

## Procedural Choreographic Programming

Cruz-Filipe, Luis; Montesi, Fabrizio

*Published in:*

Formal Techniques for Distributed Objects, Components, and Systems

*DOI:*

10.1007/978-3-319-60225-7\_7

*Publication date:*

2017

*Document version:*

Submitted manuscript

*Citation for pulished version (APA):*

Cruz-Filipe, L., & Montesi, F. (2017). Procedural Choreographic Programming. In A. Bouajjani, & A. Silva (Eds.), *Formal Techniques for Distributed Objects, Components, and Systems: 37th IFIP WG 6.1 International Conference, FORTE 2017, Held as Part of the 12th International Federated Conference on Distributed Computing Techniques, DisCoTec 2017, Neuchâtel, Switzerland, June 19-22, 2017* (pp. 92-107). Springer. [https://doi.org/10.1007/978-3-319-60225-7\\_7](https://doi.org/10.1007/978-3-319-60225-7_7)

Go to publication entry in University of Southern Denmark's Research Portal

### Terms of use

This work is brought to you by the University of Southern Denmark.  
Unless otherwise specified it has been shared according to the terms for self-archiving.  
If no other license is stated, these terms apply:

- You may download this work for personal use only.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying this open access version

If you believe that this document breaches copyright please contact us providing details and we will investigate your claim.  
Please direct all enquiries to [puresupport@bib.sdu.dk](mailto:puresupport@bib.sdu.dk)

# A Language for the Declarative Composition of Concurrent Protocols

Luís Cruz-Filipe and Fabrizio Montesi

University of Southern Denmark      {lcf, fmontesi}@imada.sdu.dk

**Abstract.** A recent study of bugs in real-world concurrent and distributed systems found that, while implementations of individual protocols tend to be robust, the composition of multiple protocols and its interplay with internal computation is the culprit for most errors. Multiparty Session Types and Choreographic Programming are methodologies for developing correct-by-construction concurrent and distributed software, based on global descriptions of communication flows. However, protocol composition is either limited or left unchecked. Inspired by these two methodologies, in this work we present a new language model for the safe composition of protocols, called Procedural Choreographies (PC). Protocols in PC are procedures, parameterised on the processes that enact them. Procedures define communications declaratively using global descriptions, and programs are written by invoking and composing these procedures. An implementation in terms of a process model is then mechanically synthesised, guaranteeing correctness and deadlock-freedom. We study PC in the settings of synchronous and asynchronous communications, and illustrate its expressivity with some representative examples.

## 1 Introduction

In the last decades, advances in multi-core hardware and large-scale networks have made concurrent and distributed systems widespread. Unfortunately, programming such systems is a notoriously error-prone activity. In [22], 105 randomly-selected concurrency bugs from real-world software projects are analysed. Of these, 31 are caused by deadlocks. Of the 74 remaining ones, 97% are caused by violations of the programmer’s intentions on atomicity or the ordering of actions.

The theory of Multiparty Session Types tackles these problems in communication protocols [18]. The idea is that developers express their intentions on communications declaratively, by writing protocol specifications from a global viewpoint using an “Alice and Bob” notation. Given such a global specification, an EndPoint Projection (EPP) mechanically synthesises the local specifications of the I/O actions for each participant. The local specifications are thus correct by construction. Then, a type system can be used to check that implementations in process models follow the generated local specifications. Multiparty session types guarantee that the implementation of each protocol, taken in isolation, is deadlock-free and faithful to its global specification – which, by its declarative nature, respects the programmer’s intentions. However, errors can still occur

due to the composition of multiple protocol executions, which is left unchecked. Different approaches for avoiding deadlocks in protocol compositions have been proposed [8,9], but these limit how protocols can be composed – e.g., protocols can be instantiated only in a certain order, or connections among processes should form a tree structure – and do not offer a means to easily and correctly translate the programmer’s intentions on ordering: the language for composing protocols is not declarative.

A more recent empirical study [21] reveals that protocol composition deserves more attention. It presents a taxonomy of 104 Distributed Concurrency (DC) bugs. The authors’ findings include a significant insight, which we quote (emphasis ours):

“Real-world DC bugs are hard to find because many of them linger in complex concurrent executions of multiple protocols. [...] *Individual protocols tend to be robust in general.* Only 18 DC bugs occur in individual protocols without any input fault condition [...]. *On the other hand, a large majority of DC bugs happen due to concurrent executions of multiple protocols [...]*”

In this quote, a protocol is not intended as an abstract specification as in multiparty session types, but rather as the concrete series of events that happen in an implementation, including internal computation. This motivates us to address the problem of protocol composition in the framework of Choreographic Programming [23]: a paradigm similar to Multiparty Session Types, where global descriptions of communications are not used as types, but rather as programs. In such programs, called choreographies, developers declaratively program communications among processes together with the internal computations that they perform. EPP is then used to synthesise an implementation in terms of a process model [5]. However, state of the art models for choreographic programming still do not support many of the features used in real-world programs, such as those in [21]. One prominent aspect, which is the focus of this paper, is that these models do not support arbitrary compositions of protocol executions, since they are not modular. In particular, the absence of full procedural abstraction disallows the development of reusable libraries that can be composed as “black boxes”.

Inspired by these observations, we propose a language model for the correct programming of concurrent and distributed systems based on message passing, called Procedural Choreographies (PC). PC is a new model for choreographic programming, where we can define a protocol execution as a procedure, parameterised on the processes that will actually enact it. Composition of protocol executions is then obtained by allowing for the arbitrary composition of procedure calls, a feature lacked by previous choreography models. Nevertheless, PC inherits all the good properties of languages based on declarative global descriptions of communications – as in choreographic programming and multiparty session types – and extends them to protocol compositions. In particular, the process implementations synthesised from choreographies are correct by construction and deadlock-free. Thus, PC contributes to bringing the current body

of work on safe concurrent programming nearer to dealing with the kinds of bugs analysed in [21].

*Example 1.* We discuss a parallel version of merge sort in PC. Although this is a toy example, it cannot be written in any previous model for choreographic programming, thus allowing us to present the key features of our work for a simple scenario. More realistic and involved examples are presented after the formal presentation of PC. We make the standard assumption that we have concurrent processes with local storage and computational capabilities. In this example, each process stores a list and can use the following local functions: `split1` and `split2`, respectively returning the first or second half of a list; `is_small`, which tests if a list has at most one element; and `merge`, which combines two sorted lists into one. The following (choreographic) procedure, `MS`, implements merge sort on the list stored at its parameter process `p`.<sup>1</sup>

```
MS(p) = if p.is_small then 0
        else p.start q1,q2; p.split1 -> q1; p.split2 -> q2;
            MS<q1>; MS<q2>; q1.* -> p; q2.* -> p.merge
```

Procedure `MS` starts by checking whether the list at process `p` is small, in which case it does not need to be sorted (0 denotes termination); otherwise, `p` starts two other processes `q1` and `q2` (`p.start q1,q2`), to which it respectively sends the first and the second half of the list (`p.split1 -> q1` and `p.split2 -> q2`). The procedure is recursively reapplied to `q1` and `q2`, which independently (concurrently) proceed to ordering their respective sub-lists. When this is done, `MS` stores the first ordered half from `q1` to `p` (`q1.* -> p`, where `*` retrieves the data stored in `q1`) and merges it with the ordered sub-list from `q2` (`q2.* -> p.merge`).

Our merge sort example showcases the key desiderata for PC:

**General recursion.** Procedure calls can be followed by arbitrary code.

**Parameterised procedures.** Procedures are parametric on their processes (`p` in `MS`), and can thus be reused with different processes (as in `MS<q1>` and `MS<q2>`).

**Process spawning.** The ability of starting new processes. There must be no bound on how many processes can be started, since this is decided at runtime. (In `MS`, the number of spawned processes depends on the size of the initial list.)

**Implicit Parallelism.** In the composition `MS<q1>; MS<q2>`, the two calls can be run in parallel because they involve separate processes and are thus non-interfering. Thus, the code synthesised from this program should be parallel.

In the remainder, we explore more sophisticated programs that require additional features, e.g., mobility of process names.

---

<sup>1</sup> In the remainder, we use a `monospaced` font for readability of our concrete examples, and other fonts for distinguishing syntactic categories in our formal arguments as usual.

## 1.1 Contributions

We summarise our development and contributions.

***Procedural Choreographies.*** We introduce Procedural Choreographies (PC), a new language model that supports all the features discussed above (§ 2). We evaluate the expressivity of PC not only with our concurrent merge sort example, but also with a more involved parallel downloader. This example makes use of additional features: mobility of process names (networks with connections that evolve at runtime) and propagation of choices among processes. It also makes heavy use of implicit parallelism to deal with the parallelisation of multiple streams.

***Asynchrony.*** PC can be endowed with both a synchronous (§ 2) and an asynchronous semantics (§ 5). All our results hold for both versions, and the two semantics enjoy a strong correspondence result (Theorem 6). This allows developers to reason about their programs in the simpler synchronous setting and then to use the asynchronous semantics to obtain a more concurrent implementation, without worrying about introducing unsafe behaviour.

***Typing.*** Mobility of process names requires careful handling, especially its interplay with procedure composition: if a procedure specifies that two processes interact, then they should be properly connected (know each other’s names). We introduce a novel typing discipline (§ 3) that prevents such errors by tracking the connections required by each procedure. It also checks that processes store data of the correct type for the local functions that use it. PC enjoys decidable type checking (Theorem 2) and type inference (Theorems 3 and 4).

***Endpoint Projection.*** We define an EndPoint Projection (EPP) that, given a choreography, synthesises a concurrent implementation in Procedural Processes (PP), our target process calculus (§ 4). PP is an abstraction of systems where concurrent processes communicate by referring to each other’s locations or identifiers, as it happens, e.g., in MPI [25] or the Internet Protocol. The synthesised code is correct by construction: it faithfully follows the behaviour of the originating choreography (Theorem 5). Also, EPP is transparent, in the sense that it does not introduce any auxiliary communications or computations. This means that a choreography always faithfully represents the actual efficiency and behaviour of the algorithm written by the programmer. All generated implementations are deadlock-free by construction (Corollary 1).

***Extensions.*** We further discuss two extensions that allow us to write more general procedures, enhancing the expressive power of PC: allowing parameters to be lists of processes and allowing procedure bodies to contain holes that can be filled at runtime with arbitrary code (§ 6). These extensions can be elegantly obtained by introducing minimal additions to the theory of PC, demonstrating its robustness, but they are nevertheless very useful in practice.

$$\begin{aligned}
C &::= \eta; C \mid I; C \mid \mathbf{0} & \eta &::= \mathbf{p}.e \rightarrow \mathbf{q}.f \mid \mathbf{p} \rightarrow \mathbf{q}[l] \mid \mathbf{p} \text{ start } \mathbf{q}^T \mid \mathbf{p} : \mathbf{q} \leftarrow r \\
\mathcal{D} &::= X(\widetilde{\mathbf{q}}^T) = C, \mathcal{D} \mid \emptyset & I &::= \text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2 \mid X(\bar{\mathbf{p}}) \mid \mathbf{0}
\end{aligned}$$

**Fig. 1.** Procedural Choreographies, Syntax.

## 2 Procedural Choreographies (PC)

We begin by introducing the language model of Procedural Choreographies (PC). We focus on the synchronous semantics in this section, as the underlying theory is a bit simpler. The asynchronous model is discussed in § 5.

**Syntax.** The syntax of PC is displayed in Figure 1. A procedural choreography is a pair  $\langle \mathcal{D}, C \rangle$ , where  $C$  is a choreography and  $\mathcal{D}$  is a set of procedure definitions. Process names  $(\mathbf{p}, \mathbf{q}, \mathbf{r}, \dots)$ , identify processes that execute concurrently. Each process is equipped with a memory cell that stores a single value of a fixed type. Specifically, we consider a fixed set  $\mathbb{T}$  of datatypes (numbers, lists, etc.); each process  $\mathbf{p}$  stores only values of type  $T_{\mathbf{p}} \in \mathbb{T}$ . Statements in a choreography can either be communication actions ( $\eta$ ) or compound instructions ( $I$ ), both of which can have continuations. Term  $\mathbf{0}$  is the terminated choreography, which we often omit in examples. We call all terms but  $\mathbf{0}; C$  *program terms*, or simply programs, since these form the syntax intended for developers to use for writing programs. Term  $\mathbf{0}; C$  is necessary only for the technical definition of the semantics, to capture termination of procedure calls with continuations, and can appear only at runtime. It is thus called a *runtime term*. This distinction plays a bigger role in § 5, which introduces more runtime terms to track the state of asynchronous communications.

Processes communicate via direct references (names) to each other.<sup>2</sup> In a value communication  $\mathbf{p}.e \rightarrow \mathbf{q}.f$ , process  $\mathbf{p}$  sends the result of evaluating expression  $e$  to  $\mathbf{q}$ . In  $e$ , the placeholder  $*$  is replaced at runtime with the data stored at process  $\mathbf{p}$ . When  $\mathbf{q}$  receives the value from  $\mathbf{p}$ , it applies to it the (total) function  $f$  and stores the result. The definition of  $f$  may also access the contents of  $\mathbf{q}$ 's memory.

In a selection term  $\mathbf{p} \rightarrow \mathbf{q}[l]$ ,  $\mathbf{p}$  communicates to  $\mathbf{q}$  its choice of label  $l$ , which is a constant. This term is intended to propagate information on which internal choice has been made by a process to another (see Remark 2 below).

In term  $\mathbf{p} \text{ start } \mathbf{q}^T$ , process  $\mathbf{p}$  spawns the new process  $\mathbf{q}$ , which stores data of type  $T$ . Process name  $\mathbf{q}$  is bound in the continuation  $C$  of  $\mathbf{p} \text{ start } \mathbf{q}^T; C$ .

Process spawning introduces the need for mobility of process names. In real-world systems, after execution of  $\mathbf{p} \text{ start } \mathbf{q}^T$ ,  $\mathbf{p}$  is the only process that knows the

<sup>2</sup> This makes PC easy to apply to mainstream settings based on actors, objects, or ranks (e.g., MPI). We could give a formulation of PC based on standard channels (from process calculi), where each pair of process names is a channel. This is evident also from the connection graph used in the semantics of PC, defined below.

name of  $q$ . Any other process wanting to communicate with  $q$  must therefore be first informed of its existence (as happens, e.g., in object- and service-oriented computing [13,16]). This is achieved with the introduction term  $p : q \leftrightarrow r$ , read “ $p$  introduces  $q$  and  $r$ ” (with  $p$ ,  $q$  and  $r$  distinct). As its double-arrow syntax suggests, this action represents *two* communications – one where  $p$  sends  $q$ ’s name to  $r$ , and another where  $p$  sends  $r$ ’s name to  $q$ . This will become explicit in § 4.

In a conditional term  $\text{if } p.e \text{ then } C_1 \text{ else } C_2$ , process  $p$  evaluates  $e$  to choose between the possible continuations  $C_1$  and  $C_2$ .

The set  $\mathcal{D}$  contains global procedures. Term  $X(\widetilde{q}^T) = C_X$  defines a procedure  $X$  with body  $C_X$ , which can be used anywhere in  $\langle \mathcal{D}, C \rangle$  – in particular, inside  $C_X$ . The names  $\widetilde{q}$  are bound to  $C_X$ , and they are exactly the free process names in  $C_X$ . Each procedure can be defined at most once in  $\mathcal{D}$ . Term  $X(\widetilde{p})$  calls (or invokes) procedure  $X$  by passing  $\widetilde{p}$  as parameters. Procedure calls inside definitions must be guarded, i.e., they can only occur after a communication action.

We work up to  $\alpha$ -equivalence in choreographies, assuming the Barendregt convention. Bound variables are renamed as needed when expanding procedure calls.

*Example 2.* Recall procedure  $\text{MS}$  from our merge sort example in the Introduction (Example 1). If we annotate the parameter  $p$  and the started processes  $q_1$  and  $q_2$  with a type, e.g.,  $\mathbf{List}(T)$  for some  $T$  (the type of lists containing elements of type  $T$ ), then  $\text{MS}$  is a valid procedure definition in PC, as long as we allow two straightforward syntactic conventions: (i)  $p \text{ start } \widetilde{q}^T$  stands for the sequence  $p \text{ start } q_1^{T_1}; \dots; p \text{ start } q_n^{T_n}$ ; (ii) a communication of the form  $p.e \rightarrow q$  stands for  $p.e \rightarrow q.\text{id}$ , where  $\text{id}$  is the identity function: it sets the content of  $q$  to the value received from  $p$ . We adopt these conventions also in the remainder.

*Remark 1 (Design choices).* We comment on two of our design choices.

The introduction action ( $p : q \leftrightarrow r$ ) requires a three-way synchronization, and essentially performs two communications. The alternative development of PC with asymmetric introduction (an action  $p : q \rightarrow r$  whereby  $p$  sends  $q$ ’s name to  $r$ , but not conversely) is essentially the same as ours. Since in our examples we always perform introductions in pairs, the current choice makes the presentation easier.

The restriction that each process stores only one value of a fixed type is, in practice, a minor constraint. As shown in Example 2, types can be tuples or lists, which mimics storing several values. Also, a process can create new processes with different types – so we can encode changing the type of  $p$  by having  $p$  create a new process  $p'$  and then continuing the choreography with  $p'$  instead of  $p$ .

