{p}] : S$ where $S \neq \text{end}$.

Theorem 3.8 (Session Progress). *If* $\cdot; \cdot \vdash^{\mathsf{f}}_{prog} (\nu s : \Delta_s)C$ *where* $\text{active}(\Delta_s)$*, then* $C \xrightarrow{\tau}{}^{*} \xrightarrow{s}$.

The proof introduces a labelled transition system for computations reduction; standard techniques such as $\top\top$-lifting [30] show the existence of a finite reduction sequence to either a value or **suspend** V for some V. Global progress follows as a consequence of an operational correspondence result between the LTS and configurations, along with similar reasoning to that of Theorem 3.6.

Let us write $\text{activeSessions}(C)$ for the set of names of sessions typable under active environments. Since (by Theorem 3.2) we can always use the structural congruence rules to hoist a session name restriction to the topmost level, global progress follows as an immediate corollary of Theorem 3.8.

Corollary 3.9 (Global Progress). *If* $\cdot; \cdot \vdash^{\mathsf{f}}_{prog} C$*, then for every* $s \in \text{activeSessions}(C)$*,* $C \equiv (\nu s)\mathcal{D}$ *for some* $\mathcal{D}$*, and* $\mathcal{D} \xrightarrow{\tau}{}^{*} \xrightarrow{s}$.

4 EXTENSIONS

We can extend Maty with various extensions: actor-level state, the ability to switch between sessions (allowing a message in one session to trigger communication in another), and the ability to support cascading failures and Erlang supervision hierarchies; due to its importance for real actor programming we concentrate on the latter. State is entirely standard so we omit it here, and we give an informal description of session switching. Full details of both can be found in Appendix A.

4.1 Switching Between Sessions

Until now we have considered scenarios where an actor is involved in *independent* sessions. Suppose we want to change our shop example from Section 1 such that we maintain a long-running session with a supplier and request delivery of an item as it runs out of stock. The key difference to our original example is that the sessions are *no longer independent*: the order message to the supplier (in the Restock session) is triggered only *as a consequence* of a buy message in the session with a customer. Whereas before we only needed to suspend an actor in a *receiving* state, this workflow requires us to also suspend an actor in a *sending* state, and switch into the session at a later stage. We call this extension $\text{Maty}_{\rightleftarrows}$. Below, we can see the extension of the shop example with the ability to switch into the restocking session; the new constructs are shaded.

```
ShopRestock ≜
  μ loop.
    Supplier ⊕ order(([ItemID] × Quantity)) .
    Supplier & ordered(Quantity) . loop

custReqHandler ≜
  handler Customer {
    getItemInfo(itemID) ↦ [...]
    checkout((itemIDs, details)) ↦
      let items = get in
      if inStock(itemIDs, items) then
        [...]
      else
        Customer ! outOfStock();
        become Restock itemIDs;
        suspend? custReqHandler
  }

shop(custAP, staffAP, restockAP) ≜
  register custAP Shop
    (registerForever(custAP, Shop, λ_. suspend? itemReqHandler) ());
  register staffAP Shop
    (registerForever(staffAP, Shop, λ_. suspend? staffReqHandler) ());
  register restockAP Shop (suspend! Restock restockHandler)

restockHandler ≜ λitemIDs .
  Supplier ! order((itemIDs, 10));
  suspend? (
    handler Supplier {
      ordered(quantity) ↦
        increaseStock(itemIDs, quantity);
        suspend! Restock restockHandler})
```

The program is implicitly parameterised by a mapping from static names like Restock to pairs of session types and payload types (in our scenario, Restock maps to (ShopRestock, [ItemID]) to show that an actor can suspend when its session type is ShopRestock, and must provide a list of ItemIDs when switching back into the session). We split the **suspend** construct into **suspend**$_?$ V (to suspend awaiting an incoming message, as previously), and **suspend**$_!$ $\underline{s}$ V (to suspend session with name $\underline{s}$ given a function V, until switched into), and introduce the **become** $\underline{s}$ V construct to switch into a suspended session. We modify the shop definition to also register with the *restockAP* access point, suspending the session (in a state that is ready to send) with the restockHandler. The restockHandler takes an item ID, sends an order message to the supplier, and suspends again.

Metatheory. $\text{Maty}_{\rightleftarrows}$ satisfies preservation. Since (by design) **become** operations are dynamic and not encoded in the protocol (for example, we might wish to queue two invocations of a send-suspended session to be executed in turn), there is no type-level mechanism of guaranteeing that a send-suspended session is invoked, so $\text{Maty}_{\rightleftarrows}$ instead enjoys a weaker version of progress where non-reducing configurations can contain send-suspended sessions (see Appendix D).

4.2 Supervision & Cascading Failure

A large reason for the success of actor languages is their support for the *let-it-crash* philosophy: if an actor encounters an error then it should crash and be restarted by a *supervisor* actor. Until now we have not considered the possibility of failure. If an actor has crashed, then it cannot send any further messages, so we need some mechanism to ensure sessions do not get 'stuck' due to a failure. Our solution is based on the *affine sessions* approach [35], in particular its adaptation to the multiparty setting [19, 28] and the asynchronous formulation introduced by Fowler et al. [14].

Syntax

Types	A, B	::=	$\cdots \mid \mathsf{Pid}$	Monitored processes	ω	::=	$\widetilde{(a, M)}$
Values	V, W	::=	$\cdots \mid a$	Configurations	$\mathcal{C}, \mathcal{D}$	::=	$\cdots \mid \langle a, \mathcal{T}, \sigma, \rho, \omega \rangle$
Computations	M, N	::=	$\cdots \mid \mathbf{suspend}\ V\ M$			$\mid$	$\lightning a \mid \lightning s[\mathsf{p}] \mid \lightning \iota$
		$\mid$	$\mathbf{monitor}\ V\ M \mid \mathbf{raise}$				

Modified typing rules for computations $\boxed{\Gamma \mid S \triangleright M : A \triangleleft T}$

$$\textsc{T-Spawn}\ \frac{\Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma \mid S \triangleright \mathbf{spawn}\ M{:}\mathsf{Pid} \triangleleft S} \qquad \textsc{T-Suspend}\ \frac{\Gamma \vdash V : \mathsf{Handler}(S^?) \quad \Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma \mid S^? \triangleright \mathbf{suspend}\ V\ M{:}A \triangleleft T}$$

$$\textsc{T-Monitor}\ \frac{\Gamma \vdash V : \mathsf{Pid} \quad \Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma \mid S \triangleright \mathbf{monitor}\ V\ M{:}\mathbf{1} \triangleleft S} \qquad \textsc{T-Raise}\ \frac{}{\Gamma \mid S \triangleright \mathbf{raise}{:}A \triangleleft T}$$

Modified configuration reduction rules $\boxed{\mathcal{C} \xrightarrow{l} \mathcal{D}}$

$$\textsc{E-React}\quad \langle a, \mathbf{idle}, \sigma[s[\mathsf{p}] \mapsto (\mathbf{handler}\ \mathsf{q}\ \{\overrightarrow{H}\}, N)], \rho, \omega \rangle \parallel s \triangleright (\mathsf{q}, \mathsf{p}, \ell(V)) \cdot \delta \xrightarrow{s} \langle a, (M\{V/x\})^{s[\mathsf{p}]}, \sigma, \rho, \omega \rangle \parallel s \triangleright \delta \quad \text{if } (\ell(V) \mapsto M) \in \overrightarrow{H}$$

E-Spawn	$\langle a, \mathcal{M}[\mathbf{spawn}\ M], \sigma, \rho, \omega \rangle$	$\xrightarrow{\tau}$	$(\nu b)(\langle a, \mathcal{M}[\mathbf{return}\ b], \sigma, \rho, \omega \rangle \parallel \langle b, M, \epsilon, \epsilon, \epsilon \rangle)$
E-Suspend	$\langle a, (\mathcal{E}[\mathbf{suspend}\ V\ M])^{s[\mathsf{p}]}, \sigma, \rho, \omega \rangle$	$\xrightarrow{\tau}$	$\langle a, \mathbf{idle}, \sigma[s[\mathsf{p}] \mapsto (V, M)], \rho, \omega \rangle$
E-Monitor	$\langle a, \mathcal{M}[\mathbf{monitor}\ b\ M], \sigma, \rho, \omega \rangle$	$\xrightarrow{\tau}$	$\langle a, \mathcal{M}[\mathbf{return}\ ()], \sigma, \rho, \omega \cup \{(b, M)\} \rangle$
E-InvokeM	$\langle a, \mathbf{idle}, \sigma, \rho, \omega \cup \{(b, M)\} \rangle \parallel \lightning b$	$\xrightarrow{\tau}$	$\langle a, M, \sigma, \rho, \omega \rangle \parallel \lightning b$
E-Raise	$\langle a, \mathcal{E}[\mathbf{raise}], \sigma, \rho, \omega \rangle$	$\xrightarrow{\tau}$	$\lightning a \parallel \lightning \sigma \parallel \lightning \rho$
E-RaiseS	$\langle a, (\mathcal{E}[\mathbf{raise}])^{s[\mathsf{p}]}, \sigma, \rho, \omega \rangle$	$\xrightarrow{\tau}$	$\lightning a \parallel \lightning s[\mathsf{p}] \parallel \lightning \sigma \parallel \lightning \rho$
E-CancelMsg	$s \triangleright (\mathsf{p}, \mathsf{q}, \ell(V)) \cdot \delta \parallel \lightning s[\mathsf{q}]$	$\xrightarrow{\tau}$	$s \triangleright \delta \parallel \lightning s[\mathsf{q}]$
E-CancelAP	$(\nu \iota)(p(\chi[\mathsf{p} \mapsto \widetilde{\iota'} \cup \{\iota\}]) \parallel \lightning \iota)$	$\xrightarrow{\tau}$	$p(\chi[\mathsf{p} \mapsto \widetilde{\iota'}])$

$$\textsc{E-CancelH}\quad \langle a, \mathbf{idle}, \sigma[s[\mathsf{p}] \mapsto (\mathbf{handler}\ \mathsf{q}\ \{\overrightarrow{H}\}, M)], \rho, \omega \rangle \parallel s \triangleright \delta \parallel \lightning s[\mathsf{q}] \xrightarrow{\tau} \langle a, M, \sigma, \rho, \omega \rangle \parallel s \triangleright \delta \parallel \lightning s[\mathsf{q}] \parallel \lightning s[\mathsf{p}] \quad \text{if } \mathsf{messages}(\mathsf{q}, \mathsf{p}, \delta) = \emptyset$$

$$\text{where } \mathsf{messages}(\mathsf{p}, \mathsf{q}, \delta) = \{\ell(V) \mid (\mathsf{r}, \mathsf{s}, \ell(V)) \in \delta \wedge \mathsf{p} = \mathsf{r} \wedge \mathsf{q} = \mathsf{s}\}$$

Structural congruence $\boxed{\mathcal{C} \equiv \mathcal{D}}$

$$(\nu s)(\lightning s[\mathsf{p}_i]_{i \in 1..n} \parallel s \triangleright \epsilon) \parallel \mathcal{C} \equiv \mathcal{C}$$
$$(\nu a)(\lightning a) \parallel \mathcal{C} \equiv \mathcal{C}$$

Syntactic sugar

$$\lightning \sigma \triangleq \lightning s_1[\mathsf{p}_1] \parallel \cdots \parallel \lightning s_n[\mathsf{p}_n] \quad (\text{where } \mathrm{dom}(\sigma) = \{s_i[\mathsf{p}_i]\}_{i \in 1..n})$$
$$\lightning \rho \triangleq \lightning \iota_1 \parallel \cdots \parallel \lightning \iota_n \quad (\text{where } \mathrm{dom}(\rho) = \{\iota_i\}_{i \in 1..n})$$

Fig. 9. Maty$_\lightning$: Modified syntax and reduction rules

The key idea behind this approach is that a role can be marked as *cancelled*, meaning that it cannot take part in any future sessions. Trying to receive from a cancelled participant, when there are no remaining messages from that participant in the session queue, raises an exception; exceptions then cause an actor to crash and propagate the failure.

Figure 9 shows the additional syntax, typing rules, and reduction rules needed for supervision and cascading failure; we call this extension Maty$_\lightning$. Concretely we make actors addressable, meaning that **spawn** will return a process identifier (PID) of type Pid. We introduce two additional constructs: **monitor** V M monitors the actor referred to by PID V and installs a callback M to be evaluated should the actor crash; and **raise**, which when evaluated causes an actor to crash and cancels all the sessions in which it is involved. We also modify the **suspend** construct to take an additional computation M to be run if the sender fails and the message is never sent; a sensible piece of syntactic sugar would be **suspend** $V \triangleq$ **suspend** V **raise** to propagate the failure.

We can make our shop actor robust by using a shopSup actor that restarts it upon failure:

$$\mathsf{shopSup}(\mathit{custAP}, \mathit{staffAP}) \triangleq \mathbf{monitor}\ (\mathbf{spawn}\ \mathsf{shop}(\mathit{custAP}, \mathit{staffAP}))\ (\mathsf{shopSup}(\mathit{custAP}, \mathit{staffAP}))$$

The shopSup actor spawns a shop actor and monitors the resulting PID. Any failure of the shop actor will be detected by the shopSup which will restart the actor and monitor it again. The restarted shop actor will re-register with the access points and can then take part in subsequent sessions.

Configurations. To capture the additional runtime behaviour we need to extend the language of configurations. The actor configuration becomes $\langle a, \mathcal{T}, \sigma, \rho, \omega \rangle$, where ω pairs monitored PIDs with callbacks to be evaluated should the actor crash. We also introduce three kinds of "zapper thread", $\lightning a$, $\lightning s[\mathsf{p}]$, $\lightning \iota$ to indicate the cancellation of an actor, role, or initialisation token respectively.

Reduction rules by example. Consider the supervised Shop example after the Customer has sent a Checkout request and is waiting for a response, and where instead of suspending to handle the Checkout message, the Shop has raised an exception. The configuration representing this scenario could be written as follows, where *shop*, *cust*, and *pp* are actors playing the Shop, Customer, and PaymentProcessor in session *s* respectively, and where the *sup* actor has monitored the *shop* actor:

$$(\nu sup)(\nu shop)(\nu cust)(\nu pp)(\nu s)\left(\begin{array}{l}\langle shop, (\mathbf{raise})^{s[\mathsf{Shop}]}, \epsilon, \epsilon, \epsilon\rangle \\ \parallel \langle cust, \mathbf{idle}, s[\mathsf{Customer}] \mapsto (\mathsf{checkoutHandler}, \mathbf{raise}), \epsilon, \epsilon\rangle \\ \parallel \langle pp, \mathbf{idle}, s[\mathsf{PaymentProcessor}] \mapsto (\mathsf{buyHandler}, \mathbf{raise}), \epsilon, \epsilon\rangle \\ \parallel s \triangleright (\mathsf{Customer}, \mathsf{Shop}, \mathsf{checkout}(([123], 510))) \\ \parallel \langle sup, \mathbf{idle}, \epsilon, \epsilon, (shop, \mathsf{shopSup}(cAP, sAP))\rangle \end{array}\right)$$

For brevity we shorten Shop, Customer, and PaymentProcessor to S, C, and PP respectively. We also define configuration contexts $\mathcal{G} ::= [\,] \mid (\nu\alpha)\mathcal{G} \mid \mathcal{G} \parallel C$; and concretely let
$\mathcal{G} = (\nu sup)(\nu shop)(\nu cust)(\nu pp)(\nu s)([\,] \parallel \langle sup, \mathbf{idle}, \epsilon, \epsilon, (shop, \mathsf{shopSup}(cAP, sAP))\rangle)$.

We can now step through the reduction sequence and see how failures propagate and are handled. Since the *shop* actor is playing role $s[\mathsf{S}]$ and raising an exception, by E-RaiseS the actor is replaced with zapper threads $\lightning shop$ and $\lightning s[\mathsf{S}]$.

$$\mathcal{G}\left[\begin{array}{l}\langle shop, (\mathbf{raise})^{s[\mathsf{S}]}, \epsilon, \epsilon, \epsilon\rangle \\ \parallel \langle cust, \mathbf{idle}, s[\mathsf{C}] \mapsto (\mathsf{checkoutHandler}, \mathbf{raise}), \epsilon, \epsilon\rangle \\ \parallel \langle pp, \mathbf{idle}, s[\mathsf{PP}] \mapsto (\mathsf{buyHandler}, \mathbf{raise}), \epsilon, \epsilon\rangle \\ \parallel s \triangleright (\mathsf{C}, \mathsf{S}, \mathsf{checkout}(([123], 510)))\end{array}\right] \longrightarrow \mathcal{G}\left[\begin{array}{l}\lightning shop \parallel \lightning s[\mathsf{S}] \\ \parallel \langle cust, \mathbf{idle}, s[\mathsf{C}] \mapsto (\mathsf{checkoutHandler}, \mathbf{raise}), \epsilon, \epsilon\rangle \\ \parallel \langle pp, \mathbf{idle}, s[\mathsf{PP}] \mapsto (\mathsf{buyHandler}, \mathbf{raise}), \epsilon, \epsilon\rangle \\ \parallel s \triangleright (\mathsf{C}, \mathsf{S}, \mathsf{checkout}(([123], 510)))\end{array}\right]$$

Next, since $s[\mathsf{S}]$ has been cancelled, the checkout message can never be received and so is removed from the queue (E-CancelMsg). Similarly since both C and PP are waiting for messages from cancelled role S, they both evaluate their failure computations, **raise** (E-CancelH). In turn this results in the cancellation of the *cust* and *pp* actors, and the $s[\mathsf{C}]$ and $s[\mathsf{PP}]$ endpoints (E-RaiseS).

$$\longrightarrow^{+} \quad \mathcal{G}\left[\begin{array}{l}\lightning shop \parallel \lightning s[\mathsf{S}] \\ \parallel \langle cust, \mathbf{idle}, (\mathbf{raise})^{s[\mathsf{C}]}, \epsilon, \epsilon\rangle \\ \parallel \langle pp, \mathbf{idle}, (\mathbf{raise})^{s[\mathsf{PP}]}, \epsilon, \epsilon\rangle \\ \parallel s \triangleright \epsilon\end{array}\right] \quad \longrightarrow^{+} \quad \mathcal{G}[\ \lightning shop \parallel \lightning s[\mathsf{S}] \parallel \lightning cust \parallel \lightning s[\mathsf{C}] \parallel \lightning pp \parallel \lightning s[\mathsf{PP}] \parallel s \triangleright \epsilon\ ]$$

At this point the session has failed and can be garbage collected, leaving the supervisor actor and the zapper thread for *shop*. Since the supervisor was monitoring *shop*, which has crashed, the monitor callback is invoked (E-InvokeM) which finally re-spawns and monitors the Shop actor.

$$\longrightarrow \ (\nu shop)(\nu sup)\left(\begin{array}{l}\lightning shop \\ \parallel \langle sup, \mathsf{shopSup}(cAP, sAP), \epsilon, \epsilon, \epsilon\rangle\end{array}\right) \longrightarrow^{+} \ (\nu shop')(\nu sup)\left(\begin{array}{l}\langle shop', \mathsf{shop}(cAP, sAP), \epsilon, \epsilon, \epsilon\rangle \\ \parallel \langle sup, \mathbf{idle}, \epsilon, \epsilon, (shop', \mathsf{shopSup}(cAP, sAP))\rangle\end{array}\right)$$

Metatheory. All metatheoretical properties continue to hold in the presence of failure (see Appendix D). We need to introduce *cancellation-aware* runtime environments that record that an endpoint has been cancelled, and extend the type LTS to take failure into account; we also need to extend the configuration typing rules. A modified version of global progress holds: in every active session, a finite number of reductions will either lead to a communication action or result in all endpoints being cancelled and garbage collected.

