

Alice or Bob?: Process polymorphism in choreographies

EVA GRAVERSEN

Department of Mathematics and Computer Science, University of Southern Denmark, Odense, Denmark
(e-mail: efgraversen@imada.sdu.dk)

ANDREW K. HIRSCH

Department of Computer Science and Engineering, University at Buffalo, SUNY, Buffalo, NY 14068, USA
(e-mail: akhirsch@buffalo.edu)

FABRIZIO MONTESI 

Department of Mathematics and Computer Science, University of Southern Denmark, Odense, Denmark
(e-mail: fmontesi@imada.sdu.dk)

Abstract

We present PolyChor λ , a language for higher-order functional *choreographic programming*—an emerging paradigm for concurrent programming. In choreographic programming, programmers write the desired cooperative behaviour of a system of processes and then compile it into an implementation for each process, a translation called *endpoint projection*. Unlike its predecessor, Chor λ , PolyChor λ has both type and *process* polymorphism inspired by System F $_{\omega}$. That is, PolyChor λ is the first (higher-order) functional choreographic language which gives programmers the ability to write generic choreographies and determine the participants at runtime. This novel combination of features also allows PolyChor λ processes to communicate *distributed values*, leading to a new and intuitive way to write delegation. While some of the functional features of PolyChor λ give it a weaker correspondence between the semantics of choreographies and their endpoint-projected concurrent systems than some other choreographic languages, we still get the hallmark end result of choreographic programming: projected programmes are deadlock-free by design.

1 Introduction

Distributed systems involve interacting processes. Usually, programmers write one programme per process and then compose those programmes in parallel. These programmes contain *send* and *receive* expressions which transmit data between processes. Predicting how the composition of programmes based on this method is challenging, so it is easy to write code that *deadlocks*, or gets stuck because patterns of sends and receives do not match. *Session types* (Honda, 1993; Honda *et al.*, 1998) can be used to describe the patterns of sends and receives in a programme, offering a foundation for static analyses aimed at preventing communication mismatches and deadlocks (Caires & Pfenning, 2010; Wadler, 2012; DeYoung *et al.*, 2012; Dardha *et al.*, 2012; Honda *et al.*, 2016; Scalas & Yoshida,



2019). Working with session types enables the programmer to ensure the communications in their system follow compatible send/receive patterns.

Alternatively, developers can use a *choreographic* language to programme the interactions that they wish to take place in the system directly from a global viewpoint (Montesi, 2023). *Choreographic programming* (Montesi, 2013) is a programming paradigm based on this idea with particularly well-explored foundations (Cruz-Filipe & Montesi, 2020; Montesi, 2023) and promising developments (see, e.g., Carbone & Montesi, 2013; López et al., 2016; Dalla Preda et al., 2017; Giallorenzo et al., 2021; Cruz-Filipe et al., 2022; Hirsch & Garg, 2022; Jongmans & van den Bos, 2022). In this paradigm, a programmer writes one programme as a choreography, which is then compiled to a programme for each process that is guaranteed to be correct by construction. Unlike session types, which only allow local code to be checked against them, choreographies compile to the local code itself. The syntax of choreographic programming languages is typically inspired by security protocol notation (Needham & Schroeder, 1978), where send and receive commands are written together as part of atomic instructions for expressing communication. This has two key advantages. First, it gives programmers the power to express the desired communication flow among processes, but without the burden of manually coding send and receive actions. Second, it ensures that there is no mismatch which can cause deadlock, a property that has become known as *deadlock-freedom by design* (Carbone & Montesi, 2013).

To see the power of this, consider the (in)famous bookseller example—a recurring example in the literature of choreographic programming and session types (Carbone & Montesi, 2013; Honda et al., 2016; Montesi, 2023). *Buyer* wants to buy a book from *Seller*. To this end, *Buyer* sends the title of the book—say, “The Importance of Being Earnest”—to *Seller*, who then sends back the price. *Buyer* then can compare the price with its budget and based on the result informs *Seller* that they want to buy the book if it is within their budget or informs them that they do not want to buy the book otherwise. We can describe this via the following choreography:

```

let  $x = \text{com}_{\text{Buyer}, \text{Seller}}$  (“The Importance of Being Earnest” @ Buyer)
in let  $y = \text{com}_{\text{Seller}, \text{Buyer}}$  (price_lookup  $x$ )
   in if  $y < \text{budget}$ 
      then  $\text{select}_{\text{Buyer}, \text{Seller}}$  Buy (() @ Seller)
      else  $\text{select}_{\text{Buyer}, \text{Seller}}$  Quit (() @ Seller)

```

(1.1)

In Listing (1.1), as in all choreographic programmes, computation takes place among multiple *processes* communicating via message passing. Values are located at processes; for example, in the first line of the choreography, the title of the book is initially located at *Buyer*. The function $\text{com}_{P, Q}$ communicates a value from the process *P* to the process *Q*. It takes a local value at *P* and returns a local value at *Q*.¹ Thus, x represents the string “The Importance of Being Earnest” at the process *Seller*, while y represents the price at the process *Buyer*. Finally, we check locally if the book’s price is in *Buyer*’s budget. Either way, we use function *select* to send a label from *Buyer* to *Seller* representing *Buyer*’s choice to

¹ Formally, we require a type annotation on $\text{com}_{P, Q}$ (see Section 3). We elide this here for clarity.

either proceed with the purchase or not. Either way, the choreography returns the dummy value () at **Seller**.

While most of the early work on choreographies focused on simple lower-order imperative programming like in the example above, recent work has shown how to develop higher-order choreographic programming languages. These languages allow a programmer to write deadlock-free code using the usual abstractions of higher-order programming, such as objects (Giallorenzo *et al.*, 2023) and higher-order functions (Hirsch & Garg, 2022; Cruz-Filipe *et al.*, 2022).

For instance, Listing (1.1) bakes in the title and the value of the book. However, we may want to use this code whenever **Buyer** wants to buy any book, and let **Buyer** use any local function to decide whether to buy the book at a price.

```

λ title : String @ Buyer.
λ buyAtPrice? : Int @ Buyer → ∅ Bool @ Buyer.
  let x = comBuyer,Seller title
  in let y = comSeller,Buyer (price_lookup x)
  in if buyAtPrice? y
    then selectBuyer,Seller Buy () @ Seller
    else selectBuyer,Seller Quit () @ Seller

```

(1.2)

Note the type of the function `buyAtPrice?`: it takes as input not just an integer, but an integer at **Buyer**; similarly, it returns a Boolean at **Buyer**. Moreover, the arrow is annotated with a set of processes, which in this case is empty (\emptyset). Other than those processes named in the input and output types of the function, these are the only processes who may participate in the computation of that function. Since that set is empty here, *no* other process may participate in the function—i.e., `buyAtPrice?` is local to **Buyer**. (Sometimes we wish for other processes to participate in the computation of a function, as we will see in Example 3.)

However, not every function with an \emptyset annotation is local. For instance, `comP,Q` is a function compatible with type $\tau \rightarrow \emptyset \tau$ for any type τ . Despite the fact that `comP,Q` is clearly not local, only **P** and **Q** are involved in the communication, leading to the \emptyset annotation. Similarly, just because the input and output of a function are at different locations does not mean that the function involves communication. For instance, it might be a constant function: the choreography $\lambda x : \text{Int} @ P. 5 @ Q$ has the same type of a communication of an integer from **P** to **Q**, i.e., $\text{Int} @ P \rightarrow \emptyset \text{Int} @ Q$.

A programmer using a higher-order choreographic language, like a programmer using any higher-order programming language, can write a programme once and use it in a large number of situations. For instance, by supplying different values of `title` and `buyAtPrice?`, the choreography in Listing (1.2) can be used to buy several different titles and **Buyer** can determine if they are willing to buy the book at the price using any method they desire.

While the move from first-order programming to higher-order programming is significant, previous work on the theoretical foundations of higher-order choreographic programming still did not account for other forms of abstraction (Hirsch & Garg, 2022; Cruz-Filipe *et al.*, 2022). In particular, they did not allow for *polymorphism*, where programmes can abstract over types as well as data, allowing them to operate in many more

settings; nor did they allow for *delegation*, where one process can ask another process to act in its stead.

These forms of abstraction are relatively standard: delegation is an important operation in concurrent calculi, and polymorphism is vital to modern programming. In choreographic programming, however, another form of abstraction becomes natural: abstraction over processes. Current higher-order choreographic languages require that code mention concrete process names. However, we often want to write more-generic code, allowing the same code to run on many processes. For example, Listing (1.2) allows **Buyer** to decide whether to buy a book from **Seller** using any local function `buyAtPrice?`. It would be more natural to write **Seller** as a book-selling *service* which different clients could interact with in the same way to buy a book.

In this paper, we tackle three new features for choreographic languages. Firstly, we show that abstraction over processes is a type of polymorphism, which we refer to as *process polymorphism*. Secondly, we extend $\text{Chor}\lambda$ —a simply-typed functional choreographic language—with polymorphism, including process polymorphism, and call this new language $\text{PolyChor}\lambda$. Thirdly, we add the ability to communicate distributed values such as functions. This gives us the ability to *delegate* (i.e., to send code to another process, which that process is then expected to run), giving a clean language to study all three forms of abstraction.

Let us examine the bookseller *service* in our extended language:

$$\begin{aligned}
 \Delta B &:: \text{Proc.} \\
 \lambda \text{ title} &: \text{String} @ B. \\
 \lambda \text{ buyAtPrice?} &: \text{Int} @ B \rightarrow \emptyset \text{ Bool} @ B. \\
 \text{let } x &= \text{com}_{B, \text{Seller}} \text{ title} \\
 \text{in let } y &= \text{com}_{\text{Seller}, B} (\text{price_lookup } x) \\
 \text{in if } &\text{buyAtPrice? } y \\
 &\text{then select}_{B, \text{Seller}} \text{ Buy } (() @ \text{Seller}) \\
 &\text{else select}_{B, \text{Seller}} \text{ Quit } (() @ \text{Seller})
 \end{aligned} \tag{1.3}$$

This programme allows a process named B to connect with **Seller** to buy a book. B then provides a string `title` and a decision function `buyAtPrice?`. Thus, we no longer have to write a separate function for every process which may want to buy a book from **Seller**.

While this addition may appear simple, it poses some unique theoretical challenges. First, the goal of a choreographic language is to compile a global programme to one local programme per process. However, since B does not represent any particular process, it is unclear how to compile the polymorphic code above. We solve this problem via a simple principle: each process knows its identity. With this principle in place, we can compile the code to a conditional in each process: one option to run if they take the role of B , and the other to run if they do not.

Notably, each process chooses dynamically which interpretation of the code to run. This flexibility is important, since we may want to allow different processes to occupy B 's place dynamically. For instance, we can imagine a situation where **Buyer₁** and **Buyer₂** work together to buy a particularly expensive book: perhaps they compare bank accounts, and whoever has more money buys the book for them to share. This can be achieved in our system with Listing (1.4), where `seller_service` is the name of the choreography from

Listing (1.3):

```

λ title : String @ Buyer1.
let x = comBuyer1, Buyer2 bank_balance1
in if x < bank_balance2
then selectBuyer2, Buyer1 Me selectBuyer2, Seller Me
    ( seller_service Buyer2 (comBuyer1, Buyer2 title)
      (λ z. z < bank_balance2) )
else selectBuyer2, Buyer1 You selectBuyer2, Seller Them
    (seller_service Buyer1 title (λ z. z < bank_balance1))
    (1.4)
    
```

Here Buyer₁ sends its bank balance, bank_balance₁ to Buyer₂, who compares the received value with its own balance, bank_balance₂. If Buyer₂ has the larger balance, then it informs Buyer₁ and Seller that Buyer₂ will be buying the book by means of the label “Me.” Buyer₁ then sends the book title to Buyer₂, which allows Buyer₂ and Seller to initiate the seller_service choreography using a buyAtPrice? function that checks whether the price is less than Buyer₂’s bank balance. If Buyer₁ has the larger balance, then Buyer₂ again informs Buyer₁ and Seller of who will be performing the role of buyer for the rest of the protocol, “You” and “Them” respectively. Then, Buyer₁ enters the seller_service choreography with similar input to the first case, except the title and buyAtPrice? are now located at Buyer₁.

A related challenge shows up in the operational semantics of our extended language. Languages like PolyChorλ generally have operational semantics which match the semantics of the compiled code by allowing *out-of-order execution*: redices in different processes might be reduced in any order. However, care must be taken with process polymorphism, since it may not be clear whether two redices are in the same or different processes.

In addition to type and process polymorphism, PolyChorλ is the first choreographic language to allow the communication of distributed values: values not located entirely at the sender. These values include full choreographies described by distributed functions, which can be used to model delegation. To see how process polymorphism and communication of distributed values enables delegation, consider Figure 1. Here, when a buyer asks for a book, the seller first checks whether it is in stock. If it is, the sale continues as normal. If not, the seller delegates to a second seller, which may sell the book to the buyer.

In more detail, after ascertaining that the book is not in stock, Seller informs B and Seller₂ that the rest of the choreography will be executed by Seller₂ in the place of Seller using two selections with label “Delegate.” Then, Seller sends first the rest of the choreography to Seller₂, followed the title of the requested book. Seller₂ uses its own lookup function to execute the code in Listing (1.2). Both Seller₂ and B need to be informed that the delegation is happening, since B needs to know that it should interact with Seller₂ rather than Seller.

In general, delegation poses a challenge: the third-party processes involved in a communicated value (processes that are neither the sender nor the receiver, such as B above) might need to change who they are going to interact with by swapping names (for instance, swapping Seller₂ and Seller above). As we will see, this challenge is relevant for both the

$$\begin{array}{l}
\Lambda B :: \text{Proc.} \\
\lambda \text{ title} : \text{String} @ B. \\
\lambda \text{ buyAtPrice?} : \text{Int} @ B \rightarrow_{\emptyset} \text{Bool} @ B. \\
\text{let } x = \text{com}_{B, \text{Seller}} \text{ title} \\
\text{in if found}(\text{price_lookup } x) \\
\text{then select}_{\text{Seller}, B} \text{ Continue} \\
\text{select}_{\text{Seller}, \text{Seller}_2} \text{ Disconnect} \\
\text{let } y = \text{com}_{\text{Seller}, B} (\text{price}(\text{price_lookup } x)) \\
\text{in if buyAtPrice? } y \\
\text{then select}_{B, \text{Seller}} \text{ Buy } (() @ B) \\
\text{else select}_{B, \text{Seller}} \text{ Quit } (() @ B) \\
\text{else select}_{\text{Seller}, B} \text{ Delegate} \\
\text{select}_{\text{Seller}, \text{Seller}_2} \text{ Delegate} \\
\text{let } F = \text{com}_{\text{Seller}, \text{Seller}_2} \left(\begin{array}{l}
\lambda \text{ title}_2 : \text{String} @ \text{Seller.} \\
\text{if found}(\text{price_lookup}_2 \text{ title}_2) \\
\text{then select}_{\text{Seller}, B} \text{ Continue} \\
\text{let } y' = \text{price}(\text{price_lookup}_2 \text{ title}_2) \\
\text{in let } y = \text{com}_{\text{Seller}, B} y' \\
\text{in if buyAtPrice? } y \\
\text{then select}_{B, \text{Seller}} \text{ Buy } (() @ B) \\
\text{else select}_{B, \text{Seller}} \text{ Quit } (() @ B) \\
\text{else select}_{\text{Seller}, B} \text{ Quit } (() @ B)
\end{array} \right) \\
\text{in let title}_2 = \text{com}_{\text{Seller}, \text{Seller}_2} x \\
\text{in } F \text{ title}_2
\end{array} \tag{1.5}$$

Fig. 1. Example of delegation.

type system and projection operation of PolyChor λ . For typing, the combination of process polymorphism and distributed value communication can make it difficult to statically determine where data are located. For projection, we need to ensure that the third-party processes involved in a communicated value perform the required changes to process names in the right places during execution.

Structure of the Paper. We begin in Section 2 by examining the system model of PolyChor λ . We then proceed with the following contributions:

- In Section 3, we describe the PolyChor λ language in detail. This language includes both type polymorphism and process polymorphism. We develop both a type system and kind system and an operational semantics for PolyChor λ .
- In Section 4, we describe the local *network language* used to describe the distributed implementation. We also detail how to obtain this implementation via *endpoint projection*, which compiles PolyChor λ programmes to a programme for each process.
- In Section 5, we describe the main theorem of this paper, the correctness of endpoint projection with respect to our operational semantics. Because of the dynamic nature

of process polymorphism, this requires significant reasoning compared to previous works on choreographies.

- In Section 6, we demonstrate how our theory can be used to model an extended example where an edge computer can delegate tasks to an external server.

Finally, we discuss related work in Section 7 and conclude in Section 8.

2 System model

We begin by discussing the assumptions we make about how PolyChor λ programmes will be run. These assumptions are as light as possible, allowing for PolyChor λ to be run in many different scenarios. In particular, we assume that we have a fixed set of processes, which can communicate via messages. These processes can each be described by a polymorphic λ -calculus, similar to System $F\omega$, but with the addition of communication primitives.

2.1 Processes

We assume that there is a fixed set \mathbf{N} of process names \mathbf{P} , \mathbf{Q} , *Alice*, et cetera. These processes can represent nodes in a distributed system, system processes, threads, or more. Process polymorphism allows us to refer to processes using type variables, which may go in or out of scope. Despite this, the set of physically-running processes remains the same.

We assume every process knows its identity. Thus, every process can choose what code to run on the basis of its identity. This assumption is reasonable for many practical settings, for instance it is common for nodes in distributed systems to know their identity. This capability is essential to our strategy for enabling process polymorphism.

2.2 Communication

We assume that processes communicate via synchronous message passing. Thus, if \mathbf{P} sends a message to \mathbf{Q} , then \mathbf{P} does not continue until \mathbf{Q} has received the message. Moreover, we assume that message passing is instantaneous and certain, so messages do not get lost.

Processes can receive two kinds of messages: values of local programmes (described below) and *labels* describing choices made during a computation. These are used to ensure that different processes stay in lock-step with each other.

2.3 Local programmes

We assume that processes run a *local* language described in Section 4. This is a functional language extended with communication features, similar to the language GV (Gay & Vasconcelos, 2010; Wadler, 2012; Lindley & Morris, 2015). Even more related to our work is FST (System F with Session Types) Lindley & Morris (2017), an extension of GV with polymorphism. As it does not have our communication of distributed values, they can base their types on System F rather than System $F\omega$.

Variables	x, y, \dots	
Type Variables	X, Y, \dots	
Integers	n	
Labels	ℓ	
Process Names	P	
Process-Name Sets	ρ	$\in 2^{\text{Type Values}}$
Kinds	K	$::= * \mid K_1 \Rightarrow K_2 \mid \text{Proc} \mid K \setminus \rho$
Types	τ	$::= v \mid \tau_1 \tau_2 \mid \tau_1 \rightarrow_{\rho} \tau_2$ $\mid \tau_1 + \tau_2 \mid \tau_1 \times \tau_2 \mid \forall X :: K. \tau \mid \lambda X :: K. \tau$
Type Values	v	$::= X \mid () @ v \mid \text{Int} @ v \mid v_1 \rightarrow_{\rho} v_2 \mid P$ $\mid v_1 + v_2 \mid v_1 \times v_2 \mid \forall X :: K. v \mid \lambda X :: K. v$
Expressions	M, N, \dots	$::= x \mid () @ v \mid n @ v \mid \lambda x : \tau. M \mid \Lambda X :: K. M$ $\mid M N \mid M \tau \mid \text{inl}_{\tau} M \mid \text{inr}_{\tau} M$ $\mid \text{case } M \text{ of inl } x \Rightarrow N_1; \text{ inr } y \Rightarrow N_2$ $\mid (M, N) \mid \text{fst } M \mid \text{snd } M$ $\mid \text{com}_{v_1, v_2}^{\tau} \mid \text{select}_{v_1, v_2} \ell M \mid f$
Values	V	$::= x \mid () @ v \mid n @ v \mid \lambda x : \tau. M \mid \Lambda x :: K. M$ $\mid \text{inl}_{\tau} V \mid \text{inr}_{\tau} V \mid (V_1, V_2)$ $\mid \text{com}_{v_1, v_2}^{\tau}$

Fig. 2. PolyChor λ syntax.

Endpoint projection translates PolyChor λ into this “Network Process” language. We have thus further extended GV with features required for our endpoint projection mechanism. For instance, in the local language described in Section 4 we provide an *Aml* expression form, which allows a process to choose which code to run based on its identity. Despite these extensions, the language should feel familiar to any reader familiar with polymorphic λ -calculi.

3 The polymorphic Chor λ language

We now turn to our first major contribution: the design of the polymorphic, choreographic λ -calculus, and PolyChor λ . This calculus extends the choreographic λ -calculus Chor λ of Cruz-Filipe *et al.* (2022) with both type and, more importantly, process polymorphism. We begin by describing the features that PolyChor λ shares with the base Chor λ before describing the new features. The syntax of PolyChor λ can be found in Figure 2.

Syntax Inherited from Chor λ . Since choreographic programmes describe the behaviour of an entire communicating network of processes, we need to reason about where terms are located. In other words, we need to know which processes store the data denoted by a term. Terms of base type, like integers, are stored by exactly one process. This is represented in our type system by matching base types with a process name. For example, integers stored by the process *Alice* are represented by the type $\text{Int} @ \text{Alice}$. Values of this type also mark the process which stores them, so a value $5 @ \text{Alice}$ (read “the integer 5 at *Alice*”) has type

Int @ Alice. In Figure 2, the only base types are $() @ P$ and $\text{Int} @ P$, but it is easy to extend the language with other base types, such as the types $\text{String} @ P$ or $\text{Bool} @ P$ used in the introduction. We will continue to freely use other base types in our examples.

While base types are located on just one process, data of more-complex types may involve multiple processes. For instance, the term $(5 @ \text{Alice}, 42 @ \text{Bob})$ involves both data stored by **Alice** and **Bob**. This is still recorded in the type: the term above has type $\text{Int} @ \text{Alice} \times \text{Int} @ \text{Bob}$. In addition to base types and product types, PolyChor λ also has sum types (written $\tau_1 + \tau_2$), along with their normal introduction and elimination forms. Note that products and coproducts in PolyChor λ may not represent a product or coproduct at the local level, since each component may be at a different process. For instance, we can represent distributed Booleans as $\text{Bool} @ \text{Alice} \times \text{Bool} @ \text{Bob} + \text{Bool} @ \text{Alice} \times \text{Bool} @ \text{Bob}$. Matching on a value with this type will cause both **Alice** and **Bob** to make the same choice.

Functions are treated more unusually: while we have standard λ and application forms, we also allow functions to be defined mutually recursively with each other. In order to do so, any PolyChor λ choreography is associated with a list, D , of bindings of functions to *function variables* f , which are also expressions. A function variable can then during execution be instantiated with its definition according to this list. As we will see in Section 3.3, PolyChor λ terms are evaluated in a context which associates each function variable with a term. Note that, while in the original Chor λ types were mutually recursive in a similar way, in PolyChor λ we do not support recursive types. To see why, note that we syntactically restrict many types to *type values*. This prevents us having to reason about processes denoted by arbitrary terms—e.g., we cannot send to the “process” $(\lambda X :: \text{Proc}. X) P$ but we can write $(\Delta Y :: \text{Proc}. \text{com}_{Q,Y}^{\tau}) ((\lambda X :: \text{Proc}. X) P)$ which, due to our call-by-value semantics, will force the type to reduce to P before Y gets instantiated. As we will see in Section 4, allowing communication between arbitrary types would make endpoint projection difficult. However, since recursive types cannot necessarily reduce to a type value, they cannot be used in many parts of the type system.