*Remark 2 (Label Selection).* We briefly motivate the need for selections ( $p \rightarrow q[l]$ ). Consider the choreography  $\text{if } p.\text{coinflip} \text{ then } (p.* \rightarrow r) \text{ else } (r.* \rightarrow p)$ . Here,  $p$  flips a coin to decide whether to send a value to  $r$  or to receive a value from  $r$ . Since processes run independently and share no data, only  $p$  knows which branch of

$$\begin{array}{c}
\frac{\mathbf{p} \xrightarrow{G} \mathbf{q} \quad e[\sigma(\mathbf{p})/*] \downarrow v \quad f[\sigma(\mathbf{q})/*](v) \downarrow w}{G, \mathbf{p}.e \rightarrow \mathbf{q}.f; C, \sigma \rightarrow_{\mathcal{D}} G, C, \sigma[\mathbf{q} \mapsto w]} \text{ [C|Com]} \\
\\
\frac{\mathbf{p} \xrightarrow{G} \mathbf{q}}{G, \mathbf{p} \rightarrow \mathbf{q}[l]; C, \sigma \rightarrow_{\mathcal{D}} G, C, \sigma} \text{ [C|Sel]} \\
\\
\frac{}{G, \mathbf{p} \text{ start } \mathbf{q}^T; C, \sigma \rightarrow_{\mathcal{D}} G \cup \{\mathbf{p} \leftrightarrow \mathbf{q}\}, C, \sigma[\mathbf{q} \mapsto \perp_T]} \text{ [C|Start]} \\
\\
\frac{\mathbf{p} \xrightarrow{G} \mathbf{q} \quad \mathbf{p} \xrightarrow{G} \mathbf{r}}{G, \mathbf{p}: \mathbf{q} \leftrightarrow \mathbf{r}; C, \sigma \rightarrow_{\mathcal{D}} G \cup \{\mathbf{q} \leftrightarrow \mathbf{r}\}, C, \sigma} \text{ [C|Tell]} \\
\\
\frac{i = 1 \text{ if } e[\sigma(\mathbf{p})/*] \downarrow \text{true}, \quad i = 2 \text{ otherwise}}{G, (\text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2); C, \sigma \rightarrow_{\mathcal{D}} G, C_i \ ; C, \sigma} \text{ [C|Cond]} \\
\\
\frac{C_1 \preceq_{\mathcal{D}} C_2 \quad G, C_2, \sigma \rightarrow_{\mathcal{D}} G', C'_2, \sigma' \quad C'_2 \preceq_{\mathcal{D}} C'_1}{G, C_1, \sigma \rightarrow_{\mathcal{D}} G', C'_1, \sigma'} \text{ [C|Struct]}
\end{array}$$

**Fig. 2.** Procedural Choreographies, Semantics.

the conditional will be executed; but this information is essential for  $r$  to decide on its behaviour. To propagate  $p$ 's decision to  $r$ , we use selections:

$$\text{if } \mathbf{p}.\text{coinflip} \text{ then } (\mathbf{p} \rightarrow \mathbf{r}[\mathbf{L}]; \mathbf{p}. * \rightarrow \mathbf{r}) \text{ else } (\mathbf{p} \rightarrow \mathbf{r}[\mathbf{R}]; \mathbf{r}. * \rightarrow \mathbf{p})$$

Now  $r$  receives a label reflecting the choice made by  $p$ , and can decide what to do.

This intuition is formalised by the definition of EndPoint Projection in § 4. The first choreography above is not projectable, whereas the second one is. See also Example 5 at the end of this section.

**Semantics.** We define a reduction semantics  $\rightarrow_{\mathcal{D}}$  for PC, parameterised over  $\mathcal{D}$ . We model the state of processes with a (total) state function  $\sigma$ , where  $\sigma(\mathbf{p})$  denotes the value stored in  $\mathbf{p}$ . We assume that each type  $T \in \mathbb{T}$  has a special value  $\perp_T$ , representing an uninitialised process state. The semantics of PC also includes a connection graph  $G$ , keeping track of which processes know each other. In the rules,  $\mathbf{p} \xrightarrow{G} \mathbf{q}$  denotes that  $G$  contains an edge between  $\mathbf{p}$  and  $\mathbf{q}$ , and  $G \cup \{\mathbf{p} \leftrightarrow \mathbf{q}\}$  denotes the graph obtained from  $G$  by adding an edge between  $\mathbf{p}$  and  $\mathbf{q}$  (if missing).

Executing a communication action  $\mathbf{p}.e \rightarrow \mathbf{q}.f$  in rule [C|Com] requires that:  $\mathbf{p}$  and  $\mathbf{q}$  are connected in  $G$ ;  $e$  is well typed; and the type of  $e$  matches that expected by the function  $f$  at the receiver. The last two conditions are encapsulated in the notation  $e \downarrow v$ , read “ $e$  evaluates to  $v$ ”. Choreographies can thus deadlock (be unable to reduce) because of errors in the programming of communications; this issue is addressed by our typing discipline in § 3.

Rule [C|Sel] defines selection as a no-op for choreographies (see Remark 2).



$$\begin{array}{c}
\frac{\text{pn}(\eta) \cap \text{pn}(\eta') = \emptyset}{\eta; \eta' \equiv_{\mathcal{D}} \eta'; \eta} \text{ [C|Eta-Eta]} \quad \frac{\text{pn}(I) \cap \text{pn}(I') = \emptyset}{I; I' \equiv_{\mathcal{D}} I'; I} \text{ [C|I-I]} \quad \frac{\text{pn}(I) \cap \text{pn}(\eta) = \emptyset}{\eta; I \equiv_{\mathcal{D}} I; \eta} \text{ [C|I-Eta]} \\
\\
\frac{}{\mathbf{0}; C \preceq_{\mathcal{D}} C} \text{ [C|End]} \quad \frac{X(\widetilde{\mathbf{q}}^T) = C_X \in \mathcal{D}}{X\langle \widetilde{\mathbf{p}} \rangle; C \preceq_{\mathcal{D}} C_X[\widetilde{\mathbf{p}}/\widetilde{\mathbf{q}}] \mathbin{\text{\textcircled{;}}} C} \text{ [C|Unfold]} \\
\\
\frac{\{\mathbf{p}, \mathbf{q}\} \cap \text{pn}(\eta) = \emptyset}{\text{if } \mathbf{p}.e \text{ then } (\eta; C_1) \text{ else } (\eta; C_2) \equiv_{\mathcal{D}} \eta; \text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2} \text{ [C|Eta-Cond]} \\
\\
\frac{}{\text{if } \mathbf{p}.e \text{ then } (C_1; \eta) \text{ else } (C_2; \eta) \equiv_{\mathcal{D}} \text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2; \eta} \text{ [C|Cond-Eta]} \\
\\
\frac{\{\mathbf{p}, \mathbf{q}\} \cap \{\mathbf{r}, \mathbf{s}\} = \emptyset}{\text{if } \mathbf{p}.e \text{ then } (\text{if } \mathbf{q}.e' \text{ then } C_1 \text{ else } C_2) \text{ else } (\text{if } \mathbf{q}.e' \text{ then } C'_1 \text{ else } C'_2)} \text{ [C|Cond-Cond]} \\
\equiv_{\mathcal{D}} \\
\text{if } \mathbf{q}.e' \text{ then } (\text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C'_1) \text{ else } (\text{if } \mathbf{p}.e \text{ then } C_2 \text{ else } C'_2)
\end{array}$$

**Fig. 3.** Procedural Choreographies, Structural precongruence  $\preceq$ .

Rule [C|Start] models the creation of a process. In the reductum, the starter and started processes are connected and can thus communicate with each other. Rule [C|Tell] captures name mobility, by creating a connection between two processes  $\mathbf{q}$  and  $\mathbf{r}$  when they are introduced by a process  $\mathbf{p}$  that is connected to both.

Rule [C|Cond] uses the auxiliary operator  $\mathbin{\text{\textcircled{;}}}$  to obtain a reductum in the syntax of PC regardless of the forms of the branches  $C_1$  and  $C_2$  and the continuation  $C$ . Operator  $\mathbin{\text{\textcircled{;}}}$  is defined by  $\eta \mathbin{\text{\textcircled{;}}} C = \eta; C$ ,  $I \mathbin{\text{\textcircled{;}}} C = I; C$  and  $(C_1; C_2) \mathbin{\text{\textcircled{;}}} C = C_1; (C_2 \mathbin{\text{\textcircled{;}}} C)$ . This operator extends the scope of bound names: any name  $\mathbf{p}$  bound in  $C$  has its scope extended also to  $C'$ . This scope extension is capture-avoiding, as the Barendregt convention guarantees that  $\mathbf{p}$  is not used in  $C'$ .

Rule [C|Struct] makes use of the structural precongruence  $\preceq_{\mathcal{D}}$ , which is defined by the rules in Figure 3. We write  $C \equiv_{\mathcal{D}} C'$  when  $C \preceq_{\mathcal{D}} C'$  and  $C' \preceq_{\mathcal{D}} C$ , and denote the set of process names (free or bound) in a choreography  $C$  by  $\text{pn}(C)$ . Rule [C|Unfold] unfolds a procedure call, again using the  $\mathbin{\text{\textcircled{;}}}$  operator defined above, and rule [C|End] is garbage collection of  $\mathbf{0}$ .

The other rules formalise the notion of implicit parallelism anticipated in § 1. Rule [C|Eta-Eta] permutes two communications performed by processes that are all distinct, modelling that processes run independently of one another. For example,  $\mathbf{p}. * \rightarrow \mathbf{q}; \mathbf{r}. * \rightarrow \mathbf{s} \equiv_{\mathcal{D}} \mathbf{r}. * \rightarrow \mathbf{s}; \mathbf{p}. * \rightarrow \mathbf{q}$  because these two communications are non-interfering, but  $\mathbf{p}. * \rightarrow \mathbf{q}; \mathbf{q}. * \rightarrow \mathbf{s} \not\equiv_{\mathcal{D}} \mathbf{q}. * \rightarrow \mathbf{s}; \mathbf{p}. * \rightarrow \mathbf{q}$ ; since the second communication causally depends on the first (both involve  $\mathbf{q}$ ).

This reasoning is extended to instructions in rule [C|I-I]; in particular, two procedure calls can be swapped if they share no arguments. This is sound because a procedure can only define actions for processes that are either passed as arguments or started inside of the procedure itself, and the latter cannot be leaked to the original call site. Thus, any actions obtained by unfolding the

first procedure call involve different processes than those obtained by unfolding the second one. As the example below shows, calls to the same procedure can be exchanged, since  $X$  and  $Y$  need not be distinct. The other rules follow similar intuitions; we omit rules  $[C|I\text{-}Cond]$  and  $[C|Cond\text{-}I]$ , analogous to  $[C|Eta\text{-}Cond]$  and  $[C|Cond\text{-}Eta]$ .

*Example 3.* In our merge sort example, rule  $[C|I\text{-}I]$  allows the recursive calls  $MS\langle q_1 \rangle$  and  $MS\langle q_2 \rangle$  to be exchanged. Furthermore, after the calls are unfolded, implicit parallelism allows their code to be interleaved in any way.

This example exhibits typical map-reduce behaviour: each new process receives its input, runs independently from all others, and then sends its result to its creator.

*Example 4.* A more refined example of implicit parallelism involves swapping communications from procedure calls that share process names. Consider the procedure

```
auth(c,a,r,l) = c.creds -> a.rCreds;
               a.chk -> r.res; a.log -> l.app
```

Client  $c$  sends its credentials to an authentication server  $a$ , which stores the result of authentication in  $r$  and appends a log of this operation at process  $l$ . In the choreography  $auth\langle c, a_1, r_1, l \rangle$ ;  $auth\langle c, a_2, r_2, l \rangle$ , a client  $c$  authenticates at two different authentication servers  $a_1$  and  $a_2$ . After unfolding the two calls, rule  $[C|Eta\text{-}Eta]$  yields the following interleaving:

```
c.creds -> a1.rCreds; c.creds -> a2.rCreds;
a2.chk -> r2.res; a1.chk -> r1.res;
a1.log -> l.app; a2.log -> l.app
```

Thus, the two authentications proceed in parallel. Observe that the logging operations cannot be swapped, since they use the same logging process  $l$ .

*Example 5.* A more sophisticated example involves modularly composing different procedures that take multiple parameters. Here, we write a choreography where a client  $c$  downloads a collection of files from a server  $s$ . The key idea is to download all files in parallel via streaming, by having the client and the server each create subprocesses to handle the transfer of each file. This allows the client to request and start downloading each file without waiting for previous downloads to finish.

```
par_download(c,s) = if c.more
  then c -> s [more]; c.start c'; s.start s';
       s: c <-> s'; c.top -> s'; pop<c>;
       c: c' <-> s'; download<c',s'>;
       par_download<c,s>; c'.file -> c.store
  else c -> s [end]
```

At the start of `par_download`, the client  $c$  checks whether it wants to download more files and informs the server  $s$  of the result via a label selection. In the

affirmative case, the client and the server start two subprocesses,  $c'$  and  $s'$  respectively, and the server introduces  $c$  to  $s'$  ( $s: c \leftrightarrow s'$ ). The client  $c$  sends to  $s'$  the name of the file to download ( $c.\text{top} \rightarrow s'$ ) and removes it from its collection, using procedure `pop` (omitted), afterwards introducing its own subprocess  $c'$  to  $s'$ . The file download is handled by  $c'$  and  $s'$  (using procedure `download`), while  $c$  and  $s$  continue operating (`par_download<c,s>`). Finally,  $c'$  waits until  $c$  is ready to store the downloaded file.

Procedure `download` has a similar structure. It implements a stream where a file is sequentially transferred in chunks from a process  $s$  to another process  $c$ .

```
download(c,s) = if s.more
  then s -> c [more]; s.next -> c.append; pop<s>; download<c,s>
  else s -> c [end]
```

The implementation of `par_download` exploits implicit parallelism considerably. All calls to `download` are made with disjoint sets of parameters (processes), thus they can be fully parallelised by our semantics: many instances of `download` run at the same time, each one implementing a (sequential) stream. By implicit parallelism, we effectively end up executing many streaming behaviours in parallel.

We can even compose `par_download` with `auth`, such that we execute the parallel download only if the client can successfully authenticate with an authentication server  $a$ . Below, we use the shortcut  $p \rightarrow \bar{q}[l]$  for  $p \rightarrow q_1[l]; \dots; p \rightarrow q_n[l]$ .

```
auth<c,a,r,l>; if r.ok then r -> c,s[ok]; par_download<c,s>
  else r -> c,s[ko]
```

### 3 Typability and Deadlock-Freedom

We give a typing discipline for PC, to check that (a) the types of functions and processes are respected by communications and (b) processes that need to communicate are first properly introduced (or connected). Regarding (b), two processes created independently can communicate only after they receive the names of each other. For instance, in Example 5, the execution of `download<c',s'>` would get stuck if  $c'$  and  $s'$  were not properly introduced in `par_download`, since our semantics requires them to be connected.

Typing judgements have the form  $\Gamma; G \vdash C \triangleright G'$ , read “ $C$  is well-typed according to the typings in  $\Gamma$ , and when executed from a connection graph that contains  $G$  it produces a connection graph that includes  $G'$ ”. Typing environments  $\Gamma$  are used to track the types of processes and procedures; they are defined as:  $\Gamma ::= \emptyset \mid \Gamma, p : T \mid \Gamma, X : G \triangleright G'$ . A typing  $p : T$  states that process  $p$  stores values of type  $T$ , and a typing  $X : G \triangleright G'$  records the effect of the body of  $X$  on graph  $G$ .

The rules for deriving typing judgements are given in Figure 4. We assume standard typing judgements for functions and expressions, and write  $* : T \vdash_{\top} e : T$  and  $* : T_1 \vdash_{\top} f : T_2 \rightarrow T_3$  meaning, respectively “ $e$  has type  $T$  assuming that  $*$  has type  $T$ ” and “ $f$  has type  $T_2 \rightarrow T_3$  assuming that  $*$  has type  $T_1$ ”.

$$\begin{array}{c}
\frac{}{\Gamma; G \vdash \mathbf{0} \triangleright G} \text{ [T|End]} \quad \frac{\mathbf{p} \xrightarrow{G} \mathbf{q} \quad \Gamma; G \vdash C \triangleright G'}{\Gamma; G \vdash \mathbf{p} \rightarrow \mathbf{q}[l]; C \triangleright G'} \text{ [T|Sel]} \quad \frac{\Gamma; G \vdash C \triangleright G'}{\Gamma; G \vdash \mathbf{0}; C \triangleright G'} \text{ [T|EndSeq]} \\
\frac{\mathbf{p} \xrightarrow{G} \mathbf{q} \quad \Gamma \vdash \mathbf{p} : T_{\mathbf{p}}, \mathbf{q} : T_{\mathbf{q}} \quad * : T_{\mathbf{p}} \vdash_{\mathbb{T}} e : T_1 \quad * : T_{\mathbf{q}} \vdash_{\mathbb{T}} f : T_1 \rightarrow T_{\mathbf{q}} \quad \Gamma; G \vdash C \triangleright G'}{\Gamma; G \vdash \mathbf{p}.e \rightarrow \mathbf{q}.f; C \triangleright G'} \text{ [T|Com]} \\
\frac{\Gamma \vdash \mathbf{p} : T \quad * : T \vdash_{\mathbb{T}} e : \text{bool} \quad \Gamma; G \vdash C_i \triangleright G_i \quad \Gamma; G_1 \cap G_2 \vdash C \triangleright G'}{\Gamma; G \vdash (\text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2); C \triangleright G'} \text{ [T|Cond]} \\
\frac{\Gamma, \mathbf{q} : T; G \cup \{\mathbf{p} \leftrightarrow \mathbf{q}\} \vdash C \triangleright G'}{\Gamma; G \vdash \mathbf{p}\text{start } \mathbf{q}^T; C \triangleright G'} \text{ [T|Start]} \quad \frac{\mathbf{p} \xrightarrow{G} \mathbf{q} \quad \mathbf{p} \xrightarrow{G} \mathbf{r} \quad \Gamma; G \cup \{\mathbf{q} \leftrightarrow \mathbf{r}\} \vdash C \triangleright G'}{\Gamma; G \vdash \mathbf{p} : \mathbf{q} \leftarrow \mathbf{r}; C \triangleright G'} \text{ [T|Tell]} \\
\frac{\Gamma \vdash X(\widetilde{\mathbf{q}}^T) : G_X \triangleright G'_X \quad \Gamma \vdash \mathbf{p}_i : T_i \quad G_X[\widetilde{\mathbf{p}}/\widetilde{\mathbf{q}}] \subseteq G \quad \Gamma; G \cup (G'_X[\widetilde{\mathbf{p}}/\widetilde{\mathbf{q}}]) \vdash C \triangleright G'}{\Gamma; G \vdash X(\widetilde{\mathbf{p}}); C \triangleright G'} \text{ [T|Call]}
\end{array}$$

**Fig. 4.** Procedural Choreographies, Typing Rules.

Verifying that communications respect the expected types is straightforward, using the connection graph  $G$  to track which processes have been introduced to each other. In rule [T|Start], we implicitly use the fact that  $\mathbf{q}$  does not appear yet in  $G$ , which is another consequence of using the Barendregt convention. The final graph  $G'$  is only used in procedure calls (rule [T|Call]). Other rules leave it unchanged.