5 IMPLEMENTATION AND EVALUATION

5.1 Implementation

Based on our formal design, we have implemented a toolchain for Maty-style event-driven actor programming in Scala. It adopts the state machine based API generation approach of Scribble [23]:

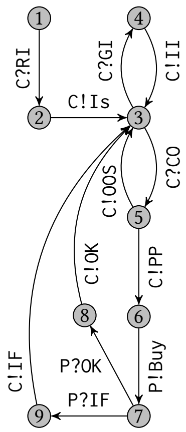

State	State types	Methods (send, suspend) or Input cases (extends state type trait)
1	S1Suspend	`suspend[D](d: D, f: (D, S1) => Done.type): Done.type`
	S1	`case class RequestItems(sid, pay, succ: S2) extends S1`
2	S2	`sendCustomerItems(pay: ItemList): S3Suspend`
3	S3Suspend	`suspend[D](d: D, f: (D, S3) => Done.type): Done.type`
	S3	`case class GetItemInfo(sid, pay, succ: S4) extends S3`
		`case class Checkout(sid, pay, succ: S5) extends S3`
4	S4	`sendCustomerItemInfo(pay): S3Suspend`
5	S5	`sendCustomerProcessingPayment(): S6`
		`sendCustomerOutOfStock(): S3Suspend`
6	S6	`sendPaymentProcessorBuy(pay): S7Suspend`
7	S7Suspend	`suspend[D](d: D, f: (D, S7) => Done.type): Done.type`
	S7	`case class OK(sid, pay, succ: S8) extends S7`
		`case class InsufficientFunds(sid, pay, succ: S9) extends S7`
8	S8	`sendCustomerOK(pay): S3Suspend`
9	S9	`sendCustomerInsufficientFunds(pay): S3Suspend`

Fig. 10. (left) CFSM for the Shop role in the Customer-Shop-PaymentProcessor protocol, and (right) summary of state types and methods in the toolchain-generated Scala API for this role.

```
// d can be used for internal, _session-specific_ actor data
def custReqHandler[T: S1orS3](d: DataS, s: T): Done.type = {
    s match {
        case c: S1 => c match {
            // pay is message payload; succ is successor state
            case RequestItems(sid, pay, succ) =>
                succ.sendCustomerItems(d.summary())
                    .suspend(d, custReqHandler[S3]) }
        case c: S3 => c match {
            case GetItemInfo(sid, pay, succ) =>
                succ.sendCustomerItemInfo(d.lookupItem(pay))
                    .suspend(d, custReqHandler[S3])
            case Checkout(sid, pay, succ) =>
                if (d.inStock(pay)) {
                    succ.sendCustomerProcessingPayment()
                        .sendPaymentProcessorBuy(d.total(pay))
                        .suspend(d, paymentResponseHandler)
```

```
            // ...continuing on from the left column
            } else {
                val sus = succ.sendCustomerOutOfStock()
                // d.staff: LOption[R1] -- this is a..
                // .."frozen" instance of state type R1
                d.staff match {
                    // R1 is the Restock protocol state type
                    case x: LSome[R1] =>
                        ibecome(d, x, restockHndlr)
                    case _: LNone =>
                        // Error handling
                        throw new RuntimeException
                }
                sus.suspend(d, custReqHandler[S3])
            }
}}}
```

Fig. 11. Example handler code from a Maty actor implemented in Scala using the toolchain-generated API

(1) The user specifies *global types* in the Scribble protocol description language [48].
(2) Our toolchain internally uses Scribble to validate global types according to the MPST-based safety conditions, *project* them to local types for each role, and construct a representation of each local type based on *communicating finite state machines* (CFSM) [6].
(3) From each CFSM, the toolchain generates a typed, *protocol-and-role-specific* API for the user to implement that role as an event-driven Maty actor in native Scala.

Typed APIs for Maty actor programming. Consider the Shop role in our running example (Fig. 2b). Fig. 10 shows the CFSM for Shop (with abbreviated message labels) and a summary of the main generated types and operations (omitting the type annotations for the `sid` and `pay` parameters, which match those in Fig. 1a). The toolchain generates Scala types for each CFSM state: non-blocking states (sends or suspends) are coloured blue, whereas blocking states (inputs) are red.

Non-blocking state types provide methods for outputs and suspend actions, with types specific to each state. The return type corresponds to the successor state type, enabling chaining of session actions: e.g., state type S2 has method `sendCustomerItems` for the transition `C!Is`. The successor state type S3Suspend includes a `suspend` method to install a handler for the input event of state 3, and to yield control back to the event loop. The `Done.type` type ensures that each handler must either complete the protocol or perform a `suspend`. Input state types are traits implemented by case classes generated for each input message. The event loop calls the user-specified handler with the corresponding case class upon each input event, with each case class carrying an instance of the successor state type. For example, `S3` (state 3) is implemented by case classes `GetItemInfo` and `Checkout` for its input transitions, which respectively carry instances of successor states `S4` and `S5`.

Fig. 11 demonstrates a user implementation of an event handler in a Shop actor, for the extended Shop+Restock protocols, using the generated APIs. This code can be compared with Fig. 4 and §4.1. The API guides the user through the protocol to construct a Maty actor with compliant handlers for every possible input event. For example, Fig. 11 handles state `S1` and could be safely supplied to the `suspend` method of `S1Suspend` immediately following a new session initiation. It *further* handles `S3` (so could also be supplied to `S3Suspend`), where the shop receives either `GetItemInfo` or `Checkout`.

For user convenience, our toolchain supports an *inline* version of **become**, as used in Fig. 11. It allows the callback for a session switching behaviour to be performed inline with the currently active handler. For this purpose, the API allows the user to "freeze" unused state type instances as a type `LOption[S]` and resume them later by an inline `ibecome`. The trade-off is this entity must be treated linearly, which our current framework checks dynamically (see below).

The runtime for our APIs executes sessions over TCP and uses the Java NIO library to run the actor event loops. It supports *fully distributed* sessions between remote Maty actors.

Discussion. Following our formal model, our generated APIs support a conventional style of actor programming where non-blocking operations are programmed in direct-style, in contrast to approaches that invert both input *and* output actions [46, 49] through the event loop.

Static Scala typing ensures that handlers safely handle all possible input events at every stage (by exhaustive matching of case classes), and that state types offer only the permitted operations at each state (by method typing). However, our API design requires *linear* usage of state type objects (e.g., `s` and `succ`). Following other works [7, 23, 39, 43, 47], we check linearity in a hybrid fashion: the `Done` return types in Fig. 10 statically require `suspend` to be invoked at least once, but our APIs rule out multiple uses *dynamically*. We exploit our formal support for failure handling (Sec. 4.2) to treat dynamic linearity errors as failures and retain safety and progress.

In summary, our toolchain enables Scala programming of Maty actors that support concurrent handling of multiple heterogeneously-typed sessions, and ensures their safe execution. A statically well-typed actor will *never* select an unavailable branch or send/receive an incompatible payload type, and an actor system will *never* become stuck due to mismatching I/O actions. As in the theory, the system will enjoy global progress provided every handler is terminating.

5.2 Evaluation

Table 1 summarises selected examples from the Savina [25] benchmark suite (lower) and larger case studies (upper); Appendix B contains sequence diagrams for the larger examples. Notably, key design features of Maty, e.g. support for handling multiple sessions per actor (mSA) and implementing multiple protocols/roles within actors (mRA), are crucial to expressing many concurrency patterns. For example, the Shop actor in both Shop examples plays the distinct Shop roles in the main Shop protocol and Shop-Staff protocol simultaneously, and handles these sessions concurrently.

The "-self" versions of Ping and Dining are versions faithful to the original Akka programs that involve internal coordination using `self ! msg` operations, i.e., self-communication. Protocols in MPST and Maty, however, focus on externally-visible I/O behaviours, and our APIs can express the formally equivalent behaviour more naturally (and simply) without needing self-communication.

Table 1. Selected case studies, examples from Savina, and key features of their Maty programs.

	MPST(s)			Maty actor programs								
	⊕/&	μ	C/P	mSA	mRA	PP	dSp	dTo	mAP	dAP	be	self
Shop (Fig. 4)	✓	✓		✓	✓	✓			✓			
ShopRestock (§ 4.1)	✓	✓		✓	✓	✓			✓		✓	
Robot [13]	✓			✓		✓	✓	(✓)				
Chat [12]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	
Ping-self [25]	✓	✓	✓	✓	✓				✓		✓	✓
Ping [25]	✓	✓										
Fib [24]				✓	✓	✓	✓	✓	✓	✓	✓	
Dining-self [25]	✓	✓	✓	✓	✓	✓	✓	(✓)	✓		✓	✓
Dining [25]	✓	✓		✓	✓	✓	✓	(✓)	✓			
Sieve [25]	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	

⊕/& = Branch type(s) μ = Recursive type(s) C/P = Concurrent/Parallel types mSA = Multiple sessions/actor
mRA = Multiple roles/actor PP = Parameterised number of actors dSp = Dynamic actor spawning
dTo = Dynamic topology mAP = Multiple APs dAP = Dynamic AP creation be = **become** self = Self communication

The (✓) distinguishes simpler forms of dynamic topologies (dTo) due to a parameterised number of clients dynamically connecting to a central server, from richer structures such as the parent-children tree topology dynamically created in Fib and the user-driven dynamic connections between clients and chat rooms in Chat; note both the latter involve dynamic access point creation (dAP).

Robot coordination. We reimplemented a real-world factory use case from Actyx AG [1], originally described by Fowler et al. [13]. In this scenario, multiple Robots access a Warehouse with a single door, with only one Robot allowed in the warehouse at a time. Concretely, each Robot actor establishes a separate session with the Door and Warehouse actors. Maty's event-driven model allows the Door and Warehouse to each be implemented as a *single* actor that can safely handle the concurrent interleavings of events across *any number* (PP, dSP) of separate Robot sessions (mSA). In contrast, standard multiparty session calculi (e.g., [9, 21]) would require us to spawn separate Door processes for each Robot client, necessitating complex state synchronisation.

Below is the straightforward user code for a Door actor to repeatedly register for an unbounded number of Robot sessions. The Door actor will safely handle all Robot sessions concurrently, coordinated by its encapsulated state (e.g., `isBusy`). The generated `ActorDoor` API provides a `register` method for the formal **register** operation, and `d1Suspend` is a user-defined handler that registers once more after every session initiation (cf. the example registerForever function in Sec. 1.3).

```
class Door(pid: Pid, port: Int, apHost: Host, apPort: Int) extends ActorDoor(pid) {
  private var isBusy = false // Shared state -- n.b. every actor is a single-threaded event loop
  def spawn(): Unit = { super.spawn(this.port); regForInit(new DataD(...)) }
  def regForInit(d: DataD) = register(this.port, apHost, apPort, d, d1Suspend)
  def d1Suspend(d: DataD, s: D1Suspend): Done.type = { regForInit(new DataD(...)); s.suspend(d, d1) }
  ... // def d1(d: DataD, s: D1): Done.type ... etc.
```

Chat server. This use case [12] involves an arbitrary number of Clients (PP) using a Registry to create new chat Rooms, and to dynamically join and leave any existing Room. We model each Client, the Registry and each Room as separate actors. Rooms are created by spawning new Room actors (dSp) with fresh access points (dAP, mAP), and we allow any Client to establish sessions with the Registry or any Room asynchronously (dTo). We decompose the Client-Registry and the Client-Room interactions into separate protocols (C/P, mAP), noting that Maty's support for event-driven processing of concurrent sessions again allows us to handle the decomposed sessions with distinct roles naturally within a single Client/Room actor (mSA, mRA). We use **become** (be) in the Room actor to broadcast chat messages to all Clients currently in that Room.

6 RELATED WORK

Event-driven session types. Several works have investigated event-driven session typing. Zhou et al. [49] introduce a multiparty session type discipline that supports statically-checked refinement types, implemented in F⋆; to avoid needing to reason about linearity, users implement callbacks for each send and receive action. This approach is used by Miu et al. [33] for session-typed web applications, and by Thiemann [46] in Agda [38]. In contrast, our approach only yields control to the event loop on actor *receives*, as in idiomatic actor programming.

Hu et al. [22] and Kouzapas et al. [27] introduced a binary session calculus with primitives used to implement an event loop; our work instead encodes an event loop directly in the semantics. Viering et al. [47] use event-driven programming in a framework for fault-tolerant session-typed distributed programming. Their model involves inversion of control on *output* as well as input events; our global progress is also stronger as it ensures possible progress on *every session in the system*. These works all concentrate on process calculi as opposed to programming language design.

Behavioural types for actor languages. Mostrous and Vasconcelos [34] were first to investigate session typing for actors, using Erlang's unique reference generation and selective receive to impose a channel-based communication model. Their approach remains unimplemented and only supports binary session types. Francalanza and Tabone [16] implement binary session typing in Elixir using pre- and post-conditions on module-level functions, but their approach can only reason about interactions between *pairs* of participants. Our approach is inspired by the model introduced by Neykova and Yoshida [37] (later implemented in Erlang [12]), but our language design supports *static* checking and is formalised. Neykova and Yoshida [36] show how causality information in global types enables *protocol-guided recovery*, leading to speedups over naïve Erlang recovery strategies. Their implementation is again dynamically-checked. Harvey et al. [19] introduce EnsembleS, which enforces session typing using a flow-sensitive effect system, focusing on supporting safe *adaptive* systems. However, each EnsembleS actor can only take part in a single session at a time.

Mailbox types [11], inspired by earlier work on typestate [10, 40], capture the expected contents of an actor mailbox as a commutative regular expression, and ensure that processes do not receive unexpected messages. Fowler et al. [13] introduce a functional language with mailbox types and show that the type discipline scales to idiomatic actor programming. Mailbox and session types both aim to ensure safe communication but address different problems: session types suit *structured* interactions among known participants, whereas mailbox types are better when participants are unknown and message ordering is unimportant. Mailbox types cannot yet handle failure.

Scalas et al. [45] introduce a behavioral typing discipline with dependent function types, allowing functions to be checked against interaction patterns written in a type-level DSL, enabling verification of properties such as liveness and termination. Their behavioural type discipline is different to session typing (e.g., supporting parameterised server interactions but not branching choice). Our session-based approach is designed for structured interactions among known participants, and it is unclear how their actor API would scale to processes handling multiple session-style interactions.

7 CONCLUSION AND FUTURE WORK

Actor languages are powerful tools for writing reliable distributed applications. This paper introduces Maty, an actor language that rules out communication mismatches and deadlocks using *multiparty session types*. Key to our approach is a novel combination of a flow-sensitive effect system and first-class message handlers; notably, Maty is *statically* typed, and actors can take part in *multiple sessions*. We have extended Maty with state and session switching, but most notably *failure handling*, allowing Maty to support Erlang-style supervision hierarchies. Finally, we have shown an implementation in Scala using an API generation approach, and evaluated our implementation on two larger case studies and a selection of examples from the Savina benchmarks.

In future it would be interesting to: implement our approach in a static typing tool for Erlang; to investigate path-dependent types in our implementation; and investigate how the approach could be integrated with more sophisticated failure handling mechanisms (e.g., [36]).

REFERENCES

[1] 2023. *Actyx AG.* https://actyx.io

[2] Gul A. Agha. 1990. *ACTORS - a model of concurrent computation in distributed systems.* MIT Press.

[3] Joe Armstrong. 2003. *Making reliable distributed systems in the presence of software errors.* Ph. D. Dissertation. Royal Institute of Technology, Stockholm, Sweden.

[4] Robert Atkey. 2009. Parameterised notions of computation. *J. Funct. Program.* 19, 3-4 (2009), 335–376.

[5] Adam D. Barwell, Alceste Scalas, Nobuko Yoshida, and Fangyi Zhou. 2022. Generalised Multiparty Session Types with Crash-Stop Failures. In *CONCUR (LIPIcs, Vol. 243).* Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 35:1–35:25.

[6] Daniel Brand and Pitro Zafiropulo. 1983. On Communicating Finite-State Machines. *J. ACM* 30, 2 (apr 1983), 323?342. https://doi.org/10.1145/322374.322380

[7] David Castro-Perez, Raymond Hu, Sung-Shik Jongmans, Nicholas Ng, and Nobuko Yoshida. 2019. Distributed programming using role-parametric session types in go: statically-typed endpoint APIs for dynamically-instantiated communication structures. *Proc. ACM Program. Lang.* 3, POPL (2019), 29:1–29:30. https://doi.org/10.1145/3290342

[8] Francesco Cesarini and Steve Vinoski. 2016. *Designing for Scalability with Erlang/OTP.* O'Reilly Media, Inc.

[9] Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. 2016. Global progress for dynamically interleaved multiparty sessions. *Math. Struct. Comput. Sci.* 26, 2 (2016), 238–302.

[10] Silvia Crafa and Luca Padovani. 2017. The Chemical Approach to Typestate-Oriented Programming. *ACM Trans. Program. Lang. Syst.* 39, 3 (2017), 13:1–13:45.

[11] Ugo de'Liguoro and Luca Padovani. 2018. Mailbox Types for Unordered Interactions. In *ECOOP (LIPIcs, Vol. 109).* Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 15:1–15:28.

[12] Simon Fowler. 2016. An Erlang Implementation of Multiparty Session Actors. In *ICE (EPTCS, Vol. 223).* 36–50.

[13] Simon Fowler, Duncan Paul Attard, Franciszek Sowul, Simon J. Gay, and Phil Trinder. 2023. Special Delivery: Programming with Mailbox Types. *Proc. ACM Program. Lang.* 7, ICFP (2023), 78–107.

[14] Simon Fowler, Sam Lindley, J. Garrett Morris, and Sára Decova. 2019. Exceptional asynchronous session types: session types without tiers. *Proc. ACM Program. Lang.* 3, POPL (2019), 28:1–28:29.

[15] Simon Fowler, Sam Lindley, and Philip Wadler. 2017. Mixing Metaphors: Actors as Channels and Channels as Actors. In *ECOOP (LIPIcs, Vol. 74).* Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 11:1–11:28.

[16] Adrian Francalanza and Gerard Tabone. 2023. ElixirST: A session-based type system for Elixir modules. *J. Log. Algebraic Methods Program.* 135 (2023), 100891.

[17] Simon J. Gay and Vasco Thudichum Vasconcelos. 2010. Linear type theory for asynchronous session types. *J. Funct. Program.* 20, 1 (2010), 19–50.

[18] Robert Harper. 2016. *Practical Foundations for Programming Languages (2nd. Ed.).* Cambridge University Press.

[19] Paul Harvey, Simon Fowler, Ornela Dardha, and Simon J. Gay. 2021. Multiparty Session Types for Safe Runtime Adaptation in an Actor Language. In *ECOOP (LIPIcs, Vol. 194).* Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 10:1–10:30.

[20] Carl Hewitt, Peter Boehler Bishop, and Richard Steiger. 1973. A Universal Modular ACTOR Formalism for Artificial Intelligence. In *IJCAI.* William Kaufmann, 235–245.

[21] Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008. Multiparty asynchronous session types. In *POPL.* ACM, 273–284.

[22] Raymond Hu, Dimitrios Kouzapas, Olivier Pernet, Nobuko Yoshida, and Kohei Honda. 2010. Type-Safe Eventful Sessions in Java. In *ECOOP 2010 - Object-Oriented Programming, 24th European Conference, Maribor, Slovenia, June 21-25, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6183)*, Theo D'Hondt (Ed.). Springer, 329–353. https://doi.org/10.1007/978-3-642-14107-2_16

[23] Raymond Hu and Nobuko Yoshida. 2016. Hybrid Session Verification Through Endpoint API Generation. In *FASE (Lecture Notes in Computer Science, Vol. 9633).* Springer, 401–418.

[24] Shams Imam. [n. d.]. Savina Actor Benchmark Suite. https://github.com/shamsimam/savina. Accessed: 2024-11-13.

[25] Shams Mahmood Imam and Vivek Sarkar. 2014. Savina - An Actor Benchmark Suite: Enabling Empirical Evaluation of Actor Libraries. In *AGERE!@SPLASH.* ACM, 67–80.

[26] Jules Jacobs, Stephanie Balzer, and Robbert Krebbers. 2022. Multiparty GV: functional multiparty session types with certified deadlock freedom. *Proc. ACM Program. Lang.* 6, ICFP (2022), 466–495.

[27] Dimitrios Kouzapas, Nobuko Yoshida, Raymond Hu, and Kohei Honda. 2016. On asynchronous eventful session semantics. *Math. Struct. Comput. Sci.* 26, 2 (2016), 303–364.

[28] Nicolas Lagaillardie, Rumyana Neykova, and Nobuko Yoshida. 2022. Stay Safe Under Panic: Affine Rust Programming with Multiparty Session Types. In *ECOOP (LIPIcs, Vol. 222)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 4:1–4:29.