Function types are also more specific than their usual construction in λ -calculus: they are written $\tau_1 \rightarrow_{\rho} \tau_2$. Here, ρ is a set of process names and type variables denoting additional participants in the function which do not have either the input or output. Thus, if **Alice** wants to communicate an integer to **Bob** directly (without intermediaries), then she should use a function of type $\text{Int} @ \text{Alice} \rightarrow_{\emptyset} \text{Int} @ \text{Bob}$. However, if she is willing to use the process **Proxy** as an intermediary, then she should use a function of type $\text{Int} @ \text{Alice} \rightarrow_{\{\text{Proxy}\}} \text{Int} @ \text{Bob}$. We will use ρ when projecting to determine that the function in question and any uses thereof must be part of the local code of **Proxy**.

In order to allow values to be communicated between processes, we provide the primitive communication function $\text{com}_{P,Q}^{\tau}$. This function takes a value of type τ at P and returns the corresponding value at Q . As mentioned in the introduction, most choreographic languages provide a communication term modelled after the “Alice-and-Bob” notation of cryptographic protocols. For instance, **Alice** \rightarrow **Bob**: 5 might represent **Alice** sending 5 to **Bob**. This is easily recovered by applying the function $\text{com}_{\text{Alice},\text{Bob}}^{\tau}$. For example, the term $\text{com}_{\text{Alice},\text{Bob}}^{\text{Int}@\text{Alice}} (5 @ \text{Alice})$ represents **Alice** sending a message containing 5 to **Bob**: it evaluates to $5 @ \text{Bob}$ and has type $\text{Int} @ \text{Bob}$.

Finally, consider the following, where M has type $\text{Int @ Alice} + \text{Int @ Alice}$:

```
case M of
  inl x ⇒ 3 @ Bob;
  inr y ⇒ 4 @ Bob
```

Clearly, **Bob** needs to know which branch is taken, since he needs to store a different return value in each branch. However, only **Alice** knows which whether M evaluates to $\text{inr}_{\text{Int@Alice}} V$ or $\text{inl}_{\text{Int@Alice}} V$ (here **inl** and **inr** are used to denote that a value is either the right or left part of a sum and annotated with the type of the other part of the sum to ensure type principality). Thus, this choreography cannot correspond to any network programme. Using the terminology found in the literature of choreographic languages, we might say that the choreography is *unrealisable* because there is insufficient *knowledge of choice* (Castagna et al., 2012; Montesi, 2023).

In order to enable programmes where a process’s behaviour differs depending on other processes data, such as how **Bob** behaved differently depending on **Alice**’s data, we provide **select** terms. These allow one process to tell another which branch has been taken, preventing knowledge from “appearing out of nowhere.” For instance, we can extend the programme above to:

```
case M of
  inl x ⇒ selectAlice,Bob Left (3 @ Bob);
  inr y ⇒ selectAlice,Bob Right (4 @ Bob)
```

This represents the same programme as above, except **Alice** tells **Bob** whether the left or the right branch has been taken. Unlike the previous version of this example, it *does* represent a (deadlock-free) network programme. In general, we allow arbitrary labels to be sent by **select** terms, so semantically-meaningful labels can be chosen.

While **com** and **select** both transfer information between two processes, they differ in what information they transfer. **com** moves a value, e.g., as an integer or a function, from the sender to the receiver. **select** on the other hand uses a label to inform the receiver of a choice made by the sender. Some choreographic languages combine the two, so both a label and a value is communicated at the same time, but like most choreographic languages PolyChor λ keeps the two separate.

Syntax Additions over Chor λ . In order to achieve (both type and process) polymorphism in PolyChor λ , we add several features based on System $F\omega$ (Girard, 1972). In particular, we add kinds and universal types $\forall X :: K. \tau$ along with type abstraction and application. From System $F\omega$, we inherit the kind $*$, which is the kind of types. We additionally inherit the kind $K_1 \Rightarrow K_2$ which represents functions from types to types.

Moreover, we inherit type-level functions $\lambda X :: K. \tau$ from System $F\omega$. These represent the definition of type constructors. We also have type-level function application $\tau_1 \tau_2$. Since types contain computation, we also define type *values*, which are simply types without application.

We use type-level functions for two primary purposes. First, we can use it to denote types which depend on process names, such as $\lambda X :: \text{Proc}. \text{Int @ } X$ and $\lambda X :: \text{Proc} \Rightarrow * . X P$. Second, we use type-level functions to type communications, as we will see in Section 3.1.

Note that the base types $() @ v$ and $\text{Int} @ v$, like local values, are *syntactically* restricted to only allow type values as subterms. This allows us to use a type variable to compute the location of a value dynamically, but not arbitrary terms, which would make it much harder to tell at time of projection where the value is located. Thus, we can write $(\lambda X :: \text{Proc. Int} @ X) (Y P)$ to compute the location of an integer dynamically ($Y P$ has to reduce to a type value before X can be instantiated), but we cannot write $\text{Int} @ (Y P)$ directly. This way, our projected calculus can tell when instantiating X (at runtime) whether it gets instantiated as P . It would be more complicated to create runtime checks for whether Y gets instantiated as a function type that outputs P or not.

In addition to the kinds $*$ and $K_1 \Rightarrow K_2$ of System $F\omega$, we also have the kind Proc of *process names*. Thus, process names are types, but they cannot be used to type any terms.

Additionally, we have *Without kinds* $K \setminus \rho$, which represents types of kind K which do not mention any of the processes in the set ρ . We also refer to this kind as having a restriction of the processes in ρ . Since we restrict the types that can be communicated based on which processes they contain, as we will see soon, the Without kind can be used to define polymorphic functions which contain communication. For instance, the term

$$\Lambda X :: \text{Proc. } \Lambda Y :: \text{Proc} \setminus \{X\}. \text{com}_{X,Y}^{\text{Int}@X} (5 @ X)$$

defines a function which, given *distinct* processes X and Y , causes X to send 5 to Y . As we will see in Section 3.2, restricting the processes involved in a type (and therefore the term being typed) is essential for typing communications. In particular, we need to ensure that a sender never tries to send something located at the receiver. Moreover, we need to ensure that every part of the communicated value located at the sender actually gets moved to the receiver, even if its location is an uninstantiated type variable.

In the rest of this section, we explore the semantics of PolyChor λ . First, we look at its static semantics, both in the form of typing and kinding. Second, we describe its operational semantics. Throughout, we will continue to give intuitions based on the concurrent interpretation of PolyChor λ , though the semantics we give here does not correspond directly to that interpretation.

3.1 Typing

We now turn to the type system for PolyChor λ . As before, our type system builds on that for Chor λ . Here, we focus on the rules that are new in this work. Thus, we focus on rules related to polymorphism, and those that have had to change due to polymorphism.

Typing judgements for PolyChor λ have the form $\Theta; \Gamma \vdash M : \tau$, where Θ is the set of process names—either names in \mathbf{N} or type variables with kind Proc —used in M or the type of M . The *typing environment* Γ is a list associating variables and function names to their types and type variables and process names to their kinds. We sometimes refer to the pair $\Theta; \Gamma$ as a *typing context*.

Selected rules for our type system can be found in Figure 3. The full collection of rules are given in Appendix A. Again, many of the rules are inherited directly from Chor λ (Cruz-Filipe *et al.*, 2022); we thus focus on the rules that have changed due to our additions. Many, if not most, of these rules are inspired by System $F\omega$. However, the addition of

$$\begin{array}{c}
\text{[TUNIT]} \frac{\Theta; \Gamma \vdash v :: \text{Proc}}{\Theta; \Gamma \vdash () @ v : () @ v} \qquad \text{[TINT]} \frac{\Theta; \Gamma \vdash v :: \text{Proc}}{\Theta; \Gamma \vdash n @ v : \text{Int} @ v} \\
\\
\text{[TAPP]} \frac{\Theta; \Gamma \vdash N : \tau_1 \rightarrow_{\rho} \tau_2 \quad \Theta; \Gamma \vdash M : \tau_1}{\Theta; \Gamma \vdash N M : \tau_2} \\
\\
\text{[TABS]} \frac{\Theta; \Gamma \vdash \tau_1 :: * \quad \Theta; \Gamma' \vdash v :: \text{Proc} \text{ for all } v \in \rho \quad \Theta \cap (\rho \cup \text{ip}(\tau_1) \cup \text{ip}(\tau_2) \cup \text{ftv}(\tau_1) \cup \text{ftv}(\tau_2)); \Gamma, x : \tau_1 \vdash M : \tau_2}{\Theta; \Gamma \vdash \lambda x : \tau_1. M : \tau_1 \rightarrow_{\rho} \tau_2} \\
\\
\text{[TSEL]} \frac{\Theta; \Gamma \vdash v_1 :: \text{Proc} \quad \Theta; \Gamma \vdash v_2 :: \text{Proc} \quad \Theta; \Gamma \vdash M : \tau}{\Theta; \Gamma \vdash \text{select}_{v_1, v_2} \ell M : \tau} \\
\\
\text{[TCOM]} \frac{\Theta; \Gamma \vdash \tau :: \text{Proc} \Rightarrow * \quad \Theta; \Gamma \vdash v_1 :: \text{Proc} \setminus (\text{mp}(\tau) \cup \text{ftv}(\tau)) \quad \Theta; \Gamma \vdash v_2 :: \text{Proc} \setminus (\text{mp}(\tau) \cup \text{ftv}(\tau))}{\Theta; \Gamma \vdash \text{com}_{v_1, v_2}^{\tau} : (\tau v_1 \rightarrow_{\emptyset} \tau v_2)} \\
\\
\text{[TAPFT]} \frac{\Theta; \Gamma \vdash M : \forall X :: K. \tau_1 \quad \Theta; \Gamma \vdash \tau_2 :: K}{\Theta; \Gamma \vdash M \tau_2 : \tau_1[X \mapsto \tau_2]} \\
\\
\text{[TABST1]} \frac{\Theta, X; \Gamma + X \& \rho \setminus \{X\}, X :: \text{Proc} \setminus \rho \vdash M : \tau}{\Theta; \Gamma \vdash \Lambda X :: \text{Proc} \setminus \rho. M : \forall X :: \text{Proc} \setminus \rho. \tau} \\
\\
\text{[TABST2]} \frac{\Theta, X; \Gamma + X, X :: \text{Proc} \vdash M : \tau}{\Theta; \Gamma \vdash \Lambda X :: \text{Proc}. M : \forall X :: \text{Proc}. \tau} \\
\\
\text{[TABST3]} \frac{\Theta; \Gamma + X \& \rho \setminus \{X\}, X :: K \setminus \rho \vdash M : \tau \quad K \neq \text{Proc}}{\Theta; \Gamma \vdash \Lambda X :: K \setminus \rho. M : \forall X :: K \setminus \rho. \tau} \\
\\
\text{[TABST4]} \frac{\Theta, X; \Gamma + X, X :: K \vdash M : \tau \quad K \neq \text{Proc} \quad \nexists K', \rho. K = K' \setminus \rho}{\Theta; \Gamma \vdash \Lambda X :: K. M : \forall X :: K. \tau} \\
\\
\text{[TEQ]} \frac{\Theta; \Gamma \vdash M : \tau_1 \quad \tau_1 \equiv \tau_2 \quad \Theta; \Gamma \vdash \tau_2 :: *}{\Theta; \Gamma \vdash M : \tau_2}
\end{array}$$

Fig. 3. Typing rules (Selected).

the kind of processes and Without kinds—i.e., kinds of the form $K \setminus \rho$ —also lead to some changes.

The rules [Tunit] and [Tint] give types to values of base types. Here, we have to ensure that the location of the term is a process. Intuitively, then, we want the location to have kind Proc . However, it might be a Without kind—that is, it might be of the form $\text{Proc} \setminus \rho$.

In this case, our subkinding system (which you can find details about in Section 3.2) still allows us to apply the rule.

We express function application and abstraction via the [Tapp] and [Tabs] rules, respectively. The application rule [Tapp] is largely standard—the only addition is the addition of a set ρ on the function type, as discussed earlier. The abstraction rule [Tabs], on the other hand, is more complicated. First, it ensures that the argument type, τ_1 , has kind $*$. Then, it ensures that every element in the set decorating the arrow is a process name—i.e., that it has kind **Proc**. Finally, it checks that, in an extended environment, the body of the function has the output type τ_2 . As is usual, this extended environment gives a type to the argument. However, it restricts the available process names to those in the set ρ and those mentioned in the types τ_1 and τ_2 .

There are two ways that a type τ can mention a process: it can either name it directly, or it can name it via a type variable. Thus, in the rule [Tabs] we allow the free variables of τ_1 and τ_2 to remain in the process context, computing them using the (standard) free-type-variable function where $\forall X :: K. M$ and $\lambda X :: K. M$ both bind X . However, we must also identify the *involved processes* in a type, which we write $\text{ip}(\tau)$ and compute as follows:

$$\begin{aligned} \text{ip}(X) &= \emptyset & \text{ip}(P) &= P & \text{ip}(() @ v) &= \text{ip}(\text{Int} @ v) = \text{ip}(v) \\ \text{ip}(v_1 \rightarrow_\rho v_2) &= \text{ip}(v_1) \cup \{P \mid P \in \rho\} \cup \text{ip}(v_2) \\ \text{ip}(\forall X :: K \setminus \rho. \tau) &= \text{ip}(\lambda X :: K \setminus \rho. \tau) = \text{ip}(\tau) \cup (\mathbf{N} \setminus \rho) \\ \text{ip}(\forall X :: K. \tau) &= \text{ip}(\lambda X :: K. \tau) = \mathbf{N} \text{ if } \nexists K', \rho. K = K' \setminus \rho \end{aligned}$$

The involved processes of other types are defined homomorphically.

The communication primitives `select` and `com` are typed with [Tsel] and [Tcom], respectively. A term `select` $_{v_1, v_2} \ell M$ behaves as M , where the process v_1 informs the process v_2 that the ℓ branch has been taken, as we saw earlier. Thus, the entire term has type τ if M does. Moreover, v_1 and v_2 must be processes.

The rule [Tcom] types `com` terms. So far we have been simplifying the type used in $\text{com}_{P, Q}^\tau$ for readability. We have been using τ to denote the input type, but as it turns out to type $\text{com}_{P, Q}^\tau$ correctly, we have to complicate things a little. Intuitively, a term $\text{com}_{v_1, v_2}^\tau M$ represents v_1 communicating the parts of M on v_1 to v_2 . Thus, we require that τ be a type *transformer* requiring a process. Moreover, v_1 and v_2 cannot be mentioned in τ ; otherwise, not every part of the type of M on v_1 in our example above would transfer to v_2 . For this, we use the following notion of *mentioned processes*:

$$\begin{aligned} \text{mp}(X) &= \emptyset & \text{mp}(P) &= P & \text{mp}(() @ v) &= \text{mp}(\text{Int} @ v) = \text{mp}(v) \\ \text{mp}(v_1 \rightarrow_\rho v_2) &= \text{mp}(v_1) \cup \{P \mid P \in \rho\} \cup \text{mp}(v_2) \\ \text{mp}(\forall X :: K \setminus \rho. \tau) &= \text{mp}(\lambda X :: K \setminus \rho. \tau) = \text{mp}(\tau) \cup \rho \\ \text{mp}(\forall X :: K. \tau) &= \text{mp}(\lambda X :: K. \tau) = \text{mp}(\tau) \text{ if } \nexists K', \rho. K = K' \setminus \rho \end{aligned}$$

Again, with other types being defined homomorphically. The difference between involved and mentioned processes is subtle. If there is no polymorphism, they are the same, but when dealing with polymorphism with restriction they are opposites: involved processes

includes every process not in the restriction (the variable could be instantiated as something involving those processes and thus they may be involved), while mentioned names includes the processes mentioned in the restriction. Mentioned names is used only when typing `com`. If we have such a type-level function, τ , and two type values ν_1 and ν_2 which are not and will not be instantiated to anything mentioned in τ then we can type `com $_{\nu_1, \nu_2}^{\tau}$` as a function from $\tau \nu_1$ to $\tau \nu_2$. Since this is direct communication, no intermediaries are necessary and we can associate this arrow with the empty set \emptyset .

It is worth noting at this point that the communication rule inspired our use of System $F\omega$ rather than plain System F, which lacks type-level computation. In `Chor λ` and other previous choreographic languages, communicated values must be local to the sender. In `PolyChor λ` , this would mean not allowing the communicated type to include type variables or processes other than the sender. Since we are introducing the idea of using communication as a means of delegation, we have slackened that restriction. This means that `PolyChor λ` programmes can communicate larger choreographies whose type may involve other processes, and importantly other type variables. We see this in the delegation example Listing (1.5), where we have the communication `com $_{Seller, Seller_2}$` . Adding in the required type annotation (which we had suppressed in the introduction), this becomes `com $_{Seller, Seller_2}^{\lambda X::Proc. String@X \rightarrow \emptyset ()@B}$` . Note that this still leaves us with a free type variable B , representing the unknown process that `Seller` is telling `Seller $_2$` to interact with! Since we cannot ban free type variables in communicated types, we must create a typing system that can handle them, and this requires type-level computation.

To see why this led us to type-level computation, consider the alternative. In `Chor λ` and other choreographic works, we would have a type communication using process *substitution* instead of communication. The annotated programme would then be `com $_{Seller, Seller_2}^{String@Seller \rightarrow \emptyset ()@B}$` . When applied to a programme of appropriate type, the result would have type

$$(\text{String @ Seller} \rightarrow_{\emptyset} () @ B)[\text{Seller} \mapsto \text{Seller}_2] = \text{String @ Seller}_2 \rightarrow_{\emptyset} () @ B$$

Note that, because B is a type variable, it was ignored by the substitution. If B is later instantiated as `Seller`, then we must substitute B with `Seller $_2$` in the output type. Thus, we need some mechanism to delay this substitution; rather than use a mechanism like explicit substitutions, we instead reached for the standard tool of System $F\omega$. The communication winds up instead being written as `com $_{Seller, Seller_2}^{\lambda X::Proc. String@X \rightarrow \emptyset ()@B}$` with X being instantiated as `Seller` in the input type and `Seller $_2$` in the output type. This seemed more elegant and less ad-hoc; moreover, it adds features which a real-world implementation of `PolyChor λ` would want anyway. To ensure that B does not get instantiated incorrectly, we use our `Without` kinds. Rule `[Tcom]` requires that both `Seller` and `Seller $_2$` are restricted on B , which, thanks to our restrictions being symmetric, means that B cannot be instantiated as either of them. The `Without` kinds here prevent nonsensical typings of `com` where, in the type, part of the output does not get moved from the sender to the receiver. This can happen if a type variable present in the type of the communicated value is instantiated during execution before the communication takes place, but has not yet been instantiated when we type the choreography. Were it not for the restrictions imposed by `Without` kinds, we would allow the choreography

$$(\Delta B :: \text{Proc. } \lambda f : \text{String @ Seller} \rightarrow_{\varnothing} () @ B. (\text{com}_{\text{Seller}, \text{Seller}_2}^{\lambda X :: \text{Proc. String @ X} \rightarrow_{\varnothing} () @ B} f)) \text{Seller}$$

to be typed as $(\text{String @ Seller} \rightarrow_{\varnothing} () @ \text{Seller}) \rightarrow_{\varnothing} (\text{String @ Seller}_2 \rightarrow_{\varnothing} () @ \text{Seller})$, which implies that part of the function is still at Seller after the communication is executed. This is not what will happen when actually executing the choreography, so the type is wrong. The Without kinds ensure that the choreography cannot be typed: the kind of B must be $\text{Proc} \setminus \{\text{Seller}, \text{Seller}_2\}$, and therefore, it cannot be instantiated as Seller.

Returning now to the typing rules of Figure 3, we next have the [TappT], [TabsT1], [TabsT2], [TabsT3] and [TabsT4] rules, which type universal quantification. The [TappT] rule is completely standard, while the others are 4 cases of what to do with a type abstraction. Each of these rules have a different definition for the typing context of M , depending on the kind of X . As is standard, we check if the body of the function has the right type when the parameter X has kind K . But first, if X is a process as in [TabsT1] and [TabsT2], then we need to extend Θ with X . In addition, we must further manipulate the context in order to ensure that the types whose kinds are restricted on X correspond to the restriction on the kind of X .

First, the new type variable X may shadow a previously defined X . Thus, we need to remove X from any Without kinds already in the context. We do this using the following operation $K + v$:

$$(K \setminus \rho) + v = (K + v) \setminus (\rho \setminus \{v\})$$

We define $+$ on other kinds homomorphically and extend this to contexts as usual:

$$\Gamma + v = \{x : \tau \mid x : \tau \in \Gamma\} \cup \{X : K + v \mid X : K \in \Gamma\}$$

Furthermore, in [TabsT1] and [TabsT4] if X itself has a Without kind—that is, X 's kind tells us it cannot be any of the processes in ρ —then we need to symmetrically add a restriction on X to every type in ρ . Otherwise, we would not be able to use the roles in ρ in any place where we cannot use X , even though we know X will not be instantiated with them. We do this with the operation $\Gamma \& \rho \setminus X$, which we define as follows:

$$\begin{aligned} \Gamma \& \rho \setminus X = & \{x : \tau \mid x : \tau \in \Gamma\} \cup \{\tau :: K \mid \tau :: K \in \Gamma \text{ and } \tau \notin \rho\} \\ & \cup \{\tau :: K \setminus (\rho_2 \cup \{X\}) \mid \tau :: K \setminus \rho_2 \in \Gamma \text{ and } \tau \in \rho\} \\ & \cup \{\tau :: K \setminus \{X\} \mid \tau :: K \in \Gamma, K \neq K_2 \setminus \rho_2, \text{ and } \tau \in \rho\} \end{aligned}$$

With these operations in place, we can now fully understand how to type the type abstractions. When K is actually a Without kind, then we must handle both shadowing and symmetrical restrictions. However, when it is not a Without kind, we must only handle shadowing. We show an example where every possible complication

Example 1 (Typing complex type abstractions). Consider the following choreography, which takes a process A and sends an integer communication with A from P to Q :

$$M = \Delta A :: \text{Proc} \setminus \{P, Q\}. \text{com}_{P, Q}^{\Delta X :: \text{Proc. Int @ X} \rightarrow_{\varnothing} \text{Int @ A}} \text{com}_{P, A}^{\Delta Y :: \text{Proc. Int @ Y}}$$

That A has a Without kind and the fact that A is a process means that we will need to use Rule [TabsT1] when typing M . In order to illustrate the necessity of shadowing, we will

include an unnecessary process P_2 in our environment. Setting $\Theta = \{P, Q, P_2\}$, we start with the following judgement:

$$\Theta; P : \text{Proc}, Q : \text{Proc}, P_2 : \text{Proc} \setminus \{A\} \vdash M : \forall A :: \text{Proc} \setminus \{P, Q\}. \text{Int} @ Q \rightarrow_{\emptyset} \text{Int} @ A$$

We need to take into account both that A is a process and that it has a Without kind in order to make the choreography typeable. First, we shadow, obtaining the following:

$$(P : \text{Proc}, Q : \text{Proc}, P_2 : \text{Proc} \setminus \{A\}) + A = P : \text{Proc}, Q : \text{Proc}, P_2 : \text{Proc}$$

so we get rid of any restrictions on previous variables called A . We then add the new symmetric restrictions necessary for typing the communication, as follows:

$$(P : \text{Proc}, Q : \text{Proc}, P_2 : \text{Proc}) \& \{P, Q\} \setminus A = P : \text{Proc} \setminus \{A\}, Q : \text{Proc} \setminus \{A\}, P_2 : \text{Proc}$$

Continuing on, we can abbreviate $K = \text{Proc} \setminus \{A\}$. Finally, we add A to the environment and Θ (writing $\Theta' = \Theta \cup \{A\}$), giving:

$$\Theta'; P : K, Q : K, P_2 : \text{Proc}, A : \text{Proc} \setminus \{P, Q\} \vdash N : \text{Int} @ Q \rightarrow_{\emptyset} \text{Int} @ A$$

where $M = \Delta A :: \text{Proc} \setminus \{P, Q\}. N$. Because of the restrictions in Rule [Tcom], N would not be typable if we had not made sure to add the symmetric restrictions. We will furthermore see in Section 3.2 that adding A to the set process names is also necessary when kinding it with the Proc kind.