To type a procedural choreography, we need to type its set of procedure definitions  $\mathcal{D}$ . We write  $\Gamma \vdash \mathcal{D}$  if: for each  $X(\widetilde{\mathbf{q}}^T) = C_X \in \mathcal{D}$ , there is exactly one typing  $X(\widetilde{\mathbf{q}}^T) : G_X \triangleright G'_X \in \Gamma$ , and this typing is such that  $\Gamma, \mathbf{q} : T, G_X \vdash C_X \triangleright G'_X$ . We say that  $\Gamma \vdash \langle \mathcal{D}, C \rangle$  if  $\Gamma, \Gamma_{\mathcal{D}}; G_C \vdash C, G'$  for some  $\Gamma_{\mathcal{D}}$  such that  $\Gamma_{\mathcal{D}} \vdash \mathcal{D}$  and some  $G'$ , where  $G_C$  is the full graph whose nodes are the free process names in  $C$ . The choice of  $G_C$  is motivated by observing that (i) all top-level processes should know each other and (ii) eventual connections between processes not occurring in  $C$  do not affect its typability.

Well-typed choreographies either terminate or diverge.<sup>3</sup>

**Theorem 1 (Deadlock freedom and Subject reduction).** *Given a choreography  $C$  and a set  $\mathcal{D}$  of procedure definitions, if  $\Gamma \vdash \mathcal{D}$  and  $\Gamma; G_1 \vdash C \triangleright G'_1$  for some  $\Gamma, G_1$  and  $G'_1$ , then either:*

- $C \preceq_{\mathcal{D}} \mathbf{0}$ ; or,
- for every  $\sigma$ , there exist  $G_2, C'$  and  $\sigma'$  such that  $G_1, C, \sigma \rightarrow_{\mathcal{D}} G_2, C', \sigma'$  and  $\Gamma'; G_2 \vdash C' \triangleright G'_2$  for some  $\Gamma' \supseteq \Gamma$  and  $G'_2$ .

(Proofs of theorems can be found in the Appendix.)

<sup>3</sup> Since we are interested in communications, we assume evaluation of functions and expressions to terminate on values with the right types (see § 7, Faults).

Checking that  $\Gamma \vdash \langle \mathcal{D}, C \rangle$  is not trivial, as it requires “guessing”  $\Gamma_{\mathcal{D}}$ . However, this set can be computed from  $\langle \mathcal{D}, C \rangle$ , entailing type inference properties for PC.

**Theorem 2.** *Given  $\Gamma$ ,  $\mathcal{D}$  and  $C$ ,  $\Gamma \vdash \langle \mathcal{D}, C \rangle$  is decidable.*

Theorem 2 may seem a bit surprising; the key idea of its proof is that type-checking may require expanding recursive definitions, but their parameters only need to be instantiated with process names from a finite set. With a similar idea, we also obtain type inference.

**Theorem 3.** *There is an algorithm that, given any  $\langle \mathcal{D}, C \rangle$ , outputs:*

- a set  $\Gamma$  such that  $\Gamma \vdash \langle \mathcal{D}, C \rangle$ , if such a  $\Gamma$  exists;
- $\text{NO}$ , if no such  $\Gamma$  exists.

**Theorem 4.** *The types of arguments in procedure definitions and the types of freshly created processes can be inferred automatically.*

*Remark 3 (Inferring introductions).* Theorems 3 and 4 allow us to omit type annotations in choreographies, if the types of functions and expressions at processes are known (from  $\vdash_{\mathbb{T}}$ ). Thus, programmers can write choreographies as in our examples.

The same reasoning could be adopted to infer missing introductions ( $\mathfrak{p} : \mathfrak{q} \leftarrow \rightarrow \mathfrak{r}$ ) in a choreography automatically, thus lifting the programmer also from having to think about connections entirely. However, while the types inferred for a choreography do not affect its behaviour, the placement of introductions does. In particular, when invoking procedures one is faced with the choice of adding the necessary introductions inside the procedure definition (weakening the conditions for its invocation) or in the code calling it (making the procedure body more efficient).

## 4 Synthesising Process Implementations

We now present our EndPoint Projection (EPP), which compiles a choreography to a concurrent implementation represented in terms of a process calculus.

### 4.1 Procedural Processes (PP)

We introduce our target process model, Procedural Processes (PP).

**Syntax.** The syntax of PP is given in Figure 5. A term  $\mathfrak{p} \triangleright_v B$  is a process, where  $\mathfrak{p}$  is its name,  $v$  is its value, and  $B$  is its behaviour. Networks, ranged over by  $N, M$ , are parallel compositions of processes, where  $\mathbf{0}$  is the inactive network. Finally,  $\langle \mathcal{B}, N \rangle$  is a procedural network, where  $\mathcal{B}$  defines the procedures that the processes in  $N$  may invoke. Values, expressions and functions are as in PC.

A process executing a send term  $\mathfrak{q}!e; B$  sends the evaluation of expression  $e$  to  $\mathfrak{q}$ , and proceeds as  $B$ . Term  $\mathfrak{p}?f; B$  is the dual receiving action: the process

$$\begin{aligned}
B ::= & \mathbf{q}!e; B \mid \mathbf{p}?f; B \mid \mathbf{q}!!r; B \mid \mathbf{p}?r; B \mid \mathbf{q} \oplus l; B \mid \mathbf{p}\&\{l_i : B_i\}_{i \in I}; B \\
& \mid \mathbf{0} \mid \mathbf{start} \mathbf{q}^T \triangleright B_2; B_1 \mid \mathbf{if} e \mathbf{then} B_1 \mathbf{else} B_2; B \mid X(\tilde{\mathbf{p}}); B \mid \mathbf{0}; B \\
\mathcal{B} ::= & X(\tilde{\mathbf{q}}) = B, \mathcal{B} \mid \emptyset \qquad N, M ::= \mathbf{p} \triangleright_v B \mid (N \mid M) \mid \mathbf{0}
\end{aligned}$$

**Fig. 5.** Procedural Processes, Syntax.

$$\begin{aligned}
& \frac{u = (f[w/*])(e[v/*])}{\mathbf{p} \triangleright_v \mathbf{q}!e; B_1 \mid \mathbf{q} \triangleright_w \mathbf{p}?f; B_2 \rightarrow_{\mathcal{B}} \mathbf{p} \triangleright_v B_1 \mid \mathbf{q} \triangleright_u B_2} \text{ [P|Com]} \\
& \frac{j \in I}{\mathbf{p} \triangleright_v \mathbf{q} \oplus l_j; B \mid \mathbf{q} \triangleright_w \mathbf{p}\&\{l_i : B_i\}_{i \in I} \rightarrow_{\mathcal{B}} \mathbf{p} \triangleright_v B \mid \mathbf{q} \triangleright_w B_j} \text{ [P|Sel]} \\
& \frac{i = 1 \text{ if } e[v/*] = \mathbf{true}, \quad i = 2 \text{ otherwise}}{\mathbf{p} \triangleright_v \mathbf{if} e \mathbf{then} B_1 \mathbf{else} B_2 \rightarrow_{\mathcal{B}} \mathbf{p} \triangleright_v B_i} \text{ [P|Cond]} \\
& \frac{\mathbf{q}' \text{ fresh}}{\mathbf{p} \triangleright_v (\mathbf{start} \mathbf{q}^T \triangleright B_2; B_1) \rightarrow_{\mathcal{B}} \mathbf{p} \triangleright_v B_1[\mathbf{q}'/\mathbf{q}] \mid \mathbf{q}' \triangleright_{\perp T} B_2} \text{ [P|Start]} \\
& \frac{}{\mathbf{p} \triangleright_v \mathbf{q}!!r; B_1 \mid \mathbf{q} \triangleright_w \mathbf{p}?r; B_2 \mid r \triangleright_u \mathbf{p}?q; B_3 \rightarrow_{\mathcal{B}} \mathbf{p} \triangleright_v B_1 \mid \mathbf{q} \triangleright_w B_2 \mid r \triangleright_u B_3} \text{ [P|Tell]} \\
& \frac{N \rightarrow_{\mathcal{B}} N'}{N \mid M \rightarrow_{\mathcal{B}} N' \mid M} \text{ [P|Par]} \qquad \frac{N \preceq_{\mathcal{B}} M \quad M \rightarrow_{\mathcal{B}} M' \quad M' \preceq_{\mathcal{B}} N'}{N \rightarrow_{\mathcal{B}} N'} \text{ [P|Struct]}
\end{aligned}$$

**Fig. 6.** Procedural Processes, Semantics.

executing it receives a value from  $\mathbf{p}$ , combines it with its value as specified by  $f$ , and then proceeds as  $B$ . Term  $\mathbf{q}!!r$  sends process name  $r$  to  $\mathbf{q}$  and process name  $\mathbf{q}$  to  $r$ , making  $\mathbf{q}$  and  $r$  “aware” of each other. The dual action is  $\mathbf{p}?r$ , which receives a process name from  $\mathbf{p}$  that replaces the bound variable  $r$  in the continuation. Term  $\mathbf{q} \oplus l; B$  sends the selection of a label  $l$  to process  $\mathbf{q}$ . Selections are received by the branching term  $\mathbf{p}\&\{l_i : B_i\}_{i \in I}$ , which can receive a selection for any of the labels  $l_i$  and proceed according to  $B_i$ . Branching terms must offer at least one branch. Term  $\mathbf{start} \mathbf{q} \triangleright B_2; B_1$  starts a new process (with a fresh name) executing  $B_2$ , and proceeds in parallel as  $B_1$ . Conditionals, procedure calls, and termination are standard. Term  $\mathbf{start} \mathbf{q} \triangleright B_2; B_1$  binds  $\mathbf{q}$  in  $B_1$ , and  $\mathbf{p}?r; B$  binds  $r$  in  $B$ .

**Semantics.** The rules defining the reduction relation  $\rightarrow_{\mathcal{B}}$  for PP are shown in Figure 6. As in PC, they are parameterised on the set of behavioural procedures  $\mathcal{B}$ . Rule [P|Com] models value communication: a process  $\mathbf{p}$  executing a send action towards a process  $\mathbf{q}$  can synchronise with a receive-from- $\mathbf{p}$  action at  $\mathbf{q}$ ; in the reductum,  $f$  is used to update the memory of  $\mathbf{q}$  by combining its contents with the value sent by  $\mathbf{p}$ . The placeholder  $*$  is replaced with the current value of  $\mathbf{p}$  in  $e$  (resp.  $\mathbf{q}$  in  $f$ ). Rule [P|Tell] establishes a three-way synchronisation,

$$\begin{array}{c}
\frac{}{p \triangleright_v \mathbf{0} \preceq_{\mathcal{B}} \mathbf{0}} \text{ [P|AZero]} \quad \frac{}{N | \mathbf{0} \preceq_{\mathcal{B}} N} \text{ [P|NZero]} \\
\frac{}{\mathbf{0}; B \preceq_{\mathcal{B}} B} \text{ [P|End]} \quad \frac{X(\widetilde{q^T}) = B_X \in \mathcal{B}}{X\langle \tilde{p} \rangle; B \preceq_{\mathcal{B}} B_X[\tilde{p}/\tilde{q}]; B} \text{ [P|Unfold]}
\end{array}$$

**Fig. 7.** Procedural Processes, Structural precongurence  $\preceq_{\mathcal{B}}$ .

allowing a process to introduce two others. Since the received names are bound at the receivers, we rely on  $\alpha$ -conversion to make the receivers agree on each other's name, as done in session types [17]. (Differently from PC, we do not assume the Barendregt convention here, in line with the tradition of process calculi.) Rule [P|Sel] is standard selection [17], where the sender process selects one of the branches offered by the receiver. In rule [P|Start], we require the name of the created process to be globally fresh. All other rules are standard. Rule [P|Struct] uses structural precongurence  $\preceq_{\mathcal{B}}$ , which is the smallest precongurence satisfying associativity and commutativity of parallel ( $|$ ) and the rules in Figure 7. Rule [P|Unfold] expands procedure calls. It uses again the  $\S$  operator, defined as for PC but with terms in the PP language.

*Remark 4.* Our three-way synchronisation in rule [P|Tell] could be easily encoded with two standard two-way communications of names, as in the  $\pi$ -calculus [27] (see also Remark 1). Our choice gives a clearer formulation of EPP.

*Example 6.* We show a process implementation of the merge sort choreography in Example 1 from § 1. All processes are annotated with type **List**( $T$ ) (omitted);  $\text{id}$  is the identity function (Example 2).

```

MSp(p) = if is_small then 0
         else start q1 ▷ (p?id; MSp<q1>; p!*);
              start q2 ▷ (p?id; MSp<q2>; p!*);
              q1!split1; q2!split2; q1?id; q2?merge

```

In the next section, we show that our EPP generates this process implementation automatically from the choreography in Example 1.

## 4.2 EndPoint Projection (EPP)

We now show how to compile programs in PC to processes in PP.

**Behaviour Projection.** We start by defining how to project the behaviour of a single process  $p$ , a partial function denoted  $\llbracket C \rrbracket_p$ . The rules defining behaviour projection are given in Figure 8. Each choreography term is projected to the local action of the process that we are projecting. For example, a communication term  $p.e \rightarrow q.f$  projects a send action for the sender  $p$ , a receive action for the receiver  $q$ , or skips to the continuation otherwise. The rules for projecting a selection or an introduction (name mobility) are similar.

$$\begin{aligned}
\llbracket p.e \rightarrow q.f; C \rrbracket_r &= \begin{cases} q!e; \llbracket C \rrbracket_r & \text{if } r = p \\ p?f; \llbracket C \rrbracket_r & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} & \llbracket p \rightarrow q[l]; C \rrbracket_r &= \begin{cases} q \oplus l; \llbracket C \rrbracket_r & \text{if } r = p \\ p\&\{l : \llbracket C \rrbracket_r\} & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket p : q \leftrightarrow r; C \rrbracket_s &= \begin{cases} q!!r; \llbracket C \rrbracket_s & \text{if } s = p \\ p?r; \llbracket C \rrbracket_s & \text{if } s = q \\ p?q; \llbracket C \rrbracket_s & \text{if } s = r \\ \llbracket C \rrbracket_s & \text{otherwise} \end{cases} & \llbracket X(\bar{p}); C \rrbracket_r &= \begin{cases} X_i(\bar{p}); \llbracket C \rrbracket_r & \text{if } r = p_i \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
& & \llbracket \mathbf{0} \rrbracket_r &= \mathbf{0} & \llbracket \mathbf{0}; C \rrbracket_r &= \llbracket C \rrbracket_r \\
\llbracket \text{if } p.e \text{ then } C_1 \text{ else } C_2; C \rrbracket_r &= \begin{cases} \text{if } e \text{ then } \llbracket C_1 \rrbracket_r \text{ else } \llbracket C_2 \rrbracket_r; \llbracket C \rrbracket_r & \text{if } r = p \\ (\llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r); \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket p \text{ start } q^T; C \rrbracket_r &= \begin{cases} \text{start } q \triangleright \llbracket C \rrbracket_q; \llbracket C \rrbracket_r & \text{if } r = p \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases}
\end{aligned}$$

**Fig. 8.** Procedural Choreographies, Behaviour Projection.

The rule for projecting a conditional uses the partial merging operator  $\sqcup$ :  $B \sqcup B'$  is isomorphic to  $B$  and  $B'$  up to branching, where the branches of  $B$  or  $B'$  with distinct labels are also included. The interesting rule defining merge is:

$$\begin{aligned}
(p\&\{l_i : B_i\}_{i \in J}; B) \sqcup (p\&\{l_i : B'_i\}_{i \in K}; B') = \\
p\&\{l_i : (B_i \sqcup B'_i)\}_{i \in J \cap K} \cup \{l_i : B_i\}_{i \in J \setminus K} \cup \{l_i : B'_i\}_{i \in K \setminus J}; (B \sqcup B')
\end{aligned}$$

The idea of merging comes from [5]. Here, we extend it to general recursion, parametric procedures, and process starts. The complete definition of merging is given in the Appendix. Merging allows the process that decides a conditional to inform other processes of its choice later on, using selections. It is found repeatedly in most choreography models [5,9,19].

Building on behaviour projection, we define how to project the set  $\mathcal{D}$  of procedure definitions. We need to consider two main aspects. The first is that, at runtime, the choreography may invoke a procedure  $X$  multiple times, but potentially passing a process  $r$  at different argument positions each time. This means that  $r$  may be called to play different “roles” in the implementation of the procedure. For this reason, we project the behaviour of each possible process parameter  $p$  as the local procedure  $X_p$ . The second aspect is: depending on the role that  $r$  is called to play by the choreography, it needs to know the names of the other processes that it is supposed to communicate with in the choreographic procedure. We deal with this by simply passing all arguments, which means that some of them may even be unknown to the process invoking the procedure. This is not a problem: we focus on typable choreographies, and typing ensures that those parameters are not actually used in the projected procedure (so they act as “dummies”). We do this for clarity, since it yields a simpler formulation of EPP. In practice, we can annotate the EPP by analysing which parameters of each



recursive definition are actually used in each of its projections, and instantiating only those (see Appendix). We can now define

$$\llbracket \mathcal{D} \rrbracket = \bigcup \left\{ \llbracket X(\tilde{\mathbf{q}}^T) = C \rrbracket \mid X(\tilde{\mathbf{q}}^T) = C \in \mathcal{D} \right\}$$

where, for  $\tilde{\mathbf{q}}^T = \mathbf{q}_1^{T_1}, \dots, \mathbf{q}_n^{T_n}$ ,

$$\llbracket X(\tilde{\mathbf{q}}^T) = C \rrbracket = \{X_1(\tilde{\mathbf{q}}) = \llbracket C \rrbracket_{\mathbf{q}_1}, \dots, X_n(\tilde{\mathbf{q}}) = \llbracket C \rrbracket_{\mathbf{q}_n}\}.$$

**Definition 1 (EPP).** *Given a procedural choreography  $\langle \mathcal{D}, C \rangle$  and a state  $\sigma$ , the endpoint projection  $\llbracket \mathcal{D}, C, \sigma \rrbracket$  is the parallel composition of the processes in  $C$  with all definitions from  $\mathcal{D}$ :*

$$\llbracket \mathcal{D}, C, \sigma \rrbracket = \langle \llbracket \mathcal{D} \rrbracket, \llbracket C, \sigma \rrbracket \rangle = \left\langle \llbracket \mathcal{D} \rrbracket, \prod_{\mathbf{p} \in \text{pn}(C)} \mathbf{p} \triangleright_{\sigma(\mathbf{p})} \llbracket C \rrbracket_{\mathbf{p}} \right\rangle$$

where  $\llbracket C, \sigma \rrbracket$ , the EPP of  $C$  wrt state  $\sigma$ , is independent of  $\mathcal{D}$ .