[29] Paul Blain Levy, John Power, and Hayo Thielecke. 2003. Modelling environments in call-by-value programming languages. *Information and Computation* 185, 2 (2003), 182–210.

[30] Sam Lindley and Ian Stark. 2005. Reducibility and TT-Lifting for Computation Types. In *TLCA (Lecture Notes in Computer Science, Vol. 3461)*. Springer, 262–277.

[31] Daniel Marino and Todd D. Millstein. 2009. A generic type-and-effect system. In *TLDI*. ACM, 39–50.

[32] Erik Meijer, Maarten M. Fokkinga, and Ross Paterson. 1991. Functional Programming with Bananas, Lenses, Envelopes and Barbed Wire. In *FPCA (Lecture Notes in Computer Science, Vol. 523)*. Springer, 124–144.

[33] Anson Miu, Francisco Ferreira, Nobuko Yoshida, and Fangyi Zhou. 2021. Communication-safe web programming in TypeScript with routed multiparty session types. In *CC*. ACM, 94–106.

[34] Dimitris Mostrous and Vasco Thudichum Vasconcelos. 2011. Session Typing for a Featherweight Erlang. In *COORDINATION (Lecture Notes in Computer Science, Vol. 6721)*. Springer, 95–109.

[35] Dimitris Mostrous and Vasco T. Vasconcelos. 2018. Affine Sessions. *Log. Methods Comput. Sci.* 14, 4 (2018).

[36] Rumyana Neykova and Nobuko Yoshida. 2017. Let it recover: multiparty protocol-induced recovery. In *CC*. ACM, 98–108.

[37] Rumyana Neykova and Nobuko Yoshida. 2017. Multiparty Session Actors. *Log. Methods Comput. Sci.* 13, 1 (2017).

[38] Ulf Norell. 2008. Dependently Typed Programming in Agda. In *Advanced Functional Programming (Lecture Notes in Computer Science, Vol. 5832)*. Springer, 230–266.

[39] Luca Padovani. 2017. A simple library implementation of binary sessions. *J. Funct. Program.* 27 (2017), e4. https://doi.org/10.1017/S0956796816000289

[40] Luca Padovani. 2018. Deadlock-Free Typestate-Oriented Programming. *Art Sci. Eng. Program.* 2, 3 (2018), 15.

[41] John C. Reynolds. 2000. *The Meaning of Types—From Intrinsic to Extrinsic Semantics*. Technical Report RS-00-32. BRICS.

[42] Alceste Scalas, Ornela Dardha, Raymond Hu, and Nobuko Yoshida. 2017. A Linear Decomposition of Multiparty Sessions for Safe Distributed Programming. In *ECOOP (LIPIcs, Vol. 74)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 24:1–24:31.

[43] Alceste Scalas and Nobuko Yoshida. 2016. Lightweight Session Programming in Scala. In *30th European Conference on Object-Oriented Programming, ECOOP 2016, July 18-22, 2016, Rome, Italy (LIPIcs, Vol. 56)*, Shriram Krishnamurthi and Benjamin S. Lerner (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 21:1–21:28. https://doi.org/10.4230/LIPICS.ECOOP.2016.21

[44] Alceste Scalas and Nobuko Yoshida. 2019. Less is more: multiparty session types revisited. *Proc. ACM Program. Lang.* 3, POPL (2019), 30:1–30:29.

[45] Alceste Scalas, Nobuko Yoshida, and Elias Benussi. 2019. Verifying message-passing programs with dependent behavioural types. In *PLDI*. ACM, 502–516.

[46] Peter Thiemann. 2023. Intrinsically Typed Sessions with Callbacks (Functional Pearl). *Proc. ACM Program. Lang.* 7, ICFP (2023), 711–739.

[47] Malte Viering, Raymond Hu, Patrick Eugster, and Lukasz Ziarek. 2021. A multiparty session typing discipline for fault-tolerant event-driven distributed programming. *Proc. ACM Program. Lang.* 5, OOPSLA (2021), 1–30.

[48] Nobuko Yoshida, Raymond Hu, Rumyana Neykova, and Nicholas Ng. 2013. The Scribble Protocol Language. In *TGC (Lecture Notes in Computer Science, Vol. 8358)*. Springer, 22–41.

[49] Fangyi Zhou, Francisco Ferreira, Raymond Hu, Rumyana Neykova, and Nobuko Yoshida. 2020. Statically verified refinements for multiparty protocols. *Proc. ACM Program. Lang.* 4, OOPSLA (2020), 148:1–148:30.

Appendices

APPENDIX CONTENTS

A Details of State and Switching Extensions 24
A.1 State 24
A.2 Session switching 24
B Details of Case Study Protocols 26
B.1 Robots 26
B.2 Chat Server 27
C Omitted Definitions and Proofs 28
C.1 Omitted Definitions 28
C.2 Preservation 28
C.3 Progress 36
C.4 Global Progress 37
D Proofs for Section 4 41
D.1 $\text{Maty}_{\rightleftarrows}$ 41
D.2 $\text{Maty}_{\lightning}$ 43

A DETAILS OF STATE AND SWITCHING EXTENSIONS

A.1 State

Figure 12 shows the full extension of Maty with state. The typing rules for functions and function applications are modified in order to record the type of the state on which the actor operates. We also modify the configuration typing rules to account for the stored state, and ensure that the currently-executing thread operates on a compatible state type.

Extended syntax

$$
\begin{array}{llcl}
\text{Types} & A, B & ::= & \cdots \mid A_1 \xrightarrow[B]{S,T} A_2 \\
\text{Computations} & M, N & ::= & \cdots \mid \textbf{get} \mid \textbf{set}\ V \mid \textbf{spawn}\ M\ V \\
\text{Configurations} & \mathcal{C}, \mathcal{D} & ::= & \langle a, \mathcal{T}, \sigma, \rho, V \rangle
\end{array}
$$

Modified value and computation typing rules $\boxed{\Gamma \vdash V : A}$ $\boxed{\Gamma \mid A \mid S \triangleright M : A \triangleleft T}$

$$
\textsc{T-Abs}\quad \frac{\Gamma, x : A_1 \mid B \mid S \triangleright M : A_2 \triangleleft T}{\Gamma \vdash \lambda x. M : A_1 \xrightarrow[B]{S,T} A_2}
\qquad
\textsc{T-App}\quad \frac{\Gamma \vdash V : A_1 \xrightarrow[B]{S,T} A_2 \qquad \Gamma \vdash W : A_1}{\Gamma \mid B \mid S \triangleright V\ W : A_2 \triangleleft T}
$$

$$
\textsc{T-Get}\quad \frac{}{\Gamma \mid A \mid S \triangleright \textbf{get} : A \triangleleft T}
\qquad
\textsc{T-Set}\quad \frac{\Gamma \vdash V : A}{\Gamma \mid A \mid S \triangleright \textbf{set}\ V : \mathbf{1} \triangleleft S}
\qquad
\textsc{T-Spawn}\quad \frac{\Gamma \mid B \mid T \triangleright M : \mathbf{1} \triangleleft \mathsf{end} \qquad \Gamma \vdash V : B}{\Gamma \mid A \mid S \triangleright \textbf{spawn}\ M\ V : \mathbf{1} \triangleleft S}
$$

Modified configuration typing rules $\boxed{\{A\}\ \Gamma; \Delta \vdash \mathcal{T}}$ $\boxed{\Gamma; \Delta \vdash \mathcal{C}}$

$$
\textsc{TT-Idle}\quad \frac{}{\{A\}\ \Gamma; \cdot \vdash \textbf{idle}}
\qquad
\textsc{TT-NoSess}\quad \frac{\Gamma \mid A \mid S \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\{A\}\ \Gamma; s[\mathsf{p}] : S \vdash (M)^{s[\mathsf{p}]}}
\qquad
\textsc{TT-NoSess}\quad \frac{\Gamma \mid A \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\{A\}\ \Gamma; \cdot \vdash M}
$$

$$
\textsc{T-Actor}\quad \frac{\{A\}\ \Gamma; \Delta_1 \vdash \mathcal{T} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \Gamma; \Delta_3 \vdash \rho \qquad \Gamma \vdash V : A}{\Gamma; \Delta_1, \Delta_2, \Delta_3 \vdash \langle a, \mathcal{T}, \sigma, \rho, V \rangle}
$$

Modified reduction rules

$$
\begin{array}{llcl}
\textsc{E-Spawn} & \langle a, \mathcal{M}[\textbf{spawn}\ M\ W], \sigma, \rho, V \rangle & \longrightarrow & \langle a, \mathcal{M}[\textbf{return}\ ()], \sigma, \rho, V \rangle \parallel \langle b, M, \epsilon, \epsilon, W \rangle \\
\textsc{E-Get} & \langle a, \mathcal{M}[\textbf{get}], \sigma, \rho, V \rangle & \longrightarrow & \langle a, \mathcal{M}[\textbf{return}\ V], \sigma, \rho, V \rangle \\
\textsc{E-Set} & \langle a, \mathcal{M}[\textbf{set}\ W], \sigma, \rho, V \rangle & \longrightarrow & \langle a, \mathcal{M}[\textbf{return}\ ()], \sigma, \rho, W \rangle
\end{array}
$$

Fig. 12. Extension of Maty with state

A.2 Session switching

Our extension to allow session switching is shown in Figure 13. We introduce a set of distinguished *session identifiers* $\underline{s}$; each session identifier is associated with a local type and a payload in an environment Σ, i.e., for each $\underline{s}$ we have $\Sigma(\underline{s}) = (S^!, A)$ for some $S^!, A$. We then split the **suspend** construct into two: $\textbf{suspend}_?\ V$ (which, as before, installs a message handler and suspends an actor) and $\textbf{suspend}_!\ \underline{s}\ V$, which suspends a session in a *send* state, installing a function taking a payload of the given type. Finally we introduce a **become** $\underline{s}\ V$ construct that queues a request for the event loop to invoke $\underline{s}$ next time the actor is available.

Metatheory. As would be expected, $\text{Maty}_{\rightleftarrows}$ satisfies preservation.

THEOREM A.1 (PRESERVATION). *Preservation (as defined in Theorem 3.2) continues to hold in* $\text{Maty}_{\rightleftarrows}$.

However, since (by design) **become** operations are dynamic and not encoded in the protocol (for example, we might wish to queue two invocations of a send-suspended session to be executed

Modified syntax

Session names	$\underline{s}, \underline{t}$		
Computations	M, N	$::=$	$\cdots \mid \textbf{suspend}_{!}\, \underline{s}\, V \mid \textbf{suspend}_{?}\, V \mid \textbf{become}\, \underline{s}\, V$
Send-suspended sessions	D	$::=$	$(s[\mathrm{p}], V)$
Handler state	σ	$::=$	$\epsilon \mid \sigma, s[\mathrm{p}] \mapsto V \mid \sigma, \underline{s} \mapsto \overrightarrow{D}$
Switch request queue	θ	$::=$	$\epsilon \mid \theta \cdot (\underline{s}, V)$
Configurations	$\mathcal{C}, \mathcal{D}$	$::=$	$\cdots \mid \langle a, \mathcal{T}, \sigma, \rho, \theta \rangle$

Modified term typing rules $\boxed{\Gamma \mid S \triangleright M : A \triangleleft T}$

T-Suspend?
$$\frac{\Gamma \vdash V : \mathsf{Handler}(S^{?})}{\Gamma \mid S^{?} \triangleright \textbf{suspend}_{?}\, V : A \triangleleft T}$$

T-Suspend!
$$\frac{\Sigma(\underline{s}) = (S^{!}, A) \qquad \Gamma \vdash V : A \xrightarrow{S^{!}, \mathsf{end}} \mathbf{1}}{\Gamma \mid S^{!} \triangleright \textbf{suspend}_{!}\, \underline{s}\, V : B \triangleleft T}$$

T-Become
$$\frac{\Sigma(\underline{s}) = (T, A) \qquad \Gamma \vdash V : A}{\Gamma \mid S \triangleright \textbf{become}\, \underline{s}\, V : \mathbf{1} \triangleleft S}$$

Modified configuration typing rules $\boxed{\Gamma; \Delta \vdash \mathcal{C}}$ $\boxed{\Gamma; \Delta \vdash \sigma}$ $\boxed{\Gamma \vdash \theta}$

T-Actor
$$\frac{\Gamma; \Delta_1 \vdash \mathcal{T} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \Gamma; \Delta_3 \vdash \rho \qquad \Gamma \vdash \theta}{\Gamma; \Delta_1, \Delta_2, \Delta_3, a \vdash \langle a, \mathcal{T}, \sigma, \rho, \theta \rangle}$$

TH-SendHandler
$$\frac{\Gamma; \Delta \vdash \sigma \qquad \Sigma(\underline{s}) = (S^{!}, A) \qquad (\Gamma \vdash V_i : A \xrightarrow{S^{!}, \mathsf{end}} \mathbf{1})_i}{\Gamma; \Delta, (s_i[\mathrm{p}_i] : S^{!})_i \vdash \sigma, \underline{s} \mapsto (s_i[\mathrm{p}_i], V_i)_i}$$

TR-Empty
$$\frac{}{\Gamma \vdash \epsilon}$$

TR-Request
$$\frac{\Gamma \vdash \theta \qquad \Sigma(\underline{s}) = (S^{!}, A) \qquad \Gamma \vdash V : A}{\Gamma \vdash \theta \cdot (\underline{s}, V)}$$

Modified reduction rules $\boxed{\mathcal{C} \longrightarrow \mathcal{D}}$

E-Suspend!-1	$\langle a, (\mathcal{E}[\textbf{suspend}_{!}\, \underline{s}\, V])^{s[\mathrm{p}]}, \sigma, \rho, \theta \rangle$	$\xrightarrow{\tau}$	$\langle a, \textbf{idle}, \sigma[\underline{s} \mapsto (s[\mathrm{p}], V)], \rho, \theta \rangle$	$(\underline{s} \notin \mathsf{dom}(\sigma))$
E-Suspend!-2	$\langle a, (\mathcal{E}[\textbf{suspend}_{!}\, \underline{s}\, V])^{s[\mathrm{p}]}, \sigma[\underline{s} \mapsto \overrightarrow{D}], \rho, \theta \rangle$	$\xrightarrow{\tau}$	$\langle a, \textbf{idle}, \sigma[\underline{s} \mapsto \overrightarrow{D} \cdot (s[\mathrm{p}], V)], \rho, \theta \rangle$	
E-Become	$\langle a, \mathcal{M}[\textbf{become}\, \underline{s}\, V], \sigma, \rho, \theta \rangle$	$\xrightarrow{\tau}$	$\langle a, \mathcal{M}[\textbf{return}\ ()], \sigma, \rho, \theta \cdot (\underline{s}, V) \rangle$	
E-Activate	$\langle a, \textbf{idle}, \sigma[\underline{s} \mapsto (s[\mathrm{p}], V) \cdot \overrightarrow{D}], \rho, (\underline{s}, W) \cdot \theta \rangle$	$\xrightarrow{\tau}$	$\langle a, (V\ W)^{s[\mathrm{p}]}, \sigma[\underline{s} \mapsto \overrightarrow{D}], \rho, \theta \rangle$	

Fig. 13. $\mathsf{Maty}_{\rightleftarrows}$: Modified syntax, typing, and reduction rules

in turn), there is no type-level mechanism of guaranteeing that a send-suspended session is ever invoked. Although all threads can reduce as before, $\mathsf{Maty}_{\rightleftarrows}$ satisfies a weaker version of progress where non-reducing configurations can contain send-suspended sessions.

Theorem A.2 (Progress ($\mathsf{Maty}_{\rightleftarrows}$)). *If $\cdot; \cdot \vdash_{prog} \mathcal{C}$, then either there exists some $\mathcal{D}$ such that $\mathcal{C} \longrightarrow \mathcal{D}$, or $\mathcal{C}$ is structurally congruent to the following canonical form:*

$$(\nu \tilde{\imath})(\nu p_{i \in 1..l})(\nu s_{j \in 1..m})(\nu a_{k \in 1..n})(p_i(\chi_i)_{i \in 1..l} \parallel (s_j \triangleright \delta_j)_{j \in 1..m} \parallel \langle a_k, \textbf{idle}, \sigma_k, \rho_k, \theta_k \rangle_{k \in 1..n})$$

where for each session s_j there exists some mapping $s_j[\mathrm{p}] \mapsto (\underline{s}, V)$ (for some role p, static session name $\underline{s}$, and callback V) contained in some σ_k where θ_k does not contain any requests for $\underline{s}$.

B DETAILS OF CASE STUDY PROTOCOLS

In this section we detail the protocols and sequence diagrams for the two case studies.

B.1 Robots

The robots protocol can be found below, both as a Scribble global type and a sequence diagram. Role `R` stands for Robot, `D` stands for Door, and `W` stands for Warehouse.

```
global protocol Robot(role R, role D, role W) {
  Want(PartNum) from R to D;
  choice at D {
    Busy() from D to R;
    Cancel() from D to W;
  } or {
    GoIn() from D to R;
    Prepare(PartNum) from D to W;
    Inside() from R to D;
    Prepared() from W to D;
    Deliver() from D to W;
    Delivered() from W to R;
    PartTaken() from R to W;
    WantLeave() from R to D;
    GoOut() from D to R;
    Outside() from R to D;
    TableIdle() from W to D;
  }
}
```

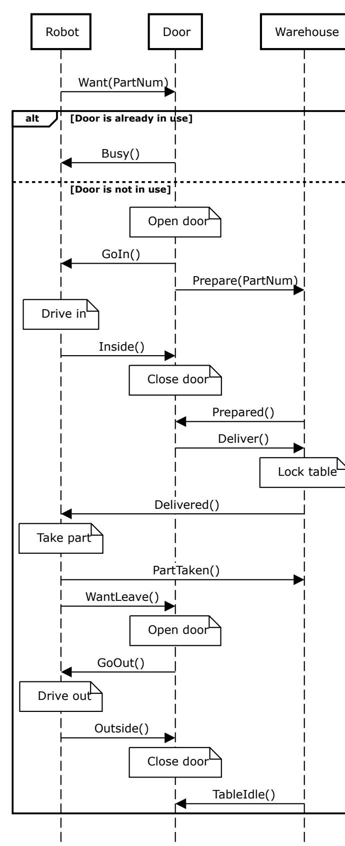

B.2 Chat Server

```
global protocol ChatServer(role C, role S) {
  choice at C {
    LookupRoom(RoomName) from C to S;
    choice at S {
      RoomPort(RoomName, Port) from S to C;
    } or {
      RoomNotFound(RoomName) from S to C;
    }
    do ChatServer(C, S);
  } or {
    CreateRoom(RoomName) from C to S;
    choice at S {
      CreateRoomSuccess(RoomName) from S to C;
    } or {
      RoomExists(RoomName) from S to C;
    }
    do ChatServer(C, S);
  } or {
    ListRooms() from C to S;
    RoomList(StringList) from S to C;
    do ChatServer(C, S);
  } or {
    Bye(String) from C to S;
  }
}

global protocol ChatSessionCtoR(role C, role R) {
  choice at C {
    OutgoingChatMessage(String) from C to R;
    do ChatSessionCtoR(C, R);
  } or {
    LeaveRoom() from C to R;
  }
}

global protocol ChatSessionRtoC(role R, role C){
  choice at R {
    IncomingChatMessage(String) from R to C;
    do ChatSessionRtoC(R, C);
  } or {
    Bye() from R to C;
  }
}
```

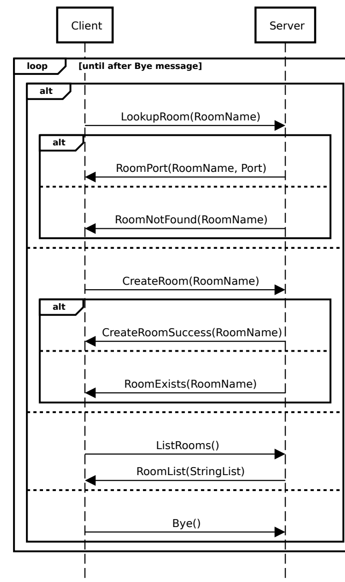

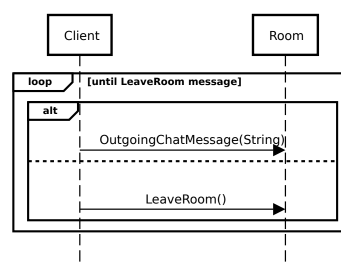

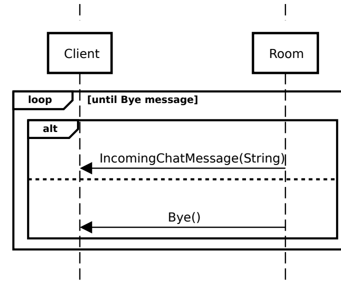

C OMITTED DEFINITIONS AND PROOFS

C.1 Omitted Definitions

Term reduction $M \longrightarrow_{\mathsf{M}} N$ is standard β-reduction:

Term reduction rules $\boxed{M \longrightarrow_{\mathsf{M}} N}$

$$\begin{array}{rcl} \textbf{let}\ x \Leftarrow \textbf{return}\ V\ \textbf{in}\ M & \longrightarrow_{\mathsf{M}} & M\{V/x\} \\ (\lambda x.\, M)\ V & \longrightarrow_{\mathsf{M}} & M\{V/x\} \\ (\textbf{rec}\ f(x).M)\ V & \longrightarrow_{\mathsf{M}} & M\{\textbf{rec}\ f(x).M/f, V/x\} \end{array}$$

$$\begin{array}{rcl} \textbf{if true then}\ M\ \textbf{else}\ N & \longrightarrow_{\mathsf{M}} & M \\ \textbf{if false then}\ M\ \textbf{else}\ N & \longrightarrow_{\mathsf{M}} & N \\ \mathcal{E}[M] \longrightarrow_{\mathsf{M}} \mathcal{E}[N] & & (\text{if}\ M \longrightarrow_{\mathsf{M}} N) \end{array}$$

C.2 Preservation

We begin with some unsurprising auxiliary lemmas.