On the other hand, although the rule looks bigger at first glance, it is much simpler to use Rule [TabsT4].

Example 2 (Typing simple type abstractions). Consider the following type abstraction, which takes a type A and applies a variable of that type to a function which also returns something of the same type:

$$\Delta A :: *. \lambda x : A. \lambda f : A \rightarrow_{\emptyset} A. f x$$

We can type this as

$$\emptyset; \emptyset \vdash \Delta A :: *. \lambda x : A. \lambda f : A \rightarrow_{\emptyset} A. f x : \forall A :: *. A \rightarrow_{\emptyset} A \rightarrow_{\emptyset} A \rightarrow_{\emptyset} A$$

Since we have no shadowing, the only way we have to manipulate our environment when entering the type abstraction is to add $A : *$ to the environment, giving us

$$\emptyset; A : * \vdash \lambda x : A. \lambda f : A \rightarrow_{\emptyset} A. f x : A \rightarrow_{\emptyset} A \rightarrow_{\emptyset} A \rightarrow_{\emptyset} A$$

Rules [TabsT2] and [TabsT3] are for cases of middling complexity. In Rule [TabsT2], we have to add the type variable to Θ , as in [TabsT1]. However, since we have no restrictions, we do not need to consider symmetric conflict. In Rule [TabsT3], we do consider symmetric conflicts, but do not add to Θ (since we are not dealing with a process).

The final addition to our type system is the rule [Teq]. This is another standard rule from System $F\omega$; it tells us that we are allowed to compute in types. More specifically, it tells us that we can replace a type with an equivalent type, using the following equivalence:

$$\begin{array}{c}
 \frac{}{\tau \equiv \tau} \qquad \frac{\tau_1 \equiv \tau_2}{\tau_2 \equiv \tau_1} \qquad \frac{\tau_1 \equiv \tau_2 \quad \tau_2 \equiv \tau_3}{\tau_1 \equiv \tau_3} \\
 \\
 \frac{\tau_1 \equiv \tau_1' \quad \tau_2 \equiv \tau_2'}{\tau_1 \rightarrow_{\rho} \tau_2 \equiv \tau_1' \rightarrow_{\rho} \tau_2'} \qquad \frac{\tau_1 \equiv \tau_1' \quad \tau_2 \equiv \tau_2'}{\tau_1 + \tau_2 \equiv \tau_1' + \tau_2'} \qquad \frac{\tau_1 \equiv \tau_1' \quad \tau_2 \equiv \tau_2'}{\tau_1 \times \tau_2 \equiv \tau_1' \times \tau_2'} \\
 \\
 \frac{\tau \equiv \tau'}{\lambda X :: K. \tau \equiv \lambda X :: K. \tau'} \qquad \frac{}{(\lambda X :: K. \tau_1) \tau_2 \equiv \tau_1[X \mapsto \tau_2]} \qquad \frac{\tau_1 \equiv \tau_1' \quad \tau_2 \equiv \tau_2'}{\tau_1 \tau_2 \equiv \tau_1' \tau_2'} \\
 \\
 \frac{\tau \equiv \tau'}{\forall X :: K. \tau \equiv \forall X :: K. \tau'}
 \end{array}$$

In addition to the rules in Figure 3 for typing choreographies, our type system needs one more rule for typing the definitions of our recursive functions. We also add an extra judgement of the form $\Theta; \Gamma \vdash D$ where $\Theta; \Gamma$ is a typing context as before, and D is a set of definitions for function variables—i.e., $D = \{f_1 = M_1, \dots, f_n = M_n\}$. We write $D(f)$ for the term associated with f in D . The only rule for this judgement is [Tdefs], which says that a set of definitions is well-formed if every variable in D is associated with a type τ in Γ , and the body of f in D can be given type τ in the context $\emptyset; \Gamma$. We require that the body of f can be typed with an empty set of roles because they are global predefined functions, and as such they should not be local to any one process.

$$\text{[TDEFS]} \frac{\forall f \in \text{domain}(D). f : \tau \in \Gamma \wedge \emptyset; \Gamma \vdash D(f) : \tau}{\Theta; \Gamma \vdash D}$$

3.2 Kinding

We finish our discussion of the static semantics of PolyChor λ by looking at our kinding system. Our kinding system uses only one judgement, $\Theta; \Gamma \vdash \tau :: K$, which says that in the typing context $\Theta; \Gamma$, the type τ has kind K . You can find the rules of our kinding system in Figure 4. These are mostly directly inherited from System $F\omega$. However, we must account for **Proc** and **Without** kinds.

For instance, the rules [Kunit] and [Kint] check that the type representing which process is storing the data indeed has the kind **Proc**. Similarly, [Kfun] ensures that all of the types in the set of possible intermediaries are processes. The rule for type variables, [Kvar], ensures that if a type variable X is assigned kind **Proc**, then X must also be in Θ .

One of the biggest differences between our kinding system and that of System $F\omega$, however, is the rule [Ksub] which tells us that our system enjoys *subkinding*. The subkinding rules come from the subset ordering on **Without** kinds. We also consider any kind equivalent to the same kind restricted on the empty set due to [SKEmpty] and [SKWithoutL].

$$\begin{array}{c}
\text{[KVAR]} \frac{X :: K \in \Gamma \quad \text{if } K \in \{\text{Proc}, \text{Proc} \setminus \rho\} \text{ then } X \in \Theta}{\Theta; \Gamma \vdash X :: K} \\
\text{[KROLE]} \frac{P :: K \in \Gamma \quad K \in \{\text{Proc}, \text{Proc} \setminus \rho\} \quad P \in \Theta \quad \text{if } K = \text{Proc} \setminus \rho \text{ then } P \notin \rho}{\Theta; \Gamma \vdash P :: K} \\
\text{[KUNIT]} \frac{\Theta; \Gamma \vdash \tau :: \text{Proc} \setminus \rho}{\Theta; \Gamma \vdash () @ \tau :: * \setminus \rho} \quad \text{[KINT]} \frac{\Theta; \Gamma \vdash \tau :: \text{Proc} \setminus \rho}{\Theta; \Gamma \vdash \text{Int} @ \tau :: * \setminus \rho} \\
\text{[KFUN]} \frac{\Theta; \Gamma \vdash \tau_1 :: * \setminus \rho_2 \quad \Theta; \Gamma \vdash \tau_2 :: * \setminus \rho_2 \quad \forall v \in \rho_1. \Theta; \Gamma \vdash v :: \text{Proc} \setminus \rho_2}{\Theta; \Gamma \vdash \tau_1 \rightarrow_{\rho_1} \tau_2 :: * \setminus \rho_2} \\
\text{[KABS]} \frac{\Theta; \Gamma, X :: K_1 \vdash \tau :: K_2}{\Theta; \Gamma \vdash \lambda X :: K_1. \tau :: K_1 \Rightarrow K_2} \quad \text{[KALL]} \frac{\Theta; \Gamma, X :: K \vdash \tau :: * \setminus \rho}{\Theta; \Gamma \vdash \forall X :: K. \tau :: * \setminus \rho} \\
\text{[KARR]} \frac{\Theta; \Gamma \vdash \tau :: (K_1 \setminus \rho) \Rightarrow (K_2 \setminus \rho)}{\Theta; \Gamma \vdash \tau :: (K_1 \Rightarrow K_2) \setminus \rho} \quad \text{[KSUB]} \frac{\Theta; \Gamma \vdash \tau :: K_1 \quad K_1 <: K_2}{\Theta; \Gamma \vdash \tau :: K_2} \\
\text{[KSUM]} \frac{\Theta; \Gamma \vdash \tau_1 :: * \setminus \rho \quad \Theta; \Gamma \vdash \tau_2 :: * \setminus \rho}{\Theta; \Gamma \vdash \tau_1 + \tau_2 :: * \setminus \rho} \\
\text{[KPROD]} \frac{\Theta; \Gamma \vdash \tau_1 :: * \setminus \rho \quad \Theta; \Gamma \vdash \tau_2 :: * \setminus \rho}{\Theta; \Gamma \vdash \tau_1 \times \tau_2 :: * \setminus \rho}
\end{array}$$

Fig. 4. Kinding rules.

The rules for subkinding are as follows:

$$\begin{array}{c}
\text{[SKREFL]} \frac{}{K <: K} \quad \text{[SKTRANS]} \frac{K_1 <: K_2 \quad K_2 <: K_3}{K_1 <: K_3} \\
\text{[SKARR]} \frac{K'_1 <: K_1 \quad K_2 <: K'_2}{K_1 \Rightarrow K_2 <: K'_1 \Rightarrow K'_2} \quad \text{[SKEMPTY]} \frac{}{K <: K \setminus \emptyset} \\
\text{[SKWITHOUTL]} \frac{K_1 <: K_2}{K_1 \setminus \rho <: K_2} \quad \text{[SKWITHOUTUNION]} \frac{K_1 <: K_2}{K_1 \setminus (\rho_1 \cup \rho_2) <: K_2 \setminus \rho_1}
\end{array}$$

Lemma 1. *Let τ be a type. If there exists a typing context $\Theta; \Gamma$ such that $\Theta; \Gamma \vdash \tau :: K$ then there exists a unique type value v such that $\tau \equiv v$.*

Proof The existence of v follows from induction on $\Theta; \Gamma \vdash \tau :: K$ and its uniqueness from induction on $\tau \equiv v$. ■

Lemma 2 (Type restriction). *Let τ be a type. If there exists a typing context $\Theta; \Gamma$ such that $\Theta; \Gamma \vdash \tau :: K \setminus \rho$ then $(ip(\tau) \cup ftv(\tau)) \cap \rho = \emptyset$.*

Proof Follows from kinding rules. ■

Theorem 1 (Kindable types). *Let M be a choreography and τ be a type such that $\Theta; \Gamma \vdash M : \tau$. Then $\Theta; \Gamma \vdash \tau :: *$.*

Proof Follows from induction on the derivation of $\Theta; \Gamma \vdash M : \tau$ and the kinding rules. ■

We also find that types have the same kinds as their equivalent type values. Due to β -expansion, a kindable type can be equivalent to an unkindable type, but not an unkindable type value.

Theorem 2 (Kind Preservation). *Let τ be a type. If there exists a typing context $\Theta; \Gamma$ such that $\Theta; \Gamma \vdash \tau :: K$, then $\Theta; \Gamma \vdash v :: K$ for any type value v such that $\tau \equiv v$.*

Proof Follows from the kinding and type equivalence rules. The only way that a kindable type τ can be equivalent to a type which is not kindable is when we have types τ_1 and τ_2 such that $\tau = \tau_1[X \mapsto \tau_2]$. In that case, if we use the rule $(\lambda X :: K. \tau_1) \tau_2 \equiv \tau_1[X \mapsto \tau_2]$ to create an unkindable $\tau' \equiv \tau$ with an extra application. However, this unkindable type is not a type value, and in fact, we must also use the same rule to remove this new type application before we get to a type value. ■

Example 3. We return to the delegation example (Listing (1.5)) and try to type it. As B appears free in the type of a value, F , being communicated between **Seller** and **Seller₂**, B must actually have the Without kind $\text{Proc} \setminus \{\text{Seller}, \text{Seller}_2\}$. The choreography therefore gets the type

$$\forall B :: \text{Proc} \setminus \{\text{Seller}, \text{Seller}_2\}. \\ \text{String} @ B \rightarrow_{\{\text{Seller}, \text{Seller}_2\}} ((\text{Int} @ B \rightarrow_{\emptyset} \text{Bool} @ B) \rightarrow_{\{\text{Seller}, \text{Seller}_2\}} ()) @ B$$

This type shows both the input, output, and involved roles of the choreography.

3.3 Operational semantics

Finally, we consider the operational semantics of PolyChor λ . In practice, the semantics of a choreographic language can be used to simulate a choreography and check if it specifies the expected collective behaviour. Its key role, however, is to prove properties about the projected local code. Specifically, we are going to prove that the projected code is compliant to the choreography (an operational correspondence result) and that as a result it is deadlock-free. The semantics of PolyChor λ are mostly a standard call-by-value reduction semantics for a typed λ calculus. However, the reduction semantics must also carry a set D of function definitions. Only a few rules are unusual or must be modified; those can be found in Figure 5. You can find the rest of the rules in Appendix B.

The rules [AppTAbS] and [MTApp1] come from System $F\omega$. The rule [AppTAbS] is similar to ordinary CBV β reduction, but tells us how to reduce a type abstraction applied

$$\begin{array}{c}
\text{[APPTABS]} \frac{\tau \equiv \nu}{(\Lambda X :: K. M) \tau \rightarrow_D M[X \mapsto \nu]} \qquad \text{[MTAPP1]} \frac{M_1 \rightarrow_D M_2}{M_1 \tau \rightarrow_D M_2 \tau} \\
\text{[DEF]} \frac{}{f \rightarrow_D D(f)} \\
\text{[SEL]} \frac{}{\text{select}_{P,Q} \ell M \rightarrow_D M} \qquad \text{[COM]} \frac{}{\text{com}_{P,Q}^\tau V \rightarrow_D V[P \mapsto Q]}
\end{array}$$

Fig. 5. Semantics of PolyChor λ (selected rules).

to a *type* value, but with the caveat that if we do not have a type value we must use type equivalence to get one before reducing. The rule [MTApp1] tells us that we can reduce a type function applied to any argument.

The rule [Def] allows us to reduce function names by looking up their definition in the set D .

Finally, we have the rules for communication. The rule [Sel] says that `select` acts as a no-op, as we stated earlier. While this may seem redundant, such terms are vital for projection, as we will see in the next section. More importantly, the [Com] rule tells us how we represent communication at the choreography level: via substitution of roles. This also helps explain some of the restrictions in [Tcom]. Since we replace *all* mentions of P with Q in V , we cannot allow other mentions of P in the type transformer of V . Otherwise, there could be some mentions of P which should not be replaced during communication, which we do not model. Unlike when typing $\text{com}_{P,Q}^\tau V$, when executing a communication we know (since we only consider choreographies without free variables) that any type variables in τ or V have already been instantiated and as such do we do not need to consider how to substitute variables which may later be instantiated to P or Q .

It may be surprising to learn that our semantics are simply call-by-value reduction semantics, especially for those readers familiar with choreographies. After all, choreographies are supposed to represent concurrent programmes, and so multiple redices should be available at any time. Indeed, previous works on choreographic programming (e.g., Carbone & Montesi, 2013; Cruz-Filipe & Montesi, 2020; Hirsch & Garg, 2022) provided a semantics with *out-of-order execution*, so that the operational semantics of the choreographies matched with all possible reductions in the concurrent interpretation. We use these simpler semantics, without out-of-order execution, instead. In exchange, our result in Section 5 will be weaker: we only promise that any value which the choreography can reduce to, so can the concurrent interpretation.

To see why we chose to obtain this weaker result, consider the choreography

$$f((\text{com}_{Q_1, Q_2}^{\lambda X :: \text{Proc. Int}@X} (3 @ Q_1)), (4 @ P))$$

Here we have a function f which needs to be instantiated with a distributed pair. P is ready to feed its part of the argument into f and start computing the result, while Q_1 and Q_2 are still working on computing their part of the argument. There are two ways we could interpret PolyChor λ concurrently: we can synchronise when all processes enter a function *or* we can allow P to enter the function early. We take the second, more practical,

route. However, this means it is not possible to reflect at least one evaluation order into the semantics of the choreography without banning distributed values or allowing us to somehow call a single value in multiple steps. This insight led to us adopting the weaker guarantee discussed above.

As is standard for call-by-value λ -calculi, we are able to show that our type system is *sound* with respect to our operational semantics, as expressed in the following two theorems:

Theorem 3 (Type Preservation). *Let M be a choreography and D a function mapping containing every function in M . If there exists a typing context $\Theta; \Gamma$ such that $\Theta; \Gamma \vdash M : \tau$ and $\Theta; \Gamma \vdash D$, then $\Theta; \Gamma \vdash M' : \tau$ for any M' such that $M \rightarrow_D M'$.*

Proof Follows from the typing and semantic rules and Theorem 2. ■

Theorem 4 (Progress). *Let M be a closed choreography and D a function mapping containing every function in M . If there exists a typing context $\Theta; \Gamma$ such that $\Theta; \Gamma \vdash M : \tau$ and $\Theta; \Gamma \vdash D$, then either $M = V$ or there exists M' such that $M \rightarrow_D M'$.*

Proof Follows from the typing and semantic rules. ■

4 Endpoint projection

We now proceed to the most important result for any choreographic programming language: *endpoint projection*. Endpoint projection gives a concurrent interpretation to our language PolyChor λ by translating it to a parallel composition of programmes, one for each process. In order to define endpoint projection, though, we must define our process language, which we refer to as a *local* language. The syntax of the local language can be found in Figure 6. There you can also find the syntax of local transition labels and network transition labels, both of which will be described when we describe the operational semantics of networks.

As in PolyChor λ , our local language inherits much of its structure from System $F\omega$. In particular, we have products, sums, functions, universal quantification, and λ types, along with their corresponding terms. In fact, some types look more like standard System $F\omega$ than PolyChor λ : function types do not need a set of processes which may participate in the function, and base types no longer need a location.

However, not everything is familiar; we have introduced new terms and new types. The terms send_v and recv_v allow terms to send and receive values, respectively. We also split select terms into two terms: an *offer* term $\&_v \{ \ell_1 : L_1, \dots, \ell_n : L_n \}$ which allows v to choose how this term will evolve. We represent such choices using *choice* terms of the form $\oplus_v \ell L$. This term informs the process represented by v that it should reduce to its subterm labelled by ℓ , and then itself reduces to the term L . While these are unusual pieces of a polymorphic language like System $F\omega$, they are familiar from process languages like π calculus. We also add undefined types and terms, written \perp and \perp , respectively. These represent terms which are ill-defined; we use them to represent data which does not exist on some process P , but which needs to be written structurally in P 's programme. For

Variables	x, y, \dots	
Type Variables	X, Y, \dots	
Process Names	P	
Local Transition Labels	μ	$::= \tau \mid P \mid \text{send}_P L L' \mid \text{recv}_P L' L$ $\mid \oplus_P \ell \mid \&_P \ell$
Network Transition Labels	μ	$::= \tau_{\mathcal{P}}$
Process labels	\mathcal{P}	$::= P \mid P, Q$
Local Types	t	$::= v \mid t_1 t_2 \mid \text{Aml } v ? t_1 \& t_2 \mid t_1 \rightarrow t_2$ $\mid t_1 + t_2 \mid t_1 \times t_2 \mid \forall X. t \mid \lambda X. t$
Local Type Values	v	$::= X \mid () \mid \text{Int} \mid v_1 \rightarrow v_2 \mid P \mid \perp$ $\mid v_1 + v_2 \mid v_1 \times v_2 \mid \forall X. v \mid \lambda X. v$
Local Expressions	B	$::= x \mid () \mid n \mid \lambda x : t. B \mid \Lambda X. B$ $\mid B_1 B_2 \mid B t \mid \text{inl}_t B \mid \text{inr}_t B$ $\mid \text{case } B \text{ of inl } x \Rightarrow B_1; \text{ inr } y \Rightarrow B_2$ $\mid (B_1, B_2) \mid \text{fst } B \mid \text{snd } B$ $\mid \text{send}_v \mid \text{recv}_v$ $\mid \&_v \{ \ell_1 : B_1, \dots, \ell_n : B_n \} \mid \oplus_v \ell B$ $\mid \text{sub}[v_1 \mapsto v_2] \mid f \mid \text{Aml } v ? B_1 \& B_2$
Local Values	L	$::= x \mid () \mid n \mid \perp \mid \lambda x : t. B \mid \Lambda X. B$ $\mid \text{inl}_t L \mid \text{inr}_t L \mid (L_1, L_2)$ $\mid \text{send}_v \mid \text{recv}_v \mid \text{sub}[v_1 \mapsto v_2]$

Fig. 6. Local language syntax.

instance, \perp is the result of sending a value without process polymorphism. We also use it as the input of *recv*, since both *send* and *recv* are functions which require an input. More generally, if a process P participates in a function but the input and/or output is located elsewhere, we will use \perp to represent that input and/or output. The type \perp is only used for the term \perp .

We also include a more unusual feature: *explicit substitutions* of processes. The term $\text{sub}[v_1 \mapsto v_2]$ is a function which, when applied, replaces the role denoted by v_1 with that denoted by v_2 in its argument. This function allows us to represent the view of communication according to third parties: the roles simply change, without any mechanism necessary. For instance, imagine that **Alice** wants to tell **Bob** to communicate an integer to **Cathy**. She can do this by sending **Bob** the function $\text{com}_{\text{Alice, Bob}}^{\lambda X :: \text{Proc. Int}@X} \text{Int}@X$. In PolyChor λ , this corresponds to the choreography

$$\text{com}_{\text{Alice, Bob}}^{\lambda X :: \text{Proc. Int}@X} \text{Int}@X \left(\text{com}_{\text{Alice, Cathy}}^{\lambda X :: \text{Proc. Int}@X} \right)$$

In order to project this choreography, we need to be able to project the communication function above even when it is not applied to any arguments. This is where we use explicit substitutions: we project the communication function to $\text{sub}[\text{Alice} \mapsto \text{Bob}]$.

Finally, we introduce our unique feature: *Aml* terms and their corresponding type. These represent the ability of a process to know its own identity and to take actions based on that knowledge. Process polymorphism requires an instantiation of a process

variable at process P to be accompanied by a conditional determining whether the variable has been instantiated as P or as some other process P may interact with. In particular, the term $\text{Aml } v ? B_1 \& B_2$ reduces to B_1 if the term is run by the process denoted by v , and B_2 otherwise. Since B_1 and B_2 may have different types, we provide types of the form $\text{Aml } v ? t_1 \& t_2$, which represent either the type t_1 (if typing a term on the process denoted by v) or t_2 (otherwise). These terms form a backbone of endpoint projection for PolyChor λ : every Λ term binding a process gets translated to include an Aml term. For instance, consider projecting the choreography

$$\Lambda X :: \text{Proc. com}_{Q,X}^{\lambda X' :: \text{Proc. Int}@X'} 4 @ Q$$

to some process P . Depending on the argument to which this function is applied, P should behave very differently: if it is applied to P itself, it should receive something from Q . However, if it is applied to any other term, it should do nothing. We therefore project the choreography above to the following programme for P :

$$\Lambda X. \text{Aml } X ? \text{recv}_Q \perp \& \perp$$

Note that the Aml construct is necessary for process polymorphism in general, unless process variables cannot be instantiated to the process they are located at. It, and the combinatorial explosion caused by having multiple process abstractions, is not caused by the choreographic language but instead the choreographic language hides it and lets programmers avoid explicitly describing both sides of the Aml separately.