Since the  $\sigma$ s are total, if  $\llbracket C, \sigma \rrbracket$  is defined for some  $\sigma$ , then  $\llbracket C, \sigma' \rrbracket$  is defined also for all other  $\sigma'$ . When  $\llbracket C, \sigma \rrbracket = N$  is defined for any  $\sigma$ , we say that  $C$  is *projectable* and that  $N$  is the projection of  $C, \sigma$ . Similar considerations apply to  $\llbracket \mathcal{D}, C, \sigma \rrbracket$ .

*Example 7.* The EPP of the choreography in Example 1 is given in Example 6.

*Example 8.* We give a more sophisticated example involving merging and introductions: the projection of procedure `par_download` (Example 5) for process  $\mathbf{s}$ . (We omit the type annotations.)

```

par_downloads(c, s) = c & {
  more: start s' ▷ (s?c; c?id; c?c'; downloads<c', s'>);
        c!!s'; par_downloads<c, s>
  end: 0
}

```

Observe that we call procedure `downloads`, as  $\mathbf{s}'$  occurs in the position of that procedure's formal argument  $\mathbf{s}$ .

**Properties.** EPP guarantees correctness by construction: the code synthesised from a choreography follows it precisely.

**Theorem 5 (EPP Theorem).** *If  $\langle \mathcal{D}, C \rangle$  is projectable,  $\Gamma \vdash \mathcal{D}$ , and  $\Gamma; G \vdash C \triangleright G^*$ , then, for all  $\sigma$ :*

- (Completeness) if  $G, C, \sigma \rightarrow_{\mathcal{D}} G', C', \sigma'$ , then  $\llbracket C, \sigma \rrbracket \rightarrow_{\llbracket \mathcal{D} \rrbracket} \succ \llbracket C', \sigma' \rrbracket$ ;
- (Soundness) if  $\llbracket C, \sigma \rrbracket \rightarrow_{\llbracket \mathcal{D} \rrbracket} N$ , then  $G, C, \sigma \rightarrow_{\mathcal{D}} G', C', \sigma'$  for some  $G', \sigma'$  such that  $\llbracket C', \sigma' \rrbracket \prec N$ .

Above, the *pruning relation*  $\prec$  from [5] eliminates the branches introduced by the merging operator  $\sqcup$  when they are not needed anymore to follow the originating choreography (we write  $N \succ N'$  when  $N' \prec N$ ). Pruning does not alter reductions, since the eliminated branches are never selected [5]. Combining Theorem 5 with Theorem 1 we get that the projections of typable PC terms never deadlock.

**Corollary 1 (Deadlock-freedom by construction).** *Let  $N = \llbracket C, \sigma \rrbracket$  for some  $C$  and  $\sigma$ , and assume that  $\Gamma; G \vdash C \triangleright G'$  for some  $\Gamma$  such that  $\Gamma \vdash \mathcal{D}$  and some  $G$  and  $G'$ . Then, either:*

- $N \preceq_{\llbracket \mathcal{D} \rrbracket} \mathbf{0}$  ( $N$  has terminated);
- or there exists  $N'$  such that  $N \rightarrow_{\llbracket \mathcal{D} \rrbracket} N'$  ( $N$  can reduce).

*Remark 5 (Amendment).* A choreography in PC can only be unprojectable because of unmergeable subterms. Thus, every choreography can be made projectable by only adding label selections. This can be formalized in an amendment algorithm, as in other choreography languages [20,11], reported in the Appendix. For example, the first (unprojectable) choreography in Remark 2 can be amended to the projectable choreography presented at the end of the same remark.

In practice, the same argument as for inferring introduction terms (Remark 3) applies. Although amendment allows us to write choreographies without worrying about label selections, it is useful to give the programmer the option to place them where most convenient. For example, consider a process  $p$  making an internal choice that affects processes  $q$  and  $r$ . If one of these two processes has to perform a slower computation in response to that choice, then it makes sense for  $p$  to send a label selection to it first, and only afterwards to notify the other process.

## 5 Asynchrony

We define an alternative semantics to PC and PP, whereby communication becomes asynchronous. We show that the results stated in the previous sections also hold for the asynchronous case.

### 5.1 Asynchronous PC (aPC)

**Syntax.** In asynchronous PC (aPC), communications are not atomically executed anymore, but consist of multiple actions. In particular, the sender of a message can proceed in its execution without waiting for the receiver to receive such message.

The intuition behind aPC is that the language is extended with new runtime terms that capture this refinement of execution steps.<sup>4</sup> At runtime communications are expanded into multiple actions. For example, a communication  $p.e \rightarrow q.f$  expands in  $p.e \xrightarrow{x} \bullet_q$  – a send action from  $p$  – and  $\bullet_p \xrightarrow{x} q.f$  – a receive action by  $q$ . The process subscripts at  $\bullet$  are immaterial for the semantics of aPC, but are useful for the definition of EPP. The tag  $x$  specifies that, in the original choreography, that message from  $p$  should reach that receive action at  $q$ . Executing  $p.e \xrightarrow{x} \bullet_q$  replaces  $x$  in the corresponding receive action with the

<sup>4</sup> Recall that runtime terms are assumed never to be used by programmers (like term  $\mathbf{0}; C$  in PC). They are used only to represent runtime states.

$$\begin{aligned} \eta ::= & \dots \mid p.e \xrightarrow{x} \bullet_q \mid \bullet_p \xrightarrow{\hat{v}} q.f \mid p \xrightarrow{x} \bullet_q[l] \mid \bullet_p \xrightarrow{\hat{l}} q[l] \mid p : \bullet_q \xrightarrow{x,y} \bullet_r \mid \bullet_p.r \xrightarrow{\hat{p}} q \\ \hat{v} ::= & x \mid v \quad \hat{l} ::= x \mid l \quad \hat{p} ::= x \mid p \end{aligned}$$

**Fig. 9.** Asynchronous PC, Syntax of New Runtime Terms.

actual value  $v$  computed from  $e$  at  $p$ , yielding  $\bullet_p \xrightarrow{v} q.f$ ; executing  $\bullet_p \xrightarrow{v} q.f$  updates the state of  $q$ .

We now define aPC formally. The new terms are given in Figure 9, and they are all runtime terms. The terms follow the intuition given for value communications, extended to selection and introduction. Label selection expands in  $p \xrightarrow{x} \bullet_q[l]$  and  $\bullet_p \xrightarrow{\hat{l}} q[l]$ , and introduction expands in  $p : \bullet_q \xrightarrow{x,y} \bullet_r$  and two  $\bullet_p.r \xrightarrow{\hat{p}} q$  actions. The tags in these actions are not essential for propagating values as above, but they make the treatment of all new actions similar. Process names for the new runtime terms are defined as follows, ignoring process names in tags and in  $\bullet$  subscripts.

$$\begin{aligned} \text{pn}(p.e \xrightarrow{x} \bullet_q) &= \text{pn}(p \xrightarrow{x} \bullet_q[l]) = \text{pn}(p : \bullet_q \xrightarrow{x,y} \bullet_r) = \{p\} \\ \text{pn}(\bullet_p \xrightarrow{\hat{v}} q.f) &= \text{pn}(\bullet_p \xrightarrow{\hat{l}} q[l]) = \text{pn}(\bullet_p.r \xrightarrow{\hat{p}} q) = \{q\} \end{aligned}$$

**Semantics.** The interplay between asynchronous communications and name mobility requires the connection graph to be directed in aPC, since the processes  $q$  and  $r$  in an introduction term  $p : q \leftrightarrow r$  may receive names at different times now. An edge from  $p$  to  $q$  in  $G$ , denoted  $q \xrightarrow{G} p$ , now means that  $p$  knows  $q$ 's name – so  $p$  is able to send messages to  $q$  or to listen for incoming messages from  $q$ .

The semantics for aPC includes:

- the rules from Figure 10, which are the asynchronous counterpart to  $[C|Com]$  in PC, and similar rules for selection and introduction;
- rules  $[C|Cond]$  and  $[C|Struct]$  from PC;
- rule  $[C|Start]$  from PC, where  $G \cup \{p \leftrightarrow q\}$  is the graph obtained by adding the two directed edges between  $p$  and  $q$  to  $G$ ;
- the rules defining  $\preceq_{\emptyset}$  (Figure 3), with  $\eta$  now ranging over all communication actions in aPC;
- the new rules for  $\preceq_{\emptyset}$  in Figure 11.

The missing new rules are given in the Appendix. By the Barendregt convention, tags introduced by the rules in Figure 11 are globally fresh. Thus, they maintain the correspondence between the value being sent and that being received.

The key to the semantics of aPC lies in the new swaps allowed by  $\preceq_{\emptyset}$ , due to the definition of  $\text{pn}$  for the new terms. This captures the concurrency that arises from asynchronous communications.

$$\frac{\frac{p \xrightarrow{G} q \quad e[\sigma(p)/*] \downarrow v}{G, p.e \xrightarrow{x} \bullet_q; C, \sigma \rightarrow_{\mathcal{D}} G, C[v/x], \sigma} \quad [C|Com-S]}{\frac{q \xrightarrow{G} p \quad f[\sigma(q)/*](v) \downarrow w}{G, \bullet_p \xrightarrow{v} q.f; C, \sigma \rightarrow_{\mathcal{D}} G, C, \sigma[q \mapsto w]} \quad [C|Com-R]}$$

**Fig. 10.** Asynchronous PC, Semantics of New Runtime Terms.

$$\frac{}{p.e \rightarrow q.f \preceq_{\mathcal{D}} p.e \xrightarrow{x} \bullet_q; \bullet_p \xrightarrow{x} q.f} \quad [C|Com-U]$$

$$\frac{}{p \rightarrow q[l] \preceq_{\mathcal{D}} p \xrightarrow{x} \bullet_q[l]; \bullet_p \xrightarrow{x} q[l]} \quad [C|Sel-U]$$

$$\frac{}{p: q \langle \rightarrow r \preceq_{\mathcal{D}} p: \bullet_q \xrightarrow{x,y} \bullet_r; \bullet_p.r \xrightarrow{x} q; \bullet_p.q \xrightarrow{y} r} \quad [C|Tell-U]$$

**Fig. 11.** Asynchronous PC, Structural Precongruence (New Additional Rules).

*Example 9.* Let  $C$  be  $p.e \rightarrow q.f; p.e' \rightarrow r.f'$ . To execute  $C$  in aPC, we expand it:

$$C \preceq_{\mathcal{D}} p.e \xrightarrow{x} \bullet_q; \bullet_p \xrightarrow{x} q.f'; p.e' \xrightarrow{y} \bullet_r; \bullet_p \xrightarrow{y} r.f'$$

Using  $\preceq_{\mathcal{D}}$ , we can swap the second term to the end (it shares no process names with the subsequent terms according to pn):

$$C \preceq_{\mathcal{D}} p.e \xrightarrow{x} \bullet_q; p.e' \xrightarrow{y} \bullet_r; \bullet_p \xrightarrow{y} r.f'; \bullet_p \xrightarrow{x} q.f'$$

Now  $p$  can send both messages immediately, and  $r$  can receive its message before  $q$ .

*Example 10.* Due to asynchrony, a process  $q$  can now send a message to another process  $r$  that does not yet know about  $q$ . However,  $r$  is still unable to receive it before learning  $q$ 's name, as expected [27]. Assume that  $p \xleftarrow{G} q$  and  $p \xleftarrow{G} r$  and consider the choreography:  $C \triangleq p: q \langle \rightarrow r; q.* \rightarrow r.f$ . We can unfold  $C$  to

$$C \preceq_{\mathcal{D}} p: \bullet_q \xrightarrow{x,y} \bullet_r; \bullet_p.r \xrightarrow{x} q; \bullet_p.q \xrightarrow{y} r; q.* \xrightarrow{z} \bullet_r; \bullet_q \xrightarrow{x} r.f$$

and, by  $\preceq_{\mathcal{D}}$ , swap the third and fourth term

$$C \preceq_{\mathcal{D}} p: \bullet_q \xrightarrow{x,y} \bullet_r; \bullet_p.r \xrightarrow{x} q; q.* \xrightarrow{z} \bullet_r; \bullet_p.q \xrightarrow{y} r; \bullet_q \xrightarrow{x} r.f$$

which corresponds to an execution path where  $q$  sends its contents to  $r$  before  $r$  is notified of  $q$ 's name. The last two actions can not be swapped, so  $r$  cannot access  $q$ 's message before receiving its name.

$$\begin{array}{c}
\frac{\mathfrak{p} \xrightarrow{G} \mathfrak{q} \quad \Gamma \vdash \mathfrak{p} : T_{\mathfrak{p}} \quad * : T_{\mathfrak{p}} \vdash_{\mathbb{T}} e : T \quad \Gamma \oplus (x : T); G \vdash C \triangleright G'}{\Gamma; G \vdash \mathfrak{p}.e \xrightarrow{x} \bullet_{\mathfrak{q}}; C \triangleright G'} \text{ [T|Com-S]} \\
\frac{\mathfrak{q} \xrightarrow{G} \mathfrak{p} \quad \Gamma \vdash \mathfrak{q} : T_{\mathfrak{q}} \quad * : T_{\mathfrak{q}} \vdash_{\mathbb{T}} f : T \rightarrow T_{\mathfrak{q}} \quad \Gamma \oplus (x : T); G \vdash C \triangleright G'}{\Gamma; G \vdash \bullet_{\mathfrak{p}} \xrightarrow{x} \mathfrak{q}.f; C \triangleright G'} \text{ [T|Com-RV]} \\
\frac{\mathfrak{q} \xrightarrow{G} \mathfrak{p} \quad \Gamma \vdash \mathfrak{q} : T_{\mathfrak{q}} \quad \vdash_{\mathbb{T}} v : T \quad * : T_{\mathfrak{q}} \vdash_{\mathbb{T}} f : T \rightarrow T_{\mathfrak{q}} \quad \Gamma; G \vdash C \triangleright G'}{\Gamma; G \vdash \bullet_{\mathfrak{p}} \xrightarrow{v} \mathfrak{q}.f; C \triangleright G'} \text{ [T|Com-RT]}
\end{array}$$

**Fig. 12.** Asynchronous PC, Typing Rules (New Runtime Terms).

To restate Theorems 1–4 for aPC, we extend our type system to the new runtime terms. In particular, we augment contexts  $\Gamma$  to contain also type declarations for tags ( $x : T$ ) and assignments of labels or process identifiers to tags ( $x = l$  or  $x = \mathfrak{p}$ ). The type system of aPC contains the rules for PC (Figure 4) together with the new rules given in Figure 12 and analogous ones for selections and introductions. Furthermore, in rule [T|End] (from PC) we require that  $\Gamma$  does not contain any assertions  $x : T$  or  $x = t$ , meaning that there should be no dangling actions.

The context  $\Gamma \oplus (x : T)$  used in the rules is defined below; the interpretation of  $\Gamma \oplus (x = t)$  is similar. The operator  $\oplus$  is partial. In particular, any judgement involving  $\Gamma \oplus (x : T)$  is false if  $\Gamma$  already declares  $x$  with a different type.

$$\Gamma \oplus (x : T) = \begin{cases} \Gamma, x : T & \text{if } x \text{ is not declared in } \Gamma \\ \Gamma \setminus \{x : T\} & \text{if } x : T \in \Gamma \\ \text{undefined} & \text{otherwise} \end{cases}$$

Thus, typing a receiving action depends on whether or not it corresponds to a message that has already been sent according to  $\Gamma$ .

If we consider only program choreographies (without runtime terms), then the typing systems for aPC and PC coincide (as the connection graph is always symmetric); in particular, Theorems 2–4 automatically hold for program choreographies in aPC. The proofs of these results, as well as that of Theorem 1, can be readily adapted to aPC, as the new cases do not pose any difficulty.

PC and aPC enjoy the following operational correspondence.

**Theorem 6.** *For any  $C$  in PC, connection graph  $G$ , and state  $\sigma$ :*

- *If  $G, C, \sigma \rightarrow_{\mathcal{O}} G', C', \sigma'$  (in PC), then  $G, C, \sigma \rightarrow_{\mathcal{O}}^* G', C', \sigma'$  (in aPC).*
- *If  $G, C, \sigma \rightarrow_{\mathcal{O}} G', C', \sigma'$  (in aPC), then there exist  $C'', G''$  and  $\sigma''$  such that  $G, C, \sigma \rightarrow_{\mathcal{O}} G'', C'', \sigma''$  (in PC) and  $G', C', \sigma' \rightarrow_{\mathcal{O}}^* G'', C'', \sigma''$  (in aPC).*

In particular, the sequence of messages from a process  $\mathfrak{p}$  to another process  $\mathfrak{q}$  is received by  $\mathfrak{q}$  in exactly the same order as it is sent by  $\mathfrak{p}$ . Thus, our semantics

$$\begin{array}{c}
\frac{\rho'_q = \rho_q \cdot \langle p, e[v/*] \rangle}{p \triangleright_v^{\rho_p} q!e; B_p \mid q \triangleright_w^{\rho_q} B_q \rightarrow_{\mathcal{B}} p \triangleright_v^{\rho_p} B_p \mid q \triangleright_w^{\rho'_q} B_q} \text{ [P|Com-S]} \\
\frac{\rho_q \preceq \langle p, v \rangle \cdot \rho'_q \quad u = (f[w/*])(v)}{q \triangleright_w^{\rho_q} p?f; B \rightarrow_{\mathcal{B}} q \triangleright_u^{\rho'_q} B} \text{ [P|Com-R]}
\end{array}$$

**Fig. 13.** Asynchronous Procedural Processes, Semantics (New Rules).

matches an interpretation of asynchrony supported by queues between processes, as we make evident in the asynchronous version of PP below.