Lemma C.1 (Substitution). *If* $\Gamma, x : B \mid S \triangleright M : A \triangleleft T$ *and* $\Gamma \vdash V : B$*, then* $\Gamma \mid S \triangleright M\{V/x\} : A \triangleleft T$.

Proof. By induction on the derivation of $\Gamma_1, x : A \mid S \triangleright M : B \triangleleft T$. □

Lemma C.2 (Subterm typability). *Suppose* $\mathbf{D}$ *is a derivation of* $\Gamma \mid S \triangleright \mathcal{E}[M] : A \triangleleft T$*. Then there exists some subderivation* $\mathbf{D}'$ *of* $\mathbf{D}$ *concluding* $\Gamma \mid S \triangleright M : B \triangleleft S'$ *for some type* B *and session type* S'*, where the position of* $\mathbf{D}'$ *in* $\mathbf{D}$ *corresponds to that of the hole in* E*.*

Proof. By induction on the structure of E. □

Lemma C.3 (Replacement). *If:*

(1) $\mathbf{D}$ *is a derivation of* $\Gamma \mid S \triangleright \mathcal{E}[M] : A \triangleleft T$
(2) $\mathbf{D}'$ *is a subderivation of* $\mathbf{D}$ *concluding* $\Gamma \mid S \triangleright M : B \triangleleft T'$ *where the position of* $\mathbf{D}'$ *in* $\mathbf{D}$ *corresponds to that of the hole in* E
(3) $\Gamma \mid S' \triangleright N : B \triangleleft T'$

then $\Gamma \mid S' \triangleright \mathcal{E}[N] : A \triangleleft T$

Proof. By induction on the structure of E. □

Since type environments are unrestricted, we also obtain a weakening result.

Lemma C.4 (Weakening). *(1) If* $\Gamma \vdash V : B$ *and* $x \notin \mathit{dom}(\Gamma)$*, then* $\Gamma, x : A \vdash V : B$.
(2) If $\Gamma \mid S \triangleright M : B \triangleleft T$ *and* $x \notin \mathit{dom}(\Gamma)$*, then* $\Gamma, x : A \mid S \triangleright M : B \triangleleft T$.
(3) If $\Gamma; \Delta \vdash \sigma$ *and* $x \notin \mathit{dom}(\Gamma)$*, then* $\Gamma, x : A; \Delta \vdash \sigma$.
(4) If $\Gamma; \Delta \vdash \rho$ *and* $x \notin \mathit{dom}(\Gamma)$*, then* $\Gamma, x : A; \Delta \vdash \rho$.
(5) If $\Gamma; \Delta \vdash C$ *and* $x \notin \mathit{dom}(\Gamma)$*, then* $\Gamma, x : A \vdash V : B$.

Proof. By mutual induction on all premises. □

Lemma C.5 (Preservation (terms)). *If* $\Gamma \mid S \triangleright M : A \triangleleft T$ *and* $M \longrightarrow_{M} N$*, then* $\Gamma \mid S \triangleright N : A \triangleleft T$.

Proof. A standard induction on the derivation of $M \longrightarrow_{\mathsf{M}} N$, noting that functional reduction does not modify the session type. □

Next, we introduce some MPST-related lemmas that are helpful for proving preservation of configuration reduction. We often make use of these lemmas implicitly.

Lemma C.6. *If* $\mathit{safe}(\Delta, \Delta')$*, then* $\mathit{safe}(\Delta)$.

Proof sketch. Splitting a context only removes potential reductions. Only by adding reductions could we violate safety. □

Lemma C.7. *If safe*(Δ_1, Δ_2) *and* $\Delta_1 \Longrightarrow \Delta'_1$*, then safe*$(\Delta'_1, \Delta_2)$.

Proof sketch. By induction on the derivation of $\Delta_1 \equiv\overset{\pi}{\Longrightarrow}\equiv \Delta'_1$.

It suffices to consider the cases where reduction could potentially make the combined environments unsafe.

In the case of Lbl-Sync-Send, the resulting reduction adds a message $(\mathsf{p}, \mathsf{q}, \ell_i(A_i))$ to a queue Q.

The only way this could violate safety is if there were some entry $s[\mathsf{q}] : \mathsf{p}\,\&\{\ell_i(A_i)\,.\,S_i\}_{i\in I}$, and $Q \equiv (\mathsf{p}, \mathsf{q}, \ell_j(A_j)) \cdot Q'$ where $j \in I$, but $(Q \cdot (\mathsf{p}, \mathsf{q}, \ell_k(A_k)) \equiv (\mathsf{p}, \mathsf{q}, \ell_k(A_k)) \cdot Q''$ with $k \notin I$. However, this is impossible since it is not possible to permute this message ahead of the existing message because of the side-conditions on queue equivalence.

A similar argument applies for Lbl-Sync-Recv. □

Lemma C.8. *If* $\Gamma;\ \Delta, s : Q \vdash s \triangleright \sigma$ *and* $\Gamma \vdash V : A$*, then* $\Gamma;\ \Delta, s : (Q \cdot (\mathsf{p}, \mathsf{q}, \ell(A))) \vdash s \triangleright \sigma \cdot (\mathsf{p}, \mathsf{q}, \ell(V))$

Proof. A straightforward induction on the derivation of $\Gamma;\ \Delta, s : Q \vdash s \triangleright \sigma$. □

Lemma C.9 (Preservation (Equivalence)). *If* $\Gamma;\ \Delta \vdash C$ *and* $C \equiv \mathcal{D}$ *then there exists some* $\Delta' \equiv \Delta$ *such that* $\Gamma;\ \Delta' \vdash \mathcal{D}$.

Proof. By induction on the derivation of $C \equiv \mathcal{D}$. The only case that causes the type environment to change is queue message reordering, which can be made typable by mirroring the change in the queue type. □

Lemma C.10 (Preservation (Configuration reduction)). *If* $\Gamma;\ \Delta \vdash C$ *with safe*(Δ) *and* $C \longrightarrow \mathcal{D}$*, then there exists some* Δ' *such that* $\Delta \Longrightarrow^{?} \Delta'$ *and* $\Gamma;\ \Delta' \vdash \mathcal{D}$.

Proof. By induction on the derivation of $C \longrightarrow \mathcal{D}$.

Case E-Send.

$$\langle a, (\mathcal{E}[\mathsf{q}\,!\,\ell(V)])^{s[\mathsf{p}]}, \sigma, \rho\rangle \parallel s \triangleright \delta \longrightarrow \langle a, (\mathcal{E}[\mathbf{return}\ ()])^{s[\mathsf{p}]}, \sigma, \rho\rangle \parallel s \triangleright \delta \cdot (\mathsf{p}, \mathsf{q}, \ell(V))$$

Assumption:

$$\dfrac{\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\mathsf{q}\,!\,\ell(V)] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : S \vdash (\mathcal{E}[\mathsf{q}\,!\,\ell(V)])^{s[\mathsf{p}]}} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma;\ s[\mathsf{p}], \Delta_2, \Delta_3 \vdash \langle a, (\mathcal{E}[\mathsf{q}\,!\,\ell(V)])^{s[\mathsf{p}]}, \sigma, \rho\rangle} \qquad \Gamma;\ s : Q \vdash s \triangleright \delta}{\Gamma;\ s[\mathsf{p}] : S, \Delta_2, \Delta_3, s : Q \vdash \langle a, (\mathcal{E}[\mathsf{q}\,!\,\ell(V)])^{s[\mathsf{p}]}, \sigma, \rho\rangle \parallel s \triangleright \delta}$$

By Lemma C.2 we have that $\Gamma \mid \mathsf{q} \oplus\{\ell_i(A_i) : T_i\}_{i\in I} \triangleright \mathsf{q}\,!\,\ell_j(V) : \mathbf{1} \triangleleft T_j$ and therefore that $S = \mathsf{q} \oplus\{\ell_i(A_i) : T_i\}_{i\in I}$.

Since $\Gamma \mid T_j \triangleright \mathbf{return}\ () : \mathbf{1} \triangleleft T_j$, we can show by Lemma C.3 we have that $\Gamma \mid T_j \triangleright \mathcal{E}[\mathbf{return}\ ()] : \mathbf{1} \triangleleft \mathsf{end}$.

By Lemma C.8, $\Gamma;\ s : Q \cdot (\mathsf{p}, \mathsf{q}, \ell_j(A_j)) \vdash s \triangleright \delta \cdot (\mathsf{p}, \mathsf{q}, \ell_j(V))$.

Therefore, recomposing:

$$\dfrac{\dfrac{\dfrac{\Gamma \mid T_j \triangleright \mathcal{E}[\mathbf{return}\ ()] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : T_j \vdash (\mathcal{E}[\mathbf{return}\ ()])^{s[\mathsf{p}]}} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma;\ s[\mathsf{p}] : T_j, \Delta_2, \Delta_3 \vdash \langle a, (\mathcal{E}[\mathbf{return}\ ()])^{s[\mathsf{p}]}, \sigma, \rho\rangle} \qquad \Gamma;\ s : Q \cdot (\mathsf{p}, \mathsf{q}, \ell_j(A_j)) \vdash s \triangleright \delta \cdot (\mathsf{p}, \mathsf{q}, \ell_j(V))}{\Gamma;\ s[\mathsf{p}] : T_j, \Delta_2, \Delta_3, s : Q \cdot (\mathsf{p}, \mathsf{q}, \ell_j(B_j)) \vdash \langle a, (\mathcal{E}[\mathbf{return}\ ()])^{s[\mathsf{p}]}, \sigma, \rho\rangle \parallel s \triangleright \delta \cdot (\mathsf{p}, \mathsf{q}, \ell_j(V))}$$

Finally,

$s[\mathsf{p}] : \mathsf{q} \oplus \{\ell_i(A_i) : T_i\}_{i \in I}, \Delta_2, \Delta_3, s : Q \Longrightarrow s[\mathsf{p}] : T_j, \Delta_2, \Delta_3, s : Q \cdot (\mathsf{p}, \mathsf{q}, \ell_j(B_j))$ by Lbl-Send as required.

Case E-React.

$$\frac{\ell(x) \mapsto M \in \overrightarrow{H}}{\langle a, \textbf{idle}, \sigma[s[\mathsf{p}] \mapsto \textbf{handler}\ \mathsf{q}\ \{\overrightarrow{H}\}], \rho\rangle \parallel s \triangleright (\mathsf{q}, \mathsf{p}, \ell(V)) \cdot \delta \longrightarrow \langle a, (M\{V/x\})^{s[\mathsf{p}]}, \sigma, \rho\rangle \parallel s \triangleright \delta}$$

For simplicity (and equivalently) let us refer to ℓ as ℓ_j.

Let **D** be the following derivation:

$$\frac{\Gamma; \cdot \vdash \textbf{idle} \qquad \dfrac{\dfrac{(\Gamma, x_i : B_i \mid S_i \triangleright M_i : \mathbf{1} \triangleleft \mathsf{end})_{i \in I}}{\Gamma \vdash \textbf{handler}\ \mathsf{q}\ \{(\ell_i(x_i) \mapsto M_i)_{i \in I}\} : \mathsf{Handler}(S^?)} \qquad \Gamma; \Delta_2 \vdash \sigma}{\Gamma; \Delta_2, s[\mathsf{p}] : S^? \vdash \sigma[s[\mathsf{p}] \mapsto \textbf{handler}\ \mathsf{q}\ \{\overrightarrow{H}\}]} \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{p}] : S^? \vdash \langle a, \textbf{idle}, \sigma[s[\mathsf{p}] \mapsto \textbf{handler}\ \mathsf{q}\ \{\overrightarrow{H}\}], \rho\rangle}$$

Assumption:

$$\frac{\mathbf{D} \qquad \dfrac{\Gamma; s : Q \vdash s \triangleright \delta}{\Gamma; s[\mathsf{p}] : S^?, s : ((\mathsf{q}, \mathsf{p}, \ell_j(A)) \cdot Q) \vdash s \triangleright (\mathsf{q}, \mathsf{p}, \ell_j(V)) \cdot \delta}}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{p}] : S^?, s : ((\mathsf{q}, \mathsf{p}, \ell_j(A)) \cdot Q) \vdash \langle a, \textbf{idle}, \sigma[s[\mathsf{p}] \mapsto \textbf{handler}\ \mathsf{q}\ \{\overrightarrow{H}\}], \rho\rangle \parallel s \triangleright (\mathsf{q}, \mathsf{p}, \ell_j(V)) \cdot \delta}$$

where $S^? = \mathsf{p}\,\&\{\ell_i(B_i).S_i\}_{i \in I}$.

Since $\mathsf{safe}(\Delta_2, \Delta_3, s[\mathsf{p}] : S^?, s : ((\mathsf{q}, \mathsf{p}, \ell_j(A)) \cdot Q))$ we have that $j \in I$ and $A = B_j$.

Similarly since $\ell_j(x_j) \mapsto M \in \overrightarrow{H}$ we have that $\Gamma, x : B_j \mid S_j \triangleright M : \mathbf{1} \triangleleft \mathsf{end}$.

By Lemma C.1, $\Gamma \mid S_j \triangleright M\{V/x_j\} : \mathbf{1} \triangleleft \mathsf{end}$.

Let **D′** be the following derivation:

$$\frac{\dfrac{\Gamma \mid S_j \triangleright M\{V/x_j\} : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : S_j \vdash (M\{V/x_j\})^{s[\mathsf{p}]}} \qquad \Gamma; \Delta_2 \vdash \sigma \Gamma; \Delta_3 \vdash \rho}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{p}] : S_j \vdash \langle a, (M\{V/x_j\})^{s[\mathsf{p}]}, \sigma, \rho\rangle}$$

Recomposing:

$$\frac{\mathbf{D'} \qquad \Gamma; s : Q \vdash s \triangleright \delta}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{p}] : S_j, s : Q \vdash \langle a, (M\{V/x_j\})^{s[\mathsf{p}]}, \sigma, \rho\rangle \parallel s \triangleright \delta}$$

Finally, we note that $\Delta_2, \Delta_3, s[\mathsf{p}] : S^?, s : ((\mathsf{q}, \mathsf{p}, \ell_j(A)) \cdot Q) \Longrightarrow \Delta_2, \Delta_3, s[\mathsf{p}] : S_j, s : Q$ by Lbl-Recv as required.

Case E-Suspend.

$$\langle a, (\mathcal{E}[\textbf{suspend}\ V])^{s[\mathsf{p}]}, \sigma, \rho\rangle \longrightarrow \langle a, \textbf{idle}, \sigma[s[\mathsf{p}] \mapsto V], \rho\rangle$$

Assumption:

$$\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\textbf{suspend } V] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : S \vdash (\mathcal{E}[\textbf{suspend } V])^{s[\mathsf{p}]}} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma;\ s[\mathsf{p}] : S, \Delta_2, \Delta_3 \vdash \langle a, (\mathcal{E}[\textbf{suspend } V])^{s[\mathsf{p}]}, \sigma, \rho \rangle}$$

By Lemma C.2 we have that:

$$\dfrac{\Gamma \vdash V : \mathsf{Handler}(S^?)}{\Gamma \mid S^? \triangleright \textbf{suspend } V : A \triangleleft T}$$

for any arbitrary A, T, and showing that $S = S^?$.
Recomposing:

$$\dfrac{\Gamma; \cdot \vdash \textbf{idle} \qquad \dfrac{\Gamma \vdash V : \mathsf{Handler}(S^?) \qquad \Gamma; \Delta_2 \vdash \sigma}{\Gamma; \Delta_2, s[\mathsf{p}] : S^? \vdash \sigma[s[\mathsf{p}] \mapsto V]} \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma;\ s[\mathsf{p}] : S^?, \Delta_2, \Delta_3 \vdash \langle a, \textbf{idle}, \sigma[s[\mathsf{p}] \mapsto V], \rho \rangle}$$

as required.
Case E-Spawn.

$$\langle a, \mathcal{M}[\textbf{spawn } M], \sigma, \rho \rangle \longrightarrow \langle a, \mathcal{M}[\textbf{return } ()], \sigma, \rho \rangle \parallel \langle a, M, \epsilon, \epsilon \rangle$$

There are two subcases based on whether the $\mathcal{M} = \mathcal{E}[-]$ or $\mathcal{M} = (\mathcal{E}[-])^{s[\mathsf{p}]}$. Both are similar so we will prove the latter case.
Assumption:

$$\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\textbf{spawn } M] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : S \vdash (\mathcal{E}[\textbf{spawn } M])^{s[\mathsf{p}]}} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{p}] : S \vdash \langle a, (\mathcal{E}[\textbf{spawn } M])^{s[\mathsf{p}]}, \sigma, \rho \rangle}$$

By Lemma C.2:

$$\dfrac{\Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma \mid S \triangleright \textbf{spawn } M : \mathbf{1} \triangleleft S}$$

By Lemma C.3, $\Gamma \mid S \triangleright \mathcal{E}[\textbf{return } ()] : \mathbf{1} \triangleleft \mathsf{end}$.
Thus, recomposing:

$$\dfrac{\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\textbf{return } ()] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : S \vdash (\mathcal{E}[\textbf{return } ()])^{s[\mathsf{p}]}} \qquad \begin{array}{c}\Gamma; \Delta_2 \vdash \sigma \\ \Gamma; \Delta_3 \vdash \rho\end{array}}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{p}] : S \vdash \langle a, (\mathcal{E}[\textbf{return } ()])^{s[\mathsf{p}]}, \sigma, \rho \rangle} \qquad \dfrac{\dfrac{\Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; \cdot \vdash M} \qquad \Gamma; \cdot \vdash \epsilon \qquad \Gamma; \cdot \vdash \epsilon}{\Gamma;\ \cdot \vdash \langle a, M, \epsilon, \epsilon \rangle}}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{p}] : S \vdash \langle a, (\mathcal{E}[\textbf{return } ()])^{s[\mathsf{p}]}, \sigma, \rho \rangle \parallel \langle a, M, \epsilon, \epsilon \rangle}$$

Case E-Reset.

$$\langle a, Q[\textbf{return}\ ()], \sigma, \rho\rangle \longrightarrow \langle a, \textbf{idle}, \sigma, \rho\rangle$$

There are two subcases based on whether $Q = [-]$ or $Q = ([-])^{s[\mathsf{p}]}$. We prove the latter case; the former is similar but does not require a context reduction.