Note that we do not have a kinding system for local programmes. In fact, we do not check the types of local programmes at all. However, because types have *computational* content, we need to project them as well. In order to preserve that computational content, we again use an equivalence of types which corresponds to β, η -equivalence. However, in order to accommodate Aml types, we must index that equivalence with a process. Then, we have two rules regarding Aml types:

$$[\text{IAM}] \frac{}{\text{Aml } P ? t_1 \& t_2 \equiv_P t_1} \qquad [\text{IAMNOT}] \frac{P \neq Q}{\text{Aml } Q ? t_1 \& t_2 \equiv_P t_2}$$

We use these equivalence rules with process annotation to ensure that processes only use equivalences indexed with their own name and do not pick the wrong branch of an Aml type. This way we project the type $(\lambda X :: \text{Proc. Int}@X) P$ as $(\lambda X. \text{Aml } X ? \text{Int} \& \perp) P$ which is equivalent to Int and P but \perp everywhere else.

Now that we have seen the syntax of the programmes which run on each process, we can look at whole networks:

Definition 1. A network \mathcal{N} is a finite map from a set of processes to local programmes. We often write $P_1[L_1] \mid \dots \mid P_n[L_n]$ for the network where process P_i has behaviour L_i .

The parallel composition of two networks \mathcal{N} and \mathcal{N}' with disjoint domains, $\mathcal{N} \mid \mathcal{N}'$, simply assigns to each process its behaviour in the network defining it. Any network is equivalent to a parallel composition of networks with singleton domain, as suggested by the syntax above.

We now consider the operational semantics of local programmes and networks. These are given via labelled-transition systems; the syntax of both sorts of label can be found in Figure 6. The network transitions are labelled with $\tau_{\mathcal{P}}$ where \mathcal{P} is the set of involved processes (either one for a local action or two for a synchronisation). The local transitions have more options for labels. The label τ denotes a normal local computation. We use the process name P as a label for an action which can only take place at P . The label $\text{send}_P L L'$ denotes sending the value L to P , leaving L' after the send—we will explain what a label left behind after the send does when we discuss the semantics of local communication in detail. The label $\text{recv}_P L' L$ is the dual: it denotes receiving L' from P , with L being the value the receiver had before receiving. Again, we explain the semantics of receiving in detail later. Finally, the label $\oplus_P \ell$ denotes sending a label ℓ to P , while the label $\&_P \ell$ denotes receiving the label ℓ from P .

Selected rules for both operational semantics can be found in Figures 7 and 8. As before, transitions are indexed by a set d of function definitions. Function variables reduce by looking up their definition in d . Since this transition involves no communication, it is labelled with the empty transition, τ .

Perhaps surprisingly, undefined arguments to functions do not immediately cause the application to be undefined. To see why, think about choreographies of the form $(\lambda x : \text{Int} @ P. M) N$ where some process $Q \neq P$ is involved in both M and N . We project this to an application on Q of the form $(\lambda x : \perp. \llbracket M \rrbracket_Q) \llbracket N \rrbracket_Q$. Note that because we know that N has type $\text{Int} @ P$, the projection $\llbracket N \rrbracket_Q$ has type \perp and eventually evaluates to \perp . Thus, if $(\lambda x : \perp. \llbracket M \rrbracket_Q) \perp$ immediately evaluated to \perp , the process Q could not participate in M , as they need to do! We therefore allow this to evaluate to $\llbracket M \rrbracket_Q$. However, when the function is also undefined, we evaluate this to \perp with the empty label τ , as you can see in the rules [NBot] and [NBott]

As mentioned earlier, the explicit substitutions $\text{sub}[P \mapsto Q]$ are functions which, when applied, perform the requested substitution in the value to which they are applied. This is implemented in the rule [NSub].

The *Aml* terms are given meaning via the rules [NAmlR] and [NAmlL]. The rule [NAmlR] says that the term $\text{Aml } P ? L_1 \& L_2$ can evaluate to L_1 with label P , while the rule [NAmlL] says that it can instead reduce to L_2 with label Q where $Q \neq P$. We will see later that in the network semantics, we only allow transitions labelled with the process performing the transition.

Choice and offer terms reduce via the rules [NCho] and [Noff]. The first, [Ncho], tells us that a choice term simply reduces to its continuation with a transition label indicating the choice that has been made. The second, [Noff], tells us that an offer term can reduce to *any* continuation, with a transition label indicating the label of the continuation it reduced to. We will see later that the semantics of networks only allows the offer term to reduce to the continuation chosen by a matching choice term.

Finally, the *send* and *recv* terms are given meaning via [NSend] and [NRecv], respectively. However, these rules behave somewhat-differently than might be expected: rather than acting as a plain send and receive, they behave more like a swap of information.

In a plain send, the sender would not have any information after the send—perhaps the term would come with a continuation, but this would not be related to the send.

$$\begin{array}{c}
 \text{[NDEF]} f \xrightarrow{\tau}_d d(f) \qquad \text{[NABSAPP]} (\lambda x : t. B) L \xrightarrow{\tau}_d B[x \mapsto L] \\
 \\
 \text{[NBABS]} \frac{t \equiv_{\mathbf{P}} v}{(\wedge X. B) t \xrightarrow{\mathbf{P}}_d B[X \mapsto v]} \qquad \text{[NBOT]} \perp \perp \xrightarrow{\tau}_d \perp \\
 \\
 \text{[NSUB]} \text{sub}[P \mapsto Q] L \xrightarrow{\tau}_d L[P \mapsto Q] \qquad \text{[NBOTT]} \perp \perp \xrightarrow{\tau}_d \perp \\
 \\
 \text{[NAMIR]} \text{Aml } P ? L_1 \& L_2 \xrightarrow{\mathbf{P}}_d L_1 \qquad \text{[NAMIL]} \frac{Q \neq \mathbf{P}}{\text{Aml } P ? L_1 \& L_2 \xrightarrow{Q}_d L_2} \\
 \\
 \text{[NCHO]} \oplus_{\mathbf{P}} \ell L \xrightarrow{\oplus_{\mathbf{P}} \ell}_d L \qquad \text{[NOFF]} \&_{\mathbf{P}} \{\ell_1 : L_1, \dots, \ell_n : L_n\} \xrightarrow{\&_{\mathbf{P}} \ell_i}_d L_i \\
 \\
 \text{[NSEND]} \text{send}_{\mathbf{P}} L_1 \xrightarrow{\text{send}_{\mathbf{P}} L_1 L_2}_d L_2 \qquad \text{[NRECV]} \text{recv}_{\mathbf{P}} L_1 \xrightarrow{\text{recv}_{\mathbf{P}} L_2 L_1}_d L_2 \\
 \\
 \text{[NAPP1]} \frac{B_1 \xrightarrow{\mu}_d B_2}{B_1 B' \xrightarrow{\mu}_d B_2 B'} \qquad \text{[NAPP2]} \frac{B \xrightarrow{\mu}_d B'}{L B \xrightarrow{\mu}_d L B'} \qquad \text{[NTAPP1]} \frac{B \xrightarrow{\mu}_d B'}{B t \xrightarrow{\mu}_d B t} \\
 \\
 \text{[NINL]} \frac{B \xrightarrow{\mu}_d B'}{\text{inl}_t B \xrightarrow{\mu}_d \text{inl}_t B'} \qquad \text{[NINR]} \frac{B \xrightarrow{\mu}_d B'}{\text{inr}_t B \xrightarrow{\mu}_d \text{inr}_t B'} \\
 \\
 \text{[NCASE]} \frac{B \xrightarrow{\mu}_d B'}{\text{case } B \text{ of inl } x \Rightarrow B_1; \text{ inr } y \Rightarrow B_2 \xrightarrow{\mu}_d \text{case } B' \text{ of inl } x \Rightarrow B_1; \text{ inr } y \Rightarrow B_2} \\
 \\
 \text{[NCASEL]} \text{case inl}_t L \text{ of inl } x \Rightarrow B_1; \text{ inr } y \Rightarrow B_2 \xrightarrow{\tau}_d B_1[x \mapsto L] \\
 \\
 \text{[NCASER]} \text{case inr}_t L \text{ of inl } x \Rightarrow B_1; \text{ inr } y \Rightarrow B_2 \xrightarrow{\tau}_d B_2[x \mapsto L] \\
 \\
 \text{[NPAIR1]} \frac{B_1 \xrightarrow{\mu}_d B'_1}{(B_1, B_2) \xrightarrow{\mu}_d (B'_1, B_2)} \qquad \text{[NPAIR2]} \frac{B_2 \xrightarrow{\mu}_d B'_2}{(B_1, B_2) \xrightarrow{\mu}_d (B_1, B'_2)} \\
 \\
 \text{[FST]} \frac{B_1 \rightarrow_D B_2}{\text{fst } B_1 \rightarrow_D \text{fst } B_2} \qquad \text{[SND]} \frac{B_1 \rightarrow_D B_2}{\text{snd } B_1 \rightarrow_D \text{snd } B_2} \qquad \text{[NPROJ1]} \text{fst } (L_1, L_2) \xrightarrow{\tau}_d L_1 \\
 \\
 \text{[NPROJ2]} \text{snd } (L_1, L_2) \xrightarrow{\tau}_d L_2
 \end{array}$$

Fig. 7. Semantics of local processes.

Moreover, the receiver would not provide any information, but merely receive the information from the sender. However, when sending a choreography with process polymorphism, the sender may need to participate in the continuation, depending on how polymorphic functions are applied. For instance, consider the following choreography, where \mathbf{P} sends a

$$\begin{array}{c}
\text{[NCOM]} \frac{L_1 \xrightarrow{\text{send}_P L (L'[Q \mapsto P])} L'_1 \quad L_2 \xrightarrow{\text{recv}_Q (L[Q \mapsto P])} L'_2}{Q[L_1] \mid P[L_2] \xrightarrow{\tau_{Q,P}} Q[L'_1] \mid P[L'_2]} \\
\text{[NSEL]} \frac{L_1 \xrightarrow{\oplus_P \ell} L'_1 \quad L_2 \xrightarrow{\&_Q \ell} L'_2}{Q[L_1] \mid P[L_2] \xrightarrow{\tau_{Q,P}} Q[L'_1] \mid P[L'_2]} \\
\text{[NPROAM]} \frac{L \xrightarrow{P} L'}{P[L] \xrightarrow{\tau_P} P[L']} \quad \text{[NPRO]} \frac{L \xrightarrow{\tau} L'}{P[L] \xrightarrow{\tau_P} P[L']} \\
\text{[NPAR]} \frac{\mathcal{N}_1 \xrightarrow{\tau_{\mathcal{P}}} \mathcal{N}_2}{\mathcal{N}_1 \mid \mathcal{N}' \xrightarrow{\tau_{\mathcal{P}}} \mathcal{N}_2 \mid \mathcal{N}'}
\end{array}$$

Fig. 8. Semantics of networks.

polymorphic function to Q , and the resulting polymorphic function is applied to P :

$$(\text{com}_{P,Q}^{\lambda Y :: \text{Proc. } \forall X :: \text{Proc. Int}@X} (\Lambda X :: \text{Proc. com}_{P,X}^{\lambda Y' :: \text{Proc. Int}@Y'} (5 @ P))) P$$

The polymorphic function that results from the com above is as follows:

$$\Lambda X :: \text{Proc.} \left(\text{com}_{Q,X}^{\lambda Y' :: \text{Proc. Int}@Y'} (5 @ Q) \right)$$

Applying this to P leads to a programme where P receives from Q . Since P needs to participate in this programme, P must have a programme remaining after sending the polymorphic function to Q .

While this explains why send terms cannot simply, for instance, return unit, it does not explain why send and recv terms swap results. To see this, consider what happens when a term is sent from a process P to another process Q . We know from our type system that Q is not mentioned in the type of the term being sent, and we know that after the send all mentions of P are changed to mentions of Q . Hence, after the send, P 's version of the term should be the view of a process not involved in the term. This is exactly what Q 's version of the term is *before* the send. Thus, sends and recvs behaving as swaps leads to the correct behaviour.

Example 4 (Send And Receive). We now show the local projection (formalised in Section 4.1) and desired behaviour of

$$(\text{com}_{P,Q}^{\lambda Y :: \text{Proc. } \forall X :: \text{Proc. Int}@X} (\Lambda X :: \text{Proc. com}_{P,X}^{\lambda Y' :: \text{Proc. Int}@Y'} (5 @ P))) P$$

This choreography generates the network:

$$\begin{array}{l}
P[(\text{send}_Q (\Lambda X. \text{Aml } X ? (\lambda x : \lambda Y'. \text{Aml } Y' ? \text{Int} \& \perp P. x) \& (\text{send}_X 5))) P] \\
Q[(\text{recv}_P (\Lambda X. \text{Aml } X ? (\text{recv}_P \perp) \& (\perp))) P]
\end{array}$$

Using our semantics, we get the following reductions:

$$\begin{aligned}
& P[(\text{send}_Q (\Lambda X. \text{Aml } X ? (\lambda x : \lambda Y'. \text{Aml } Y' ? \text{Int} \& \perp P.x) \& (\text{send}_X 5))) P] \\
& Q[(\text{recv}_P (\Lambda X. \text{Aml } X ? (\text{recv}_P \perp) \& (\perp))) P] \\
& \xrightarrow{\tau_{P,Q}} \emptyset \\
& P[(\Lambda X. \text{Aml } X ? (\text{recv}_Q \perp) \& (\perp)) P] \\
& Q[(\Lambda X. \text{Aml } X ? (\lambda x : \lambda Y'. \text{Aml } Y' ? \text{Int} \& \perp Q.x) \& (\text{send}_X 5)) P] \\
& \xrightarrow{\tau_P} \emptyset \\
& P[(\text{Aml } P ? (\text{recv}_Q \perp) \& (\perp))] \\
& Q[(\Lambda X. \text{Aml } X ? (\lambda x : \lambda Y'. \text{Aml } Y' ? \text{Int} \& \perp Q.x) \& (\text{send}_X 5)) P] \\
& \xrightarrow{\tau_P} \emptyset \\
& P[\text{recv}_Q \perp] \\
& Q[(\Lambda X. \text{Aml } X ? (\lambda x : \lambda Y'. \text{Aml } Y' ? \text{Int} \& \perp Q.x) \& (\text{send}_X 5)) P] \\
& \xrightarrow{\tau_Q} \emptyset \\
& P[\text{recv}_Q \perp] \\
& Q[\text{Aml } P ? (\lambda x : \lambda Y'. \text{Aml } Y' ? \text{Int} \& \perp Q.x) \& (\text{send}_X 5)] \\
& \xrightarrow{\tau_Q} \emptyset \\
& P[\text{recv}_Q \perp] | Q[(\text{send}_X 5)] \\
& \xrightarrow{\tau_{Q,P}} \emptyset \\
& P[5] | Q[\perp]
\end{aligned}$$

Now that we have discussed the semantics of local programmes, we discuss the semantics of networks. Each transition in the network semantics has a silent label indexed with the processes participating in that reduction: $\tau_{\mathcal{P}}$, where \mathcal{P} consists of either one process name (for local actions at that process) or two process names (for interactions involving these two processes). We treat \mathcal{P} as a set, implicitly allowing for exchange.

For instance, the rule [NCom] describes communication. Here, one local term must reduce with a *send* label, while another reduces with a *recv* label. These labels must match, in the sense that the value received by the *recv* must be the value sent by the *send*—though with the receiver in place of the sender—and vice-versa. Then, a network in which the local terms are associated with the appropriate processes, *Q* and *P*, can reduce with the label $\tau_{Q,P}$. Similarly, the rules [NSel] reduces matching choice and select terms, resulting in the label $\tau_{Q,P}$.

While [NCom] and [NSel] describe communication, the rest of the rules describe how a single process's term can evolve over time in a network. Particularly interesting is [NProam], which says that a *Aml* term can reduce only according to the process it is associated with. We can see here that the resulting label is τ_P , indicating that this reduction step only involves *P*.

The rules [NPro] tells us how to lift steps with an empty label τ . Such steps make no assumptions about the network, and so such terms can be associated with any process *P*. When such a reduction takes place in a network, we label the resulting transition τ_P .

Finally, the rule [NPar] says that if one part of a network can reduce with a label $\tau_{\mathcal{P}}$, then the entire network can reduce with that same label. This allows the other rules, which assume minimal networks, to be applied in larger networks.

In the future, we will use \rightarrow^* and \rightarrow^+ to denote respectively a sequence and a sequence of at least one action with arbitrary labels.

4.1 Projection

We can now define the endpoint projection (EPP) of choreographies. This describes a single process's view of the choreography; the concurrent interpretation of a choreography is given by composing the projection to every process in parallel. Endpoint projection to a particular process P is defined as a recursive function over typing derivations $\Theta; \Gamma \vdash M : \tau$. For readability, however, we write it as a recursive function over the term M , and use the notation $\text{typeof}(N)$ to refer to the types assigned to any term N in the implicit typing derivation. Similarly, we use $\text{kindof}(\tau)$ to refer to the kind of a type τ in the implicit typing derivation. We write $\llbracket M \rrbracket_P$ to denote the projection of the term M (implicitly a typing derivation for M , proving that it has *some* type) to the process P .

Intuitively, projection translates a choreography term to its corresponding local behaviour. For example, a communication action projects to a send (for the sender), a receive (for the receiver), a substitution (for the other processes in the type of the value being communicated) or an empty process (for the remaining processes). However, this is more complicated for `case` statements. For instance, consider the following choreography, which matches on a sum type which is either an integer on `Alice` or a unit on `Alice`. If it is an integer, then `Bob` receives that integer from `Alice` and the choreography returns the integer now located at `Bob`. Otherwise, the choreography returns the default value 42 also located at `Bob`. `Alice` informs `Bob` of which branch she has taken using `select` terms.

$$\left(\begin{array}{l} \lambda z : (\text{Int @ Alice}) + (() @ Alice). \\ \text{case } z \text{ of} \\ \text{inl } x \Rightarrow \text{select}_{\text{Alice, Bob}} \text{ Just } (\text{com}_{\text{Alice, Bob}}^{\lambda X. \text{Int}@X} x); \\ \text{inr } y \Rightarrow \text{select}_{\text{Alice, Bob}} \text{ Nothing } (42 @ \text{Bob}) \end{array} \right) \text{inl}_{()@Alice} (3 @ \text{Alice})$$

Imagine projecting this to `Bob`'s point of view. He does not have any of the information in the sum, so he cannot participate in choosing which branch of the `case` expression gets evaluated. Instead, he has to wait for `Alice` to tell him which branch he is in. If we naïvely translate just the first branch of the case expression, `Bob` waits for `Alice` to send him the label `Just` and then waits for `Alice` to send him an integer. Similarly, in the second branch `Bob` waits for `Alice` to send him the label `Nothing` before returning the default value 42. Somehow, we need to combine these so that `Bob` waits for either label, and then takes the corresponding action.

We do this by *merging* `Bob`'s local programmes for each branch (Carbone *et al.*, 2012; Honda *et al.*, 2016; Cruz-Filipe & Montesi, 2020). Merging is a *partial* function which combines two compatible local programmes, combining choice statements. In other words, the key property of merging is:

$$\&_P \{ \ell_i : B_i \}_{i \in I} \sqcup \&_P \{ \ell_j : B'_j \}_{j \in J} = \&_P \left(\{ \ell_k : B_k \sqcup B'_k \}_{k \in I \cup J} \cup \{ \ell_i : B_i \}_{i \in I \setminus J} \cup \{ \ell_j : B'_j \}_{j \in J \setminus I} \right)$$

Merging is defined homomorphically on other terms, though it is undefined on incompatible terms. Thus, for example, $\text{inl}_t B \sqcup \text{inl}_t B' = \text{inl}_t (B \sqcup B')$, but $\text{inl}_{t_1} B \sqcup \text{inr}_{t_2} B'$ is undefined.

We can then use this to project the choreography above to **Bob** as:

$$(\lambda z : \perp. \&_{\text{Alice}} \{ \text{Just} : (\text{recv}_{\text{Alice}} \perp), \text{Nothing} : 42 \}) \perp$$

where \perp represents a part of the choreography executed by **Alice**.

Definition 2. The EPP of a choreography M for process P is defined by the rules in Figures 9, 10 and 11.

To project a network from a choreography, we therefore project the choreography for each process and combine the results in parallel: $\llbracket M \rrbracket = \prod_{P \in \text{ip}(M)} P \llbracket \llbracket M \rrbracket_P \rrbracket$.

Intuitively, projecting a choreography to a process that is not involved in it returns a \perp . More-complex choreographies, though, may involve processes that are not shown in their type. This explains the first clause for projecting an application: even if P does not appear in the type of M , it may participate in interactions inside M . A similar observation applies to the projection of *case*, where merging is also used.

Selections and communications follow the intuition given above, with one interesting detail: self-selections are ignored, and self-communications project to the identity function. This is different from many other choreography calculi, where self-communications are not allowed—we do not want to impose this in PolyChor λ , since we have process polymorphism and therefore do not want to place unnecessary restrictions on which processes a choreography can be instantiated with.

Any process P must prepare two behaviours for a process abstraction $\Lambda X :: \text{Proc}. M$: one for when X is instantiated with P itself, and one for when X is instantiated with another process. To do this, we use *Aml* terms, which allow P to use its knowledge of its identity to select which behaviour takes place. (This also holds when X has a *Without* kind, as long as the base kind is *Proc*, though if P is excluded from the type of X and P does not participate in M then we simply project to \perp .) However, type abstractions $\Lambda X :: K. M$ do not use *Aml* terms if K is not a kind of processes, since P cannot instantiate X .

When projecting an application, we may project both the function and its argument, either one but not the other, or neither. While it may seem simple—just project both sides, and get rid of any \perp s or \perp s that come up—it turns out to be somewhat complicated. In order to ensure every process performs actions in the same order and avoid communication mismatches, we must project an abstraction for any process involved in the computation, even if they do not have the input (Cruz-Filipe *et al.*, 2023, Example 6). To see why this causes complications, consider $M = \lambda x : \text{Int} @ P. 5 @ Q$. When M gets projected to Q , it becomes $\lambda x. 5$. However, applying M to an argument—say, $M 2 @ P$ —needs to lead to a function application on Q ! Thus, we project this to $(\lambda x. 5) \perp$, allowing Q to instantiate its function. We use the type system to identify the cases where we need to keep \perp or \perp and those where we should only project the function part of an (type) application.