## 5.2 Asynchronous PP (aPP)

**Syntax.** The asynchronous variant of PP, aPP, is easier to define, as the underlying language is almost unchanged. The only syntactic difference is that processes now have the form  $p \triangleright_v^\rho B$ , where  $\rho$  is a queue of incoming messages. A message is a pair  $\langle q, m \rangle$ , where  $q$  is the sender process and  $m$  is a value, label, or process identifier.

**Semantics.** The semantics of aPP consists of rules  $[P|Cond]$ ,  $[P|Par]$ ,  $[P|Cond]$ , and  $[P|Start]$  from PP (Figure 6) together with the rules given in Figure 13 and analogous variants for selection and introduction (see Appendix). In the reduction of  $[P|Start]$ , the newly created process is initialized with an empty queue. Structural precongruence for aPP is defined exactly as for PP. We write  $\rho \cdot \langle q, m \rangle$  to denote the queue obtained by appending message  $\langle q, m \rangle$  to  $\rho$ , and  $\langle q, m \rangle \cdot \rho$  for the queue with  $\langle q, m \rangle$  at the head and  $\rho$  as tail. We simulate having one separate queue for each other process by allowing incoming messages from different senders to be exchanged, which we represent using the congruence  $\rho \preceq \rho'$  defined by the rule  $\langle p, m \rangle \cdot \langle q, m' \rangle \preceq \langle q, m' \rangle \cdot \langle p, m \rangle$  if  $p \neq q$ .

All behaviours of PP are valid also in aPP.

**Theorem 7.** *Let  $N$  be a PP network. If  $N \rightarrow_{\mathcal{B}} N'$  (in PP), then  $N_{\square} \rightarrow_{\mathcal{B}}^* N'_{\square}$  (in aPP), where  $N_{\square}$  denotes the asynchronous network obtained by adding an empty queue to each process.*

The converse is not true, so the relation between PP and aPP is not so strong as in Theorem 6 for PC and aPC. This is because of deadlocks: in PP, a communication action can only take place when the sender and receiver are ready to synchronize; in aPP, a process can send a message to another process, even though the intended recipient is not yet able to receive it. For example, the network  $p \triangleright_v q! * \mid q \triangleright_w \mathbf{0}$  is deadlocked in PP but not in aPP.

## 5.3 EndPoint Projection (EPP) from aPC to aPP

Defining an EPP from aPC to aPP requires extending the previous definition with clauses for the new runtime terms. The most interesting part is gener-

ating the local queues for each process: when compiling  $\bullet_p.v \rightarrow q.f$ , for example, we need to add  $\langle p, v \rangle$  at the top of  $q$ 's queue. However, if we follow this intuition, Theorem 5 no longer holds for arbitrary aPC choreographies. This happens because we can write choreographies that use runtime terms in a “wrong” way, and these are not correctly compiled to aPP. Consider the choreography  $C = p.1 \rightarrow q.f; \bullet_p \xrightarrow{2} q.f$ . If we naively project it for some  $\sigma$ , we obtain  $p \triangleright_{\sigma(p)}^{\square} q!1 \mid q \triangleright_{\sigma(q)}^{\langle p, 2 \rangle} p?f; p?f$ , where  $\square$  is the empty queue, and  $q$  will receive 2 before 1.

To avoid this undesired behaviour, we recall that runtime terms should not be used by the programmer and restrict ourselves to *well-formed* choreographies: those that can arise from executing a choreography that does not contain runtime terms (i.e., a program).

**Definition 2 (Well-formedness).** *A choreography  $C$  in aPC is well-formed if  $C^R \circ C^{PC} \preceq C$  where:*

- $\preceq$  is structural precongruence without rule  $\lfloor C \rfloor \text{Unfold}$ ;
- $C^R$  only contains unmatched and instantiated receive actions:  $\bullet_p \xrightarrow{v} q.f$ ,
- $\bullet_p \xrightarrow{l} q[l]$ ,  $\bullet_{p.r} \xrightarrow{r} q$ ;
- $C^{PC}$  is a PC choreography.

A procedure  $X(\widetilde{q^T}) = C$  is well-formed if  $C$  does not contain any runtime actions. A procedural choreography  $\langle \mathcal{D}, C \rangle$  is well-formed if all procedures in  $\mathcal{D}$  are well-formed and  $C$  is well-formed.

Well-formedness is decidable, since the set of choreographies equivalent up to  $\preceq$  is decidable. More efficiently, one can check that  $C$  is well-formed by swapping all runtime actions to the beginning and folding all paired send/receive terms. Furthermore, choreography execution preserves well-formedness. The problematic choreography described above is not well-formed.

**Definition 3 (EPP from aPC to aPP).** *Let  $\langle \mathcal{D}, C \rangle$  be in aPC and  $\sigma$  be a state. The EPP of  $\langle \mathcal{D}, C \rangle$  and  $\sigma$  is defined as*

$$\llbracket \mathcal{D}, C, \sigma \rrbracket = \langle \llbracket \mathcal{D} \rrbracket, \llbracket C, \sigma \rrbracket \rangle = \left\langle \llbracket \mathcal{D} \rrbracket, \prod_{p \in \text{pn}(C)} p \triangleright_{\sigma(p)}^{\langle C \rangle_p} \llbracket C \rrbracket_p \right\rangle$$

where  $\llbracket C \rrbracket_p$  is defined as for PC (Figure 8) with the extra rules in Figure 14 and  $\langle C \rangle_p$  is defined by the rules in Figure 15. The similar rules for selection and introduction are defined in the Appendix.

In the last case of the definition of  $\langle C \rangle_p$  (bottom-right),  $\eta$  and  $I$  range over all cases that are not covered previously. The rule for the conditional may seem a bit surprising: in the case of projectable choreographies, mergeability and well-formedness together imply that unmatched receive actions at a process must occur in the same order in both branches. Note that projection of choreographies with ill-formed runtime terms (e.g.,  $\bullet_p \xrightarrow{l'} q[l]$  with  $l \neq l'$ , which cannot occur in a well-formed choreography) is not defined.

With this definition, we can restate Theorem 5 for aPC and aPP.

$$\llbracket \mathbf{p}.e \xrightarrow{x} \bullet_{\mathbf{q}}; C \rrbracket_r = \begin{cases} \mathbf{q}!e; \llbracket C \rrbracket_r & \text{if } r = \mathbf{p} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \quad \llbracket \bullet_{\mathbf{p}} \xrightarrow{\hat{v}} \mathbf{q}.f; C \rrbracket_r = \begin{cases} \mathbf{p}?f; \llbracket C \rrbracket_r & \text{if } r = \mathbf{q} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases}$$

**Fig. 14.** Asynchronous PC, Behaviour Projection (New Rules).

$$\llbracket \bullet_{\mathbf{p}} \xrightarrow{v} \mathbf{q}.f; C \rrbracket_r = \begin{cases} \langle \mathbf{p}, v \rangle \cdot \llbracket C \rrbracket_r & \text{if } r = \mathbf{q} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \quad \llbracket \text{if } \mathbf{p} \leq \mathbf{q} \text{ then } C_1 \text{ else } C_2; C \rrbracket_r = \llbracket C_1 \rrbracket_r \cdot \llbracket C \rrbracket_r \\ \llbracket \eta; C \rrbracket_r = \llbracket I; C \rrbracket_r = \llbracket C \rrbracket_r$$

**Fig. 15.** Asynchronous PC, State Projection.

**Theorem 8 (Asynchronous EPP Theorem).** *If  $\langle \mathcal{D}, C \rangle$  in aPC is projectable and well-formed,  $\Gamma \vdash \mathcal{D}$ , and  $\Gamma; G \vdash C \triangleright G^*$ , then, for all  $\sigma$ :*

- (Completeness) *if  $G, C, \sigma \rightarrow_{\mathcal{D}} G', C', \sigma'$ , then  $\llbracket C, \sigma \rrbracket \rightarrow_{\llbracket \mathcal{D} \rrbracket} \llbracket C', \sigma' \rrbracket$ ;*
- (Soundness) *if  $\llbracket C, \sigma \rrbracket \rightarrow_{\llbracket \mathcal{D} \rrbracket} N$ , then  $G, C, \sigma \rightarrow_{\mathcal{D}} G', C', \sigma'$  for some  $G', \sigma'$  such that  $\llbracket C', \sigma' \rrbracket \prec N$ .*

As a consequence, Corollary 1 applies also to the asynchronous case: the processes projected from aPC into aPP are deadlock-free.

## 6 Language Extensions

We sketch two simple extensions that enhance the expressivity of (synchronous and asynchronous) PC without affecting its underlying theory: they only require simple modifications, and all the results from the previous sections still hold.

**Parameter Lists.** Our first extension allows procedure parameters to be lists of processes, on which procedures can then act uniformly by recursion. Formally, a procedure parameter can now be either a process or a list of processes all with the same type. We restrict the usage of such lists to the arguments of procedure calls; however, they may be manipulated by means of pure total functions that take a list as their only argument. For example, if procedure  $X$  takes a list  $P$  as an argument, then  $X$ 's body may call other procedures on  $P$ ,  $\text{head}(P)$  or  $\text{tail}(P)$ , but it may not include a communication involving  $\text{head}(P)$  and another process.

The semantics is extended with a new garbage collection rule:  $X \langle [] \rangle \preceq_{\mathcal{D}} \mathbf{0}$ , i.e., calling a procedure with an empty list is equivalent to termination ( $\mathbf{0}$ ). In order to type procedures that take process lists as arguments, in a typing  $X : G \triangleright G'$  we allow the vertices of  $G$  and  $G'$  to be not only processes, but also (formal) lists. This requires adjusting the premise of  $\llbracket \text{T|Call} \rrbracket$  to the case when some of the arguments are lists. For when a concrete list is given as argument, we update the premise  $G_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}] \subseteq G$  to check that: if  $G_X$  contains an edge between two process lists  $P$  and  $Q$ , then  $G$  must contain edges between  $\mathbf{p}$  and  $\mathbf{q}$  for each  $\mathbf{p}$



in the list supplied for  $P$  and each  $q$  in the list supplied for  $Q$ . The interpretation of an edge between an argument process and an argument list is similar.

We extend EPP to this enriched language by similarly extending the syntax of PP and adding the rule  $q \triangleright_v X(\tilde{p}) \preceq_{\mathcal{D}} q \triangleright_v \mathbf{0}$  if  $q \notin \tilde{p}$ . Although we cannot call procedures with parameter lists based on the result of a conditional without changing PP (in particular, due to the definition of merge), in this language we can write more sophisticated examples. In [10], we show how to implement Quick-sort, Gaussian elimination and the Fast Fourier Transform in PC, and discuss how implicit parallelism yields efficient implementations of these algorithms.

**Procedures with Holes.** The second extension allows the presence of holes in procedure definitions. Holes are denoted  $\square_h$ , where  $h$  is a unique name for the hole in the procedure. Procedure calls are then allowed to specify a choreography to inject in each hole by means of the syntax  $X(\tilde{p})$  with  $[h_i \mapsto C_i]_i$  where each  $C_i$  is a choreography (in particular, it can be another procedure call) with  $\text{pn}(C_i) \subseteq \tilde{p}$ . We assume unspecified holes to be filled with  $\mathbf{0}$ .

The typing of procedure definitions is extended to  $X(\tilde{q}^T) : G \triangleright G'$  with  $\{h_i : G_i\}_i$ , where  $G_i$  is the guaranteed connection graph when reaching hole  $\square_{h_i}$ . The remaining adaptations (syntax and semantics of PC and PP, and EPP) are trivial.

*Example 11.* Holes enable higher-order composition of choreographies. Consider the following `exchange` procedure:

```
exchange(p,q,r,s) = p.item -> q.get;  $\square_{h1}$ ; r.item -> s.get;  $\square_{h2}$ 
```

We can either leave the holes empty, or use them to perform some extra actions, e.g., inserting a payment:

```
exchange<p,q,r,s> with h1  $\mapsto$  pay<p,q>, h2  $\mapsto$  pay<r,s>
```

We can even insert new causal dependencies between the two communications in `exchange`, e.g., using  $q$  as a broker:

```
exchange<p,q,r,s> with h1  $\mapsto$  q.item -> r.get
```

Using holes, we can also define a general-purpose iterator (see Appendix).

## 7 Discussion, Future and Related Work

**Design of PC.** PC is part of a long-term research effort motivated by empirical studies on concurrency bugs, such as the taxonomies [22] and [21]. They distinguish between deadlock and non-deadlock bugs. In distributed systems, both are mostly caused by the wrong composition of different protocols, rather than the wrong implementation of single protocols. Deadlocks are prevented in PC (Corollary 1). Non-deadlock bugs are more subtle, and typically due to the informal specification (or lack thereof) of the programmer's intentions wrt protocol composition, which leads to unexpected bad executions. In PC, the programmer's

intentions are formalised as a choreography and the composition of procedures in code. Synthesis (EPP) guarantees that these intentions are respected (Theorem 5). Typical non-deadlock bugs include lack of causality between messages, or between an internal computation and a message [21]. For example, a client may forget to wait for an authentication protocol to complete and include the resulting security key in its request to a server in another protocol. We can formalise the correct scenario in PC using our `auth` procedure from Example 4 (`c` is the client and `s` is the server):

```
auth<c,a,r,l>;
if r.ok then r -> c,s[ok];
    r.key -> c.addKey; request<c,s>
else r -> c,s[ko]
```

There is still ground to be covered to support all the features (e.g., pre-emption [21]) used in some of the programs included in these empirical studies. Thus, this kind of studies will also be important for future developments of PC.

**Multiparty Session Types.** In Multiparty Session Types (MPST) [18], global types are specifications of single protocols, used for verifying the code of manually-written implementations in process models. Global types are similar to a simplified fragment of PC, obtained (among other restrictions) by replacing expressions and functions with constants (representing types), removing process creation (the processes are fixed), and restricting recursion to parameterless tail recursion.

MPST leaves the composition of protocols as an implementation detail to the implementors of processes. As a consequence, protocol compositions may lead to deadlocks, unlike in PC. We illustrate this key difference with an example, using our syntax. Consider the protocols  $X(r,s) = r.e \rightarrow s.f$  and  $Y(r',s') = r'.e' \rightarrow s'.f'$ , and assume that we want to compose their instantiations  $X\langle p,q \rangle$  and  $Y\langle q,p \rangle$ . According to MPST, a valid implementation (paraphrased in PP) is  $p \triangleright_v q?f'; q!e \mid q \triangleright_v p?f; p!e'$ . Although this network is obviously deadlocked, this is not detected by MPST because the interleaving of the two protocols is not checked. In PC, we can only obtain correct implementations, because compositions are defined at the level of choreographies, e.g.,  $X\langle p,q \rangle; Y\langle q,p \rangle$  or  $Y\langle q,p \rangle; X\langle p,q \rangle$ .

The authors of [9] augment MPST with a type system that prevents protocol compositions that may lead to deadlocks, but their approach is too restrictive when actions from different protocols are interleaved. Consider the following example, reported as untypable but deadlock-free in [9] (adapted to PC).

```
X(p,q) = p.* -> q; □h1; q.* -> p; p.* -> q; □h2
X<p,q> with h1 ↦ q.* -> p; p.* -> q, h2 ↦ q.* -> p
```

This code is both typable and projectable in PC, yielding a correct implementation.

Another technique for deadlock-freedom in MPST and similar models is to restrict connections among processes participating in different protocols to form

a tree [4,6,8]. PC is more expressive, since connections can form an arbitrary graph.

Recent work investigated how to extend MPST to capture protocols where the number of participants in a session is fixed only at runtime [30], or can grow during execution [15]. These results use ad-hoc primitives and “middleware” terms in the process model, e.g., for tracking the number of participants in a session [15]. Such machinery is not needed in PC. The authors of [14] propose a theory of nested MPST, which partially recalls our notion of parametric procedures. Differently from PC, procedures are invoked by a coordinator (requiring extra communications), and compositions of such nested types can deadlock.

***Choreographic Programming.*** None of the examples in this work can be written in previous models for choreographic programming, which lack full procedural abstraction. As far as we know, this is also the first work exploring choreographic programming in a general setting, instead of in the context of web services [5,7,29].

The models in [5,7] are based on choreographic programming and thus support deadlock-freedom even when multiple protocols are used. However, these models do not support general recursion and parametric procedures, both of which are novel in this work. Composition and reusability in these models are thus limited: procedures cannot have continuations; there can only be a limited number of protocols running at any time (modulo dangling asynchronous actions); and the process names used in a procedure are statically determined, preventing reuse. In PC, all these limitations are lifted.

Asynchronous communications in choreographic programming were addressed in [7] using an ad-hoc transition rule. Our approach in aPC, which builds on implicit parallelism, is simpler and yields an EPP Theorem (Theorem 8) that states a lockstep correspondence between the programmer’s choreography and the synthesised implementation. Instead, in [7], EPP guarantees only a weaker and more complex confluence correspondence.

Holes in PC recall adaptation scopes for choreographies in [12], used to pinpoint where runtime adaptation of code can take place. However, that work deals only with choreographies with a finite number of processes.

A major distinguishing feature of PC is the management of connections among processes, using graphs that are manipulated at runtime. Channel delegation in choreographies [7] behaves in a similar way, but it is less expressive: a process that introduces two other processes cannot communicate with them thenceforth without requiring a new auxiliary channel.