Assumption:

$$\frac{\dfrac{\Gamma \mid \mathsf{end} \triangleright \textbf{return}\ () : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : \mathsf{end} \vdash (\textbf{return}\ ())^{s[\mathsf{p}]}} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{p}] : \mathsf{end} \vdash \langle a, (\textbf{return}\ ())^{s[\mathsf{p}]}, \sigma, \rho\rangle}$$

We can show that $\Delta_2, \Delta_3, s[\mathsf{p}] : \mathsf{end} \xRightarrow{\mathsf{end}(s, \mathsf{p})} \Delta_2, \Delta_3$, so we can reconstruct:

$$\frac{\Gamma; \cdot \vdash \textbf{idle} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma;\ \Delta_2, \Delta_3 \vdash \langle a, \textbf{idle}, \sigma, \rho\rangle}$$

as required.

Case E-NewAP.

$$\frac{c \text{ fresh}}{\langle a, \mathcal{M}[\textbf{newAP}_{(\mathsf{p}_i : S_i)_{i \in I}}], \sigma, \rho\rangle \longrightarrow (\nu p)(\langle a, \mathcal{M}[\textbf{return}\ p], \sigma, \rho\rangle \parallel p((\mathsf{p}_i \mapsto \epsilon)_{i \in 1..n}))}$$

As usual we prove the case where $\mathcal{M} = (\mathcal{E}[-])^{s[\mathsf{p}]}$; the case where $\mathcal{M} = (\mathcal{E}[-])$ is similar.

Assumption:

$$\frac{\dfrac{\Gamma \mid T \triangleright \mathcal{E}[\textbf{newAP}_{(\mathsf{p}_i : S_i)_{i \in I}}] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : T \vdash (\mathcal{E}[\textbf{newAP}_{(\mathsf{p}_i : S_i)_{i \in I}}])^{s[\mathsf{p}]}} \qquad \begin{array}{c}\Gamma; \Delta_2 \vdash \sigma \\ \Gamma; \Delta_3 \vdash \rho\end{array}}{\Gamma;\ \Delta_2, \Delta_3 \vdash \langle a, (\mathcal{E}[\textbf{newAP}_{(\mathsf{p}_i : S_i)_{i \in I}}])^{s[\mathsf{p}]}, \sigma, \rho\rangle}$$

By Lemma C.2:

$$\frac{\varphi \text{ is a safety property} \qquad \varphi((\mathsf{p}_i : S_i)_{i \in I})}{\Gamma \mid T \triangleright \textbf{newAP}_{(\mathsf{p}_i : S_i)_{i \in I}} : \mathsf{AP}((\mathsf{p}_i : S_i)_{i \in I}) \triangleleft T}$$

By Lemma C.3, $\Gamma, c : \mathsf{AP}((\mathsf{p}_i : S_i)_{i \in I}) \mid T \triangleright \mathcal{E}[\textbf{return}\ c] : \mathbf{1} \triangleleft \mathsf{end}$.
Let $\Gamma' = \Gamma, c : \mathsf{AP}((\mathsf{p}_i : S_i)_{i \in I})$.
By Lemma C.4, since c is fresh we have that $\Gamma'; \Delta_2 \vdash \sigma$ and $\Gamma'; \Delta_3 \vdash \rho$.
Recomposing:

$$\frac{\dfrac{\dfrac{\dfrac{\Gamma' \mid T \triangleright \mathcal{E}[\textbf{return}\ c] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma'; s[\mathsf{p}] : T \vdash (\mathcal{E}[\textbf{return}\ c])^{s[\mathsf{p}]}} \qquad \begin{array}{c}\Gamma'; \Delta_2 \vdash \sigma \\ \Gamma'; \Delta_3 \vdash \rho\end{array}}{\Gamma';\ \Delta_2, \Delta_3, s[\mathsf{p}] : T \vdash \langle a, (\mathcal{E}[\textbf{return}\ c])^{s[\mathsf{p}]}, \sigma, \rho\rangle} \qquad \dfrac{\begin{array}{cc} c : \mathsf{AP}((\mathsf{p}_i : S_i)_{i \in 1..n}) \in \Gamma & (\cdot \vdash \epsilon : S_i)_{i \in 1..n} \\ \varphi((\mathsf{p}_i : S_i)_{i \in 1..n}) & \varphi \text{ is a safety property}\end{array}}{\Gamma';\ c : \mathsf{AP} \vdash c((\mathsf{p}_i \mapsto \epsilon)_{i \in 1..n})}}{\Gamma';\ \Delta_2, \Delta_3, s[\mathsf{p}] : T, c : \mathsf{AP} \vdash \langle a, (\mathcal{E}[\textbf{return}\ c])^{s[\mathsf{p}]}, \sigma, \rho\rangle \parallel p((\mathsf{p}_i \mapsto \epsilon)_{i \in 1..n})}}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{p}] : T \vdash (\nu c)(\langle a, (\mathcal{E}[\textbf{return}\ c])^{s[\mathsf{p}]}, \sigma, \rho\rangle \parallel p((\mathsf{p}_i \mapsto \epsilon)_{i \in 1..n}))}$$

as required.

Case E-Register.

$$\dfrac{\iota \text{ fresh}}{\langle a, \mathcal{M}[\mathbf{register}\ p\ \mathsf{p}\ M], \sigma, \rho\rangle \parallel p(\chi[\mathsf{p} \mapsto \widetilde{\iota'}]) \longrightarrow (\nu\iota)(\langle a, \mathcal{M}[\mathbf{return}\ ()], \sigma, \rho[\iota \mapsto M]\rangle \parallel p(\chi[\mathsf{p} \mapsto \widetilde{\iota'} \cup \{\iota\}]))}$$

Again, we prove the case where $\mathcal{M} = (\mathcal{E}[-])^{s[\mathsf{q}]}$ and let $\mathsf{p} = \mathsf{p}_j$ for some j.
Let $\Delta = \Delta_2, \Delta_3, \Delta_4, \widetilde{\iota_j^- : S_j}, s[\mathsf{p}] : T$.
Let **D** be the following derivation:

$$\dfrac{\dfrac{\Gamma \mid T \triangleright \mathcal{E}[\mathbf{register}\ p\ \mathsf{p}_j\ M] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{q}] : T \vdash (\mathcal{E}[\mathbf{register}\ p\ \mathsf{p}_j\ M])^{s[\mathsf{q}]}} \qquad \begin{array}{l}\Gamma; \Delta_2 \vdash \sigma \\ \Gamma; \Delta_3 \vdash \rho\end{array}}{\Gamma; \Delta_2, \Delta_3, s[\mathsf{q}] : T \vdash \langle a, (\mathcal{E}[\mathbf{register}\ p\ \mathsf{p}_j\ M])^{s[\mathsf{q}]}, \sigma, \rho\rangle}$$

Assumption:

$$\dfrac{\mathbf{D} \qquad \dfrac{\dfrac{\{(\mathsf{p}_i : S_i)_{i \in 1..n}\}\ \Delta_4 \vdash \chi}{\{(\mathsf{p}_i : S_i)_{i \in 1..n}\}\ \Delta_4, \overrightarrow{\iota'^-_j : S_j} \vdash \chi[\mathsf{p}_j \mapsto \widetilde{\iota'}]} \qquad \begin{array}{l} c : \mathsf{AP}((\mathsf{p}_i : S_i)_{i \in 1..n}) \in \Gamma \\ \varphi((\mathsf{p}_i : S_i)_{i \in 1..n}) \\ \varphi \text{ is a safety property}\end{array}}{\Gamma;\ \Delta_4, \overrightarrow{\iota'^-_j : S_j}, p : \mathsf{AP} \vdash p(\chi[\mathsf{p}_j \mapsto \widetilde{\iota'}])}}{\Gamma;\ \Delta \vdash \langle a, (\mathcal{E}[\mathbf{register}\ p\ \mathsf{p}_j\ M])^{s[\mathsf{q}]}, \sigma, \rho\rangle \parallel p(\chi[\mathsf{p}_j \mapsto \widetilde{\iota'}])}$$

By Lemma C.2:

$$\dfrac{\Gamma \vdash c : \mathsf{AP}((\mathsf{p}_i : S_i)_i) \qquad \Gamma \mid S_j \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma \mid T \triangleright \mathbf{register}\ c\ \mathsf{p}_j\ M : \mathbf{1} \triangleleft T}$$

By Lemma C.3, $\Gamma \mid T \triangleright \mathcal{E}[\mathbf{return}\ ()] : \mathbf{1} \triangleleft \mathsf{end}$.
Now, let **D′** be the following derivation:

$$\dfrac{\dfrac{\Gamma \mid T \triangleright \mathcal{E}[\mathbf{return}\ ()] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{q}] : T \vdash (\mathcal{E}[\mathbf{return}\ ()])^{s[\mathsf{q}]}} \qquad \Gamma; \Delta_2 \vdash \sigma \qquad \dfrac{\Gamma \mid S_j \triangleright M : \mathbf{1} \triangleleft \mathsf{end} \qquad \Gamma; \Delta_3 \vdash \rho}{\Gamma; \Delta_3, \iota^+ : S_j \vdash \rho[\iota^+ \mapsto M]}}{\Gamma;\ \Delta_2, \Delta_3, s[\mathsf{q}] : S, \iota^+ : S_j \vdash \langle a, (\mathcal{E}[\mathbf{return}\ ()])^{s[\mathsf{q}]}, \sigma, \rho\rangle}$$

Finally, we can recompose:

$$\dfrac{\dfrac{\mathbf{D} \qquad \dfrac{\dfrac{\{(\mathsf{p}_i : S_i)_{i \in 1..n}\}\ \Delta_4 \vdash \chi}{\{(\mathsf{p}_i : S_i)_{i \in 1..n}\}\ \Delta_4, \overrightarrow{\iota'^-_j : S_j}, \iota^- : S_j \vdash \chi[\mathsf{p}_j \mapsto \widetilde{\iota'} \cup \{\iota\}]} \qquad \begin{array}{l} c : \mathsf{AP}((\mathsf{p}_i : S_i)_{i \in 1..n}) \in \Gamma \\ \varphi((\mathsf{p}_i : S_i)_{i \in 1..n}) \\ \varphi \text{ is a safety property}\end{array}}{\Gamma;\ \Delta_4, \overrightarrow{\iota'^-_j : S_j}, \iota^- : S_j, p : \mathsf{AP} \vdash p(\chi[\mathsf{p}_j \mapsto \widetilde{\iota'} \cup \{\iota\}])}}{\Gamma;\ \Delta, \iota^+ : S_j, \iota^- : S_j \vdash \langle a, (\mathcal{E}[\mathbf{return}\ ()])^{s[\mathsf{q}]}, \sigma, \rho\rangle \parallel p(\chi[\mathsf{p}_j \mapsto \widetilde{\iota'} \cup \{\iota\}])}}{\Gamma;\ \Delta \vdash (\nu\iota)(\langle a, (\mathcal{E}[\mathbf{return}\ ()])^{s[\mathsf{q}]}, \sigma, \rho\rangle \parallel p(\chi[\mathsf{p}_j \mapsto \widetilde{\iota'} \cup \{\iota\}]))}$$

as required.
Case E-Init.

$$\frac{s \text{ fresh}}{\begin{array}{c}(\nu \iota_{\mathsf{p}_i})_{i\in 1..n}(p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}} \cup \{\iota_{\mathsf{p}_i}\})_{i\in 1..n}) \parallel \langle a_i, \mathbf{idle}, \sigma_i, \rho_i[\iota_{\mathsf{p}_i} \mapsto M_i]\rangle_{i\in 1..n}) \xrightarrow{\tau} \\ (\nu s)(p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}})_{i\in 1..n}) \parallel s \triangleright \epsilon \parallel \langle a_i, (M_i)^{s[\mathsf{p}_i]}, \sigma_i, \rho_i\rangle_{i\in 1..n})\end{array}}$$

For each actor composed in parallel we have:

$$\frac{\Gamma;\cdot \vdash \mathbf{idle} \qquad \Gamma;\Delta_{i_2} \vdash \sigma_i \qquad \dfrac{\Gamma \mid S_i \triangleright M_i : \mathbf{1} \triangleleft \mathsf{end} \qquad \Gamma;\Delta_{i_3} \vdash \rho}{\Gamma;\Delta_{i_3}, \iota_i^+ : S_i \vdash \rho_i[\iota_{\mathsf{p}_i} \mapsto M_i]}}{\Gamma;\ \Delta_{i_2}, \Delta_{i_3}, \iota_i^+ : S_i, a_i \vdash \langle a_i, \mathbf{idle}, \sigma_i, \rho_i\rangle}$$

Let:

- $\Delta_{tok+} = \iota_1^+ : S_1, \ldots, \iota_n^+ : S_n$
- $\Delta_{tok-} = \iota_1^- : S_1, \ldots, \iota_n^- : S_n$
- $\Delta_a = \Delta_{1_2}, \Delta_{1_3}, \ldots, \Delta_{n_2}, \Delta_{n_3}, a_1, \ldots, a_n$
- $\Delta_b = \Delta_a, \Delta_{tok+}$

Then by repeated use of TC-Par we have that $\Gamma;\ \Delta_a, \Delta_{tok+} \vdash (\langle a, \mathbf{idle}, \sigma_i, \rho_i[\iota_{\mathsf{p}_i} \mapsto M_i]\rangle)_{i\in 1..n}$
Assumption (given some Δ):

$$\frac{\dfrac{\dfrac{\begin{array}{l} \qquad c : \mathsf{AP}((\mathsf{p}_i : S_i)_i) \in \Gamma \\ \{\mathsf{p}_i : S_i\}\ \Delta, \Delta_{tok-} \vdash (\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}} \cup \{\iota_{\mathsf{p}_i}\})_{i\in 1..n} \\ \varphi((\mathsf{p}_i : S_i)_{i\in 1..n}) \qquad \varphi \text{ is a safety property}\end{array}}{\Gamma;\ \Delta, \Delta_{tok-} \vdash p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}} \cup \{\iota_{\mathsf{p}_i}\})_{i\in 1..n})} \qquad \Gamma;\ \Delta_a, \Delta_{tok+} \vdash (\langle a, \mathbf{idle}, \sigma_i, \rho_i[\iota_{\mathsf{p}_i} \mapsto M_i]\rangle)_{i\in 1..n}}{\Gamma;\ \Delta, \Delta_a, \Delta_{tok+}, \Delta_{tok-} \vdash p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}} \cup \{\iota_{\mathsf{p}_i}\})_{i\in 1..n}) \parallel (\langle a, \mathbf{idle}, \sigma_i, \rho_i[\iota_{\mathsf{p}_i} \mapsto M_i]\rangle)_{i\in 1..n}}}{\Gamma;\ \Delta, \Delta_a \vdash (\nu \iota_1) \cdots (\nu \iota_n)(p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}} \cup \{\iota_{\mathsf{p}_i}\})_{i\in 1..n}) \parallel (\langle a, \mathbf{idle}, \sigma_i, \rho_i[\iota_{\mathsf{p}_i} \mapsto M_i]\rangle)_{i\in 1..n})}$$

Through the access point typing rules we can show that we can remove each ι_{p_i} from the access point: $\Gamma;\ \Delta \vdash p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}})_{i\in 1..n})$.

Similarly, for each actor composed in parallel we can construct:

$$\frac{\dfrac{\Gamma \mid S_i \triangleright M_i : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}_i] : S_i \vdash (M_i)^{s[\mathsf{p}_i]}} \qquad \Gamma;\Delta_{i_2} \vdash \sigma_i \qquad \Gamma;\Delta_{i_3} \vdash \rho_i}{\Gamma;\ \Delta_{i_2}, \Delta_{i_3}, s[\mathsf{p}_i] : S_i \vdash \langle a, (M_i)^{s[\mathsf{p}_i]}, \sigma_i, \rho_i\rangle}$$

Let $\Delta_s = s[\mathsf{p}_1] : S_1, \ldots, s[\mathsf{p}_n] : S_n$
Then by repeated use of TC-Par we have that $\Gamma;\ \Delta_a, \Delta_s \vdash \langle a, (M_i)^{s[\mathsf{p}_i]}, \sigma_i, \rho_i\rangle_{i\in 1..n}$.
Recomposing:

$$\frac{\dfrac{\begin{array}{l}\varphi(\Delta_s) \\ \varphi \text{ is a safety property} \\ \Gamma;\ \Delta \vdash p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}})_{i\in 1..n})\end{array} \qquad \dfrac{\dfrac{}{\Gamma;\ s : \epsilon \vdash s \triangleright \epsilon} \qquad \Gamma;\ \Delta_a, \Delta_s \vdash (\langle a, (M_i)^{s[\mathsf{p}_i]}, \sigma_i, \rho_i\rangle)_{i\in 1..n}}{\Gamma;\ \Delta_a, \Delta_s, s : \epsilon \vdash s \triangleright \epsilon \parallel (\langle a, (M_i)^{s[\mathsf{p}_i]}, \sigma_i, \rho_i\rangle)_{i\in 1..n}}}{\Gamma;\ \Delta, \Delta_a, \Delta_s, s : \epsilon \vdash p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}})_{i\in 1..n}) \parallel s \triangleright \epsilon \parallel (\langle a, (M_i)^{s[\mathsf{p}_i]}, \sigma_i, \rho_i\rangle)_{i\in 1..n}}}{\Gamma;\ \Delta, \Delta_a \vdash (\nu s)(p((\mathsf{p}_i \mapsto \widetilde{\iota'_{\mathsf{p}_i}})_{i\in 1..n}) \parallel s \triangleright \epsilon \parallel (\langle a, (M_i)^{s[\mathsf{p}_i]}, \sigma_i, \rho_i\rangle)_{i\in 1..n})}$$

as required.

Case E-Lift.
Immediate by Lemma C.5.

Case E-Nu.
Immediate by the IH, noting that by the definition of safety, reduction of a safe context results in another safe context.

Case E-Par.
Immediate by the IH and Lemma C.7.

Case E-Struct.
Immediate by the IH and Lemma C.9.

□

Theorem 3.2 (Preservation). *Typability is preserved by structural congruence and reduction.*

(≡) *If* $\Gamma;\, \Delta \vdash C$ *and* $C \equiv \mathcal{D}$ *then there exists some* $\Delta' \equiv \Delta$ *such that* $\Gamma;\, \Delta' \vdash \mathcal{D}$.

(⟶) *If* $\Gamma;\, \Delta \vdash C$ *with* $\mathit{safe}(\Delta)$ *and* $C \longrightarrow \mathcal{D}$, *then there exists some* Δ' *such that* $\Delta \Longrightarrow^{?} \Delta'$ *and* $\Gamma;\, \Delta' \vdash \mathcal{D}$.

Proof. Immediate from Lemmas C.9 and C.10. □

C.3 Progress

Let Ψ be a type environment containing only references to access points:

$$\Psi ::= \cdot \mid \Psi, p : \mathsf{AP}((\mathsf{p}_i : S_i)_i)$$

Functional reduction satisfies progress.

Lemma C.11 (Term Progress). *If $\Psi \mid S_1 \triangleright M : A \triangleleft S_2$ then either:*

- *$M = \textbf{return}\ V$ for some value V; or*
- *there exists some N such that $M \longrightarrow_M N$; or*
- *M can be written $\mathcal{E}[M']$ where M' is a communication or concurrency construct, i.e.*
    - *$M = \textbf{spawn}\ N$ for some N; or*
    - *$M = \mathsf{p}\,!\,m(V)$ for some role p and message $m(V)$; or*
    - *$M = \textbf{suspend}\ V$ or some V; or*
    - *$M = \textbf{newAP}_{(\mathsf{p}_i : T_i)}$ for some collection of participants $(\mathsf{p}_i : T_i)$*
    - *$M = \textbf{register}\ V\ \mathsf{p}$ for some value V and role p*

Proof. A standard induction on the derivation of $\Psi \mid S_1 \triangleright M : A \triangleleft S_2$; there are β-reduction rules for all STLC terms, leaving only values and communication / concurrency terms. □

The key *thread progress* lemma shows that each actor is either idle, or can reduce; the proof is by inspection of $\mathcal{T}$, noting there are reduction rules for each construct; the runtime typing rules ensure the presence of any necessary queues or access points.