Type applications work a bit differently. Since there is no chance of communication happening while computing a type, we can project only the body of a type abstraction without the actual abstraction to P when we know the argument is not located at P . In addition, we

$$\begin{aligned}
\llbracket x \rrbracket_{\mathbf{P}} &= \begin{cases} \perp & \text{if } \llbracket \text{typeof}(x) \rrbracket_{\mathbf{P}} = \perp \\ x & \text{otherwise} \end{cases} & \llbracket f \rrbracket_{\mathbf{P}} &= f \\
\llbracket () @ v \rrbracket_{\mathbf{P}} &= \begin{cases} () & \text{if } \llbracket v \rrbracket_{\mathbf{P}} = \mathbf{P} \\ \perp & \text{otherwise} \end{cases} & \llbracket n @ v \rrbracket_{\mathbf{P}} &= \begin{cases} n & \text{if } \llbracket v \rrbracket_{\mathbf{P}} = \mathbf{P} \\ \perp & \text{otherwise} \end{cases} \\
\llbracket \lambda x : \tau. M \rrbracket_{\mathbf{P}} &= \begin{cases} \perp & \text{if } \llbracket M \rrbracket_{\mathbf{P}} = \perp \\ & \text{and } \llbracket \tau \rrbracket_{\mathbf{P}} = \perp \\ \lambda x : \llbracket \tau \rrbracket_{\mathbf{P}}. \llbracket M \rrbracket_{\mathbf{P}} & \text{otherwise} \end{cases} \\
\llbracket M N \rrbracket_{\mathbf{P}} &= \begin{cases} \perp & \text{if } \llbracket M \rrbracket_{\mathbf{P}} = \llbracket N \rrbracket_{\mathbf{P}} = \perp \\ \llbracket M \rrbracket_{\mathbf{P}} \llbracket N \rrbracket_{\mathbf{P}} & \text{if } \mathbf{P} \in \text{ip}(\text{typeof}(M)) \\ & \text{or } \llbracket M \rrbracket_{\mathbf{P}} \neq \perp \neq \llbracket N \rrbracket_{\mathbf{P}} \\ \llbracket M \rrbracket_{\mathbf{P}} & \text{if } \llbracket N \rrbracket_{\mathbf{P}} = \perp \\ \llbracket N \rrbracket_{\mathbf{P}} & \text{otherwise} \end{cases} \\
\llbracket \Delta X :: K. M \rrbracket_{\mathbf{P}} &= \begin{cases} \Delta X. \text{Aml } X & \text{if } K \in \{\text{Proc}, \text{Proc} \setminus \rho\} \\ \quad ? \llbracket M[X \mapsto \mathbf{P}] \rrbracket_{\mathbf{P}} \\ \quad \& \llbracket M \rrbracket_{\mathbf{P}} \\ \llbracket M \rrbracket_{\mathbf{P}} & \text{if } K = K' \setminus \{\mathbf{P}\} \cup \rho \\ \Delta X. \llbracket M \rrbracket_{\mathbf{P}} & \text{otherwise} \end{cases} \\
\llbracket M \tau \rrbracket_{\mathbf{P}} &= \begin{cases} \perp & \text{if } \llbracket M \rrbracket_{\mathbf{P}} = \llbracket \tau \rrbracket_{\mathbf{P}} = \perp \\ \llbracket M \rrbracket_{\mathbf{P}} & \text{if } \llbracket \tau \rrbracket_{\mathbf{P}} = \perp \\ & \text{and } \mathbf{P} \notin \text{ip}(\text{typeof}(M)) \\ \llbracket M \rrbracket_{\mathbf{P}} \llbracket \tau \rrbracket_{\mathbf{P}} & \text{otherwise} \end{cases} \\
\llbracket \text{inl}_{\tau} M \rrbracket_{\mathbf{P}} &= \begin{cases} \perp & \text{if } \llbracket M \rrbracket_{\mathbf{P}} = \perp \text{ and} \\ & \text{kindof}(\tau) = K \setminus \rho \\ \llbracket M_1 \rrbracket_{\mathbf{P}} & \text{if } \llbracket \text{typeof}(M) \rrbracket_{\mathbf{P}} = \perp \\ \text{inl}_{\llbracket \tau \rrbracket_{\mathbf{P}}} \llbracket M_1 \rrbracket_{\mathbf{P}} & \text{otherwise} \end{cases} \\
\llbracket \text{inr}_{\tau} M \rrbracket_{\mathbf{P}} &= \begin{cases} \perp & \text{if } \llbracket M \rrbracket_{\mathbf{P}} = \perp \text{ and} \\ & \text{kindof}(\tau) = K \setminus \rho \\ \llbracket M_1 \rrbracket_{\mathbf{P}} & \text{if } \llbracket \text{typeof}(M) \rrbracket_{\mathbf{P}} = \perp \\ \text{inr}_{\llbracket \tau \rrbracket_{\mathbf{P}}} \llbracket M_1 \rrbracket_{\mathbf{P}} & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 9. Projection of PolyChor λ programmes.

$$\begin{aligned}
 \llbracket \text{case } M \text{ of inl } x \Rightarrow N_1; \text{ inr } y \Rightarrow N_2 \rrbracket_P = & \begin{cases} \text{case } \llbracket M \rrbracket_P \text{ of inl } x \Rightarrow \llbracket N_1 \rrbracket_P; \text{ inr } y \Rightarrow \llbracket N_2 \rrbracket_P & \text{if } P \in \text{ip}(\text{typeof}(M)) \\ \perp & \text{if } \llbracket M \rrbracket_P = \llbracket N_1 \rrbracket_P = \llbracket N_2 \rrbracket_P = \perp \\ \llbracket M \rrbracket_P & \text{if } \llbracket N_1 \rrbracket_P = \llbracket N_2 \rrbracket_P = \perp \\ \llbracket N_1 \rrbracket_P \sqcup \llbracket N_2 \rrbracket_P & \text{if } \llbracket M \rrbracket_P = \perp \\ (\lambda z : \perp. (\llbracket N_1 \rrbracket_P \sqcup \llbracket N_2 \rrbracket_P)) \llbracket M \rrbracket_P \text{ (z fresh)} & \text{otherwise} \end{cases} \\
 \llbracket (M_1, M_2) \rrbracket_P = & \begin{cases} \perp & \text{if } \llbracket M_1 \rrbracket_P = \llbracket M_2 \rrbracket_P = \perp \\ (\llbracket M_1 \rrbracket_P, \llbracket M_2 \rrbracket_P) & \text{otherwise} \end{cases} \\
 \llbracket \text{fst } M \rrbracket_P = & \begin{cases} \perp & \text{if } \llbracket M \rrbracket_P = \perp \\ \llbracket M_1 \rrbracket_P & \text{if } \llbracket \text{typeof}(M) \rrbracket_P = \perp \\ \text{fst } \llbracket M_1 \rrbracket_P & \text{otherwise} \end{cases} \\
 \llbracket \text{snd } M \rrbracket_P = & \begin{cases} \perp & \text{if } \llbracket M \rrbracket_P = \perp \\ \llbracket M_1 \rrbracket_P & \text{if } \llbracket \text{typeof}(M) \rrbracket_P = \perp \\ \text{snd } \llbracket M_1 \rrbracket_P & \text{otherwise} \end{cases} \\
 \llbracket \text{select}_{Q_1, Q_2} \ell M \rrbracket_P = & \begin{cases} \bigoplus_{Q'} \ell \llbracket M \rrbracket_P & \text{if } P = Q_1 \neq Q_2 \\ \&_S \{ \ell : \llbracket M \rrbracket_P \} & \text{if } P = Q_2 \neq Q_1 \\ \llbracket M \rrbracket_P & \text{otherwise} \end{cases} \\
 \llbracket \text{com}^\tau_{Q_1, Q_2} \rrbracket_P = & \begin{cases} \lambda x : \llbracket \tau P \rrbracket_P . x & \text{if } P = Q_1 = Q_2 \\ \text{send}_{Q_2} & \text{if } P = Q_1 \neq Q_2 \\ \text{recv}_{Q_1} & \text{if } P = Q_2 \neq Q_1 \\ \text{sub}[Q_1 \mapsto Q_2] & \text{if } \llbracket \tau \rrbracket_P \neq \perp \\ \perp & \text{otherwise} \end{cases}
 \end{aligned}$$

Fig. 10. Projection of PolyChorλ programmes (ctd.).

do not have a case for projecting only the argument, since the context surrounding a type abstraction will not expect a type.

In general, projecting a type yields \perp at any process not used in that type. We use the restrictions on kinds to avoid projecting type variables and type abstractions when we know we do not need to and project all process names to themselves, but otherwise the projection of type constructs is similar to that of corresponding process terms.

Finally, to execute a projected choreography, we need to project the set of definitions of choreographic functions to a set of definitions of local functions. Since these functions are all parametrised over every involved process, this is as simple as projecting the definitions

$$\begin{aligned}
\llbracket X \rrbracket_P &= \begin{cases} \perp & \text{if } \text{kindof}(X) = K \setminus (\{P\} \cup \rho) \text{ for } K \neq \text{Proc} \\ X & \text{otherwise} \end{cases} & \llbracket Q \rrbracket_P &= Q \\
\llbracket () @ Q \rrbracket_P &= \begin{cases} () & \text{if } P = Q \\ \perp & \text{otherwise} \end{cases} & \llbracket \text{Int} @ Q \rrbracket_P &= \begin{cases} \text{Int} & \text{if } P = Q \\ \perp & \text{otherwise} \end{cases} \\
\llbracket \tau_1 \times \tau_2 \rrbracket_P &= \begin{cases} \perp & \text{if } \llbracket \tau_1 \rrbracket_P = \llbracket \tau_2 \rrbracket_P = \perp \\ \llbracket \tau_1 \rrbracket_P \times \llbracket \tau_2 \rrbracket_P & \text{otherwise} \end{cases} \\
\llbracket \tau_1 + \tau_2 \rrbracket_P &= \begin{cases} \perp & \text{if } \llbracket \tau_1 \rrbracket_P = \llbracket \tau_2 \rrbracket_P = \perp \\ \llbracket \tau_1 \rrbracket_P + \llbracket \tau_2 \rrbracket_P & \text{otherwise} \end{cases} \\
\llbracket \tau_1 \rightarrow_\rho \tau_2 \rrbracket_P &= \begin{cases} \llbracket \tau_1 \rrbracket_P \rightarrow \llbracket \tau_2 \rrbracket_P & \text{if } P \in \rho \text{ or } \llbracket \tau_1 \rrbracket_P \neq \perp \neq \llbracket \tau_2 \rrbracket_P \\ \perp & \text{otherwise} \end{cases} \\
\llbracket \forall X :: K. \tau \rrbracket_P &= \\
&\begin{cases} \perp & \text{if } \llbracket \tau \rrbracket_P = \perp \text{ and } K = K' \setminus (\{P\} \cup \rho) \\ \forall X. \text{Aml } X ? [\llbracket \tau[X \mapsto P] \rrbracket_P] \& [\llbracket \tau \rrbracket_P] & \text{if } K \in \{\text{Proc}, \text{Proc} \setminus \rho\} \\ \forall X. [\llbracket \tau \rrbracket_P] & \text{otherwise} \end{cases} \\
\llbracket \tau_1 \tau_2 \rrbracket_P &= \begin{cases} \perp & \text{if } \llbracket \tau_1 \rrbracket_P = \llbracket \tau_2 \rrbracket_P = \perp \\ \llbracket \tau_1 \rrbracket_P & \text{if } \llbracket \tau_2 \rrbracket_P = \perp \text{ and } \text{kindof}(\tau_2) = K \setminus (\{P\} \cup \rho) \\ \llbracket \tau_2 \rrbracket_P & \text{if } \llbracket \tau_1 \rrbracket_P = \perp \\ \llbracket \tau_1 \rrbracket_P \llbracket \tau_2 \rrbracket_P & \text{otherwise} \end{cases} \\
\llbracket \lambda X :: K. \tau \rrbracket_P &= \\
&\begin{cases} \perp & \text{if } \llbracket \tau \rrbracket_P = \perp \text{ and } K = K' \setminus (\{P\} \cup \rho) \\ \lambda X. \text{Aml } X ? [\llbracket \tau[X \mapsto P] \rrbracket_P] \& [\llbracket \tau \rrbracket_P] & \text{if } K \in \{\text{Proc}, \text{Proc} \setminus \rho\} \\ \lambda X. [\llbracket \tau \rrbracket_P] & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 11. Projection of PolyChor λ types.

onto an arbitrarily chosen process name.

$$\llbracket D \rrbracket = \{f \mapsto \llbracket D(f) \rrbracket_P \mid f \in \text{domain}(D)\}$$

Note that function names always get projected everywhere. This means that if we have a function which does not terminate when applied to some value in any process, then it diverges when applied to that value in the choreography and in every other process.

Example 5. We will now show how to project the bookseller service example Equation (1.3). As in that example we use $\text{let } x = B \text{ in } B'$ as syntactic sugar for $\lambda x : t. B' B$ for some t and $\text{if } B_1 \text{ then } B_2 \text{ else } B_3$ as syntactic sugar for $\text{case } B_1 \text{ of inl } x \Rightarrow B_2; \text{ inr } x \Rightarrow B_3$ for some $x \notin (\text{fv}(B_2) \cup \text{fv}(B_3))$. We project for **Seller** and get the following process:


```

 $\Delta B.$ 
Aml  $B$ 
?  $\lambda$  title.
   $\lambda$  buyAtPrice?.
    let  $x = (\lambda y. y)$  title
    in let  $y = (\lambda z. z)$  (price_lookup  $x$ )
      in if buyAtPrice?  $y$ 
        then ()
        else ()
&  $\lambda$  title.
   $\lambda$  buyAtPrice?.
    let  $x = \text{recv}_B \perp$ 
    in let  $y = \text{send}_B$  (price_lookup  $x$ )
      in  $\&_B$  {Buy : (), Quit : ()}

```

Here we can see that if the buyer B turns out to be **Seller** itself, then all the communications become identity functions, and the seller does not inform itself of its choice. Otherwise, we get a function which, after being instantiated with a buyer, also needs to be instantiated with two \perp s representing values existing at B . It then waits for B to send a title, returns the price of this title, and waits for B to decide whether to buy or not. It might seem strange to have a function parametric on two values that are located at B and will therefore here be instantiated with \perp s, but this example actually illustrates why when projecting we cannot in cases like this remove the first two λ s from the process without causing a deadlock. Consider that $\text{let } y = \text{send}_B(\text{price_lookup } x) \text{ in } \&_B \{\text{Buy} : (), \text{Quit} : ()\}$ is syntactic sugar for $(\lambda y. \&_B \{\text{Buy} : (), \text{Quit} : ()\}) (\text{send}_B(\text{price_lookup } x))$. Here we need to have the abstraction on y even though it gets instantiated as \perp after **Seller** sends the result of $\text{price_lookup } x$ to B . If instead we only had $(\&_B \{\text{Buy} : (), \text{Quit} : ()\}) (\text{send}_B(\text{price_lookup } x))$, then the first part of the application would not be a value, and would be waiting for B to choose between Buy and Quit while B has the abstraction on y and therefore considers the first part of the application a function which must wait to be instantiated. B therefore expects to receive the result of $\text{price_lookup } x$, and we get a deadlock in our system. This is why we never want to project a value to a non-value term and need to keep any abstractions guarding a part of the choreography involving **Seller**.

5 The correctness of endpoint projection

We now show that there is a close correspondence between the executions of choreographies and of their projections. Intuitively, this correspondence states that a choreography can execute an action if, and only if, its projection can execute the same action, and both transition to new terms in the same relation. However, this is not completely true: if a choreography M reduces by rule [CaseL], then the result has fewer branches than the network obtained by performing the corresponding reduction in the projection of C .

In order to capture this we revert to the branching relation (Cruz-Filipe & Montesi, 2020; Montesi, 2023), defined by $M \sqsupseteq N$ iff $M \sqcup N = M$. Intuitively, if $M \sqsupseteq N$, then M offers the

same and possibly more behaviours than N . This notion extends to networks by defining $\mathcal{N} \sqsupseteq \mathcal{N}'$ to mean that, for any role P , $\mathcal{N}(P) \sqsupseteq \mathcal{N}'(P)$.

Using this, we can show that the EPP of a choreography can do all that (completeness) and only what (soundness) the original choreography does. For traditional imperative choreographic languages, this correspondence takes the form of one action in the choreography corresponding to one action in the projected network. We instead have a correspondence between one action in the choreography and multiple actions in the network due to allowing choreographies to manipulate distributed values in one action such as in $\lambda x : \text{Int} @ \text{Bob} \times \text{Int} @ \text{Alice}. M (3 @ \text{Bob}, 3 @ \text{Alice})$ where both **Bob** and **Alice** independently take the first part of the pair.

Theorem 5 (Completeness). *Given a closed choreography M , if $M \rightarrow_D M'$, $\Theta; \Gamma \vdash D$, $\Theta; \Gamma \vdash M : \tau$, and $\llbracket M \rrbracket$ is defined, then there exists network \mathcal{N} and choreography M'' such that: $\llbracket M \rrbracket \rightarrow_{\llbracket D \rrbracket}^+ \mathcal{N}$ and $\mathcal{N} \sqsupseteq \llbracket M'' \rrbracket$.*

Proof We prove this by structural induction on $M \rightarrow_D M'$. We take advantage of the fact that type values project to \perp at processes not involved in them, while choreographic values correspondingly project to \perp at processes not involved in their type. See Appendix C for full details. ■

Theorem 6 (Soundness). *Given a closed choreography M and a function mapping D , if $\Theta; \Gamma \vdash M : \tau$, $\Theta; \Gamma \vdash D$, and $\llbracket M \rrbracket \rightarrow_{\llbracket D \rrbracket}^* \mathcal{N}$ for some network \mathcal{N} , then there exist a choreography M' and a network \mathcal{N}' such that: $M \rightarrow_D^* M'$, $\mathcal{N} \rightarrow_{\llbracket D \rrbracket}^* \mathcal{N}'$, and $\mathcal{N}' \sqsupseteq \llbracket M' \rrbracket$.*

Proof We prove this by structural induction on M in the accompanying technical report, taking advantage of the fact that thanks to projecting function names everywhere, a choreography that diverges at one role diverges at every role. See Appendix D for full details. ■

From Theorems 3 to 6, we get the following corollary, which states that a network derived from a well-typed closed choreography can continue to reduce until all roles contain only local values.

Corollary 1. *Given a closed choreography M and a function environment D containing all the function names of M , if $\Theta; \Gamma \vdash M : T$ and $\Theta; \Gamma \vdash D$, then: whenever $\llbracket M \rrbracket \rightarrow_{\llbracket D \rrbracket}^* \mathcal{N}$ for some network \mathcal{N} , either there exists \mathcal{P} such that $\mathcal{N} \xrightarrow{\tau \mathcal{P}}_{\llbracket D \rrbracket} \mathcal{N}'$ or $\mathcal{N} = \prod_{P \in \text{ip}(M)} P[V_P]$.*

6 Case study

We now show how our language can be used in an extended example (Figure 12). This example involves three processes: a client **C**, an edge computer **E**, and a server **S**. Intuitively, **C** wants to request that **E** does some computation. However, **E** may not have the resources to perform the computation; in this case, it will forward the request to **S**.

```

λ HandleHere : (Int @ E →∅ Int @ E) →∅ Bool @ E.
(
  λ x : Int @ E →∅ Int @ E.
  LogRequest E S C x
  (
    case HandleHere x of
    inl y ⇒ (
      selectE,C here
      selectE,S here
      comE,CλX::Proc\{S}.Int@X (Compute x)
    );
    inr y ⇒ selectE,C atS
      selectE,S you
      (
        comE,CλY::Proc\{S}.Int@Y→∅Bool@S
        (
          comS,EλY::Proc\{C}.∀X::Proc.Int@Y→∅Bool@X
          (
            Λ Z :: Proc. λ z : Int @ S. Authenticate
            comS,ZλX::Proc.Int@X z
          )
        )
      )
      AuthKey
      (
        selectS,E go
        selectS,C go
        comS,CλX::Proc\{E}.Int@X
        Compute (comE,SλX::Proc\{C}.Int@X→∅Int@X x)
      );
    inr y ⇒ selectS,E no
      selectS,C no
      0 @ C
      comC,EλX::Proc\{C}.Int@X→∅Int@X task
  )
) of

```

Fig. 12. Code of the case study on edge computing.

Whenever **S** receives a request, then **C** must first perform an authentication protocol. Whether or not the task is outsourced to **S**, **S** logs the request.

Here we assume the following data:

- A task (of type $\text{Int} @ C \rightarrow_{\emptyset} \text{Int} @ C$) located at **C**
- For each of **E** and **S**, a local function `Compute` which executes a task
- An authentication choreography `Authenticate` between **S** and another process **Z**. This choreography takes a key `AuthKey` and checks if the holder of that key is authorised to run a task on **S**.
- A key `AuthKey` for **C**
- A logging choreography `LogRequest` involving two roles, provocatively called **E** and **S**. This choreography takes a client, a task, and the result of executing the task (at **C**) as input. It then creates a log entry at **S**.
- A local function `HandleHere`, which **E** uses to determine whether it can handle a task locally. If `HandleHere` returns false, then the task must be shipped to **S**. Unlike other data, this is represented as input to the choreography.

The choreography begins with **C** sending the task to **E**; we call the resulting task x . (Note that x is the name of the task *on E*, not the name of the task on **C**.) **E** then checks whether it can handle x using `HandleHere`. If so, **E** informs **S** and **C** that it is computing the task. After performing the task, **E** sends the result to **C**. It furthermore informs **S** so that it knows that it needs to log the task.

If **E** cannot handle the task, then it again informs **C** and **S**. **S** then makes a decision on whether **C** has authorisation to request a task from **S**. To make this decision, **S** sends an authentication protocol to **E**. Because communications swap the sender and the receiver in the communicated value, we write this authentication protocol with **S** playing the role of the client. The protocol is therefore parameterised on the authenticator. Once **E** receives the authentication protocol, it can instantiate the authenticator as **S**. **E** finally sends the (now complete) protocol to **C**; running this protocol will have **C** send its key to **S**, possibly among other actions required for authentication. If the authentication procedure succeeds, then **S** informs **E** and **C** of this. **E** can then send the task to **S**, who computes it and returns the result to **C**. If authentication fails, then **S** informs **E** and **C** of this and the task fails, resulting in a 0 on **C**. Either way, we finish the choreography by logging the task and its result using the function `LogRequest`.

For **S** to send an authentication protocol which it is itself involved with requires a bit of trickery. Usually, we would expect every part of the sent value located at **S** to be moved to the receiver (first **E** and then after another communication **C**) but obviously that would mean **S** cannot be involved. We therefore send an authentication protocol that is parametric on the authenticator, **Z**, and only instantiate **Z** as **S** after the first communication from **S** to **E** has taken place.

Projecting this protocol to **C** leads to the following code:

$$\begin{aligned}
 & \lambda \text{ HandleHere} : \perp. \\
 & \left(\lambda x : \perp. \right. \\
 & \quad \text{LogRequest } E \ S \ C \ \perp \\
 & \quad \left. \left. \begin{array}{l} \text{here} : \text{recv}_E \ \perp, \\ \&_E \text{ atS:} \left(\begin{array}{l} \left(\lambda y : \perp. \&_S \{ \text{go} : \text{recv}_S \ \perp, \text{no} : 0 \} \right. \right. \\ \left. \left. \left(\begin{array}{l} \text{recv}_E \ \text{sub}[S \mapsto E] \\ \text{AuthKey} \end{array} \right) \left(\begin{array}{l} \wedge Z :: \text{Proc.} \\ \text{Aml } Z \\ ? \lambda z : \perp. \text{Authenticate} \\ \text{recv}_S \ \perp \\ \& \text{sub}[S \mapsto Z] \ \perp \end{array} \right) \right) \end{array} \right) \right) \end{array} \right) \\
 & \quad \left. \left. \begin{array}{l} \text{send}_E \ \text{task} \\ S \end{array} \right) \right) \\
 & \end{aligned} \tag{6.1}
 \end{aligned}$$

We see that the second case gets projected as an application of a new abstraction on the new variable y , with **C**'s part of the condition as the right side of the application. Since the condition contains a delegation, we get some process substitutions representing values with unknown locations being communicated between other processes. Because of the **Aml** in the second branch, none of the substituted processes are ever reached. Therefore, these substitutions do nothing. Since **C** is not involved in **E**'s decision to delegate to **S** (or not), we do not see any of the code involved in the decision here. Instead, we get a straightforward offer as the result of merging the projection of each branch of the involved conditional.