***Other paradigms.*** Reo is another approach that, like PC, disentangles computation from communication and focuses on the programming of interactions, rather than the actions that implement them. In Reo, protocols are obtained by composing (graphical) connectors, which mediate communications among components [1]. As such, Reo produces rather different and more abstract artifacts than PC, which require additional machinery to be implemented (a compiler from Reo to Java uses constraint automata [2]). Programs in PC can be implemented rather directly via standard send and receive primitives (which is

typical of choreographies [7]). PC also supports dynamic creation of processes and guarantees deadlock-freedom by construction, through its simple type system and EPP. Reo supports many interesting connectors and compositions (e.g., multicast), which would be interesting to investigate in PC.

In a different direction, there have been proposals of integrating communication protocols (given as session types) with functional programming [28]. These approaches allow parallel computations of functions to synchronise and exchange values. As for MPST, protocol compositions are left to the programmer, and the same limitations discussed for MPST apply.

**Local concurrency.** Some choreography models support an explicit parallel operator, e.g.,  $C \mid C'$ . We chose not to present it here to keep our model simple, as most behaviours of  $C \mid C'$  are captured by implicit parallelism and asynchrony in PC. For example,  $X\langle p, q \rangle \mid Y\langle r, s \rangle$  is equivalent to  $X\langle p, q \rangle; Y\langle r, s \rangle$  in PC (Rule [C|I-I], Figure 3). However, when some processes names are shared in the parallel composition, then the parallel operator is more expressive, because it allows for *local concurrency*. As an example, consider  $X\langle p, q \rangle \mid Y\langle p, s \rangle$ : here,  $p$  is doing more actions concurrently. In PC, we can encode this by spawning new processes, e.g. our example would become  $p \text{ start } r; p: r \leftrightarrow s; X\langle p, q \rangle; Y\langle r, s \rangle$ . The price we pay is the extra communications to introduce  $r$  ( $p$ 's new process) to  $s$  (plus, we may need to send  $p$ 's value to  $r$ , and maybe send  $r$ 's value to  $p$  after  $Y$  terminates).

If local concurrency in PC is desired, there are different options for introducing it. A straightforward one is to adapt the parallel operator from [5,18], which, however, abstracts from how local concurrency is implemented. A more interesting strategy would be to support named local threads in processes. Threads would share the process location, such that they can receive at the same name without the need to be introduced. Then, implicit parallelism would allow for swapping actions at threads (not just processes) with distinct names. Threads at the same process should not receive from the same sender process at the same time, to avoid races, which can be prevented via typing (cf. [5]). Integration with separation logic [26] would also be interesting, to allow threads also to safely share data.

**Sessions, Connections, and Mobility.** All recent theories based on session types (e.g., [5,7,8,9,18]) assume that all processes in a session (a protocol execution) have a private full-duplex channel to communicate with each other in that session. This amounts to requiring that the graph of connections among processes in a protocol is always complete. This assumption is not necessary in PC, since our semantics and typing require only the processes that actually communicate to be connected. This makes PC a suitable model for reasoning about different kinds of network topologies. In the future, it would be interesting to see whether our type system and connection graphs can be used to enforce pre-defined network structures (e.g., hypercubes or butterflies), making PC a candidate for the programming of choreographies that account for hardware restrictions.

Another important aspect of sessions is that each new protocol execution requires the creation of a new session (and all the connections among its par-

ticipants!). By contrast, protocol executions (running procedures) in PC can reuse all available connections – allowing for more efficient implementations. We use this feature in our parallel downloader example (Example 5), where  $c'$  can still communicate with  $c$  even after the latter introduced it to  $s'$ . Also, aPC supports for the first time the asynchronous establishment of new connections among processes.

The standard results of communication safety found in session-typed calculi can be derived from our EPP Theorem (Theorem 5), as discussed in [7].

**Faults.** Some interesting distributed bugs [21] are triggered by unexpected fault conditions at nodes, making such faults an immediate candidate for the future developments of PC. Useful inspiration to this aim may be provided by [3].

We have also abstracted from faults and divergence of internal computations: in PC, we assume that all internal computations terminate successfully. If we relax these conditions, deadlock-freedom can still be achieved simply by using timeouts and propagating faults through communications.

**Integration.** Sometimes, the code synthesised from a choreography has to be used in combination with legacy process code. For example, we may want to use an existing authentication server in our `auth` procedure in Example 4. This issue is addressed in [24] by using a type theory based on sessions. We leave an adaptation of this idea in our setting to future work.

## References

1. Farhad Arbab. Reo: a channel-based coordination model for component composition. *Mathematical Structures in Computer Science*, 14(3):329–366, 2004.
2. Farhad Arbab, Christel Baier, Jan J. M. M. Rutten, and Marjan Sirjani. Modeling component connectors in reo by constraint automata: (extended abstract). *Electr. Notes Theor. Comput. Sci.*, 97:25–46, 2004.
3. Sara Capecchi, Elena Giachino, and Nobuko Yoshida. Global escape in multiparty sessions. *Mathematical Structures in Computer Science*, 26(2):156–205, 2016.
4. Marco Carbone and Søren Debois. A graphical approach to progress for structured communication in web services. In *Proc. of ICE'10*, 2010.
5. Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.*, 34(2):8, 2012.
6. Marco Carbone, Sam Lindley, Fabrizio Montesi, Carsten Schürmann, and Philip Wadler. Coherence generalises duality: a logical explanation of multiparty session types. In *CONCUR*, 2016. To appear.
7. Marco Carbone and Fabrizio Montesi. Deadlock-freedom-by-design: multiparty asynchronous global programming. In R. Giacobazzi and R. Cousot, editors, *POPL*, pages 263–274. ACM, 2013.
8. Marco Carbone, Fabrizio Montesi, Carsten Schürmann, and Nobuko Yoshida. Multiparty session types as coherence proofs. In L. Aceto and D. de Frutos-Escrig, editors, *CONCUR*, volume 42 of *LIPICs*, pages 412–426. Schloss Dagstuhl, 2015.
9. Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. Global progress for dynamically interleaved multiparty sessions. *Mathematical Structures in Computer Science*, 26(2):238–302, 2016.
10. Luís Cruz-Filipe and Fabrizio Montesi. Choreographies in practice. In Elvira Albert and Ivan Lanese, editors, *FORTE*, volume 9688 of *LNCS*, pages 1–10. Springer, 2016.
11. Luís Cruz-Filipe and Fabrizio Montesi. A core model for choreographic programming. 2016. Accepted for publication.
12. Mila Dalla Preda, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro. Dynamic choreographies – safe runtime updates of distributed applications. In T. Holvoet and M. Viroli, editors, *COORDINATION*, volume 9037 of *LNCS*, pages 67–82. Springer, 2015.
13. Frank S. de Boer, Mohammad Mahdi Jaghoori, Cosimo Laneve, and Gianluigi Zavattaro. Decidability problems for actor systems. *Logical Methods in Computer Science*, 10(4), 2014.
14. Romain Demangeon and Kohei Honda. Nested protocols in session types. In *CONCUR*, pages 272–286, 2012.
15. Pierre-Malo Deniérou and Nobuko Yoshida. Dynamic multirole session types. In T. Ball and M. Sagiv, editors, *POPL*, pages 435–446. ACM, 2011.
16. Maurizio Gabbrielli, Saverio Giallorenzo, and Fabrizio Montesi. Applied choreographies. *CoRR*, abs/1510.03637, 2015.
17. Kohei Honda, Vasco Vasconcelos, and Makoto Kubo. Language primitives and type disciplines for structured communication-based programming. In C. Hankin, editor, *ESOP*, volume 1381 of *LNCS*, pages 122–138. Springer, 1998.
18. Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty Asynchronous Session Types. *J. ACM*, 63(1):9, 2016.

19. Ivan Lanese, Claudio Guidi, Fabrizio Montesi, and Gianluigi Zavattaro. Bridging the gap between interaction- and process-oriented choreographies. In A. Cerone and S. Gruner, editors, *SEFM*, pages 323–332. IEEE, 2008.
20. Ivan Lanese, Fabrizio Montesi, and Gianluigi Zavattaro. Amending choreographies. In *Proceedings 9th International Workshop on Automated Specification and Verification of Web Systems, WWV 2013, Florence, Italy, 6th June 2013.*, pages 34–48, 2013.
21. Tanakorn Leesatapornwongsa, Jeffrey F. Lukman, Shan Lu, and Haryadi S. Gunawi. TaxDC: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In *ASPLOS*, pages 517–530. ACM, 2016.
22. Shan Lu, Soyeon Park, Eunsoo Seo, and Yuanyuan Zhou. Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. *ACM SIGARCH Computer Architecture News*, 36(1):329–339, 2008.
23. Fabrizio Montesi. *Choreographic Programming*. Ph.D. Thesis, IT University of Copenhagen, 2013. <http://fabriziomontesi.com/files/choreographic-programming.pdf>.
24. Fabrizio Montesi and Nobuko Yoshida. Compositional choreographies. In P.R. D’Argenio and H.C. Melgratti, editors, *CONCUR*, volume 8052 of *LNCS*, pages 425–439. Springer, 2013.
25. MPI Forum. *MPI: A Message-Passing Interface Standard*. High-Performance Computing Center Stuttgart, 2015. Version 3.1.
26. John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*, pages 55–74. IEEE Computer Society, 2002.
27. Davide Sangiorgi and David Walker. *The  $\pi$ -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
28. Vasco Thudichum Vasconcelos, Simon J. Gay, and António Ravara. Type checking a multithreaded functional language with session types. *Theor. Comput. Sci.*, 368(1–2):64–87, 2006.
29. W3C WS-CDL Working Group. Web services choreography description language version 1.0. <http://www.w3.org/TR/2004/WD-ws-cdl-10-20040427/>, 2004.
30. Nobuko Yoshida, Pierre-Malo Denielou, Andi Bejleri, and Raymond Hu. Parameterised multiparty session types. In C.-H. Luke Ong, editor, *FOSSACS*, volume 6014 of *LNCS*, pages 128–145. Springer, 2010.

## A Detailed definitions and proofs

We report on definitions and proofs that were omitted from the main part of this paper, as well as some more detailed examples.

### A.1 Type checking and type inference

We start with some technical lemmas about typing.

**Lemma 1 (Monotonicity).** *Let  $\Gamma$  and  $\Gamma'$  be typing contexts with  $\Gamma \subseteq \Gamma'$ ,  $G_1$ ,  $G'_1$  and  $G$  be connection graphs such that  $G_1 \subseteq G$ , and  $C$  be a choreography. If  $\Gamma; G_1 \vdash C \triangleright G'_1$ , then  $\Gamma'; G \vdash C \triangleright G \cup G'_1$ .*

*Proof.* Straightforward by induction on the derivation of  $\Gamma; G_1 \vdash C \triangleright G'_1$ .

**Lemma 2 (Sequentiality).** *Let  $\Gamma$  be a typing context,  $G_1$ ,  $G'_1$ ,  $G_2$  and  $G'_2$  be connection graphs such that  $G_2 \subseteq G'_1$ , and  $C_1$ ,  $C_2$  be choreographies. If  $\Gamma; G_1 \vdash C_1 \triangleright G'_1$  and  $\Gamma; G_2 \vdash C_2 \triangleright G'_2$ , then  $\Gamma; G_1 \vdash C_1 \mathbin{\text{;}} C_2 \triangleright G'_1 \cup G'_2$ .*

*Proof.* Straightforward by induction on the derivation of  $\Gamma; G_2 \vdash C_2 \triangleright G'_2$ .

**Lemma 3 (Substitution).** *Let  $\Gamma$  be a typing context,  $G$  and  $G'$  be connection graphs, and  $C$  be a choreography. Let  $\bar{p}$  be a set of process names that are free in  $C$  and  $\bar{q}$  be a set of process names that do not occur (free or bound) in  $C$ . If  $\Gamma; G \vdash C \triangleright G'$ , then  $\Gamma[\bar{p}/\bar{q}]; G[\bar{p}/\bar{q}] \vdash C[\bar{p}/\bar{q}] \triangleright G'[\bar{p}/\bar{q}]$ .*

*Proof.* Straightforward by induction on the derivation of  $\Gamma; G \vdash C \triangleright G'$ , as all typing rules are valid when substitutions are applied.

We are now ready to start proving Theorem 1. The following lemma takes care of the base cases, and is required for one of the inductive steps.

**Lemma 4.** *Let  $\Gamma$  be a set of typing judgements,  $\mathcal{D}$  a set of procedure definitions,  $G_1$  and  $G'_1$  connection graphs, and  $C$  a choreography that does not start with  $\mathbf{0}$  or a procedure call. Assume that  $\Gamma \vdash \mathcal{D}$  and  $\Gamma; G_1 \vdash C \triangleright G'_1$ . For every state  $\sigma$ , there exist  $\Gamma'$ ,  $\sigma'$ ,  $C'$ ,  $G_2$  and  $G'_2$  such that  $G_1, C, \sigma \rightarrow_{\mathcal{D}} G_2, C', \sigma'$  and  $\Gamma'; G_2 \vdash C' \triangleright G'_2$ .*

*Proof.* By case analysis on the last step of the proof of  $\Gamma; G_1 \vdash C \triangleright G'_1$ . By hypothesis, this proof cannot end with an application of rules [T|End], [T|EndSeq] or [T|Call]; we detail all cases for completeness, but the only non-trivial one is the last.

– [T|Start]: then  $C$  is  $\mathbf{p\,start\,q}^T; C^\circ$  and by hypothesis

$$\Gamma, q : T; G_1 \cup \{p \leftrightarrow q\} \vdash C^\circ \triangleright G'_1.$$

Since  $G_1, \mathbf{p\,start\,q}^T; C, \sigma \rightarrow_{\mathcal{D}} G_1 \cup \{p \leftrightarrow q\}, C^\circ, \sigma[q \mapsto \perp_T]$  by rule [C|Start], taking  $\Gamma' = \Gamma, q : T$ ,  $\sigma' = \sigma[q \mapsto \perp_T]$ ,  $C' = C^\circ$ ,  $G_2 = G_1 \cup \{p \leftrightarrow q\}$  and  $G'_2 = G'_1$  establishes the thesis.



- [T|Com]: then  $C$  is  $\mathbf{p}.e \rightarrow \mathbf{q}.f; C^\circ$  and by hypothesis  $p \xrightarrow{G_1} q, f[\sigma(\mathbf{q})/*](e[\sigma(\mathbf{p})/*])$  is a valid expression of type  $T_q$ , and  $\Gamma; G_1 \vdash C \triangleright G'_1$ . Then all the preconditions of [C|Com] are met, so taking  $\Gamma' = \Gamma, \sigma' = \sigma[\mathbf{q} \mapsto f[\sigma(\mathbf{q})/*](e[\sigma(\mathbf{p})/*])], C' = C^\circ, G_2 = G_1$  and  $G'_2 = G'_1$  establishes the thesis.
- [T|Sel]: then  $C$  is  $\mathbf{p} \rightarrow \mathbf{q}[l]; C^\circ$  and by hypothesis  $p \xrightarrow{G_1} q$  and  $\Gamma; G_1 \vdash C \triangleright G'_1$ . By [C|Sel],  $G_1, \mathbf{p} \rightarrow \mathbf{q}[l]; C^\circ, \sigma \rightarrow_{\mathcal{D}} G_1, C^\circ, \sigma$ , so taking  $\Gamma' = \Gamma, \sigma' = \sigma, C' = C^\circ, G_2 = G_1$  and  $G'_2 = G'_1$  again establishes the thesis.
- [T|Tell]: then  $C$  is  $\mathbf{p}: \mathbf{q} \leftrightarrow \mathbf{r}; C^\circ$  and by hypothesis both  $p \xrightarrow{G_1} q, p \xrightarrow{G_1} r$ , and  $\Gamma; G_1 \cup \{q \leftrightarrow r\} \vdash C^\circ \triangleright G'_1$ . Since the preconditions of rule [C|Tell] are met, by taking  $\Gamma' = \Gamma, \sigma' = \sigma, C' = C^\circ, G_2 = G_1 \cup \{q \leftrightarrow r\}$  and  $G'_2 = G'_1$  establishes the thesis.
- [T|Cond]: then  $C$  is  $\text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2; C^\circ$  and by hypothesis  $e[\sigma(\mathbf{p})/*]$  is a valid Boolean expression,  $\Gamma; G_1 \vdash C_i \triangleright G_i^\circ$  and  $\Gamma; G_1^\circ \cap G_2^\circ \vdash C^\circ \triangleright G'_1$ . Suppose  $e[\sigma(\mathbf{p})/*] = \text{true}$  (the other case is similar). Then  $G_1, \text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2; C^\circ, \sigma \rightarrow_{\mathcal{D}} G_1, C_1 \mathbin{\text{;}} C, \sigma$ . Since  $G_1^\circ \cap G_2^\circ \subseteq G_1^\circ$ , Lemma 2 allows us to conclude that  $\Gamma; G_1 \vdash C_1 \mathbin{\text{;}} C \triangleright G'_1 \cup G_1^\circ$ , whence the thesis follows by taking  $\Gamma' = \Gamma, C' = C_1 \mathbin{\text{;}} C, G_2 = G_1$  and  $G'_2 = G'_1 \cup G_1^\circ$ .

*Proof (Theorem 1).* If  $C \preceq_{\mathcal{D}} \mathbf{0}$ , then the first case holds. Assume that  $C \not\preceq_{\mathcal{D}} \mathbf{0}$ ; we show that the second case holds by induction on the proof of  $\Gamma; G_1 \vdash C \triangleright G'_1$ . By hypothesis, the last rule applied in this proof cannot be [T|End]; the cases where the last rule applied is [T|Start], [T|Com], [T|Sel], [T|Tell] or [T|Cond] follow immediately from Lemma 4, while the case of rule [T|EndSeq] is straightforward from the induction hypothesis.