Lemma C.12 (Thread Progress). *Let $\mathcal{C} = \mathcal{G}[\langle a, \mathcal{T}, \sigma, \rho \rangle]$. If $\cdot ; \cdot \vdash \mathcal{C}$ then either $\mathcal{T} = \textbf{idle}$, or there exist $\mathcal{G}', \mathcal{T}', \sigma', \rho'$ such that $\mathcal{C} \longrightarrow \mathcal{G}'[\langle a, \mathcal{T}', \sigma', \rho' \rangle]$.*

Proof. If $\mathcal{T} = \textbf{idle}$ then the theorem is satisfied, so consider the cases where $\mathcal{T} = M$ or $\mathcal{T} = (M)^{s[\mathsf{p}]}$. By Lemma C.11, either M can reduce (and the configuration can reduce via E-Lift), M is a value (and the thread can reduce by E-Reset), or M is a communication or concurrency construct. Of these:

- **spawn** N can reduce by E-Spawn
- **suspend** V can reduce by E-Suspend
- $\textbf{newAP}_{(\mathsf{p}_i : S_i)_i}$ can reduce by E-NewAP

Next, consider **register** p p M. Since we begin with a closed environment, it must be the case that p is ν-bound so by T-APName and T-AP there must exist some subconfiguration $p(\chi)$ of $\mathcal{G}$; the configuration can therefore reduce by E-Register.

Finally, consider $M = \mathsf{q}\,!\,\ell(V)$. It cannot be the case that $\mathcal{T} = \mathsf{q}\,!\,\ell(V)$ since by T-Send the term must have an output session type as a precondition, whereas TT-NoSess assigns a precondition of end. Therefore, it must be the case that $\mathcal{T} = (\mathsf{q}\,!\,\ell(V))^{s[\mathsf{p}]}$ for some s, p. Again since the initial runtime typing environment is empty, it must be the case that s is ν-bound and so by T-SessionName and T-EmptyQueue/T-ConsQueue there must be some session queue $s \triangleright \delta$. The thread must therefore be able to reduce by E-Send. □

Theorem 3.6 (Progress). *If $\cdot ; \cdot \vdash_{prog} \mathcal{C}$, then either there exists some $\mathcal{D}$ such that $\mathcal{C} \longrightarrow \mathcal{D}$, or $\mathcal{C}$ is structurally congruent to the following canonical form:*

$$(\nu \tilde{\iota})(\nu p_{i \in 1..m})(\nu a_{j \in 1..n})(p_1(\chi_1)_{i \in 1..m} \parallel \langle a_j, \textbf{idle}, \epsilon, \rho_j \rangle_{j \in 1..n})$$

Proof. By Proposition 3.4 $\mathcal{C}$ can be written in canonical form:

$$(\nu \tilde{\iota})(\nu p_{i \in 1..l})(\nu s_{j \in 1..m})(\nu a_{k \in 1..n})(p_i(\chi_i)_{i \in 1..l} \parallel (s_j \triangleright \delta_j)_{j \in 1..m} \parallel \langle a_k, \mathcal{T}_k, \sigma_k, \rho_k \rangle_{k \in 1..n})$$

Additional Syntax

$$\text{Labels} \quad \mathscr{L} \quad ::= \quad \tau \mid \mathsf{spawn}(M) \mid \mathsf{send}(\mathtt{p}, \ell, V) \mid \mathsf{newAP}(a) \mid \mathsf{register}(p, \mathtt{p}, M)$$

Labelled Transition Semantics for Computations $\boxed{M \xhookrightarrow{\mathscr{L}} N}$

LTS-Spawn
$$\frac{}{\mathbf{spawn}\ M \xhookrightarrow{\mathsf{spawn}(M)} \mathbf{return}\ ()}$$

LTS-Send
$$\frac{}{\mathtt{p}\,!\,\ell(V) \xhookrightarrow{\mathsf{send}(\mathtt{p},\ell,V)} \mathbf{return}\ ()}$$

LTS-NewAP
$$\frac{}{\mathbf{newAP}_{(\mathtt{p}_i:S_i)_i} \xhookrightarrow{\mathsf{newAP}(a)} \mathbf{return}\ a}$$

LTS-Register
$$\frac{}{\mathbf{register}\ p\ \mathtt{p}\ M \xhookrightarrow{\mathsf{register}(p,\mathtt{p},M)} \mathbf{return}\ ()}$$

LTS-Beta
$$\frac{M \longrightarrow_{\mathsf{M}} N}{M \xhookrightarrow{\tau} N}$$

LTS-Lift
$$\frac{M \xhookrightarrow{\mathscr{L}} N}{\mathcal{E}[M] \xhookrightarrow{\mathscr{L}} \mathcal{E}[N]}$$

Fig. 14. LTS Semantics for Computations

By repeated applications of Lemma C.12, either the configuration can reduce or all threads are idle:

$$(\nu\tilde{\iota})(\nu p_{i\in 1..l})(\nu s_{j\in 1..m})(\nu a_{k\in 1..n})(p_i(\chi_i)_{i\in 1..l} \parallel (s_j \triangleright \delta_j)_{j\in 1..m} \parallel \langle a_k, \mathbf{idle}, \sigma_k, \rho_k\rangle_{k\in 1..n})$$

By the linearity of runtime type environments Δ, each role endpoint $s[\mathtt{p}]$ must be contained in precisely one actor. There are two ways an endpoint can be used: either by TT-Sess in order to run a term in the context of a session, or by TH-Handler to record a receive session type as a handler. Since all threads are idle, it must be the case the only applicable rule is TH-Handler and therefore each role must have an associated stored handler.

Since the types for each session must satisfy progress, the collection of local types must reduce. Since all session endpoints must have a receive session type, the only type reductions possible are through Lbl-Sync-Recv. Since all threads are idle we can pick the top message from any session queue and reduce the actor with the associated stored handler by E-React.

The only way we could not do such a reduction is if there were to be no sessions, leaving us with a configuration of the form:

$$(\nu\tilde{\iota})(\nu p_{i\in 1..m})(\nu a_{j\in 1..n})(p_i(\chi_i)_{i\in 1..m} \parallel \langle a_j, \mathbf{idle}, \sigma_j, \rho_j\rangle_{j\in 1..n})$$

□

C.4 Global Progress

The overview of the global progress proof is as follows:

- We design a labelled transition system semantics for term reduction (Figure 14).
- We argue that our LTS is strongly-normalising up to **suspend** (Proposition C.13).
- We prove an operational correspondence between the LTS reduction and configuration reduction, specifically that reductions in the LTS semantics can drive configuration reduction, and that every configuration reduction affecting an actor term can be reflected by the LTS.
- Finally we can use this result to show that any session can eventually reduce.

C.4.1 LTS Semantics.

Figure 14 shows the labelled transition semantics for our language of computations. Standard β-reductions are reflected as τ-transitions, and communication and concurrency actions reduce in a single step and are accounted for using labelled reductions.

Proposition C.13 (Strong Normalisation (LTS)). *If $\Psi \mid S \triangleright^{\mathsf{f}} M : A \triangleleft \mathbf{end}$ then there exists some finite reduction sequence such that either:*

- $M \overset{\mathcal{L}_1}{\hookrightarrow} \cdots \overset{\mathcal{L}_n}{\hookrightarrow} V$ *for some* V*; or*
- $M \overset{\mathcal{L}_1}{\hookrightarrow} \cdots \overset{\mathcal{L}_n}{\hookrightarrow} \mathcal{E}[\textbf{suspend}\ V]$ *for some* $\mathcal{E}$ *and* V

Proof sketch. Fine-grain call-by-value is strongly normalising; this can be shown using techniques such as $\top\top$-lifting [30], which also extends to exceptions. These results extend to our LTS as all additional constructs reduce immediately. □

Lemma C.14 (Reduction under contexts). *If $\Gamma; \Delta \vdash \mathcal{G}[C]$ and $C \longrightarrow \mathcal{D}$ then there exists some $\mathcal{G}'$ such that $\mathcal{G}[C] \longrightarrow \mathcal{G}'[\mathcal{D}]$.*

Proof. By induction on the structure of $\mathcal{G}$. □

We also have a special case for straightforward β-reduction:

Lemma C.15 (Term reduction under contexts). *Let $C = \mathcal{G}[\langle a, \mathcal{Q}[\mathcal{E}[M]], \sigma, \rho\rangle]$. If $\Gamma; \Delta \vdash C$ and $M \overset{\tau}{\hookrightarrow} N$, then $C \longrightarrow \mathcal{G}[\langle a, \mathcal{Q}[\mathcal{E}[N]], \sigma, \rho\rangle]$.*

Proof. By induction on the structure of $\mathcal{G}$. □

Lemma C.16 (Simulation). *Suppose $\Psi; \cdot \vdash^{\mathsf{f}} C$ where $C = \mathcal{G}[\langle a, \mathcal{Q}[M], \sigma, \rho\rangle]$. If $M \overset{\mathcal{L}}{\hookrightarrow} N$, then there exist some $\mathcal{G}'$, σ', and ρ' such that $C \longrightarrow \mathcal{G}'[\langle a, \mathcal{Q}[M'], \sigma', \rho'\rangle]$.*

Proof. By induction on the derivation of $M \overset{\mathcal{L}}{\hookrightarrow} N$.

Case LTS-Spawn.

Assumption:

$$\frac{}{\textbf{spawn}\ M \overset{\mathsf{spawn}(M)}{\hookrightarrow} \textbf{return}\ ()}$$

By Lemma C.14, E-Spawn, and E-Lift,

$$\mathcal{G}[\langle a, \mathcal{Q}[\textbf{spawn}\ M], \sigma, \rho\rangle] \longrightarrow \mathcal{G}'[\langle a, \mathcal{Q}[\textbf{return}\ ()], \sigma, \rho\rangle \parallel \langle b, M, \sigma, \rho\rangle]$$

which we can write as $\mathcal{G}''[\langle a, \mathcal{Q}[\textbf{return}\ ()], \sigma, \rho\rangle]$ as required.

Case LTS-Send.

Assumption:

$$\frac{}{\mathsf{p} ! \ell(V) \overset{\mathsf{send}(\mathsf{p}, \ell, V)}{\hookrightarrow} \textbf{return}\ ()}$$

Since $\Psi; \cdot \vdash^{\mathsf{f}} C$ where $C = \mathcal{G}[\langle a, \mathcal{Q}[\mathsf{p} ! \ell(V)], \sigma, \rho\rangle]$, by T-Session and the linearity of runtime environments, there must exist some $\mathcal{G}'$ such that $C \equiv \mathcal{G}'[\langle a, \mathcal{Q}[\mathsf{p} ! \ell(V)], \sigma, \rho\rangle] \parallel s \triangleright \delta$ which can reduce by E-Send to $\mathcal{G}'[\langle a, \mathcal{Q}[\textbf{return}\ ()], \sigma, \rho\rangle] \parallel s \triangleright \delta \cdot (\mathsf{p}, \mathsf{q}, \ell(V))$ as required.

Case LTS-NewAP.

Assumption:

$$\frac{}{\mathbf{newAP}_{(\mathsf{p}_i:S_i)_i} \xrightarrow{\mathsf{newAP}(a)} \mathbf{return}\ a}$$

We have that $\Psi; \cdot \vdash^{\mathsf{f}} C$ where $C = \mathcal{G}[\langle a, \mathcal{Q}[\mathbf{newAP}_{(\mathsf{p}_i:S_i)_i}], \sigma, \rho\rangle]$; the result follows from reduction by E-NewAP.

Case LTS-Register.

Assumption:

$$\frac{}{\mathbf{register}\ p\ \mathsf{p}\ M \xrightarrow{\mathsf{register}(p,\mathsf{p},M)} \mathbf{return}\ ()}$$

Since $\Psi; \cdot \vdash^{\mathsf{f}} C$ where $C = \mathcal{G}[\langle a, \mathcal{Q}[\mathbf{register}\ p\ \mathsf{p}\ M], \sigma, \rho\rangle]$, by T-Session and the linearity of runtime environments, there must exist some $\mathcal{G}'$ such that $C \equiv \mathcal{G}'[\langle a, \mathcal{Q}[\mathbf{register}\ p\ \mathsf{p}\ M], \sigma, \rho\rangle] \parallel p(\chi[\mathsf{p} \mapsto \tilde{\iota}])$ which can reduce by E-Register to $(\nu\iota')(\mathcal{G}'[\langle a, \mathbf{return}\ (), \sigma, \rho[\iota' \mapsto M]\rangle] \parallel p(\chi[\mathsf{p} \mapsto \tilde{\iota} \cup \{\iota'\}]))$ as required.

Case LTS-Beta.

Immediate by Lemma C.15.

Case LTS-Lift.

Immediate by the IH and E-LiftM. □

Lemma C.17 (Determinism (Term Reduction)). *Suppose* $\Psi; \Delta \vdash^{\mathsf{f}} C$ *where* $C = \mathcal{G}[\langle a, \mathcal{Q}[M], \sigma, \rho\rangle]$. *If:*

- $C \longrightarrow \mathcal{G}_1[\langle a, \mathcal{Q}[N_1], \sigma_1, \rho_1\rangle]$, *where* $M \neq N_1$
- $C \longrightarrow \mathcal{G}_2[\langle a, \mathcal{Q}[N_2], \sigma_2, \rho_2\rangle]$, *where* $M \neq N_2$

then up to the identities of fresh variables, $\mathcal{G}_1 = \mathcal{G}_2$, *and* $N_1 = N_2$, *and* $\sigma_1 = \sigma_2$, *and* $\rho_1 = \rho_2$.

Proof. Since $M \neq N_1$ and $M \neq N_2$ the overall reduction must be driven by the reduction from M into N_1 or N_2 respectively. The result then follows by inspection on the reduction rules, noting that β-reduction is deterministic, as are the relevant rules E-Send, E-Spawn, E-NewAP, and E-Register. □

Lemma C.18 (Reflection). *Suppose* $\Psi; \Delta \vdash^{\mathsf{f}} C$ *where* $C = \mathcal{G}[\langle a, \mathcal{Q}[M], \sigma, \rho\rangle]$.
If $C \longrightarrow \mathcal{G}'[\langle a, \mathcal{Q}[N], \sigma', \rho'\rangle]$ *for some* $\mathcal{G}'$, N, σ' *and* ρ' *where* $M \neq N$, *then there exists some* $\mathscr{L}$ *such that* $M \xrightarrow{\mathscr{L}} N$.

Proof. Since $M \neq N$, by Lemma C.17, the reduction from C must be unique, and will be the one specified by Lemma C.16. □

Proposition C.19 (Operational Correspondence).

Suppose $\Psi; \Delta \vdash^{\mathsf{f}} C$ *where* $C = \mathcal{G}[\langle a, \mathcal{Q}[M], \sigma, \rho\rangle]$.

- *If* $M \xrightarrow{\mathscr{L}} N$, *then there exist some* $\mathcal{G}'$, σ', *and* ρ' *such that* $C \longrightarrow \mathcal{G}'[\langle a, \mathcal{Q}[N], \sigma', \rho'\rangle]$.
- *If* $C \longrightarrow \mathcal{G}'[\langle a, \mathcal{Q}[N], \sigma', \rho'\rangle]$ *for some* $\mathcal{G}'$, N, σ' *and* ρ' *where* $M \neq N$, *then there exists some* $\mathscr{L}$ *such that* $M \xrightarrow{\mathscr{L}} N$.

Proof. Follows as a consequence of Lemmas C.16 and C.18. □

Lemma C.20. *If* $\cdot; \cdot \vdash (vs : \Delta)C$ *and* $C \xrightarrow{\tau} \mathcal{D}$, *then* $\cdot; \cdot \vdash (vs : \Delta)\mathcal{D}$.

Proof. A straightforward corollary of Theorem 3.2. □

Theorem 3.8 (Session Progress). *If* $\cdot; \cdot \vdash^{\mathsf{f}}_{prog} (\nu s : \Delta_s) C$ *where active*(Δ_s), *then* $C \xrightarrow{\tau}{}^{*} \xrightarrow{s}$.

Proof. By T-SessionName we have that $\cdot; s[\mathsf{p}_1] : S_1, \ldots, s[\mathsf{p}_n] : S_n \vdash^{\mathsf{f}} C$ and thus by the linearity of Δ_s alongside rule T-Actor we have some set of actors:

$$\{\langle a_i, \mathcal{T}_i, \sigma_i, \rho_i \rangle\}_{i \in 1..m}$$

such that for each role p_j for $j \in 1..n$, either:

- there exists some $\mathcal{T}_k$ such that $\mathcal{T}_k = (M)^{s[\mathsf{p}_j]}$ for some M
- $s[\mathsf{p}_j] \in \mathsf{dom}(\sigma_k)$ for some $k \in 1..m$

Consider the subset of actors where $\mathcal{T}_i \neq \textbf{idle}$, i.e., $\mathcal{T}_i = N_i$ or $\mathcal{T}_i = (N_i)^{s'[\mathsf{p}_j]}$ for some N_i. In this case, for each actor, by Proposition C.13 we have that $N_i \xhookrightarrow{\mathscr{L}_1} \cdots \xhookrightarrow{\mathscr{L}_n} N'_i$ where either $N'_i = \textbf{return}\ ()$, or $N'_i = E[\textbf{suspend}\ V]$ for some value V. By Proposition C.19, we can simulate each reduction sequence as a configuration reduction (and moreover, by the reflection direction, each term can *only* follow this reduction sequence). At this point we can revert each actor to idle by either E-Suspend or E-Reset.

If any labelled reduction, simulated as a configuration reduction, is labelled with session s then we can conclude. Otherwise we have that $C \xrightarrow{\tau}{}^{*} \mathcal{D}$ where again by typing we have some subset of actors such that:

$$\{\langle a_i, \textbf{idle}, \sigma_i, \rho_i \rangle\}_{i \in 1..m'}$$

By Lemma C.20 we have that $\cdot; \cdot \vdash^{\mathsf{f}}_{\mathrm{prog}} (\nu s : \Delta_s)\mathcal{D}$ and thus it remains the case that $\Delta \Longrightarrow$. Thus by similar reasoning to Theorem 3.6 it must be the case that some actor a_j (where $j \in 1..m'$) can reduce by E-React as required. □

D PROOFS FOR SECTION 4

This appendix details the proofs of the metatheoretical properties enjoyed by Maty$_{\rightleftarrows}$ and Maty$_{\lightning}$; we omit the proofs for Maty with state, which is entirely standard.

D.1 Maty$_{\rightleftarrows}$

D.1.1 Preservation.

THEOREM A.1 (PRESERVATION). *Preservation (as defined in Theorem 3.2) continues to hold in Maty$_{\rightleftarrows}$.*

PROOF. Preservation of typing under structural congruence follows straightforwardly.

For preservation of typing under reduction, we proceed by induction on the derivation of $\mathcal{C} \longrightarrow \mathcal{D}$.

Case E-SUSPEND$_!$-1.
Similar to E-SUSPEND$_!$-2.