We now show the projection for **E**:

7 Related work

7.1 Choreographies

Choreographies are inspired by the “Alice-and-Bob” notation for security protocols by Needham & Schroeder (1978), which included a term for expressing a communication from a participant to another. The same idea inspired later several graphical notations for modelling interactions among processes, including sequence diagrams and message sequence charts (International Telecommunication Union, 1996; Object Management Group, 2017).

A systematic introduction to theory of choreographic languages and their historical development can be found in Montesi (2023). We recap and discuss relevant recent developments. The first sophisticated languages for expressing choreographies were invented to describe interactions between web services. The Web Services Choreography Description Language (WS-CDL) by The World Wide Web Consortium (W3C) (2004) is a choreographic language which describes the expected observable interactions between web services from a global point of view (Zongyan et al., 2007). Carbone et al. (2012) later formalised endpoint projection for a theory of choreographies based on WS-CDL, and in particular introduced the merging operator (which we adjusted and extended to our setting). This inspired more work on choreographies and projection and eventually the birth of choreographic programming—where choreographies are programmes compiled to executable code—and the first choreographic programming language, Chor (Montesi, 2013). As choreographic programming languages became more complex, Cruz-Filipe & Montesi (2020) developed a *core calculus of choreographies* (CC). Montesi (2023) revisited and generalised CC in his text on foundations of choreographic languages. Cruz-Filipe et al. (2021) then formalised this new version and its properties in the Coq theorem prover (The Coq Development Team, 2004). Later, Pohjola et al. (2022) developed a certified end-to-end compiler from another variation of CC to CakeML by using the HOL theorem prover.

One of the primary design goals of all of choreographic programming languages is *deadlock-freedom by design* (Carbone & Montesi, 2013)—the operational correspondence between the choreography and the distributed network ensures deadlock-freedom for the network. PolyChor λ achieves this goal. Montesi (2023) discusses restrictions for a procedural imperative choreographic language in order to obtain a stronger liveness property (starvation-freedom). The idea is to prove that processes eventually get involved in transitions at the choreographic level and then leverage the correctness of endpoint projection to obtain the same result about choreography projections. This idea might work for PolyChor λ as well, but whether and how the technical devices for starvation-freedom in Montesi (2023) can be adapted to PolyChor λ is not clear due to the different nature of our language (functional instead of imperative). Alternatively, one could adapt static analyses for lock-freedom—like that in Kobayashi (2006)—to choreographies. We leave explorations of liveness properties other than deadlock-freedom in PolyChor λ to future work.

The first choreographic language with limited process polymorphism was Procedural Choreographies (PC) (Cruz-Filipe & Montesi, 2017b). In PC, a choreography comes with an environment of predefined *procedures* parametric on process names which may be

called by the choreography. These procedures have a number of limitations compared to the process polymorphism of $\text{PolyChor}\lambda$: they cannot contain any free processes, they cannot be partially instantiated, and they are lower-order—that is, they must be defined in the environment rather than as part of a larger choreography. These limitations allow the projection function to know how the procedure will be instantiated, whereas in $\text{PolyChor}\lambda$ we may need to compute the processes involved first. This has major implications for projection: in PC, it is easy to tell when projecting a procedure call which processes are involved and therefore need a local copy of the call. However, $\text{PolyChor}\lambda$'s full process polymorphism allows the function and process names to each be enclosed in a context. While this allows greater flexibility for programmers, it forces us to project a process-polymorphic functions to every process and let each process determine at runtime whether it is involved.

Recently, there has been a fair amount of interest in higher-order and functional programming for choreographies (Hirsch & Garg, 2022; Cruz-Filipe *et al.*, 2022; Shen *et al.*, 2023; Giallorenzo *et al.*, 2023). The first higher-order choreographic programming language, Choral (Giallorenzo *et al.*, 2023) is an object-oriented choreographic language compiled to Java code. Thus, Choral choreographies can depend on other choreographies, allowing programmers to reuse code. Choral was also the first choreographic language to treat $\text{com}_{P,Q}^{\tau}$ as a first-class function.

While Choral gave a practical implementation of higher-order choreographies, it did not establish their theoretical properties. Two different—but independently developed—works filled this gap, including $\text{Chor}\lambda$, the basis of $\text{PolyChor}\lambda$. $\text{Chor}\lambda$ is a functional choreographic calculus based on the λ -calculus. In this work, we extended $\text{Chor}\lambda$ with type and process polymorphism and the ability to send non-local values such as choreographies. $\text{Chor}\lambda$, and hence $\text{PolyChor}\lambda$, provides a core language for higher-order choreographies, thus allowing us to establish their properties. Since the original $\text{Chor}\lambda$ has parametric procedures like PC and Choral, it lacks $\text{PolyChor}\lambda$'s property that a choreography diverging in one process must diverge in every process. This forces $\text{Chor}\lambda$ to have both a complex notion of out-of-order execution and a more lax correspondence between actions in the network and the choreography.

The other work establishing the theoretical properties of higher-order choreographic programming is Pirouette (Hirsch & Garg, 2022), which is also a functional choreographic programming language based on simply-typed λ calculus. Unlike $\text{Chor}\lambda$ (and thus $\text{PolyChor}\lambda$), Pirouette does not allow processes to send messages written in Pirouette. Instead, it takes inspiration from lower-order choreographic programming languages in which (the computations to produce) messages are written in their own separate language. Like other choreographic languages (Cruz-Filipe *et al.*, 2021; Montesi, 2023), Pirouette's design is parametrised by the language for writing messages. Thus, Pirouette can describe communication patterns between processes that draw from a large swath of languages for their local computations. Nevertheless, this design means that Pirouette fundamentally cannot allow programmes to send choreographic functions, unlike $\text{PolyChor}\lambda$.

Moreover, unlike $\text{Chor}\lambda$ and $\text{PolyChor}\lambda$, Pirouette forces every process to synchronise when applying a function. This allows Pirouette to establish a bisimulation relation with its network programming language, a result formalised in Coq. This result allows a traditional—and verified—proof of deadlock-freedom by construction. However, this constant synchronisation would be a bottleneck in real-world systems; $\text{PolyChor}\lambda$'s choice

to obtain a weaker—but strong-enough—connection between the languages allows it to avoid this high cost.

7.2 Concurrent functional programming

Functional concurrent programming has a long history, starting with attempts to parallelise executions of functional programmes (Burton, 1987). The first language for functional programming with communications on channels was Facile (Giacalone et al., 1989) which, unlike later choreographic languages, had an abstraction over *process IDs* very similar to process polymorphism. A more recent work, which more closely resembles choreographic programming, is Links (Cooper et al., 2006), with the RPC calculus (Cooper & Wadler, 2009) as its core language. Links and the RPC calculus, like choreographies, allow a programmer to describe programmes where different parts of the computation takes place at different locations and then compile it to separate code for each location. Interestingly, though Links has explicit communication, in the RPC calculus the top level does not, and communications are created only when projecting a function located at a different process. Moreover, the RPC calculus does not feature multiple threads of computation; instead, on communication the single thread of computation moves to a new location while other locations block. The RPC calculus was later extended with location polymorphism, very similar to our and Facile's process polymorphism (Choi et al., 2020). However, as the RPC calculus only deals with systems of 2 processes, a client and a server, they project a process abstraction as a pair and then the location as picking the correct part of this pair. This solution creates a simpler network language but is not suitable for a framework with an arbitrary number of participants such as PolyChor λ . Moreover, the RPC calculus—like PolyChor λ but unlike traditional choreographic languages—does not have out-of-order execution at the top level.

Session types were applied to a concurrent functional calculus with asynchronous communication by Gay & Vasconcelos (2010). Though initially this language did not guarantee deadlock-freedom, only runtime safety, later versions of GV (Wadler, 2012; Lindley & Morris, 2015) did. Jacobs et al. (2022) extended GV with *global types* (Honda et al., 2016), which generalise session types to protocols with multiple participants. Similarly to choreographic programming, global types offer a global viewpoint on interactions. However, they are intended as specifications and thus cannot express computation. Global types are typically projected onto *local types*, which manually-written programmes can later be checked against. In choreographic programming, by contrast, choreographies are projected directly to programmes. Some works mix the approaches (e.g., Scalas et al., 2017): given a global type, a compiler produces typestate-oriented libraries (Aldrich et al., 2009) that help the users with following the global type correctly (but not with performing the right computations at the right time).

Session types have also been used to study global higher-order programming outside of functional settings. Mostrous & Yoshida (2007) describe the challenges associated with obtaining subject reduction when sessions can pass other sessions between them. Based on this, Mostrous & Yoshida (2009) define the higher-order π -Calculus with asynchronous sessions, a calculus combining elements of the π -calculus and λ -calculus.

Castellani *et al.* (2020) propose a notion of *session types with delegation*. They write delegation by enclosing a part of the global type in brackets. During the execution of such a part, one process acts on another's behalf by temporarily taking its name. This means that they do not need to inform other participants in the delegated computation that delegation is happening. However, nested delegations can cause deadlocks.

ML5 (Murphy VII *et al.*, 2007; Licata & Harper, 2010) is a functional concurrent programming language based on the semantics of modal logic. However, instead of the `send` and `recv` terms of choreographic languages, they have a primitive `get[w] M`, which makes another process w evaluate M and return the result. Since M may include other gets, this construct gives ML5 something resembling PolyChor λ 's ability to send a full choreography. However, the result of evaluating this "choreography" must be at the receiver and then returned to the sender.

Multitier programming languages, like ScalaLoc (Weisenburger & Salvaneschi, 2020), offer another paradigm for describing global views of distributed systems. Like Choral, ScalaLoc is built on top of an existing object-oriented language: in this case, Scala. In ScalaLoc and other *tierless* languages, an object describes a whole system containing multiple processes and functions. Differently from choreographic programming, multitier programming does not allow for modelling the intended flow of communications. Rather, communication happens implicitly through remote function calls and the concrete protocol to be followed is largely left to be decided by a runtime middleware. For a more detailed comparison of choreographic and multitier languages, see the work of Giallorenzo *et al.* (2021).

8 Conclusion

In this paper, we presented PolyChor λ , the first higher-order choreographic programming language with process polymorphism. PolyChor λ has a type and kind system based on System $F\omega$, but extended such that process names are types of kind `Proc`. Moreover, we showed how to obtain a concurrent interpretation of PolyChor λ programmes in a process language by using a new construct corresponding by the ability of a process to know its identity. We found that this construct was necessary if process variables are able to be instantiated as the process they are located at, but using a choreographic language abstracts from this necessity. Our explorations of process polymorphism also allowed PolyChor λ to describe a communication of a non-local value from P to Q as sending the part of the message owned by P to Q . These non-local values include full choreographies, creating a simple and flexible way to describe delegation by communicating a distributed function describing the delegated task. This innovation required a new notion of communication as an exchange in which the delegator rather than being left with an empty value after sending a choreography is left with a function which will allow it to potentially take part in the delegated task, e.g., by receiving a result at the end.

Process polymorphism fills much of the gap between previous works on the theory of higher-order choreographies and practical languages. However, there is still more work to do. For instance, currently PolyChor λ does not support recursive types.

Our current results rely on types normalising to a type value, which recursive types may not do. System $F\omega$ does not have our restriction of type abstractions only being instantiated with type values. However, PolyChor λ needs to ensure that communications

are only undertaken between processes, rather than complicated type expressions resulting in processes. Thus, we need to treat our type system as call-by-name, leading to the restriction above.

In order to support recursive types, we would need to either make endpoint projection capable of projecting to a possibly-nonterminating description of a process, or limit recursive types ability to make type computations fail to terminate.

Furthermore, one can imagine allowing processes to send types and process names as well as values. This would, for example, allow us to programme a server to wait to receive the name of a client which it will have to interact with. Since this form of delegation is common in practice, understanding how to provide this capability in a choreographic language, while retaining the guarantees of choreographic programming, would enable programmers to apply their usual programming patterns to choreographic code.

We project local type despite lacking a typing system for local processes. Our unusual network communication semantics have made it difficult to define local typing rules for `sends` and `recvs`, and we therefore leave local typing (or alternatively type erasure) as future work.

Certain, instant, and synchronous communication is convenient for theoretical study, but such assumptions do not match real-world distributed systems. Cruz-Filipe & Montesi (2017a) model asynchronous communication in choreographies via runtime terms representing messages in transit. We could adapt this method to PolyChor λ by having the communication primitive reduce in two steps: first to a runtime term and then to the delivered value. However, this extension would be nontrivial, since it is not obvious how to represent messages in transit when those messages are non-local values such as choreographies. In addition, the way we represent a communication at the local level (swapping values rather than only moving a value from sender to receiver) might require additional machinery (e.g., new runtime terms) to capture its asynchronous execution.

We also leave practical implementation of PolyChor λ 's new features to future work. This could be achieved by extending Choral (Giallorenzo *et al.*, 2023), the original inspiration for Chor λ . Communication primitives in Choral are user-defined—not fixed by any middleware or compiler—so it is possible to define new communication abstractions involving multiple roles. However, we need to manipulate roles at runtime in our local semantics, while the Choral compiler erases roles when projection code to Java. We may be able to overcome this issue by reifying roles in projected code or by using reflection.

While these gaps between theory and practice persist, process polymorphism in PolyChor λ brings us much closer to realistic choreographic languages for distributed systems. Choreographic programmes promise to provide easier and cleaner concurrent and distributed programming with better guarantees. With higher-order choreographic programming and process polymorphism, the fulfilment of that promise is nearly within reach.

Acknowledgements

We thank the anonymous reviewers for their useful comments.

Funding

This work was supported by Villum Fonden (F.M., grant numbers 29518, 50079) and Independent Research Fund Denmark (F.M., grant number 0135-00219).

Conflicts of interest

The authors report no conflict of interest.

References

- Aldrich, J., Sunshine, J., Saini, D. & Sparks, Z. (2009) Typestate-oriented programming. In Companion to the 24th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2009, October 25–29, 2009, Orlando, Florida, USA. ACM, pp. 1015–1022.
- Burton, F. (1987) Functional programming for concurrent and distributed programming. *Comput. J.* **30**(5), 437–450.
- Caires, L. & Pfenning, F. (2010) Session types as intuitionistic linear propositions. In Concurrency Theory (CONCUR).
- Carbone, M., Honda, K. & Yoshida, N. (2012) Structured communication-centered programming for web services. *Trans. Program. Lang. Syst. (TOPLAS)* **34**(2), 1–78.
- Carbone, M. & Montesi, F. (2013) Deadlock-freedom-by-design: Multiparty asynchronous global programming. In Principles of Programming Languages (POPL).
- Castagna, G., Dezani-Ciancaglini, M. & Padovani, L. (2012) On global types and multi-party session. *Log. Methods Comput. Sci.* **8**(1), 1–45.
- Castellani, I., Dezani-Ciancaglini, M., Giannini, P. & Horne, R. (2020) Global types with internal delegation. *TCS* **807**, 128–153.
- Choi, K., Cheney, J., Fowler, S. & Lindley, S. (2020) A polymorphic RPC calculus. *SCP* **197**, 102499.
- Cooper, E., Lindley, S., Wadler, P. & Yallop, J. (2006) Links: Web programming without tiers. In *Formal Methods for Components and Objects (FMCO)*.
- Cooper, E. E. & Wadler, P. (2009) The RPC calculus. In Principles and Practice of Declarative Programming (PPDP).
- Cruz-Filipe, L., Graversen, E., Lugovic, L., Montesi, F. & Peressotti, M. (2022) Functional choreographic programming. In ICTAC.
- Cruz-Filipe, L., Graversen, E., Lugovic, L., Montesi, F. & Peressotti, M. (2023) Modular compilation for higher-order functional choreographies. In ECOOP.
- Cruz-Filipe, L. & Montesi, F. (2017a) On asynchrony and choreographies. In Interaction and Concurrency Experience (ICE).
- Cruz-Filipe, L. & Montesi, F. (2017b) Procedural choreographic programming. In Formal Techniques for Distributed Objects, Components, and Systems (FORTE).
- Cruz-Filipe, L. & Montesi, F. (2020) A core model for choreographic programming. *Theor. Comput. Sci.* **802**, 38–66.
- Cruz-Filipe, L., Montesi, F. & Peressotti, M. (2021) Formalizing a turing-complete choreographic language in coq. In Interactive Theorem Proving (ITP).
- Dalla Preda, M., Gabbriellini, M., Giallorenzo, S., Lanese, I. & Mauro, J. (2017) Dynamic choreographies: Theory and implementation. *Logical Methods Comput. Sci. (LMCS)* **13**(2), 1–57.
- Dardha, O., Giachino, E. & Sangiorgi, D. (2012) Session types revisited. In Principles and Practice of Declarative Programming (PPDP).
- DeYoung, H., Caires, L., Pfenning, F. & Toninho, B. (2012) Cut reduction in linear logic as asynchronous session-typed communication. In Computer Science Logic (CSL).

- Gay, S. J. & Vasconcelos, V. T. (2010) Linear type theory for asynchronous session types. *J. Funct. Program. (JFP)* **20**(1), 19–50.
- Giacalone, A., Mishra, P. & Prasad, S. (1989) Facile: A symmetric integration of concurrent and functional programming. *Int. J. Parallel Program.* **18**, 121–160.
- Giallorenzo, S., Montesi, F. & Peressotti, M. (2023) Choral: Object-oriented choreographic programming. *ACM Trans. Program. Lang. Syst.* Just Accepted.
- Giallorenzo, S., Montesi, F., Peressotti, M., Richter, D., Salvaneschi, G. & Weisenburger, P. (2021) Multiparty languages: The choreographic and multitier cases (pearl). In European Conference on Object-Oriented Programming (ECOOP).
- Girard, J.-Y. (1972) *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. PhD Thesis. Université Paris 7.
- Hirsch, A. K. & Garg, D. (2022) Pirouette: Higher-order typed functional choreographies. In Principles of Programming Languages (POPL).
- Honda, K. (1993) Types for dyadic interaction. In Concurrency Theory (CONCUR).
- Honda, K., Vasconcelos, V. T. & Kubo, M. (1998) Language primitives and type discipline for structured communication-based programming. In European Symposium on Programming (ESOP).
- Honda, K., Yoshida, N. & Carbone, M. (2016) Multiparty asynchronous session types. *J. ACM* **63**(1), 1–67.
- International Telecommunication Union. (1996) Recommendation Z.120: Message sequence chart.
- Jacobs, J., Balzer, S. & Krebbers, R. (2022) Multiparty GV: Functional multiparty session types with certified deadlock freedom. *Proc. ACM Program. Lang.* **6**(ICFP), 466–495.
- Jongmans, S. & van den Bos, P. (2022) A predicate transformer for choreographies - computing preconditions in choreographic programming. In European Symposium on Programming (ESOP).
- Kobayashi, N. (2006) A new type system for deadlock-free processes. In CONCUR 2006 - Concurrency Theory, 17th International Conference, CONCUR 2006, Bonn, Germany, August 27–30, 2006, Proceedings. Springer, pp. 233–247.
- Licata, D. R. & Harper, R. (2010) A monadic formalization of ML5. In Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP).
- Lindley, S. & Morris, J. (2017) Lightweight functional session types. In *Behavioural Types: from Theory to Tools*, Gay, S. & Ravara, A. (eds). River Publishers. chapter 12.
- Lindley, S. & Morris, J. G. (2015) A semantics for propositions as sessions. In *Programming Languages and Systems*.
- López, H. A., Nielson, F. & Nielson, H. R. (2016) Enforcing availability in failure-aware communicating systems. In Formal Techniques for Distributed Objects, Components, and Systems (FORTE).
- Montesi, F. (2013) *Choreographic Programming*. PhD Thesis. IT University of Copenhagen.
- Montesi, F. (2023) *Introduction to Choreographies*. Cambridge University Press.
- Mostrous, D. & Yoshida, N. (2007) Two session typing systems for higher-order mobile processes. In TLCA. Springer, pp. 321–335.
- Mostrous, D. & Yoshida, N. (2009) Session-based communication optimisation for higher-order mobile processes. In TLCA. Springer, pp. 203–218.
- Murphy VII, T., Crary, K. & Harper, R. (2007) Type-safe distributed programming with ML5. In Trustworthy Global Computer (TGC).
- Needham, R. M. & Schroeder, M. D. (1978) Using encryption for authentication in large networks of computers. *Commun. ACM* **21**(12), 993–999.
- Object Management Group. (2017) Unified modelling language, version 2.5.1.
- Pohjola, J. Å., Gómez-Londoño, A., Shaker, J. & Norrish, M. (2022) Kalas: A verified, end-to-end compiler for a choreographic language. In 13th International Conference on Interactive Theorem Proving, ITP 2022, August 7–10, 2022, Haifa, Israel. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 27:1–27:18.
- Scalas, A., Dardha, O., Hu, R. & Yoshida, N. (2017) A linear decomposition of multiparty sessions for safe distributed programming. In 31st European Conference on Object-Oriented Programming,

- ECOOP 2017, June 19–23, 2017, Barcelona, Spain. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 24:1–24:31.
- Scalas, A. & Yoshida, N. (2019) Less is more: Multiparty session types revisited. In *Principles of Programming Languages (POPL)*.
- Shen, G., Kashiwa, S. & Kuper, L. (2023) Haschor: Functional choreographic programming for all (functional pearl). CoRR. abs/2303.00924.
- The Coq Development Team. (2004) *The Coq proof assistant reference manual*. LogiCal Project. Version 8.0.
- The World Wide Web Consortium (W3C). (2004) *Ws choreography model overview*. Accessed January 29, 2021.
- Wadler, P. (2012) Propositions as sessions. In *International Conference on Functional Programming (ICFP)*.
- Weisenburger, P. & Salvaneschi, G. (2020) Implementing a language for distributed systems: Choices and experiences with type level and macro programming in scala. *ASEP* 4(3), 17.
- Zongyan, Q., Xiangpeng, Z., Chao, C. & Hongli, Y. (2007) Towards the theoretical foundation of choreography. In *The Web Conference (WWW)*.