We focus on the case of rule [T|Call]. In this case,  $C$  has the form  $X(\tilde{\mathbf{p}}); C^\circ$ , and we know that  $\Gamma \vdash X(\tilde{\mathbf{q}}^T) : (G_X \triangleright G'_X), \Gamma \vdash \tilde{\mathbf{p}} : T, G_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}] \subseteq G_1$  and  $\Gamma; G_1 \cup (G'_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}]) \vdash C^\circ \triangleright G'_1$ . From the hypothesis that  $\Gamma \vdash \mathcal{D}$  we also know that  $\Gamma_X; G_X \vdash C_X \triangleright G'_X$ , where  $C_X$  is the body of  $X$  as defined in  $\mathcal{D}$ . By Lemma 3,  $\Gamma_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}]; G_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}] \vdash C_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}] \triangleright G'_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}]$ , whence by Lemma 1 also  $\Gamma; G_1 \vdash C_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}] \triangleright G'_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}] \cup G_1$ . By applying rule [C|Unfold], we conclude that  $X(\tilde{\mathbf{p}}); C^\circ \preceq_{\mathcal{D}} C_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}] \mathbin{\text{;}} C^\circ$ , and Lemma 2 allows us to conclude that  $\Gamma; G_1 \vdash C_X[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}] \mathbin{\text{;}} C^\circ \triangleright G'_1$ . Since procedure calls are guarded,  $C_X$  does not begin with a procedure call, and Lemma 4 establishes the thesis.

*Proof (Theorem 2).* The proof of this result proceeds in several stages. We first observe that deciding whether  $\Gamma; G \vdash C \triangleright G'$  is completely mechanical, as the typing rules are deterministic. Furthermore, those rules can also be used to construct  $G'$  from  $G$  and  $C$ ; therefore, the key step of this proof is showing, given  $\Gamma$  and  $(\mathcal{D}, C)$ , how to find a “canonical typing” for the recursive definitions, the set  $\Gamma_{\mathcal{D}}$ , such that  $\Gamma_{\mathcal{D}} \vdash \mathcal{D}$  and  $\Gamma, \Gamma_{\mathcal{D}}; G_C \vdash C \triangleright G'$  (with  $G'$  inferred) iff  $\Gamma, \Gamma'; G_C \vdash C \triangleright G''$  for some  $\Gamma'$  and  $G''$ . More precisely, we need to find graphs  $G_X$  and  $G'_X$  for each procedure  $X$  defined in  $\mathcal{D}$ .

Our proof proceeds in three steps. First, for each  $X$  we compute an underapproximation  $G_X^\circ$  of the output graph  $G'_X$ , containing all the relevant connections that executing  $X$  can add. Using this, we are able to compute the input graph

$G_X$  and the output graph  $G'_X = G_X \cup G_X^\circ$ . Both these steps are achieved by computing a minimal fixpoint of a monotonic operator in the set of all graphs whose vertices are the parameters of  $X$ . Finally, we argue that the typing  $X : G_X \triangleright G_X$  is minimal, and therefore the set  $\Gamma_{\mathcal{D}}$  of all such typings fulfills the property we require.

Throughout the remainder of this proof, we assume  $\mathcal{D} = \{X_i(\tilde{\mathbf{q}}^i) = C_i \mid i = 1, \dots, n\}$ .

1. In order to compute  $G_{X_i}^\circ$ , we define an auxiliary function `fwd` with intended meaning as follows:  $\text{fwd}_{C_j}^{\tilde{G}_i}(G)$  computes the communication graph obtained from  $G$  after one execution of the body of  $X_j$ , assuming that  $X_i(\tilde{\mathbf{q}}^i) : \emptyset \triangleright G_i$  for all  $i$  and ignoring newly created processes. We use a conditional union operator  $\uplus$  where  $G \uplus \{e\}$  denotes  $G \cup \{e\}$  if  $e$  is an edge connecting two vertices in  $G$ , and  $G$  otherwise. The function `fwd` is defined as follows.

$$\begin{aligned}
\text{fwd}_{\mathbf{0}}^{\tilde{G}_i}(G) &= G \\
\text{fwd}_{\mathbf{0};C}^{\tilde{G}_i}(G) &= \text{fwd}_{C}^{\tilde{G}_i}(G) \\
\text{fwd}_{p.e \rightarrow q.f;C}^{\tilde{G}_i}(G) &= \text{fwd}_{C}^{\tilde{G}_i}(G) \\
\text{fwd}_{p \rightarrow q[l];C}^{\tilde{G}_i}(G) &= \text{fwd}_{C}^{\tilde{G}_i}(G) \\
\text{fwd}_{p \text{ start } q^T;C}^{\tilde{G}_i}(G) &= \text{fwd}_{C}^{\tilde{G}_i}(G) \\
\text{fwd}_{p:q \leftrightarrow r;C}^{\tilde{G}_i}(G) &= \text{fwd}_{C}^{\tilde{G}_i}(G \uplus \{q \leftrightarrow r\}) \\
\text{fwd}_{\text{if } p.e \text{ then } C_1 \text{ else } C_2;C}^{\tilde{G}_i}(G) &= \text{fwd}_{C}^{\tilde{G}_i}(\text{fwd}_{C_1}^{\tilde{G}_i}(G) \cap \text{fwd}_{C_2}^{\tilde{G}_i}(G)) \\
\text{fwd}_{X_i(\tilde{\mathbf{p}});C}^{\tilde{G}_i}(G) &= \text{fwd}_{C}^{\tilde{G}_i}(G \uplus G_i[\tilde{\mathbf{p}}/\tilde{\mathbf{q}}^i])
\end{aligned}$$

Using `fwd`, we define an operator  $\mathcal{F}_{\text{fwd}}$  over the set  $\mathcal{G}$  of tuples of graphs over the parameters of  $X_i$ , i.e.  $\mathcal{G} = \{\tilde{G}_i \mid G_i \text{ is a graph over } \tilde{\mathbf{q}}^i\}$ . Observe that  $\mathcal{G}$  is a complete lattice wrt componentwise inclusion.

$$\mathcal{F}_{\text{fwd}}(\tilde{G}_i) = \widetilde{\text{fwd}_{C_i}^{\tilde{G}_i}(G_i)}$$

This operator is monotonic, since `fwd` only adds edges to its argument, and thus has a least fixpoint that can be computed by iterating  $\mathcal{F}_{\text{fwd}}$  from the tuple of empty graphs over the right sets of vertices. Furthermore, since  $\mathcal{G}$  is finite (each graph has a finite number of vertices) this fixpoint corresponds to a finite iterate, and can thus be computed in finite time. We denote this fixpoint by  $\tilde{G}_{X_i}^\circ$ .

2. The construction of the input graphs  $G_{X_i}$  follows the same idea: we go through the  $C_i$ s noting the edges that are required for all communications to be able to take place. It is however slightly more complicated, because we have to keep track of edges that the choreography adds to the graph; we therefore need a function `bck` that manipulates two graphs instead of one.

More precisely,  $\text{bck}_C^{\tilde{G}_i}(G, G)$  returns the graph extending  $G$  that is needed for correctly executing  $C$  (ignoring newly created processes); the first argument keeps track of the edges that need to be added to  $G$ , and the second argument keeps track of edges added by executing  $C$ . This function uses the graphs  $G_{X_i}^\circ$  computed earlier, which explains why it has to be defined afterwards. We use the same notational conventions as above, and let  $\text{fst}\langle a, b \rangle = a$  and  $\text{snd}\langle a, b \rangle = b$ . The definition of  $\text{bck}$  is given in Figure 16; it is not the simplest

$$\begin{aligned}
& \text{bck}_0^{\tilde{G}_i}(G_a, G_b) = \langle G_a, G_b \rangle \\
& \text{bck}_{0;C}^{\tilde{G}_i}(G_a, G_b) = \text{bck}_C^{\tilde{G}_i}(G_a, G_b) \\
& \text{bck}_{p.e \rightarrow q.f;C}^{\tilde{G}_i}(G_a, G_b) = \begin{cases} \text{bck}_C^{\tilde{G}_i}(G_a, G_b) & \text{if } p \leftrightarrow q \in G_b \\ \text{bck}_C^{\tilde{G}_i}(G_a \uplus \{p \leftrightarrow q\}, G_b \uplus \{p \leftrightarrow q\}) & \text{otherwise} \end{cases} \\
& \text{bck}_{p \rightarrow q[l];C}^{\tilde{G}_i}(G_a, G_b) = \begin{cases} \text{bck}_C^{\tilde{G}_i}(G_a, G_b) & \text{if } p \leftrightarrow q \in G_b \\ \text{bck}_C^{\tilde{G}_i}(G_a \uplus \{p \leftrightarrow q\}, G_b \uplus \{p \leftrightarrow q\}) & \text{otherwise} \end{cases} \\
& \text{bck}_{p \text{ start } q T;C}^{\tilde{G}_i}(G_a, G_b) = \text{bck}_C^{\tilde{G}_i}(G_a, G_b) \\
& \text{bck}_{p:q \leftarrow r;C}^{\tilde{G}_i}(G_a, G_b) = \begin{cases} \text{bck}_C^{\tilde{G}_i}(G_a, G_b \uplus \{q \leftrightarrow r\}) & \text{if } p \leftrightarrow q, p \leftrightarrow r \in G_b \\ \text{bck}_C^{\tilde{G}_i}(G_a \uplus \{p \leftrightarrow q\}, G_b \uplus \{p \leftrightarrow q, q \leftrightarrow r\}) & \text{if } p \leftrightarrow q \notin G_b, p \leftrightarrow r \in G_b \\ \text{bck}_C^{\tilde{G}_i}(G_a \uplus \{p \leftrightarrow r\}, G_b \uplus \{p \leftrightarrow r, q \leftrightarrow r\}) & \text{if } p \leftrightarrow q \in G_b, p \leftrightarrow r \notin G_b \\ \text{bck}_C^{\tilde{G}_i}(G_a \uplus \{p \leftrightarrow q, p \leftrightarrow r\}, G_b \uplus \{p \leftrightarrow q, p \leftrightarrow r, q \leftrightarrow r\}) & \text{if } p \leftrightarrow q, p \leftrightarrow r \notin G_b \end{cases} \\
& \text{bck}_{\text{if } p.e \text{ then } C_1 \text{ else } C_2;C}^{\tilde{G}_i}(G_a, G_b) = \text{bck}_C^{\tilde{G}_i}(\text{fst}(\text{bck}_{C_1}^{\tilde{G}_i}(G_a, G_b)) \cup \text{fst}(\text{bck}_{C_2}^{\tilde{G}_i}(G_a, G_b)), \text{snd}(\text{bck}_{C_1}^{\tilde{G}_i}(G_a, G_b)) \cap \text{snd}(\text{bck}_{C_2}^{\tilde{G}_i}(G_a, G_b))) \\
& \text{bck}_{X_i(\tilde{p});C}^{\tilde{G}_i}(G_a, G_b) = \text{bck}_C^{\tilde{G}_i}(G_a \uplus (G_i[\tilde{p}/\tilde{q}^i] \setminus G_b), G_b \uplus G_i[\tilde{p}/\tilde{q}^i] \uplus G_i^\circ[\tilde{p}/\tilde{q}^i])
\end{aligned}$$

**Fig. 16.** Definition of  $\text{bck}$  (case 2 in the proof of Theorem 2).

possible, but the formulation given is sufficient for our purposes. Again we define a monotonic operator over the same  $\mathcal{G}$  as above.

$$\mathcal{T}_{\text{bck}}(\tilde{G}_i) = \widetilde{\text{fst}(\text{bck}_{C_i}^{\tilde{G}_i}(G_i, G_i))}$$

We do not need to recompute  $G_i^\circ$ , since these graphs contain all edges that can possibly be added by executing  $C_i$ . The least fixpoint of  $\mathcal{T}_{\text{bck}}$  can again be computed by finitely iterating this operator, and it is precisely  $\widetilde{G_{X_i}}$ . We then define  $G'_{X_i} = G_{X_i} \cup G_{X_i}^\circ$ .

3. We now show that  $\Gamma_{\mathcal{D}} = \{X_i(\tilde{q}_i) : G_{X_i} \triangleright G'_{X_i}\}$  is a minimal typing of  $\mathcal{D}$ , in the sense explained earlier. Observe that it is possible that  $\Gamma_{\mathcal{D}} \not\vdash \mathcal{D}$ , in particular if the  $X_i$  are ill-formed choreographies.

Suppose that  $\Gamma, \Gamma'; G_C \vdash C \triangleright G$  for some  $\Gamma'$  and  $G$ . We argue that  $\Gamma, \Gamma_{\mathcal{D}}; G_C \vdash C \triangleright G'$ , where  $G'$  is inferred from the typing rules. For each procedure  $X_i(\tilde{q}_i) = C_i$ , there must be a unique typing  $X_i(\tilde{q}_i) : G_{X_i}^* \triangleright G_{X_i}^{**}$  in  $\Gamma'$ . By a simple inductive argument one can show that  $G_{X_i}^\circ \subseteq G^{**}$  (since  $\emptyset \subseteq G_{X_i}^{**}$  and  $\mathcal{T}_{\text{bck}}$  preserves inclusion in  $G_{X_i}^{**}$ ). Similarly, one shows that  $G_{X_i} \subseteq G_{X_i}^*$ .

and that  $G_{X_i}^{**} \setminus G_{X_i}^* \subseteq G'_{X_i} \setminus G_{X_i}$ . As a consequence, the typing derivation for  $\Gamma, \Gamma'; G_C \vdash C \triangleright G$  can be used for  $\Gamma, \Gamma_{\mathcal{D}}; G_C \vdash C \triangleright G'$ , as all applications of rule [T|Call] are guaranteed to be valid (their preconditions hold) and to produce the same results (they change the communication graph in the same way).

A consequence of this result is that we can obtain a type inference algorithm (Theorem 3), as we only need to “guess” types for the processes in the choreography. As a corollary of the proof, we can also infer the types for parameters of procedural definitions and freshly created processes.

*Proof (Theorem 3).* Construct  $\Gamma$  by going through  $C$  and adding  $\mathfrak{p} : T_{\mathfrak{p}}$  every time there is an action that depends on  $\mathfrak{p}$ 's type (i.e.  $\mathfrak{p}$  is a sender or receiver in a communication, or an argument of a procedure call). If  $\Gamma$  contains two different types for any process, then output  $\text{NO}$ , else output  $\Gamma$ . This algorithm will not necessarily assign a type to all processes in  $C$ , in case  $C$  contains processes whose memory is never accessed.

*Proof (Theorem 4).* Inferring the types of freshly created processes is analogous to the previous proof.

As for parameters of procedure definitions, we omit the details of the proof, as it repeats ideas previously used. Define an operator  $\mathcal{S}_{\mathbb{T}}$  over tuples of typing contexts (one for each  $X_i$  defined in  $\mathcal{D}$ ) that generates a typing context for each  $X_i$  in the same way as in the previous proof. If any contradictions are found, then fail. Iterate  $\mathcal{S}_{\mathbb{T}}$  until either failure occurs (in which case the  $X_i$ s are not properly defined) or a fixpoint is reached. Finally, assign a random type (e.g.  $\mathbb{N}$ ) to each process variable that has not received a type during this procedure. The algorithm readily extends to infer the types of processes created inside procedure definitions.

## A.2 EndPoint Projection

*Merging.* The full definition of merging is given in Figure 17.

*Projections of procedure definitions.* When projecting procedure definitions, we simply kept all arguments, and relied on typing to guarantee that no process will ever attempt to communicate with another process it does not know. However, it is “cleaner” to refine the definition of projection so that such parameters are not even formally used.

This is technically not challenging, as we can use typing information to decide which parameters should be kept in each projection. In other words, when computing  $\llbracket X(\widetilde{\mathfrak{q}}^T) = C \rrbracket$ , the projected procedure  $X_i$  will only contain as arguments those  $\mathfrak{q}_j$  such that  $\mathfrak{q}_i \xrightarrow{G_X} \mathfrak{q}_j$ , where  $G_X$  is given by typing. (This definition is non-deterministic, but it can be made deterministic by using the minimal graph  $G_X$  computed by the type inference algorithm.) A simple annotation can then ensure that projected procedure calls only keep the arguments

$$\begin{aligned}
(\mathbf{q}!e; B) \sqcup (\mathbf{q}!e; B') &= \mathbf{q}!e; (B \sqcup B') & (\mathbf{p}?f; B) \sqcup (\mathbf{p}?f; B') &= \mathbf{p}?f; (B \sqcup B') \\
(\mathbf{q}!!r; B) \sqcup (\mathbf{q}!!r; B') &= \mathbf{q}!!r; (B \sqcup B') & (\mathbf{q}?r; B) \sqcup (\mathbf{q}?r; B') &= \mathbf{q}?r; (B \sqcup B') \\
(\mathbf{q} \oplus l; B) \sqcup (\mathbf{q} \oplus l; B') &= \mathbf{q} \oplus l; (B \sqcup B') & (X\langle \bar{p} \rangle; B) \sqcup (X\langle \bar{p} \rangle; B') &= X\langle \bar{p} \rangle; (B \sqcup B') \\
(\mathbf{start} \mathbf{q} \triangleright B_2; B_1) \sqcup (\mathbf{start} \mathbf{q} \triangleright B'_2; B'_1) &= \mathbf{start} \mathbf{q} \triangleright (B_2 \sqcup B'_2); (B_1 \sqcup B'_1) \\
B_1 \sqcup B_2 &= B'_1 \sqcup B'_2 \quad (\text{if } B_1 \preceq B'_1 \text{ and } B_2 \preceq B'_2) & (\mathbf{0}; B) \sqcup B' &= B \sqcup B' \\
(\text{if } e \text{ then } B_1 \text{ else } B_2); B \sqcup (\text{if } e \text{ then } B'_1 \text{ else } B'_2); B' &= (\text{if } e \text{ then } (B_1 \sqcup B'_1) \text{ else } (B_2 \sqcup B'_2)); (B \sqcup B') \\
(\mathbf{p}\&\{l_i : B_i\}_{i \in J}; B) \sqcup (\mathbf{p}\&\{l_i : B'_i\}_{i \in K}; B') &= \mathbf{p}\&(\{l_i : (B_i \sqcup B'_i)\}_{i \in J \cap K} \cup \{l_i : B_i\}_{i \in J \setminus K} \cup \{l_i : B'_i\}_{i \in K \setminus J}); (B \sqcup B')
\end{aligned}$$

**Fig. 17.** Merging operator in PP.

in the corresponding positions, and typing guarantees that all those arguments are known at runtime by the process invoking the recursive call.