Case E-SUSPEND$_!$-2.

$$\langle a, (\mathcal{E}[\mathbf{suspend}_! \ \underline{s} \ V])^{s[\mathsf{p}]}, \sigma[\underline{s} \mapsto \overrightarrow{D}], \rho, \theta\rangle \xrightarrow{\tau} \langle a, \mathbf{idle}, \sigma[\underline{s} \mapsto \overrightarrow{D} \cdot (s[\mathsf{p}], V)], \rho, \theta\rangle$$

Assumption:

$$\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\mathbf{suspend}_! \ \underline{s} \ V] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : S \vdash (\mathcal{E}[\mathbf{suspend}_! \ \underline{s} \ V])^{s[\mathsf{p}]}} \qquad \dfrac{\Gamma; \Delta_1 \vdash \sigma \qquad \Sigma(\underline{s}) = (S^!, A) \qquad (\Gamma \vdash W_i : A \xrightarrow{S^!, \mathsf{end}} \mathbf{1})_i}{\Gamma; \Delta_1, (s_i[\mathsf{q}_i] : S^!)_i \vdash \sigma[\underline{s} \mapsto (s_i[\mathsf{q}_i], W_i)_i]} \qquad \begin{array}{c}\Gamma; \Delta_2 \vdash \rho \\ \Gamma \vdash \Delta_3 \theta\end{array}}{\Gamma;\ \Delta_1, \Delta_2, \Delta_3, s[\mathsf{p}] : S, (s_i[\mathsf{q}_i] : S^!)_i, a \vdash \langle a, \mathbf{idle}, \sigma[\underline{s} \mapsto (s_i[\mathsf{q}_i], W_i)_i], \rho, \theta\rangle}$$

Consider the subderivation $\Gamma \mid S \triangleright \mathcal{E}[\mathbf{suspend}_! \ \underline{s} \ V] : \mathbf{1} \triangleleft \mathsf{end}$. By Lemma C.2 there exists a subderivation:

$$\dfrac{\Sigma(\underline{s}) = (S^!, A) \qquad \Gamma \vdash V : A \xrightarrow{S^!, \mathsf{end}} \mathbf{1}}{\Gamma \mid S^! \triangleright \mathbf{suspend}_! \ \underline{s} \ V : \mathbf{1} \triangleleft \mathsf{end}}$$

Therefore we have that $S = S^!$.
Recomposing:

$$\dfrac{\dfrac{}{\Gamma; \cdot \vdash \mathbf{idle}} \qquad \dfrac{\Gamma; \Delta_1 \vdash \sigma \qquad \Sigma(\underline{s}) = (S^!, A) \qquad (\Gamma \vdash W_i : A \xrightarrow{S^!, \mathsf{end}} \mathbf{1})_i \qquad \Gamma \vdash V : A \xrightarrow{S^!, \mathsf{end}} \mathbf{1}}{\Gamma; \Delta_1, (s_i[\mathsf{q}_i] : S^!)_i, s[\mathsf{p}] : S^! \vdash \sigma[\underline{s} \mapsto (s_i[\mathsf{q}_i], W_i)_i \cdot (s[\mathsf{p}], V)]} \qquad \begin{array}{c}\Gamma; \Delta_2 \vdash \rho \\ \Gamma \vdash \Delta_3 \theta\end{array}}{\Gamma;\ \Delta_1, \Delta_2, \Delta_3, s[\mathsf{p}] : S, (s_i[\mathsf{q}_i] : S^!)_i, a \vdash \langle a, (\mathcal{E}[\mathbf{suspend}_! \ \underline{s} \ V])^{s[\mathsf{p}]}, \sigma[\underline{s} \mapsto (s_i[\mathsf{q}_i], W_i)_i \cdot (s[\mathsf{p}], V)], \rho, \theta\rangle}$$

as required.

Case E-BECOME.

$$\langle a, \mathcal{M}[\mathbf{become} \ \underline{s} \ V], \sigma, \rho, \theta\rangle \xrightarrow{\tau} \langle a, \mathcal{M}[\mathbf{return} \ ()], \sigma, \rho, \theta \cdot (\underline{s}, V)\rangle$$

Assumption (considering the case that $\mathcal{M} = \mathcal{E}[-]$ for some $\mathcal{E}$; the case in the context of a session is identical:

$$\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\mathbf{become} \ \underline{s} \ V] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; \cdot \vdash \mathcal{E}[\mathbf{become} \ \underline{s} \ V]} \qquad \Gamma; \Delta_1 \vdash \sigma \qquad \Gamma; \Delta_2 \vdash \rho \qquad \Gamma \vdash \Delta_3 \theta}{\Gamma;\ \Delta_1, \Delta_2, \Delta_3, a \vdash \langle a, \mathcal{T}, \sigma, \rho, \theta\rangle}$$

By Lemma C.2 we have:

$$\dfrac{\Sigma(\underline{s}) = (T, A) \qquad \Gamma \vdash V : A}{\Gamma \mid S \triangleright \textbf{become}\ \underline{s}\ V : \mathbf{1} \triangleleft S}$$

By Lemma C.3 we can show that $\Gamma \mid S \triangleright \mathcal{E}[\textbf{return}\ ()] : \mathbf{1} \triangleleft \mathsf{end}$.
Recomposing:

$$\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\textbf{return}\ ()] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; \cdot \vdash \mathcal{E}[\textbf{return}\ ()]} \qquad \Gamma; \Delta_1 \vdash \sigma \qquad \Gamma; \Delta_2 \vdash \rho \qquad \dfrac{\Gamma \vdash \Delta\theta \qquad \Sigma(\underline{s}) = (S^!, A) \qquad \Gamma \vdash V : A}{\Gamma \vdash \Delta_3 \theta \cdot (\underline{s}, V)}}{\Gamma;\ \Delta_1, \Delta_2, \Delta_3, a \vdash \langle a, \mathcal{E}[\textbf{return}\ ()], \sigma, \rho, \theta \cdot (\underline{s}, V) \rangle}$$

as required.

Case E-Activate.

$$\langle a, \textbf{idle}, \sigma[\underline{s} \mapsto (s[\mathsf{p}], V) \cdot \overrightarrow{D}], \rho, (\underline{s}, W) \cdot \theta \rangle \xrightarrow{\tau} \langle a, (V\ W)^{s[\mathsf{p}]}, \sigma[\underline{s} \mapsto \overrightarrow{D}], \rho, \theta \rangle$$

Assumption:

$$\dfrac{\dfrac{}{\Gamma; \cdot \vdash \textbf{idle}} \qquad \dfrac{\begin{array}{c}\Gamma; \Delta_1 \vdash \sigma \\ \Sigma(\underline{s}) = (S^!, A) \qquad \Gamma \vdash V : A \xrightarrow{S^!, \mathsf{end}} \mathbf{1} \qquad (\Gamma \vdash V_i : A \xrightarrow{S^!, \mathsf{end}} \mathbf{1})_i\end{array}}{\Gamma; \Delta_1, s[\mathsf{p}] : S^!, (s_i[\mathsf{p}_i] : S^!)_i \vdash \sigma, \underline{s} \mapsto (s[\mathsf{p}], V) \cdot (s_i[\mathsf{p}_i], V_i)_i} \qquad \Gamma; \Delta_2 \vdash \rho \qquad \dfrac{\begin{array}{c}\Gamma \vdash \Delta_3 \theta \\ \Sigma(\underline{s}) = (S^!, A) \\ \Gamma \vdash W : A\end{array}}{\Gamma \vdash \Delta_3 (\underline{s}, W) \cdot \theta}}{\Gamma;\ \Delta_1, \Delta_2, s[\mathsf{p}] : S^!, (s_i[\mathsf{p}_i] : S^!)_i, \Delta_3, a \vdash \langle a, \textbf{idle}, \sigma[\underline{s} \mapsto (s[\mathsf{p}], V) \cdot (s_i[\mathsf{p}_i], V_i)_i], \rho, (\underline{s}, W) \cdot \theta \rangle}$$

Recomposing:

$$\dfrac{\dfrac{\dfrac{\Gamma \vdash V : A \xrightarrow{S^!, \mathsf{end}} \mathbf{1} \qquad \Gamma \vdash W : A}{\Gamma \mid S^! \triangleright V\ W : \mathbf{1} \triangleleft \mathsf{end}}}{\Gamma; s[\mathsf{p}] : S^! \vdash (V\ W)^{s[\mathsf{p}]}} \qquad \dfrac{\begin{array}{c}\Gamma; \Delta_1 \vdash \sigma \\ \Sigma(\underline{s}) = (S^!, A) \qquad (\Gamma \vdash V_i : A \xrightarrow{S^!, \mathsf{end}} \mathbf{1})_i\end{array}}{\Gamma; \Delta_1, (s_i[\mathsf{p}_i] : S^!)_i \vdash \sigma, \underline{s} \mapsto (s_i[\mathsf{p}_i], V_i)_i} \qquad \Gamma; \Delta_2 \vdash \rho \qquad \Gamma \vdash \Delta_3 \theta}{\Gamma;\ \Delta_1, \Delta_2, s[\mathsf{p}] : S^!, (s_i[\mathsf{p}_i] : S^!)_i, \Delta_3, a \vdash \langle a, (V\ W)^{s[\mathsf{p}]}, \sigma[\underline{s} \mapsto (s_i[\mathsf{p}_i], V_i)_i], \rho, \theta \rangle}$$

as required. □

D.1.2 Progress.

Theorem A.2 (Progress (Maty$_{\rightleftarrows}$)). *If $\cdot; \cdot \vdash_{prog} \mathcal{C}$, then either there exists some $\mathcal{D}$ such that $\mathcal{C} \longrightarrow \mathcal{D}$, or $\mathcal{C}$ is structurally congruent to the following canonical form:*

$$(\nu \tilde{\iota})(\nu p_{i \in 1..l})(\nu s_{j \in 1..m})(\nu a_{k \in 1..n})(p_i(\chi_i)_{i \in 1..l} \parallel (s_j \triangleright \delta_j)_{j \in 1..m} \parallel \langle a_k, \textbf{idle}, \sigma_k, \rho_k, \theta_k \rangle_{k \in 1..n})$$

where for each session s_j there exists some mapping $s_j[\mathsf{p}] \mapsto (\underline{s}, V)$ (for some role p, static session name $\underline{s}$, and callback V) contained in some σ_k where θ_k does not contain any requests for $\underline{s}$.

Proof. The proof follows that of Theorem 3.6. Thread progress (Lemma C.12) holds as before, since we can always evaluate **become** by E-Become, and we can always evaluate $\textbf{suspend}_!$ by E-Suspend-$!_1$ or E-Suspend-$!_2$.

Following the same reasoning as Theorem 3.6 we can write $\mathcal{C}$ in canonical form, where all threads are idle:

$$(\nu \tilde{\iota})(\nu p_{i \in 1..l})(\nu s_{j \in 1..m})(\nu a_{k \in 1..n})(p_i(\chi_i)_{i \in 1..l} \parallel (s_j \triangleright \delta_j)_{j \in 1..m} \parallel \langle a_k, \textbf{idle}, \sigma_k, \rho_k \rangle_{k \in 1..n})$$

However, there are now *three* places each role endpoint $s[\mathsf{p}]$ can be used: either by TT-Sess to run a term in the context of a session or by TH-Handler to record a receive-suspended session

Runtime syntax

Cancellation-aware runtime envs.	Φ	::=	$\cdot \mid \Phi, p \mid \Phi, \iota^{\pm} : S \mid \Phi, s[\mathsf{p}] : S \mid \Phi, s[\mathsf{p}] : \lightning \mid \Phi, s : Q$
Labels	γ	::=	$\cdots \mid \lightning s[\mathsf{p}] \mid s : \mathsf{p} \lightning \mathsf{q}{::}\ell \mid s : \mathsf{p} \lightning \mathsf{q}$

Modified typing rules for configurations $\boxed{\Gamma; \Phi \vdash C}$ $\boxed{\Gamma; \Phi \vdash \sigma}$

T-ActorName
$$\frac{\Gamma, a : \mathsf{Pid}; \Phi, a \vdash C}{\Gamma; \Phi \vdash (\nu a)C}$$

T-ZapActor
$$\frac{}{\Gamma; a \vdash \lightning a}$$

T-ZapRole
$$\frac{}{\Gamma; s[\mathsf{p}] : \lightning \vdash \lightning s[\mathsf{p}]}$$

T-ZapTok
$$\frac{}{\Gamma; \iota^{+} : S \vdash \lightning \iota}$$

T-Actor
$$\frac{\Gamma; \Phi_1 \vdash \mathcal{T} \qquad \Gamma; \Phi_2 \vdash \sigma \qquad \Gamma; \Phi_3 \vdash \rho \qquad \forall (b, M) \in \omega.\ \Gamma \vdash b : \mathsf{Pid} \wedge \Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; \Phi_1, \Phi_2, \Phi_3, a \vdash \langle a, \mathcal{T}, \sigma, \rho, \omega \rangle}$$

TH-Handler
$$\frac{\Gamma \vdash V : \mathsf{Handler}(S^{?}) \qquad \Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end} \qquad \Gamma; \Phi \vdash \sigma}{\Gamma; \Phi, s[\mathsf{p}] : S^{?} \vdash \sigma[s[\mathsf{p}] \mapsto (V, M)]}$$

Additional LTS rules $\boxed{\Phi \xrightarrow{\gamma} \Phi'}$ $\boxed{\Phi \overset{s[\mathsf{p}]}{\rightsquigarrow} \Phi}$

Lbl-ZapMsg $\quad \Phi, s[\mathsf{q}] : \lightning, s : (\mathsf{p}, \mathsf{q}, \ell(A)) \cdot Q \xrightarrow{s:\mathsf{p}\lightning\mathsf{q}::\ell} \Phi, s[\mathsf{q}] : \lightning, s : Q$

Lbl-ZapRecv $\quad \Phi, s[\mathsf{p}]{:}\mathsf{q}\,\&\{\ell_i(A_i).S_i\}_{i \in I}, s[\mathsf{q}]{:}\lightning, s{:}Q \xrightarrow{s:\mathsf{p}\lightning\mathsf{q}} \Phi, s[\mathsf{p}]{:}\lightning, s[\mathsf{q}]{:}\lightning, s{:}Q \quad$ (if $\mathsf{messages}(\mathsf{q}, \mathsf{p}, Q) = \emptyset$)

Lbl-Zap $\quad \Phi, s[\mathsf{p}] : S \overset{s[\mathsf{p}]}{\rightsquigarrow} \Phi, s[\mathsf{p}] : \lightning$

Fig. 15. Maty$_{\lightning}$: Modified configuration typing rules and type LTS

type as before, but now also by TH-SendHandler to record a send-suspended session type. As before, the former is impossible as all threads are idle, so now we must consider the cases for TH-Handler.

Following the same reasoning as Theorem 3.6, we can reduce any handlers that have waiting messages. Thus we are finally left with the scenario where the session type LTS can reduce, but not the configuration: this can only happen when the sending reduction is send-suspended, as required. □

D.2 Maty$_{\lightning}$

Figure 15 shows the necessary modifications to the configuration typing rules and type LTS. We extend runtime type environments to *cancellation-aware* environments Φ that include an additional entry of the form $s[\mathsf{p}] : \lightning$, denoting that endpoint $s[\mathsf{p}]$ has been cancelled. We also need to extend the type LTS to account for failure propagation; we take a similar approach to Barwell et al. [5]. Rule Lbl-Zap accounts for the possibility that in any given reduction step, a role may be cancelled (for example, as a result of E-RaiseS), but it is a separate relation since it is unnecessary for determining behavioural properties of types.

All metatheoretical results continue to hold.

D.2.1 Preservation. First, it is useful to show that safety is preserved even if several roles are cancelled; we use this lemma implicitly throughout the preservation proof.

Let us write $\mathsf{roles}(\Delta) = \{\mathsf{p} \mid s[\mathsf{p}] : S \in \Phi\}$ to retrieve the roles from an environments.

Let us also define the operation $\mathsf{zap}(\Phi, \widetilde{\mathsf{p}})$ that cancels any role in the given set, i.e., $\mathsf{zap}(s[\mathsf{p}_1] : S_1, s[\mathsf{p}_2] : S_2, a, \{\mathsf{p}_1\}) = s[\mathsf{p}_1] : \lightning, s[\mathsf{p}_2] : S_2, a$.

Lemma D.1. *If* $\mathit{safe}(\Phi)$ *then* $\mathit{safe}(\mathit{zap}(\Phi, \widetilde{\mathsf{p}}))$ *for any* $\widetilde{\mathsf{p}} \subseteq \mathit{roles}(\Phi)$.

Proof. Zapping a role does not affect safety; the only way to violate safety is by *adding* further unsafe communication reductions. □

We need a slightly modified preservation theorem in order to account for cancelled roles; specifically we write $\Rrightarrow$ for the relation $\Longrightarrow^{?}\rightsquigarrow^{*}$. The safety property is unchanged for cancellation-aware environments.

Theorem D.2 (Preservation ($\longrightarrow$, $\text{Maty}_{\lightning}$)). *If* $\Gamma;\ \Phi \vdash C$ *with* $\mathit{safe}(\Phi)$ *and* $C \longrightarrow \mathcal{D}$*, then there exists some* Φ' *such that* $\Phi \Rrightarrow \Phi'$ *and* $\Gamma;\ \Phi' \vdash \mathcal{D}$*.*

Proof. Preservation of typability by structural congruence is straightforward, so we concentrate on preservation of typability by reduction. We proceed by induction on the derivation of $C \longrightarrow \mathcal{D}$, concentrating on the new rules rather than the adapted rules (which are straightforward changes to the existing proof).

Case E-Monitor.

$$\langle a, \mathcal{M}[\mathbf{monitor}\ b\ M], \sigma, \rho, \omega\rangle \xrightarrow{\tau} \langle a, \mathcal{M}[\mathbf{return}\ ()], \sigma, \rho, \omega \cup \{(b, M)\}\rangle$$

We consider the case where $\mathcal{M} = \mathcal{E}[-]$ for some $\mathcal{E}$; the case in the context of a session is similar.

Assumption:

$$\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\mathbf{monitor}\ b\ M] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; \cdot \vdash \mathcal{E}[\mathbf{monitor}\ b\ M]} \qquad \Gamma; \Phi_1 \vdash \sigma \qquad \Gamma; \Phi_2 \vdash \rho}{\Gamma;\ \Phi_1, \Delta_2, a \vdash \langle a, \mathcal{E}[\mathbf{monitor}\ b\ M], \sigma, \rho, \omega\rangle}$$

where $\forall (b, N) \in \omega.\ \Gamma \vdash b : \mathsf{Pid}\ \wedge\ \Gamma \mid \mathsf{end} \triangleright N : \mathbf{1} \triangleleft \mathsf{end}$.

By Lemma C.2, we know:

$$\dfrac{\Gamma \vdash b : \mathsf{Pid} \qquad \Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma \mid S \triangleright \mathbf{monitor}\ b\ M : \mathbf{1} \triangleleft S}$$

By Lemma C.3 we know $\Gamma \mid S \triangleright \mathcal{E}[\mathbf{return}\ ()] : \mathbf{1} \triangleleft \mathsf{end}$.

Recomposing:

$$\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\mathbf{return}\ ()] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; \cdot \vdash \mathcal{E}[\mathbf{return}\ ()]} \qquad \Gamma; \Phi_1 \vdash \sigma \qquad \Gamma; \Phi_2 \vdash \rho}{\Gamma;\ \Phi_1, \Delta_2, a \vdash \langle a, \mathcal{E}[\mathbf{return}\ ()], \sigma, \rho, \omega \cup (b, N)\rangle}$$

noting that $\omega \cup (b, N)$ is safe since $\Gamma \vdash b : \mathsf{Pid}$ and $\Gamma \mid S \triangleright \mathcal{E}[\mathbf{return}\ ()] : \mathbf{1} \triangleleft \mathsf{end}$, as required.

Case E-InvokeM.

$$\langle a, \mathbf{idle}, \sigma, \rho, \omega \cup \{(b, M)\}\rangle \parallel \lightning b \xrightarrow{\tau} \langle a, M, \sigma, \rho, \omega\rangle \parallel \lightning b$$

Assumption:

$$\dfrac{\dfrac{\Gamma; \cdot \vdash \mathbf{idle} \qquad \Gamma; \Phi_1 \vdash \sigma \qquad \Gamma; \Phi_2 \vdash \rho}{\Gamma;\ \Phi_1, \Phi_2, a \vdash \langle a, \mathcal{T}, \sigma, \rho, \omega \cup \{(b, M)\}\rangle} \qquad \dfrac{}{\Gamma;\ b \vdash \lightning b}}{\Gamma;\ \Phi_1, \Phi_2, a, b \vdash \langle a, \mathcal{T}, \sigma, \rho, \omega \cup \{(b, M)\}\rangle \parallel \lightning b}$$

where $\forall (a', M) \in \omega \cup \{(b, N)\}.\ \Gamma \vdash b : \mathsf{Pid}\ \wedge\ \Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}$.

Recomposing:

$$\dfrac{\dfrac{\dfrac{\Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; \cdot \vdash M} \qquad \Gamma; \Phi_1 \vdash \sigma \qquad \Gamma; \Phi_2 \vdash \rho}{\Gamma;\ \Phi_1, \Phi_2, a \vdash \langle a, M, \sigma, \rho, \omega \rangle} \qquad \dfrac{}{\Gamma;\ b \vdash \lightning b}}{\Gamma;\ \Phi_1, \Phi_2, a, b \vdash \langle a, M, \sigma, \rho, \omega \rangle \parallel \lightning b}$$

as required.

Case E-Raise.