A Full PolyChor λ typing rules

$$\begin{array}{c}
\text{[TUNIT]} \frac{\Theta; \Gamma \vdash v :: \mathbf{Proc}}{\Theta; \Gamma \vdash () @ v : () @ v} \qquad \text{[TINT]} \frac{\Theta; \Gamma \vdash v :: \mathbf{Proc}}{\Theta; \Gamma \vdash n @ v : \mathbf{Int} @ v} \\
\\
\text{[TAPP]} \frac{\Theta; \Gamma \vdash N : \tau_1 \rightarrow_{\rho} \tau_2 \quad \Theta; \Gamma \vdash M : \tau_1}{\Theta; \Gamma \vdash NM : \tau_2} \\
\\
\text{[TABS]} \frac{\Theta; \Gamma \vdash \tau_1 :: * \quad \Theta; \Gamma' \vdash v :: \mathbf{Proc} \text{ for all } v \in \rho \quad \Theta \cap (\rho \cup \text{ip}(\tau_1) \cup \text{ip}(\tau_2) \cup \text{ftv}(\tau_1) \cup \text{ftv}(\tau_2)); \Gamma, x : \tau_1 \vdash M : \tau_2}{\Theta; \Gamma \vdash \lambda x : \tau_1. M : \tau_1 \rightarrow_{\rho} \tau_2} \\
\\
\text{[TSEL]} \frac{\Theta; \Gamma \vdash v_1 :: \mathbf{Proc} \quad \Theta; \Gamma \vdash v_2 :: \mathbf{Proc} \quad \Theta; \Gamma \vdash M : \tau}{\Theta; \Gamma \vdash \text{select}_{v_1, v_2} \ell M : \tau} \\
\\
\text{[TCOM]} \frac{\Theta; \Gamma \vdash v_1 :: \mathbf{Proc} \setminus (\text{ip}(\tau) \cup \text{ftv}(\tau)) \quad \Theta; \Gamma \vdash v_2 :: \mathbf{Proc} \setminus (\text{ip}(\tau) \cup \text{ftv}(\tau)) \quad \Theta; \Gamma \vdash \tau :: \mathbf{Proc} \Rightarrow *}{\Theta; \Gamma \vdash \text{com}_{v_1, v_2}^{\tau} : (\tau v_1 \rightarrow \emptyset \tau v_2)} \\
\\
\text{[TAPPT]} \frac{\Theta; \Gamma \vdash M : \forall X :: K. \tau_1 \quad \Theta; \Gamma \vdash \tau_2 :: K}{\Theta; \Gamma \vdash M \tau_2 : \tau[X \mapsto \tau_2]} \\
\\
\text{[TABST]} \frac{\Theta'; \Gamma', X :: K \vdash M : \tau \quad \text{if } \exists K', \rho. K = K' \setminus \rho \text{ then } \Gamma' = (\Gamma + X) \& \rho \setminus X \text{ else } \Gamma' = \Gamma + X \quad \text{if } K = \mathbf{Proc} \text{ or } \exists \rho. K = \mathbf{Proc} \setminus \rho \text{ then } \Theta' = \Theta, X \text{ else } \Theta' = \Theta}{\Theta; \Gamma \vdash \Lambda X :: K. M : \forall X :: K. \tau} \\
\\
\text{[TEQ]} \frac{\Theta; \Gamma \vdash M : \tau_1 \quad \tau_1 \equiv \tau_2 \quad \Theta; \Gamma \vdash \tau_2 :: *}{\Theta; \Gamma \vdash M : \tau_2} \\
\\
\text{[TDEFS]} \frac{\forall f \in \text{domain}(D). f : \tau \in \Gamma \wedge \emptyset; \Gamma \vdash D(f) : \tau}{\Theta; \Gamma \vdash D} \qquad \text{[TVAR]} \frac{x : \tau \in \Gamma}{\Theta; \Sigma; \Gamma \vdash x : \tau} \\
\\
\text{[TCASE]} \frac{\Gamma \vdash N : \tau_1 + \tau_2 \quad \Theta; \Gamma, x : T_1 \vdash M' : \tau \quad \Theta; \Gamma, x' : T_2 \vdash M'' : \tau}{\Theta; \Gamma \vdash \text{case } N \text{ of } \text{inl } x \Rightarrow M'; \text{inr } x' \Rightarrow M'' : \tau} \\
\\
\text{[TFUN]} \frac{f : \tau \in \Gamma}{\Theta; \Gamma \vdash f : \tau} \qquad \text{[TPAIR]} \frac{\Theta; \Gamma \vdash M : \tau_1 \quad \Theta; \Gamma \vdash N : \tau_2}{\Theta; \Gamma \vdash (M, N) : \tau_1 \times \tau_2} \\
\\
\text{[TPROJ1]} \frac{\Theta; \Gamma \vdash M : \tau_1 \times \tau_2}{\Theta; \Gamma \vdash \text{fst } M : \tau_1} \qquad \text{[TPROJ2]} \frac{\Theta; \Gamma \vdash M : \tau_1 \times \tau_2}{\Theta; \Gamma \vdash \text{snd } M : \tau_2} \\
\\
\text{[TINL]} \frac{\Theta; \Sigma; \Gamma \vdash M : \tau_1}{\Theta; \Sigma; \Gamma \vdash \text{inl}_M \tau_2 : \tau_1 + \tau_2} \qquad \text{[TINR]} \frac{\Theta; \Sigma; \Gamma \vdash M : \tau_2}{\Theta; \Sigma; \Gamma \vdash \text{inr}_M \tau_2 : \tau_1 + \tau_2}
\end{array}$$

B Full PolyChor λ operational semantics

$$\begin{array}{c}
\text{[APPABS]} \frac{}{(\lambda x : \tau. M) V \rightarrow_D M[x \mapsto V]} \\
\text{[APP1]} \frac{M_1 \rightarrow_D M_2}{M_1 N \rightarrow_D M_2 N} \qquad \text{[APP2]} \frac{N_1 \rightarrow_D N_2}{V N_1 \rightarrow_D V N_2} \\
\text{[APPTABS]} \frac{\tau \equiv \nu}{(\Lambda X :: K. M) \tau \rightarrow_D M[X \mapsto \nu]} \qquad \text{[MTAPP1]} \frac{M_1 \rightarrow_D M_2}{M_1 \tau \rightarrow_D M_2 \tau} \\
\text{[MTAPP2]} \frac{\tau_1 \rightarrow_D \tau_2}{V \tau_1 \rightarrow_D V \tau_2} \\
\text{[INL]} \frac{M_1 \rightarrow_D M_2}{\text{inl}_\tau M_1 \rightarrow_D \text{inl}_\tau M_2} \qquad \text{[INR]} \frac{M_1 \rightarrow_D M_2}{\text{inr}_\tau M_1 \rightarrow_D \text{inr}_\tau M_2} \\
\text{[CASE]} \frac{N_1 \rightarrow_D N_2}{\text{case } N_1 \text{ of inl } x \Rightarrow M_1; \text{ inr } y \Rightarrow M_2 \rightarrow_D \text{case } N_2 \text{ of inl } x \Rightarrow M_1; \text{ inr } y \Rightarrow M_2} \\
\text{[CASEL]} \frac{}{\text{case inl}_\tau V \text{ of inl } x \Rightarrow M_1; \text{ inr } y \Rightarrow M_2 \rightarrow_D M_1[x \mapsto V]} \\
\text{[CASER]} \frac{}{\text{case inr}_\tau V \text{ of inl } x \Rightarrow M_1; \text{ inr } y \Rightarrow M_2 \rightarrow_D M_1[x \mapsto V]} \\
\text{[PAIR1]} \frac{M_1 \rightarrow_D M_2}{(M_1, N) \rightarrow_D (M_2, N)} \qquad \text{[PAIR2]} \frac{N_1 \rightarrow_D N_2}{(V, N_1) \rightarrow_D (V, N_2)} \\
\text{[FST]} \frac{M_1 \rightarrow_D M_2}{\text{fst } M_1 \rightarrow_D \text{fst } M_2} \qquad \text{[SND]} \frac{M_1 \rightarrow_D M_2}{\text{snd } M_1 \rightarrow_D \text{snd } M_2} \\
\text{[PROJ1]} \frac{}{\text{fst } (V_1, V_2) \rightarrow_D V_1} \qquad \text{[PROJ2]} \frac{}{\text{snd } (V_1, V_2) \rightarrow_D V_2} \\
\text{[DEF]} \frac{}{f \rightarrow_D D(f)} \\
\text{[SEL]} \frac{}{\text{select}_{P,Q} \ell M \rightarrow_D M} \qquad \text{[COM]} \frac{}{\text{com}_{P,Q}^\tau V \rightarrow_D V[P \mapsto Q]}
\end{array}$$

C Proof of Theorem 5

In the foregoing, we use L to denote local expressions and U to denote local values in order to make the proofs more readable, as we will be switching back and forth between layers a lot.

Before we can prove completeness, we need a few lemmas. First, we show that choreographic values always project to local values.

Lemma 3. *For any choreographic value V and process P , if $\Theta; \Gamma \vdash V : \tau$ then $\llbracket V \rrbracket_P$ is either a value or undefined.*

Proof Straightforward from the projection rules. ■

We then prove the same for type values.

Lemma 4. *Given a type value v , if $\Theta; \Gamma \vdash v :: K$ then for any process P in $ip(v)$, $\llbracket v \rrbracket_P = v$.*

Proof Straightforward from the projection rules. ■

We then show that type values are projected to \perp at uninvolved processes.

Lemma 5. *Given a type value $v \neq P$, for any process $Q \notin ip(v)$, $\llbracket v \rrbracket_Q = \perp$.*

Proof Straightforward from induction on v . ■

And similarly, we show that choreographic values project to \perp at processes not involved in their type.

Lemma 6. *Given a value V , if $\Theta; \Gamma \vdash V : \tau$ then for any process $P \notin ip(\tau)$, we have $\llbracket V \rrbracket_P = \perp$ or $\llbracket V \rrbracket_P$ is undefined.*

Proof Follows from Lemmas 3 and 5 and the projection rules. ■

Finally, we show that equivalent types are projected to equivalent local types.

Lemma 7. *Given a closed type τ , if $\tau \equiv \tau'$ and $\Theta; \Gamma \vdash \tau :: K$, then for any process P , $\llbracket \tau \rrbracket_P \equiv_P \llbracket \tau' \rrbracket_P$.*

Proof We prove this by structural induction on $\tau \equiv \tau'$. All but one case follow by simple induction.

The one interesting case is if $\tau = \lambda X. K \tau_1 v$ and $\tau' = \tau_1[X := v]$. Then (1) if $K = K' \setminus (\{R\} \cup \rho)$ and $\llbracket \tau_1 \rrbracket_P = \perp$, we have $\llbracket \tau \rrbracket_P = \llbracket \tau' \rrbracket_P = \perp$. (2) If $K \in \{\text{Proc}, \text{Proc} \setminus \rho\}$ then $\llbracket \tau \rrbracket_P = (\forall X. \text{Aml } X ? \llbracket \tau_1[X := P] \rrbracket_P \& \llbracket \tau_1 \rrbracket_P) \llbracket v \rrbracket_P$ and $\llbracket \tau' \rrbracket_P = \llbracket \tau_1[X := v] \rrbracket_P$. Since τ is a closed type and $\Theta; \Gamma \vdash v :: K$, v must be a process. If $v = P$ then obviously $(\forall X. \text{Aml } X ? \llbracket \tau_1[X := P] \rrbracket_P \& \llbracket \tau_1 \rrbracket_P) \llbracket v \rrbracket_P \equiv_P \llbracket \tau_1[X := v] \rrbracket_P$. If $v \neq P$ then $(\forall X. \text{Aml } X ? \llbracket \tau_1[X := P] \rrbracket_P \& \llbracket \tau_1 \rrbracket_P) \llbracket v \rrbracket_P \equiv_P \llbracket \tau_1 \rrbracket_P[X := \llbracket v \rrbracket_P]$, but since v is a process Q , $\llbracket v \rrbracket_P = Q$ and $\llbracket X \rrbracket_P = X$, and therefore we get $\llbracket \tau \rrbracket_P \equiv_P \llbracket \tau' \rrbracket_P$. And (3) otherwise we have $\llbracket \tau \rrbracket_P = (\lambda X. \llbracket \tau_1 \rrbracket_P) \llbracket v \rrbracket_P$ and $\llbracket \tau' \rrbracket_P = \llbracket \tau_1 \rrbracket_P[X := \llbracket v \rrbracket_P]$. We therefore get $\llbracket \tau \rrbracket_P \equiv_P \llbracket \tau' \rrbracket_P$. ■

We also need to prove that performing a substitution before and after projection yield the same result.

Lemma 8. *Given a choreography M with a free variable x , a value V , and a type τ such that $\Theta; \Gamma \vdash \lambda x : \tau. M \ V : \tau'$ and $\llbracket \lambda x : \tau. M \ V \rrbracket$ is defined, we get $\llbracket M[x := V] \rrbracket_{\mathcal{P}} = \llbracket M \rrbracket_{\mathcal{P}} [x := \llbracket V \rrbracket_{\mathcal{P}}]$.*

Proof If $\tau = \perp$ then by definition $\llbracket x \rrbracket_{\mathcal{P}} = \perp$ and by Lemma 6, $\llbracket V \rrbracket_{\mathcal{P}} = \perp$. If $\tau \neq \perp$ then $\llbracket x \rrbracket_{\mathcal{P}} = x$ and since we use α -conversion when substituting, we can guarantee that $\text{typeof}(V) = \tau$ anywhere it gets substituted into M , meaning the projection will always be the same as it does not depend on context, only on syntax and type. We therefore get $\llbracket M[x := V] \rrbracket_{\mathcal{P}} = \llbracket M \rrbracket_{\mathcal{P}} [x := \llbracket V \rrbracket_{\mathcal{P}}]$. ■

We are now ready to prove completeness.

Proof [Proof of Theorem 5] We prove this by structural induction on $M \rightarrow_D M'$.

- Assume $M = (\lambda x : \tau. N) \ V$ and $M' = N[x := V]$. Then for any process \mathcal{P} such that $\llbracket N \rrbracket_{\mathcal{P}} \neq \perp$ or $\llbracket \tau \rrbracket_{\mathcal{P}} \neq \perp$, we have $\llbracket M \rrbracket_{\mathcal{P}} = (\lambda x : \llbracket \tau \rrbracket_{\mathcal{P}}. \llbracket N \rrbracket_{\mathcal{P}}) \ \llbracket V \rrbracket_{\mathcal{P}}$ and $\llbracket M' \rrbracket_{\mathcal{P}} = \llbracket N \rrbracket_{\mathcal{P}} [x := \llbracket V \rrbracket_{\mathcal{P}}]$, and for any other \mathcal{P}' , we have $\mathcal{P}' \notin \text{ip}(\text{typeof}(V))$ and therefore $\llbracket M \rrbracket_{\mathcal{P}'} = \llbracket M' \rrbracket_{\mathcal{P}'} = \perp$. We therefore get $\mathcal{P} \llbracket \llbracket M \rrbracket_{\mathcal{P}} \rrbracket \xrightarrow{[D]} \llbracket \llbracket M' \rrbracket_{\mathcal{P}} \rrbracket$ for all $\mathcal{P} \in \text{ip}(\text{typeof}(\lambda x : \tau. N))$ and define $\mathcal{N} = \prod_{\mathcal{P} \in \text{ip}(\text{typeof}(\lambda x : \tau. N))} \mathcal{P} \llbracket \llbracket M' \rrbracket_{\mathcal{P}} \rrbracket \mid \prod_{\mathcal{P}' \notin \text{ip}(\text{typeof}(\lambda x : \tau. N))} \mathcal{P}' \llbracket \perp \rrbracket$ and the result follows.

- Assume $M = (\Delta X :: K. N) \ \tau$, $\tau \equiv v$, and $M' = N[X := v]$. Then if $K \in \{\mathcal{P}, \mathcal{P} \setminus \rho\}$, for any process \mathcal{P} , $\llbracket M \rrbracket_{\mathcal{P}} = (\Delta X :: K. \text{Aml } X \ ? \llbracket N[t := \mathcal{P}] \rrbracket_{\mathcal{P}} \ \& \ \llbracket N \rrbracket_{\mathcal{P}}) \ \llbracket \tau \rrbracket_{\mathcal{P}}$ and the result follows Lemmas 4 and 7, and Rules [NBabs], [Nlaml], and [Nlaml]. If $K \notin \{\text{Proc}, \text{Proc} \setminus \rho\}$ then for any process \mathcal{P} such that $\llbracket N \rrbracket_{\mathcal{P}} = \perp$ and $K = K' \setminus (\{\mathcal{P}\} \cup \rho)$, we have $\llbracket M \rrbracket_{\mathcal{P}} = \llbracket M' \rrbracket_{\mathcal{P}} = \perp$, for any other \mathcal{P}' we have $\llbracket M \rrbracket_{\mathcal{P}'} = (\Delta X. \llbracket N \rrbracket_{\mathcal{P}'}) \ \llbracket \tau \rrbracket_{\mathcal{P}'}$ and $\llbracket M' \rrbracket_{\mathcal{P}'} = \llbracket N \rrbracket_{\mathcal{P}'} [X := \llbracket v \rrbracket_{\mathcal{P}'}]$. We therefore get $\mathcal{P} \llbracket \llbracket M \rrbracket_{\mathcal{P}} \rrbracket \rightarrow_{[D]}^* \llbracket \llbracket M' \rrbracket_{\mathcal{P}} \rrbracket$ for all \mathcal{P} and the result follows.

- Assume $M = N \ M''$, $M' = N' \ M''$, and $N \rightarrow_D N'$. Then, for any process \mathcal{P} such that $\llbracket N \rrbracket_{\mathcal{P}} = \llbracket M'' \rrbracket_{\mathcal{P}} = \perp$, by induction we have $\llbracket N' \rrbracket_{\mathcal{P}} = \perp$, and therefore $\llbracket M \rrbracket_{\mathcal{P}} = \llbracket M' \rrbracket_{\mathcal{P}} = \perp$. For any process \mathcal{P}' such that $\mathcal{P}' \in \text{ip}(\text{typeof}(N))$ or $\llbracket N \rrbracket_{\mathcal{P}'} \neq \perp \neq \llbracket M'' \rrbracket_{\mathcal{P}'}$, $\llbracket M \rrbracket_{\mathcal{P}'} = \llbracket N \rrbracket_{\mathcal{P}'} \ \llbracket M'' \rrbracket_{\mathcal{P}'}$ and $\llbracket M' \rrbracket_{\mathcal{P}'} = \llbracket N' \rrbracket_{\mathcal{P}'} \ \llbracket M'' \rrbracket_{\mathcal{P}'}$. For any other process \mathcal{P}'' such that $\llbracket N \rrbracket_{\mathcal{P}''} = \perp$, by induction we get $\llbracket N' \rrbracket_{\mathcal{P}''} = \perp$ and therefore $\llbracket M \rrbracket_{\mathcal{P}''} = \llbracket M' \rrbracket_{\mathcal{P}''} = \llbracket M'' \rrbracket_{\mathcal{P}''}$. For any other process \mathcal{P}''' such that $\llbracket M'' \rrbracket_{\mathcal{P}'''} = \perp$, we get $\llbracket M \rrbracket_{\mathcal{P}'''} = \llbracket N \rrbracket_{\mathcal{P}'''} = \perp$ and $\llbracket M' \rrbracket_{\mathcal{P}'''} = \llbracket N' \rrbracket_{\mathcal{P}'''} = \perp$. And by induction $\llbracket N \rrbracket \rightarrow_{[D]}^* \mathcal{N}_N$ and $\mathcal{N}_N \sqsupseteq \llbracket N' \rrbracket$. For any process \mathcal{P} we therefore get $\llbracket N \rrbracket_{\mathcal{P}} \xrightarrow{[D]}^{\mu_0} \xrightarrow{[D]}^{\mu_1} \dots L_{\mathcal{P}}$ for $L_{\mathcal{P}} \sqsupseteq \llbracket N' \rrbracket_{\mathcal{P}}$ for some sequences of transitions $\xrightarrow{[D]}^{\mu_0} \xrightarrow{[D]}^{\mu_1} \dots$, and from the network semantics we get

$$\llbracket M \rrbracket \rightarrow_{[D]}^* \left(\prod_{\substack{\llbracket N \rrbracket_{\mathcal{P}} = \llbracket M'' \rrbracket_{\mathcal{P}} = \perp \\ \llbracket M \rrbracket_{\mathcal{P}''} = \llbracket M'' \rrbracket_{\mathcal{P}''}}} \prod_{\mathcal{P}' \in \text{ip}(\text{typeof}(N)) \text{ or } \llbracket N \rrbracket_{\mathcal{P}'} \neq \perp \neq \llbracket M'' \rrbracket_{\mathcal{P}'}} \mathcal{P}' \llbracket \perp \rrbracket \mid \prod_{\substack{\llbracket M \rrbracket_{\mathcal{P}''} = \llbracket M'' \rrbracket_{\mathcal{P}''}}} \prod_{\llbracket M \rrbracket_{\mathcal{P}''} = \llbracket N \rrbracket_{\mathcal{P}''}} \mathcal{P}'' \llbracket \llbracket M'' \rrbracket_{\mathcal{P}''} \rrbracket \mid \prod_{\llbracket M \rrbracket_{\mathcal{P}''} = \llbracket N \rrbracket_{\mathcal{P}''}} \mathcal{P}'' \llbracket L_{\mathcal{P}''} \rrbracket \right) = \mathcal{N}$$

and $M \rightarrow_D N' \ M''$. And since $\llbracket N \rrbracket \rightarrow_{[D]}^* \mathcal{N}'$ and $\llbracket N' \rrbracket \rightarrow_{[D]}^* \mathcal{N}'_N$, we know these sequences of transitions can synchronise when necessary, and if $\llbracket N \rrbracket_{\mathcal{P}'''} \neq \llbracket N' \rrbracket_{\mathcal{P}'''} = \perp$ then we can do the extra application to get rid of this unit.

- Assume $M = V N$, $M' = V N'$, and $N \rightarrow_D N'$. This is similar to the previous case, using Lemma 3 to ensure every process can use Rule [NApp2].
- Assume $M = \text{case } N \text{ of } \text{inl } x \Rightarrow N'; \text{ inr } x' \Rightarrow N''$, $M' = \text{case } M' \text{ of } \text{inl } x \Rightarrow N'; \text{ inr } x' \Rightarrow N''$, and $N \rightarrow_D M'$. Then for any process P such that $P \in \text{ip}(\text{typeof}(N))$, we have $\llbracket M \rrbracket_P = \text{case } \llbracket N \rrbracket_P \text{ of } \text{inl } x \Rightarrow \llbracket N' \rrbracket_P; \text{ inr } x' \Rightarrow \llbracket N'' \rrbracket_P$ and $\llbracket M' \rrbracket_P = \text{case } \llbracket M' \rrbracket_P \text{ of } \text{inl } x \Rightarrow \llbracket N' \rrbracket_P; \text{ inr } x' \Rightarrow \llbracket N'' \rrbracket_P$. For any other process P' such that $\llbracket N \rrbracket_{P'} = \llbracket N' \rrbracket_{P'} = \llbracket N'' \rrbracket_{P'} = \perp$, by induction we get $\llbracket M' \rrbracket_{P'} = \perp$, and therefore $\llbracket M \rrbracket_{P'} = \llbracket M' \rrbracket_{P'} = \perp$. For any other process P'' such that $\llbracket N \rrbracket_{P''} = \perp$, we get $\llbracket M \rrbracket_{P''} = \llbracket M' \rrbracket_{P''} = \llbracket N' \rrbracket_{P''} \sqcup \llbracket N'' \rrbracket_{P''}$. For any other processes P''' such that $\llbracket N' \rrbracket_{P'''} = \llbracket N'' \rrbracket_{P'''} = \perp$, we have $\llbracket M \rrbracket_{P'''} = \llbracket N \rrbracket_{P'''} = \perp$ and $\llbracket M' \rrbracket_{P'''} = \llbracket M'' \rrbracket_{P'''} = \perp$. For any other process P'''' , we have $\llbracket M \rrbracket_{P''''} = (\lambda x. \perp. \llbracket N' \rrbracket_{P''''} \sqcup \llbracket N'' \rrbracket_{P''''}) \llbracket N \rrbracket_{P''''}$ and $\llbracket M' \rrbracket_{P''''} = (\lambda x. \llbracket N' \rrbracket_{P''''} \sqcup \llbracket N'' \rrbracket_{P''''}) \llbracket M'' \rrbracket_{P''''}$ for $x \notin \text{fv}(N') \cup \text{fv}(N'')$. The rest follows by simple induction similar to the second case.
- Assume $M = \text{case inl}_\tau V \text{ of } \text{inl } x \Rightarrow N; \text{ inr } x' \Rightarrow N'$ and $M' = N[x := V]$. Then for any process $P \in \text{ip}(\text{typeof}(\text{inl}_\tau V))$, we have $\llbracket M \rrbracket_P = \text{case inl}_{[\tau]} \llbracket V \rrbracket_P \text{ of } \text{inl } x \Rightarrow \llbracket N \rrbracket_P; \text{ inr } x' \Rightarrow \llbracket N' \rrbracket_P$ and $\llbracket M' \rrbracket_P = \llbracket N[x := \llbracket V \rrbracket_P] \rrbracket_P$. By Lemma 6, $\llbracket N[x := \llbracket V \rrbracket_P] \rrbracket_P = \llbracket N \rrbracket_P[x := \llbracket V \rrbracket_P]$. For any other process $P' \notin \text{ip}(\text{typeof}(\text{inl}_\tau V))$, $\llbracket \text{inl}_V \tau P' \rrbracket = \perp$, and therefore $\llbracket M \rrbracket_{P'} = \llbracket N \rrbracket_{P'} \sqcup \llbracket N' \rrbracket_{P'} \supseteq \llbracket N \rrbracket_{P'} = \llbracket M' \rrbracket_{P'}$. The result follows.
- Assume $M = \text{case inr}_\tau V \text{ of } \text{inl } x \Rightarrow N; \text{ inr } x' \Rightarrow N'$ and $M' = N'[x' := V]$. This case is similar to the previous.
- Assume $M = \text{com}_{P,Q}^\tau V$ and $M' = V[Q := P]$. Then if $Q \neq P$, $\llbracket M \rrbracket_P = \text{recv}_Q \llbracket V \rrbracket_P$, $\llbracket M' \rrbracket_P = \llbracket V[Q := P] \rrbracket_P = \llbracket V \rrbracket_Q[Q := P]$, $\llbracket M \rrbracket_Q = \text{send}_P \llbracket V \rrbracket_Q$, $\llbracket M' \rrbracket_Q = \llbracket V[Q := P] \rrbracket_Q = \llbracket V \rrbracket_P[Q := P]$, and for any P' such that $\llbracket \tau \rrbracket_{P'} \neq \perp$, we have $\llbracket M \rrbracket_{P'} = \text{sub}[Q \mapsto P] \llbracket V \rrbracket_{P'}$ and $\llbracket M' \rrbracket_{P'} = \llbracket V[Q := P] \rrbracket_{P'} = \llbracket V \rrbracket_{P'}[Q := P]$, and for any other P'' , $\llbracket M \rrbracket_{P''} = \llbracket M' \rrbracket_{P''} = \perp$. We therefore get $\llbracket M \rrbracket_P \xrightarrow{\text{recv}_Q \llbracket V \rrbracket_Q[Q := P] \llbracket V \rrbracket_P} \llbracket M' \rrbracket_P$, $\llbracket M \rrbracket_Q \xrightarrow{\text{send}_P \llbracket V \rrbracket_P[Q := P] \llbracket V \rrbracket_Q} \llbracket M' \rrbracket_Q$, and $\llbracket M \rrbracket_{P'} \xrightarrow{\tau} \llbracket M' \rrbracket_{P'}$. We define $\mathcal{N} = \llbracket M' \rrbracket$ and the result follows. If $Q = P$, then $\llbracket M \rrbracket_P = (\lambda x.x) \llbracket V \rrbracket_P$ and $\llbracket M' \rrbracket_P = \llbracket V \rrbracket_P$ and $\mathcal{N} = \llbracket M' \rrbracket$ and the result follows.
- Assume $M = \text{select}_{Q,P} \ell M'$. Then $\llbracket M \rrbracket_Q = \bigoplus_P \ell \llbracket M' \rrbracket_Q$, $\llbracket M \rrbracket_P = \&_Q \{ \ell : \llbracket M' \rrbracket_P \}$, and for any $P' \notin \{Q, P\}$, $\llbracket M \rrbracket_{P'} = \llbracket M' \rrbracket_{P'}$. We therefore get $\llbracket M \rrbracket \xrightarrow{\tau_{P,Q}} \llbracket M \rrbracket \setminus \{P, Q\} \mid P[\llbracket M' \rrbracket_P] \mid Q[\llbracket M' \rrbracket_Q]$ and the result follows.
- Assume $M = (N, N')$, $N \rightarrow_D N''$, and $M' = (N'', N')$. Then, the result follows from simple induction.
- Assume $M = (V, N)$, $N \rightarrow_D N'$, and $M' = (V, N')$. Then, the result follows from simple induction.
- Assume $M = \text{fst}(V, V')$ and $M' = V$. Then for any process P such that $\llbracket V \rrbracket_P \neq \perp$ or $\llbracket V' \rrbracket_P \neq \perp$, $\llbracket M \rrbracket_P = \text{fst}(\llbracket M' \rrbracket_P, \llbracket V' \rrbracket_P)$ and for any other process $P' \notin \text{ip}(\text{typeof}((M', V')))$, we have $\llbracket M \rrbracket_{P'} = \perp$ and $\llbracket M' \rrbracket_{P'} = \perp$. We define $\mathcal{N} = \llbracket M' \rrbracket$ and the result follows.
- Assume $M = \text{snd}(V, V')$ and $M' = V'$. Then, the case is similar to the previous.
- Assume $M = f$ and $M' = D(f)$. Then, the result follows from the definition of $\llbracket D \rrbracket$.

■

D Proof of Theorem 6

As with completeness, we need some ancillary lemmas before we can prove soundness. For this, we need a notion of removing processes from a network.

Definition 3. Given a network $\mathcal{N} = \prod_{P \in \rho} P[L_P]$, we have $\mathcal{N} \setminus \rho' = \prod_{P \in (\rho \setminus \rho')} P[L_P]$.

First we show that actions in a network do not affect the roles not mentioned in the transition label.

Lemma 9. For any process P and network \mathcal{N} , if $\mathcal{N} \xrightarrow{\tau_{\mathcal{P}}} \mathcal{N}'$ and $P \notin \mathcal{P}$ then $\mathcal{N}(P) = \mathcal{N}'(P)$.

Proof Straightforward from the network semantics. ■

Then we show that removing processes from a network does not prevent it from performing actions involving different processes.

Lemma 10. For any set of processes ρ and network \mathcal{N} , if $\mathcal{N} \xrightarrow{\tau_{\mathcal{P}}} \mathcal{N}'$ and $\mathcal{P} \cap \rho = \emptyset$ then $\mathcal{N} \setminus \rho \xrightarrow{\tau_{\mathcal{P}}} \mathcal{N}' \setminus \rho$.

Proof Straightforward from the network semantics. ■

We finally show that if the projection of a choreographic type is equivalent to a local type value, then the original choreographic type is equivalent to a choreographic type value.

Lemma 11. Given a closed type τ_1 and process P , if $\Theta; \Gamma \vdash \tau_1 :: K$ and $\llbracket \tau_1 \rrbracket_P \equiv_P v$, then there exist a type v' such that: $\tau_1 \equiv v'$ and $\llbracket v' \rrbracket_P = v$.

Proof We prove this by structural induction on τ_1 . All but one case follows from simple induction.

Assume $\tau_1 = \tau_2 \tau_3$. Then if $\llbracket \tau_2 \rrbracket_P = \llbracket \tau_2 \rrbracket_P = \perp$, we have $\llbracket \tau_1 \rrbracket_P = \perp$, and therefore, $v = \perp = v'$. Otherwise, if $\llbracket \tau_3 \rrbracket_{P'} = \perp$ and $\text{kindof}(\tau_3) = K \setminus (\{P\} \cup \rho)$, we get $\llbracket \tau_1 \rrbracket_P = \llbracket \tau_2 \rrbracket_P$ and the result follows from induction. Otherwise if $\llbracket \tau_2 \rrbracket_P = \perp$, we get $\llbracket \tau_1 \rrbracket_P = \llbracket \tau_3 \rrbracket_P$ and the result follows from induction. Otherwise, we get $\llbracket \tau \rrbracket_P = \llbracket \tau_2 \rrbracket_P \llbracket \tau_3 \rrbracket_P$. By induction, $\llbracket \tau_2 \rrbracket_P \equiv_P v_2$ and there exists v'_2 such that $\tau_2 \equiv v'_2$ and $\llbracket v'_2 \rrbracket_P = v_2$ and $\llbracket \tau_3 \rrbracket_P \equiv_P v_3$ and there exists v'_3 such that $\tau_3 \equiv v'_3$ and $\llbracket v'_3 \rrbracket_P = v_3$. Because τ_1 is kindable, we have a kind K' such that $\Theta; \Gamma \vdash \tau_2 :: K' \Rightarrow K$ and $\Theta; \Gamma \vdash \tau_3 :: K'$. This means that $v'_2 = \lambda X :: K'. v_4$ and if $K' \in \{\text{Proc}, \text{Proc} \setminus \rho\}$ then $v_2 = \lambda X. \text{Aml } X ? \llbracket v_4[X := P] \rrbracket_P \& \llbracket v_4 \rrbracket_P$, otherwise $v_2 = \lambda X. \llbracket v_4 \rrbracket_P$. We then get $v \equiv_P \llbracket v'_4[X := v_3] \rrbracket_P$ and $v' \equiv v'_4[X := v'_3]$, and since X and v'_3 are both base types, so are $\llbracket v'_4[X := v_3] \rrbracket_P$ and $v'_4[X := v'_3]$. ■

We are then ready to prove soundness.

Proof [Proof of Theorem 6] We prove this by structural induction on M .

- Assume $M = V$. Then for any process P , $\llbracket M \rrbracket_P = U$, and therefore, $\llbracket M \rrbracket \not\xrightarrow{\tau_{\mathcal{P}}}$.

- Assume $M = N_1 N_2$. Then for any process P such that $\llbracket N_1 \rrbracket_P = \llbracket N_2 \rrbracket_P = \perp$, we have $\llbracket M \rrbracket_P = \perp$. For any process P' such that $P' \in \text{ip}(\text{typeof}(N_1))$ or $\llbracket N_1 \rrbracket_{P'} \neq \perp \neq \llbracket N_2 \rrbracket_{P'}$, $\llbracket M \rrbracket_{P'} = \llbracket N_1 \rrbracket_{P'} \llbracket N_2 \rrbracket_{P'}$. For any other process P'' such that $\llbracket N_2 \rrbracket_{P''} = \perp$, we get $\llbracket M \rrbracket_{P''} = \llbracket N_1 \rrbracket_{P''}$. For any other process P''' , we get $\llbracket M \rrbracket_{P'''} = \llbracket N_2 \rrbracket_{P'''}$. We then have 2 cases.
 - Assume $N_2 = V$. Then, $\llbracket N_2 \rrbracket_P = U$ by Lemma 3, and for any P' such that $P' \notin \text{ip}(\text{typeof}(N_2)) \subseteq \text{ip}(\text{typeof}(N_1))$, by Lemma 6, $\llbracket N_2 \rrbracket_{P'} = \perp$, and therefore, $\llbracket M \rrbracket_{P'} = \llbracket N_1 \rrbracket_{P'}$, and we have 5 cases.
 - * Assume $N_1 = \lambda x : \tau. N_3$. Then for any process P such that $\llbracket N_3 \rrbracket_P \neq \perp$ or $\llbracket \tau \rrbracket_P \neq \perp$, $\llbracket N_1 \rrbracket_P = \lambda x : \llbracket \tau \rrbracket_P. \llbracket N_3 \rrbracket_P$. And for any other process, $\llbracket N_1 \rrbracket_P = \perp$. The only transition available at any process, would then use Rule [NAbsApp]. This means for any transition $M \xrightarrow{\mathcal{P}}$, there exists P'' such that $\mathcal{P} = P''$. We then get $\llbracket M \rrbracket \xrightarrow{\mathcal{P}} \llbracket M \rrbracket \setminus \{P''\} \mid P''[\llbracket N_3 \rrbracket_{P''}[x := \llbracket N_2 \rrbracket_{P''}]]$. We say that $M' = N_3[x := N_2]$, and the result follows from using Rule [NAbsApp] in every process P such that $\llbracket M \rrbracket_P \neq \perp$ and induction.
 - * Assume $N_1 = \text{com}_{Q,P}^{\tau}$. Then if $Q \neq P$, $\llbracket M \rrbracket_Q = \text{send}_P \llbracket N_2 \rrbracket_Q$, $\llbracket M \rrbracket_P = \text{recv}_P \perp$, for any $P' \in \text{ip}(\tau)$, $\llbracket M \rrbracket_{P'} = \text{sub}[Q \mapsto P] \llbracket V \rrbracket_{P'}$, and for any other process P'' , $\llbracket N_1 \rrbracket_{P''} = \perp = \llbracket M \rrbracket_{P''}$. And if $Q = P$ then $\llbracket N_1 \rrbracket_P = \lambda x.x$. If $\mathcal{P} = Q, P$ then $\mathcal{N} = \llbracket M \rrbracket \setminus \{Q, P\} \mid Q[\llbracket N_2 \rrbracket_P [Q := P]] \mid P[\llbracket N_2 \rrbracket_Q [Q := P]]$. Because $\llbracket N_2 \rrbracket_P = \perp$ and $\llbracket N_2 \rrbracket_Q = U$, $N_2 = V$. Therefore, $M \xrightarrow{\mathcal{P}}_D V[Q := P]$ and for any $P' \in \text{ip}(\tau)$, by Rule [NSub], $\mathcal{N}(P') \xrightarrow{\tau}_{[D]} \llbracket V[Q := P] \rrbracket_{P'}$ and the result follows from induction. If $\mathcal{P} = P$ then either $Q = P$ or $\llbracket N_1 \rrbracket_P = \text{sub}[Q \mapsto P]$. If $Q = P$ then $\mathcal{N} = \llbracket M \rrbracket \setminus \{P\} \mid P[\llbracket N_2 \rrbracket_P]$ and the rest is similar to above. If $\llbracket N_1 \rrbracket_P = \text{sub}[Q \mapsto P]$ then the case is similar to one of the other two.
 - * Otherwise, $N_1 \neq V$ and either $\mathcal{P} = P$ or $\mathcal{P} = P, Q$. If $\mathcal{P} = P$ then either $\llbracket N_1 \rrbracket_P \xrightarrow{\tau}_{[D]} L$ and $P \in \text{ip}(\text{typeof}(N_1))$, $\mathcal{N} = \llbracket M \rrbracket \setminus \{P\} \mid P[\llbracket N_2 \rrbracket_P]$. We therefore have $\llbracket N_1 \rrbracket \xrightarrow{\tau_P}_{[D]} \llbracket N_1 \rrbracket \setminus \{P\} \mid P[L]$, and by induction, $N_1 \rightarrow_D^* N'_1$ such that $\llbracket N_1 \rrbracket \setminus \{P\} \mid P[L] \rightarrow_D^* \mathcal{N} \sqsupseteq \llbracket N'_1 \rrbracket$. Since all these transitions can be propagated past N_2 in the network and $\llbracket N_2 \rrbracket_{P'}$ in any process P' involved, we get the result for $M' = N'_1 N_2$. If $\mathcal{P} = P, Q$ then the case is similar.
 - If $N_2 \neq V$ then we have 2 cases.
 - * If $\mathcal{P} = P$ then either $\llbracket N_1 \rrbracket_P \xrightarrow{\tau}_{[D]} L$ or $\llbracket N_1 \rrbracket_P = U$ and $\llbracket N_2 \rrbracket_P \xrightarrow{\tau}_{[D]} L$ and the case is similar to the previous.
 - * If $\mathcal{P} = Q, P$ then there exists U such that either $\llbracket N_1 \rrbracket_Q \xrightarrow{\text{send}_P U}_{[D]} L_Q$ or $\llbracket N_2 \rrbracket_Q \xrightarrow{\text{send}_P U}_{[D]} L_Q$ and $\llbracket N_1 \rrbracket_P \xrightarrow{\text{recv}_Q U[Q:=P]}_{[D]} L_P$ or $\llbracket N_2 \rrbracket_P \xrightarrow{\text{recv}_Q U[Q:=P]}_{[D]} L_P$.

If $\llbracket N_1 \rrbracket_Q \xrightarrow{\text{send}_P U} \llbracket D \rrbracket_Q L_Q$ then $\llbracket N_1 \rrbracket_Q \neq U'$ and therefore $\llbracket N_1 \rrbracket_P \xrightarrow{\text{recv}_Q U[Q:=P]} \llbracket D \rrbracket_P L_P$ and the case is similar to the previous. If $\llbracket N_2 \rrbracket_Q \xrightarrow{\text{send}_P U} \llbracket D \rrbracket_Q L_Q$ then $\llbracket N_1 \rrbracket_Q = U'$, and therefore $\llbracket N_2 \rrbracket_P \xrightarrow{\text{recv}_Q U[Q:=P]} \llbracket D \rrbracket_P L_P$ and the case is similar to the previous.

- Assume $M = N \tau$. Then for any process P such that $\llbracket N \rrbracket_P = \llbracket \tau \rrbracket_P = \perp$, we have $\llbracket M \rrbracket_P = \perp$. For any process P' such that $\llbracket \tau \rrbracket_{P'} = \perp$ and $\text{kindof}(\tau) = K \setminus (\{P\} \cup \rho)$, $\llbracket M \rrbracket_{P'} = \llbracket N \rrbracket_{P'}$. For any other process P'' such that $\llbracket \tau \rrbracket_{P''} = \perp$, we get $\llbracket M \rrbracket_{P''} = \llbracket N \rrbracket_{P''}$. For any other process P''' , we get $\llbracket M \rrbracket_{P'''} = \llbracket N \rrbracket_{P'''} \llbracket \tau \rrbracket_{P'''}$. This case is similar to the previous unless $N = \Delta X :: K. N'$.

If $N = \Delta X :: K. N'$ and $\tau \equiv \nu$ then we have two cases. Either $K \in \{\text{Proc}, \text{Proc} \setminus \rho\}$ or not. If $K \in \{\text{Proc}, \text{Proc} \setminus \rho\}$ then for any P' , $\llbracket M \rrbracket_{P'} = \Delta X. \text{Aml } X ? \llbracket N'[t := P'] \rrbracket_P \& \llbracket N' \rrbracket_{P'} \llbracket \nu \rrbracket_{P'}$. As $\llbracket \nu \rrbracket_{P'} = P$ for some P , the only available transition is using Rule [NBabs], and we therefore get $\mathcal{P} = P''$ for some P'' and $\mathcal{N} = \llbracket M \rrbracket \setminus \{P''\} \mid P''[\text{Aml } P ? \llbracket N'[X := P''] \rrbracket_{P''} \& \llbracket N' \rrbracket_{P''}]$. We then define $M' = N'[X := \nu]$ and see that the result follows from using Rules [Nlaml] and [NProam] on P'' if $P'' = P$ and otherwise using Rules [Nlaml] and [NProam], at all other processes using Rule [NBabs] and then either Rules [Nlaml] and [NProam] or Rules [Nlaml] and [NProam] and the result follows from induction.

If $K \notin \{\text{Proc}, \text{Proc} \setminus \rho\}$ then the case is similar to $N_1 = \lambda X : \tau. N_3$ above.

- Assume $M = \text{fst } N$. Then either $N \neq V$ and the result follows from induction, or $N = (V, V')$ and for any process $P \in \text{ip}(\text{typeof}((V, V')))$, $\llbracket M \rrbracket_P = \text{fst}(\llbracket V \rrbracket_P, \llbracket V' \rrbracket_P)$ and for any other process $P' \notin \text{ip}(\text{typeof}((V, V')))$, by Theorem 6 we have $\llbracket M \rrbracket_{P'} = \llbracket N \rrbracket_{P'} = \perp$, and therefore $\llbracket M \rrbracket_{P'} \not\rightarrow \llbracket D \rrbracket_{P'}$.

If $\mathcal{P} = P \in \text{ip}(\text{typeof}((V, V')))$ then $\mathcal{N} = \llbracket M \rrbracket \setminus \{P\} \mid P[\llbracket V \rrbracket_P]$ and $M \xrightarrow{\tau \mathcal{P}}_D V$. The result follows by use of Rule [NProj1] and Theorem 6 and induction.

- Assume $N_1 = \text{snd } N_2$. This case is similar to the previous.
- Assume $M = (M_1, M_2)$. Then, the result follows from simple induction.
- Assume $M = \text{case } N \text{ of inl } x \Rightarrow N'; \text{ inr } x' \Rightarrow N''$. Then for any process P such that $P \in \text{ip}(\text{typeof}(N))$, we have $\llbracket M \rrbracket_P = \text{case } \llbracket N \rrbracket_P \text{ of inl } x \Rightarrow \llbracket N' \rrbracket_P; \text{ inr } x' \Rightarrow \llbracket N'' \rrbracket_P$. For any other process P' such that $\llbracket N \rrbracket_{P'} = \llbracket N' \rrbracket_{P'} = \llbracket N'' \rrbracket_{P'} = \perp$, $\llbracket M \rrbracket_{P'} = \perp$. For any other process P'' such that $\llbracket N \rrbracket_{P''} = \perp$, we get $\llbracket M \rrbracket_{P''} = \llbracket N' \rrbracket_{P''} \sqcup \llbracket N'' \rrbracket_{P''}$. For any other processes P''' such that $\llbracket N' \rrbracket_{P'''} = \llbracket N'' \rrbracket_{P'''} = \perp$, we have $\llbracket M \rrbracket_{P'''} = \llbracket N \rrbracket_{P'''}$. For any other process P'''' , we have $\llbracket M \rrbracket_{P''''} = (\lambda x : \perp. \llbracket N' \rrbracket_{P''''} \sqcup \llbracket N'' \rrbracket_{P''''}) \llbracket N \rrbracket_{P''''}$. We have two cases.

– Assume $\mathcal{P} = P \in \text{ip}(\text{typeof}(N))$. Then, we have three cases.

- * Assume $N = \text{inl}_\tau V$. Then, $\llbracket N \rrbracket_P = \text{inl}_{\llbracket \tau \rrbracket_P} \llbracket V \rrbracket_P$ and $\mathcal{N} = \llbracket M \rrbracket \setminus \{P\} \mid P[\llbracket N'[x := \llbracket V \rrbracket_P] \rrbracket_P]$. We define $M' = N'$ and since $\llbracket N' \rrbracket_{P'} \sqsupseteq \llbracket N' \rrbracket_{P'} \sqcup \llbracket N'' \rrbracket_{P'}$ the result follows from using Rules [NAbsApp] and [NCasel] and induction.
- * Assume $N = \text{inr}_\tau V$. Then, the case is similar to the previous.

- * Otherwise, we use Rule [NCase] and we have a transition $\llbracket N \rrbracket_{\mathbf{P}} \xrightarrow{\tau} \llbracket D \rrbracket_{\mathbf{P}} L$ such that

$$\mathcal{N} = \llbracket M \rrbracket \setminus \{\mathbf{P}\} \mid \mathbf{P}[\text{case } L \text{ of } \text{inl } x \Rightarrow \llbracket N' \rrbracket_{\mathbf{P}}; \text{inr } x' \Rightarrow \llbracket N'' \rrbracket_{\mathbf{P}}]$$

and the result follows from induction similar to the last application case.

- Assume $\mathcal{P} = \mathbf{Q}, \mathbf{P}$. Then, the logic is similar to the third subcases of the previous case.
 - Assume $M = \text{select}_{\mathbf{Q}, \mathbf{P}} \ell N$. This is similar to the $N_1 = \text{com}_{\mathbf{Q}, \mathbf{P}}^{\tau}$ case above.
 - Assume $M = f$. Then for any process \mathbf{P} , $\llbracket M \rrbracket_{\mathbf{P}} = f$. We therefore have some process \mathbf{P} such that $\mathcal{P} = \mathbf{P}$ and $\mathcal{N} = (\llbracket M \rrbracket \setminus \mathbf{P}) \mid \mathbf{P}[\llbracket D \rrbracket(f)]$. We then define the required choreography $M' = D(f)$ and network $\mathcal{N}' = \llbracket M' \rrbracket$, and the result follows. ■

■