Similar to E-RaiseS.

Case E-RaiseS.

$$\langle a, (\mathcal{E}[\mathbf{raise}])^{s[\mathsf{p}]}, \sigma, \rho, \omega \rangle \xrightarrow{\tau} \lightning a \parallel \lightning s[\mathsf{p}] \parallel \lightning \sigma \parallel \lightning \rho$$

$$\dfrac{\dfrac{\Gamma \mid S \triangleright \mathcal{E}[\mathbf{raise}] : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; s[\mathsf{p}] : S \vdash (\mathcal{E}[\mathbf{raise}])^{s[\mathsf{p}]}} \qquad \Gamma; \Phi_1 \vdash \sigma \qquad \Gamma; \Phi_2 \vdash \rho}{\Gamma;\ \Phi_1, \Phi_2, s[\mathsf{p}] : S, a \vdash \langle a, (\mathcal{E}[\mathbf{raise}])^{s[\mathsf{p}]}, \sigma, \rho, \omega \rangle}$$

where $\forall (b, M) \in \omega.\ \Gamma \vdash b : \mathsf{Pid}\ \wedge\ \Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}$.

Let us write $\lightning \Phi = \{ s[\mathsf{p}] : \lightning \mid s[\mathsf{p}] : S \in \Phi \}$. It follows that for a given environment, $\Phi \rightsquigarrow^* \lightning \Phi$.

The result follows by noting that due to TH-Handler and TI-Callback we have that $\mathsf{fn}(\Phi_1) = \mathsf{fn}(\sigma)$ and $\mathsf{fn}(\Phi_2) = \mathsf{fn}(\rho)$. Thus:

- $\Gamma;\ \lightning \Phi_1 \vdash \lightning \sigma$,
- $\Gamma;\ \lightning \Phi_2 \vdash \lightning \rho$,
- $\Gamma;\ \lightning \Phi_1, \lightning \Phi_2, s[\mathsf{p}] : \lightning, a \vdash \lightning a \parallel \lightning s[\mathsf{p}] \parallel \lightning \sigma \parallel \lightning \rho$

with the environment reduction:

$$\Phi_1, \Phi_2, s[\mathsf{p}] : S, a \rightsquigarrow^+ \lightning \Phi_1, \lightning \Phi_2, s[\mathsf{p}] : \lightning, a$$

as required.

Case E-CancelMsg.

$$s \triangleright (\mathsf{p}, \mathsf{q}, \ell(V)) \cdot \delta \parallel \lightning s[\mathsf{q}] \xrightarrow{\tau} s \triangleright \delta \parallel \lightning s[\mathsf{q}]$$

Assumption:

$$\dfrac{\dfrac{\Gamma \vdash V : A \qquad \Gamma;\ s : Q \vdash s \triangleright \delta}{\Gamma;\ s : (\mathsf{p}, \mathsf{q}, \ell(A)) \cdot Q \vdash s \triangleright (\mathsf{p}, \mathsf{q}, \ell(V)) \cdot \delta} \qquad \Gamma;\ s[\mathsf{q}] : \lightning \vdash \lightning s[\mathsf{q}]}{\Gamma;\ s[\mathsf{q}] : \lightning, s : (\mathsf{p}, \mathsf{q}, \ell(V)) \cdot Q \vdash s \triangleright (\mathsf{p}, \mathsf{q}, \ell(V)) \cdot \delta \parallel \lightning s[\mathsf{q}]}$$

Recomposing, we have:

$$\dfrac{\Gamma;\ s : Q \vdash s \triangleright \delta \qquad \Gamma;\ s[\mathsf{q}] : \lightning \vdash \lightning s[\mathsf{q}]}{\Gamma;\ s[\mathsf{q}] : \lightning, s : Q \vdash s \triangleright \delta \parallel \lightning s[\mathsf{q}]}$$

with

$$s[\mathsf{q}] : \lightning, s : (\mathsf{p}, \mathsf{q}, \ell(V)) \cdot Q \xrightarrow{s:\mathsf{p} \lightning \mathsf{q}::\ell} s[\mathsf{q}] : \lightning, s : Q$$

as required.

Case E-CancelAP.

$$(\nu\iota)(p(\chi[\mathsf{p} \mapsto \widetilde{\iota'} \cup \{\iota\}]) \parallel \lightning\,\iota) \xrightarrow{\tau} p(\chi[\mathsf{p} \mapsto \widetilde{\iota'}])$$

Assumption:

$$\frac{\dfrac{\dfrac{p : \mathsf{AP}(\mathsf{p}_i : S_i)_i \qquad \dfrac{\{(\mathsf{p}_i : S_i)_i\} \ \Phi \vdash \chi}{\{(\mathsf{p}_i : S_i)_i\} \ \Phi, \widetilde{\iota'^- : S_j}, \iota^- : S_j \vdash \chi[\mathsf{p}_j \mapsto \widetilde{\iota'} \cup \{\iota\}]}}{\Gamma;\ \Phi, \widetilde{\iota'^- : S_j}, \iota^- : S_j \vdash p(\chi[\mathsf{p}_j \mapsto \widetilde{\iota'} \cup \{\iota\}])} \qquad \dfrac{}{\Gamma;\ \iota^+ : S_j \vdash \lightning\,\iota}}{\Gamma;\ \Phi, \widetilde{\iota'^- : S_j}, \iota^+ : S_j, \iota^- : S_j, p \vdash p(\chi[\mathsf{p}_j \mapsto \widetilde{\iota'} \cup \{\iota\}]) \parallel \lightning\,\iota}}{\Gamma;\ \Phi, \widetilde{\iota'^- : S_j}, p \vdash (\nu\iota)(p(\chi[\mathsf{p}_j \mapsto \widetilde{\iota'} \cup \{\iota\}]) \parallel \lightning\,\iota)}$$

Recomposing:

$$\frac{p : \mathsf{AP}(\mathsf{p}_i : S_i)_i \qquad \dfrac{\{(\mathsf{p}_i : S_i)_i\} \ \Phi \vdash \chi}{\{(\mathsf{p}_i : S_i)_i\} \ \Phi, \widetilde{\iota'^- : S_j} \vdash \chi[\mathsf{p}_j \mapsto \widetilde{\iota'}]}}{\Gamma;\ \Phi, \widetilde{\iota'^- : S_j}, p \vdash p(\chi[\mathsf{p}_j \mapsto \widetilde{\iota'}])}$$

as required.
Case E-CancelH.

$$\begin{array}{l}\langle a, \mathbf{idle}, \sigma[s[\mathsf{p}] \mapsto (V, M), \rho, \omega\rangle \parallel s \triangleright \delta \parallel \lightning\, s[\mathsf{q}] \xrightarrow{\tau} \\ \quad \langle a, M, \sigma, \rho, \omega\rangle \parallel s \triangleright \delta \parallel \lightning\, s[\mathsf{q}] \parallel \lightning\, s[\mathsf{p}] \quad \text{if } \mathsf{messages}(\mathsf{q}, \mathsf{p}, \delta) = \emptyset\end{array}$$

Let $\mathbf{D}$ be the following derivation:

$$\frac{\Gamma; \cdot \vdash \mathbf{idle} \qquad \dfrac{\begin{array}{cc} T = \mathsf{q}\,\&\{\ell_i(x_i) \mapsto S_i\}_i & \Gamma \vdash V : \mathsf{Handler}(T) \\ \Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end} & \Gamma; \Phi_1 \vdash \sigma \end{array}}{\Gamma; \Phi_1, s[\mathsf{p}] : T \vdash \sigma[s[\mathsf{p}] \mapsto (V, M)]} \qquad \Gamma; \Phi_2 \vdash \rho}{\Gamma;\ \Phi_1, \Phi_2, s[\mathsf{p}] : T, a \vdash \langle a, \mathbf{idle}, \sigma[s[\mathsf{p}] \mapsto (V, M)], \rho, \omega\rangle}$$

Assumption:

$$\frac{\mathbf{D} \qquad \dfrac{\Gamma;\ s : Q \vdash s \triangleright \delta \qquad \Gamma;\ s[\mathsf{p}] : \lightning \vdash \lightning\, s[\mathsf{p}]}{\Gamma;\ s : Q, s[\mathsf{p}] : \lightning \vdash s \triangleright \delta \parallel \lightning\, s[\mathsf{p}]}}{\Gamma;\ \Phi_1, \Phi_2, s[\mathsf{p}] : T, s : Q, s[\mathsf{q}] : \lightning, a \vdash \langle a, \mathbf{idle}, \sigma[s[\mathsf{p}] \mapsto (V, M)], \rho, \omega\rangle \parallel s \triangleright \delta \parallel \lightning\, s[\mathsf{p}]}$$

We can recompose as follows. Let $\mathbf{D'}$ be the following derivation:

$$\frac{\dfrac{\Gamma \mid \mathsf{end} \triangleright M : \mathbf{1} \triangleleft \mathsf{end}}{\Gamma; \cdot \vdash M} \qquad \Gamma; \Phi_1 \vdash \sigma \qquad \Gamma; \Phi_2 \vdash \rho}{\Gamma;\ \Phi_1, \Phi_2, a \vdash \langle a, M, \sigma, \rho, \omega\rangle}$$

Then we can construct the remaining derivation:

$$\frac{\mathbf{D} \qquad \dfrac{\Gamma;\ s : Q \vdash s \triangleright \delta \qquad \dfrac{\dfrac{}{\Gamma;\ s[\mathsf{p}] : \lightning \vdash \lightning\, s[\mathsf{p}]} \qquad \dfrac{}{\Gamma;\ s[\mathsf{q}] : \lightning \vdash \lightning\, s[\mathsf{q}]}}{\Gamma;\ s[\mathsf{p}] : \lightning, s[\mathsf{q}] : \lightning \vdash \lightning\, s[\mathsf{p}] \parallel \lightning\, s[\mathsf{q}]}}{\Gamma;\ s : Q, s[\mathsf{p}] : \lightning, s[\mathsf{q}] : \lightning \vdash s \triangleright \delta \parallel \lightning\, s[\mathsf{p}] \parallel \lightning\, s[\mathsf{q}]}}{\Gamma;\ \Phi_1, \Phi_2, s : Q, s[\mathsf{p}] : \lightning, s[\mathsf{q}] : \lightning, a \vdash \langle a, M, \sigma, \rho, \omega\rangle \parallel s \triangleright \delta \parallel \lightning\, s[\mathsf{p}] \parallel \lightning\, s[\mathsf{q}]}$$

Finally, we need to show environment reduction:

$$\Phi_1, \Phi_2, s[\mathsf{p}] : T, s : Q, s[\mathsf{q}] : ↯, a \xrightarrow{s:\mathsf{p}↯\mathsf{q}} \Phi_1, \Phi_2, s : Q, s[\mathsf{p}] : ↯, s[\mathsf{q}] : ↯, a$$

as required. □

D.2.2 Progress. Maty$_↯$ enjoys a similar progress property since E-CancelMsg discards messages that cannot be received, and E-CancelMsg invokes the failure continuation whenever a message will never be sent due to a failure; monitoring is orthogonal. The one change is that zapper threads for actors may remain if the actor name is free in an existing monitoring or initialisation callback.

We require a slightly-adjusted progress property on environments to account for session failure.

Definition D.3 (Progress (Cancellation-aware environments)). A runtime environment Φ *satisfies progress*, written $\mathsf{prog}_↯(\Phi)$, if $\Phi \Longrightarrow^* \Phi' \not\Longrightarrow$ implies that either $\Phi' = s : \epsilon$ or $\Phi' = (s[\mathsf{p}_i] : ↯)_i, s : \epsilon$.

We first need to define a canonical form that takes zapper threads into account.

Definition D.4 (Canonical form (Maty$_↯$)). A Maty$_↯$ configuration C is in *canonical form* if it can be written:

$$(\nu\tilde{\imath})(\nu p_{i\in 1..l})(\nu s_{j\in 1..m})(\nu a_{k\in 1..n})(p_i(\chi_i)_{i\in 1..l} \parallel (s_j \triangleright \delta_j)_{j\in 1..m} \parallel \langle a_k, \mathcal{T}_k, \sigma_k, \rho_k, \omega_k \rangle_{k\in 1..n'-1} \parallel \widetilde{↯\alpha})$$

with $↯a_{k k\in n'..n}$ contained in $\widetilde{↯\alpha}$.

As before, all well-typed configurations can be written in canonical form; as usual the proof relies on the fact that structural congruence is type-preserving.

Lemma D.5. *If $\Gamma; \Phi \vdash C$ then there exists a $\mathcal{D} \equiv C$ where $\mathcal{D}$ is in canonical form.*

It is also useful to see that the progress property on environments is preserved even if some roles become cancelled.

Lemma D.6. *If $\mathsf{prog}_↯(\Phi)$ then $\mathsf{prog}_↯(\mathit{zap}(\Phi, \widetilde{\mathsf{p}}))$ for any $\widetilde{\mathsf{p}} \subseteq \mathit{roles}(\Phi)$.*

Proof. Zapping a role may prevent Lbl-Recv from firing, but in this case would enable either a Lbl-ZapRecv and Lbl-ZapMsg reduction. □

Thread progress needs to change to take into account the possibility of an exception due to E-Raise or E-RaiseExn:

Lemma D.7 (Thread Progress). *Let $C = \mathcal{G}[\langle a, \mathcal{T}, \sigma, \rho \rangle]$. If $\cdot; \cdot \vdash C$ then either:*

- *$\mathcal{T} =$ **idle**, or*
- *there exist $\mathcal{G}', \mathcal{T}', \sigma', \rho'$ such that $C \longrightarrow \mathcal{G}'[\langle a, \mathcal{T}', \sigma', \rho' \rangle]$, or*
- *$C \longrightarrow \mathcal{G}'[↯a \parallel ↯\sigma \parallel ↯\rho]$ if $\mathcal{T} = \mathcal{E}[\textbf{raise}]$, or*
- *$C \longrightarrow \mathcal{G}'[↯a \parallel ↯s[\mathsf{p}] \parallel ↯\sigma \parallel ↯\rho]$ if $\mathcal{T} = (\mathcal{E}[\textbf{raise}])^{s[\mathsf{p}]}$.*

Proof. As with Lemma C.12 but taking into account that:

- **monitor** $b\ M$ can always reduce by E-Monitor;
- **raise** can always reduce by either E-Raise or E-RaiseS.

□

Theorem D.8 (Progress (Maty$_↯$)). *If $\cdot; \cdot \vdash_{\mathit{prog}_↯} C$, then either there exists some $\mathcal{D}$ such that $C \longrightarrow \mathcal{D}$, or C is structurally congruent to the following canonical form:*

$$(\nu\tilde{\imath})(\nu p_{i\in 1..m})(\nu a_{j\in 1..n})(p_1(\chi_1)_{i\in 1..m} \parallel \langle a_j, \textbf{idle}, \epsilon, \rho_j, \omega_j \rangle_{j\in 1..n'-1} \parallel (↯a_j)_{j\in n'..n})$$

Proof. The reasoning is similar to that of Theorem 3.6. By Lemma D.5, C can be written in canonical form:

$$(\nu\tilde{\iota})(\nu p_{i\in 1..l})(\nu s_{j\in 1..m})(\nu a_{k\in 1..n})(p_i(\chi_i)_{i\in 1..l} \parallel (s_j \triangleright \delta_j)_{j\in 1..m} \parallel \langle a_k, \mathcal{T}_k, \sigma_k, \rho_k, \omega_k\rangle_{k\in 1..n'-1} \parallel \widetilde{\lightning\alpha})$$

with $(\lightning a_k)_{k\in n'..n}$ contained in $\widetilde{\lightning\alpha}$.

By repeated applications of Lemma D.7, either the configuration can reduce or all threads are idle:

$$(\nu\tilde{\iota})(\nu p_{i\in 1..l})(\nu s_{j\in 1..m})(\nu a_{k\in 1..n})(p_i(\chi_i)_{i\in 1..l} \parallel (s_j \triangleright \delta_j)_{j\in 1..m} \parallel \langle a_k, \mathbf{idle}, \sigma_k, \rho_k, \omega_k\rangle_{k\in 1..n'-1} \parallel \widetilde{\lightning\alpha})$$

By the linearity of runtime type environments Δ, each role endpoint $s[\mathsf{p}]$ must either be contained in an actor, or exist as a zapper thread $\lightning s[\mathsf{p}] \in \widetilde{\lightning\alpha}$. Let us first consider the case that the endpoint is contained in an actor; we know by previous reasoning that each role must have an associated stored handler.

Since the types for each session must satisfy progress, the collection of local types must reduce. There are two potential reductions: either Lbl-Sync-Recv in the case that the queue has a message, or Lbl-ZapRecv if the sender is cancelled and the queue does not have a message. In the case of Lbl-Sync-Recv, since all actors are idle we can reduce using E-React as usual. In the case of Lbl-ZapRecv typing dictates that we have a zapper thread for the sender and so can reduce by E-CancelH.

It now suffices to reason about the case where all endpoints are zapper threads (and thus by linearity, where all handler environments are empty). In this case we can repeatedly reduce with E-CancelMsg until all queues are cleared, at which point we have a configuration of the form:

$$(\nu\tilde{\iota})(\nu p_{i\in 1..l})(\nu s_{j\in 1..m})(\nu a_{k\in 1..n})(p_i(\chi_i)_{i\in 1..l} \parallel (s_j \triangleright \epsilon)_{j\in 1..m} \parallel \langle a_k, \mathbf{idle}, \epsilon, \rho_k, \omega_k\rangle_{k\in 1..n'-1} \parallel \widetilde{\lightning\alpha})$$

We must now account for the remaining zapper threads. If there exists a zapper thread $\lightning a$ where a is contained within some monitoring environment ω then we can reduce with E-InvokeM. If a does not occur free in any initialisation callback or monitoring callback then we can eliminate it using the garbage collection congruence $(\nu a)(\lightning a) \parallel C \equiv C$.

Next, we eliminate all zapper threads for initialisation tokens using E-CancelAP.

Finally, we can eliminate all failed sessions $(\nu s)(\lightning s[\mathsf{p}_1] \parallel \cdots \parallel \lightning s[\mathsf{p}_n] \parallel s \triangleright \epsilon)$, and we are left with a configuration of the form:

$$(\nu\tilde{\iota})(\nu p_{i\in 1..m})(\nu a_{j\in 1..n})(p_1(\chi_1)_{i\in 1..m} \parallel \langle a_j, \mathbf{idle}, \epsilon, \rho_j, \omega_j\rangle_{j\in 1..n'-1} \parallel (\lightning a_j)_{j\in n'..n})$$

as required. □

D.2.3 Global Progress.

Lemma D.9 (Session Progress (Maty$_{\lightning}$)). *If* $\cdot;\cdot \vdash^{\mathsf{f}}_{prog} (\nu s : \Delta_s)C$, *then there exists some* $\mathcal{D}_1$ *such that* $C \xrightarrow{\tau}{}^{*} \mathcal{D}_1$ *and either* $\mathcal{D}_1 \xrightarrow{s}$, *or* $(\nu s)\mathcal{D}_1 \equiv \mathcal{D}_2$ *for some* $\mathcal{D}_2$ *where* $s \notin activeSessions(\mathcal{D}_2)$.

Proof. The proof is as with Theorem 3.8, except we must account for failed sessions arising as a consequence of reduction. □

A modified version of global progress holds: for every active session, in a finite number of reductions, either the session can make a communication action, or all endpoints become cancelled and can be garbage collected.

Theorem D.10 (Global Progress (Maty$_{\lightning}$)). *If* $\cdot;\cdot \vdash^{\mathsf{f}}_{prog_{\lightning}} C$*, then for every* $s \in \mathit{activeSessions}(C)$*, then there exist* $\mathcal{D}$ *and* $\mathcal{D}_1$ *such that* $C \equiv (\nu s)\mathcal{D}$ *where* $\mathcal{D} \xrightarrow{\tau}{}^{*} \mathcal{D}_1$ *and either* $\mathcal{D}_1 \xrightarrow{s}$*, or* $\mathcal{D}_1 \equiv \mathcal{D}_2$ *for some* $\mathcal{D}_2$ *where* $s \notin \mathit{activeSessions}(\mathcal{D}_2)$*.*

Proof. Arises as a corollary of Lemma D.9. □