An implementation of Quicksort showcasing this reasoning was previously published in [10].

*EPP.* We sketch the proof of the EPP Theorem.

*Proof (Theorem 5 (sketch)).* The structure of the proof is standard, from [23], so we only show the most interesting differing details. In particular, we need to be careful about how we deal with connections, which is a new key ingredient in PC wrt to previous work. We demonstrate this point for the direction of (*Completeness*); the direction for (*Soundness*) is proven similarly. The proof proceeds by induction on the derivation of  $G, C, \sigma \rightarrow_{\mathcal{D}} G'', C', \sigma'$ . The interesting cases are reported below.

– [C|Tell]: From the definition of EPP we get:

$$\begin{aligned}
& \llbracket \mathbf{p} : \mathbf{q} \langle - \rangle r; C^\circ, \sigma \rrbracket \preceq \\
& \quad \mathbf{p} \triangleright_{\sigma(\mathbf{p})} \mathbf{q}!!r; \llbracket C^\circ \rrbracket_{\mathbf{p}} \mid \mathbf{q} \triangleright_{\sigma(\mathbf{q})} \mathbf{p}?r; \llbracket C^\circ \rrbracket_{\mathbf{q}} \mid r \triangleright_{\sigma(r)} \mathbf{p}?q; \llbracket C^\circ \rrbracket_r \mid N
\end{aligned}$$

By [P|Tell] we get:

$$\begin{aligned}
& \llbracket \mathbf{p} : \mathbf{q} \langle - \rangle r; C^\circ, \sigma \rrbracket \rightarrow \\
& \quad \mathbf{p} \triangleright_{\sigma'(\mathbf{p})} \llbracket C^\circ \rrbracket_{\mathbf{p}} \mid \mathbf{q} \triangleright_{\sigma'(\mathbf{q})} \llbracket C^\circ \rrbracket_{\mathbf{q}} \mid r \triangleright_{\sigma'(r)} \llbracket C^\circ \rrbracket_r \mid N
\end{aligned}$$

which proves the thesis, since we can assume that the projection of  $C'$  remains unchanged for the other processes ( $N$  stays the same).

– [C|Start]: This is the most interesting case. From the definition of EPP we get:

$$\llbracket \mathbf{p} \mathbf{start} \mathbf{q}^T; C^\circ, \sigma \rrbracket \preceq \mathbf{p} \triangleright_{\sigma(\mathbf{p})} \mathbf{start} \mathbf{q}^T \triangleright \llbracket C^\circ \rrbracket_{\mathbf{q}}; \llbracket C^\circ \rrbracket_{\mathbf{p}} \mid N$$

From the semantics of PP we get:

$$\llbracket \mathbf{p} \text{ start } \mathbf{q}^T; C^\circ, \sigma \rrbracket \rightarrow \mathbf{p} \triangleright_{\sigma'(\mathbf{p})} (\llbracket C^\circ \rrbracket_{\mathbf{p}})[\mathbf{q}'/\mathbf{q}] \mid \mathbf{q}' \triangleright_{\perp_T} \llbracket C^\circ \rrbracket_{\mathbf{q}'} \mid N$$

The Barendregt convention implies  $C \preceq (\mathbf{p} \text{ start } \mathbf{q}^T; C^\circ)[\mathbf{q}'/\mathbf{q}]$ . We now have to prove that:

$$\llbracket C^\circ[\mathbf{q}'/\mathbf{q}], \sigma \rrbracket \preceq \mathbf{p} \triangleright_{\sigma'(\mathbf{p})} (\llbracket C^\circ \rrbracket_{\mathbf{p}})[\mathbf{q}'/\mathbf{q}] \mid \mathbf{q}' \triangleright_{\perp_T} \llbracket C^\circ \rrbracket_{\mathbf{q}'} \mid N$$

We observe that this is true only if process  $\mathbf{q}$  does not occur free in  $N$ , i.e.,  $\mathbf{q}$  appear in  $N$  only inside the scope of a binder. The latter must be of the form  $r? \mathbf{q}; B$ . This is guaranteed by the fact that  $C$  is well-typed, since the typing rules prevent other processes in  $N$  to communicate with  $\mathbf{q}$  without being first introduced.  $\square$

*Choreography Amendment.* We now define precisely an amendment function that makes every choreography projectable. As mentioned earlier, we follow ideas previously used in other choreography models [20,11]. Observe that, by definition, the only choreography construct that can lead to unprojectability is the conditional.

**Definition 4 (Amendment).** *Given a choreography  $C$ , the transformation  $\text{Amend}(C)$  repeatedly applies the following procedure until it is no longer possible, starting from the inner-most subterms in  $C$ . For each conditional subterm if  $\mathbf{p}.e$  then  $C_1$  else  $C_2$  in  $C$ , let  $\tilde{r} \subseteq (\text{pn}(C_1) \cup \text{pn}(C_2))$  be such that  $\llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r$  is undefined for all  $r \in \tilde{r}$ ; then if  $\mathbf{p} \leq \mathbf{q}$  then  $C_1$  else  $C_2$  in  $C$  is replaced with:*

$$\begin{aligned} &\text{if } \mathbf{p}.e \text{ then } \mathbf{p} \rightarrow r_1[L]; \dots; \mathbf{p} \rightarrow r_n[L]; C_1 \\ &\text{else } \mathbf{p} \rightarrow r_1[R]; \dots; \mathbf{p} \rightarrow r_n[R]; C_2 \end{aligned}$$

By the definitions of  $\text{Amend}$  and  $\text{EPP}$  and the semantics of  $\text{PC}$ , we get the following properties, where  $\rightarrow^*$  is the transitive closure of  $\rightarrow$ .

**Theorem 9 (Amendment).** *Let  $C$  be a choreography. Then:*

**(Completeness)**  $\text{Amend}(C)$  is defined;

**(Projectability)** for all  $\sigma$ ,  $\llbracket \text{Amend}(C), \sigma \rrbracket$  is defined;

**(Correspondence)** for all  $G, \sigma$  and  $\mathcal{D}$ :

- if  $G, C, \sigma \rightarrow_{\mathcal{D}} G', C', \sigma$ , then  $G, \text{Amend}(C), \sigma \rightarrow_{\mathcal{D}}^* G', \text{Amend}(C'), \sigma'$ ;
- if  $G, \text{Amend}(C), \sigma \rightarrow_{\mathcal{D}} G', C', \sigma'$ , then there exist  $C''$  and  $\sigma''$  such that  $G, C, \sigma \rightarrow_{\mathcal{D}} G', C'', \sigma''$  and  $G', C', \sigma' \rightarrow_{\mathcal{D}}^* G', C'', \sigma''$ .

*Proof (Theorem 9).* (Completeness) and (Projectability) are immediate by definition of  $\text{Amend}$ . For (Correspondence), the proof is by analysis of the possible transitions. The only interesting cases occur when the transition consumes a conditional. In the case  $G, C, \sigma \rightarrow G', C', \sigma$ , then  $\text{Amend}(C)$  also has to consume the label selections introduced by amendment in the branch taken in order to reach  $\text{Amend}(C')$ . Conversely, if  $G, \text{Amend}(C), \sigma$  makes a transition that consumes a conditional, then  $C'$  needs to consume the label selections introduced by amendment in order to match the corresponding move by  $C$ .

$$\begin{array}{c}
\frac{p \xrightarrow{G} q \quad e[\sigma(p)/*] \downarrow v}{G, p.e \xrightarrow{x} \bullet_q; C, \sigma \rightarrow_{\mathcal{D}} G, C[v/x], \sigma} \quad [\text{C|Com-S}] \\
\\
\frac{q \xrightarrow{G} p \quad f[\sigma(q)/*](v) \downarrow w}{G, \bullet_p \xrightarrow{v} q.f; C, \sigma \rightarrow_{\mathcal{D}} G, C, \sigma[q \mapsto w]} \quad [\text{C|Com-R}] \\
\\
\frac{p \xrightarrow{G} q}{G, p \xrightarrow{x} \bullet_q[l]; C, \sigma \rightarrow_{\mathcal{D}} G, C[l/x], \sigma} \quad [\text{C|Sel-S}] \\
\\
\frac{q \xrightarrow{G} p}{G, \bullet_p \xrightarrow{l} q[l]; C, \sigma \rightarrow_{\mathcal{D}} G, C, \sigma} \quad [\text{C|Sel-R}] \\
\\
\frac{p \xrightarrow{G} q \quad p \xrightarrow{G} r}{G, p: \bullet_q \xrightarrow{x,y} \bullet_r; C, \sigma \rightarrow_{\mathcal{D}} G, C[q/x, r/y], \sigma} \quad [\text{C|Tell-S}] \\
\\
\frac{q \xrightarrow{G} p}{G, \bullet_p.r \xrightarrow{r} q; C, \sigma \rightarrow_{\mathcal{D}} G \cup \{q \rightarrow r\}, C, \sigma} \quad [\text{C|Tell-R}]
\end{array}$$

**Fig. 18.** Asynchronous PC, Semantics of New Runtime Terms.

This result also holds if we replace  $\mathcal{D}$  by  $\text{Amend}(\mathcal{D})$  in the relevant places.

The first choreography in Remark 2, which is unprojectable, can be amended to the projectable choreography presented at the end of the same remark.

### A.3 Asynchrony

We detail the whole sets of rules for the semantics of aPC (Figure 18) and its type system (Figure 19), as well as for the semantics of aPP (Figure 20). The full definition of EPP is given in Figures 21 and 22.

The proofs of the relationships between PC/PP and their asynchronous counterparts are mechanical.

*Proof (Theorem 6).* Straightforward by case analysis on the possible transitions of  $C$ .

*Proof (Theorem 7).* Straightforward by case analysis on the possible transitions of  $N$ .

### A.4 Procedures with holes

We illustrate the use of holes to define a general-purpose iterator, procedure `Iter` below. Using `Iter`, we define a procedure `Reverse` for reversing the list of a process  $q$ . We omit the straightforward procedures for decrementing a counter and popping a list.

$$\begin{array}{c}
\frac{p \xrightarrow{G} q \quad \Gamma \vdash p : T_p \quad * : T_p \vdash_{\top} e : T \quad \Gamma \oplus (x : T); G \vdash C \triangleright G'}{\Gamma; G \vdash p.e \xrightarrow{x} \bullet_q; C \triangleright G'} \text{ [T|Com-S]} \\
\frac{q \xrightarrow{G} p \quad \Gamma \vdash q : T_q \quad * : T_q \vdash_{\top} f : T \rightarrow T_q \quad \Gamma \oplus (x : T); G \vdash C \triangleright G'}{\Gamma; G \vdash \bullet_p \xrightarrow{x} q.f; C \triangleright G'} \text{ [T|Com-RV]} \\
\frac{q \xrightarrow{G} p \quad \Gamma \vdash q : T_q \quad \vdash_{\top} v : T \quad * : T_q \vdash_{\top} f : T \rightarrow T_q \quad \Gamma; G \vdash C \triangleright G'}{\Gamma; G \vdash \bullet_p \xrightarrow{v} q.f; C \triangleright G'} \text{ [T|Com-RT]} \\
\frac{p \xrightarrow{G} q \quad \Gamma \oplus (x = l); G \vdash C \triangleright G'}{\Gamma; G \vdash p \xrightarrow{x} \bullet_q[l]; C \triangleright G'} \text{ [T|Sel-S]} \\
\frac{q \xrightarrow{G} p \quad \Gamma \oplus (x = l); G \vdash C \triangleright G'}{\Gamma; G \vdash \bullet_p \xrightarrow{x} q[l]; C \triangleright G'} \text{ [T|Sel-RV]} \quad \frac{q \xrightarrow{G} p \quad \Gamma; G \vdash C \triangleright G'}{\Gamma; G \vdash \bullet_p \xrightarrow{l} q[l]; C \triangleright G'} \text{ [T|Sel-RT]} \\
\frac{p \xrightarrow{G} q \quad p \xrightarrow{G} r \quad \Gamma \oplus (x = q) \oplus (y = r) \vdash C \triangleright G'}{\Gamma; G \vdash p : \bullet_q \xleftrightarrow{x,y} \bullet_r; C \triangleright G'} \text{ [T|Tell-S]} \\
\frac{q \xrightarrow{G} p \quad \Gamma; G \cup \{q \rightarrow r\} \vdash C \triangleright G'}{\Gamma; G \vdash \bullet_{p,r} \xrightarrow{r} q; C \triangleright G'} \text{ [T|Tell-RV]} \\
\frac{q \xrightarrow{G} p \quad \Gamma \oplus (x = r); G \cup \{q \rightarrow r\} \vdash C \triangleright G'}{\Gamma; G \vdash \bullet_{p,r} \xrightarrow{x} q; C \triangleright G'} \text{ [T|Tell-RT]}
\end{array}$$

**Fig. 19.** Asynchronous PC, Typing Rules (New Runtime Terms).

```

Iter(p,q,r) = if p.is_zero
then p -> q,r[stop]
else dec<p>; p -> q,r[cont];  $\square_h$ ; Iter<p,q,r>

Reverse(q) =
q starts p,r; q: p<->r;
q.size -> p; q.empty -> r;
Iter<p,q,r> with h  $\mapsto$  q.top -> r.append; pop<q>;
r.* -> q

```



$$\begin{array}{c}
\frac{\rho'_q = \rho_q \cdot \langle \mathbf{p}, e[v/*] \rangle}{\mathbf{p} \triangleright_v^{\rho_p} \mathbf{q}!e; B_p \mid \mathbf{q} \triangleright_w^{\rho_q} B_q \rightarrow_{\mathcal{B}} \mathbf{p} \triangleright_v^{\rho_p} B_p \mid \mathbf{q} \triangleright_w^{\rho'_q} B_q} \text{ [P|Com-S]} \\
\frac{\rho_q \preceq \langle \mathbf{p}, v \rangle \cdot \rho'_q \quad u = (f[w/*])(v)}{\mathbf{q} \triangleright_w^{\rho_q} \mathbf{p}?f; B \rightarrow_{\mathcal{B}} \mathbf{q} \triangleright_w^{\rho'_q} B} \text{ [P|Com-R]} \\
\frac{\rho'_q = \rho_q \cdot \langle \mathbf{p}, \mathbf{r} \rangle \quad \rho'_r = \rho_r \cdot \langle \mathbf{p}, \mathbf{q} \rangle}{\mathbf{p} \triangleright_v^{\rho_p} \mathbf{q}!!\mathbf{r}; B_p \mid \mathbf{q} \triangleright_w^{\rho_q} B_q \mid \mathbf{r} \triangleright_z^{\rho_r} B_r \rightarrow_{\mathcal{B}} \mathbf{p} \triangleright_v^{\rho_p} B_p \mid \mathbf{q} \triangleright_w^{\rho'_q} B_q \mid \mathbf{r} \triangleright_z^{\rho'_r} B_r} \text{ [P|Tell-S]} \\
\frac{\rho_q \preceq \langle \mathbf{p}, \mathbf{r} \rangle \cdot \rho'_q}{\mathbf{q} \triangleright_w^{\rho_q} \mathbf{p}?\mathbf{r}; B \rightarrow_{\mathcal{B}} \mathbf{q} \triangleright_w^{\rho'_q} B} \text{ [P|Tell-R]} \\
\frac{\rho'_q = \rho_q \cdot \langle \mathbf{p}, l \rangle}{\mathbf{p} \triangleright_v^{\rho_p} \mathbf{q} \oplus l; B \mid \mathbf{q} \triangleright_w^{\rho_q} B_q \rightarrow_{\mathcal{B}} \mathbf{p} \triangleright_v^{\rho_p} B \mid \mathbf{q} \triangleright_w^{\rho'_q} B_q} \text{ [P|Sel-S]} \\
\frac{\rho_q \preceq \langle \mathbf{p}, l_j \rangle \cdot \rho'_q \quad j \in I}{\mathbf{q} \triangleright_w^{\rho_q} \mathbf{p}\&\{l_i : B_i\}_{i \in I} \rightarrow_{\mathcal{B}} \mathbf{q} \triangleright_w^{\rho'_q} B_j} \text{ [P|Sel-R]}
\end{array}$$

**Fig. 20.** Asynchronous Procedural Processes, Semantics (New Rules).

$$\begin{array}{l}
\llbracket \mathbf{p}.e \xrightarrow{x} \bullet_{\mathbf{q}}; C \rrbracket_r = \begin{cases} \mathbf{q}!e; [C]_r & \text{if } r = \mathbf{p} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \quad \llbracket \bullet_{\mathbf{p}} \xrightarrow{\hat{v}} \mathbf{q}.f; C \rrbracket_r = \begin{cases} \mathbf{p}?f; [C]_r & \text{if } r = \mathbf{q} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \mathbf{p} \xrightarrow{x} \bullet_{\mathbf{q}}[l]; C \rrbracket_r = \begin{cases} \mathbf{q} \oplus l & \text{if } r = \mathbf{p} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \quad \llbracket \bullet_{\mathbf{p}} \xrightarrow{\hat{l}} \mathbf{q}[l]; C \rrbracket_r = \begin{cases} \mathbf{p}\&\{l : [C]_r\} & \text{if } r = \mathbf{q} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \mathbf{p} : \bullet_{\mathbf{q}} \xrightarrow{x,y} \bullet_{\mathbf{r}}; C \rrbracket_s = \begin{cases} \mathbf{q}!!\mathbf{r}; [C]_s & \text{if } s = \mathbf{p} \\ \llbracket C \rrbracket_s & \text{otherwise} \end{cases} \quad \llbracket \bullet_{\mathbf{p}}.r \xrightarrow{\hat{p}} \mathbf{q}; C \rrbracket_s = \begin{cases} \mathbf{p}?\mathbf{r}; [C]_s & \text{if } s = \mathbf{q} \\ \llbracket C \rrbracket_s & \text{otherwise} \end{cases}
\end{array}$$

**Fig. 21.** Asynchronous PC, Behaviour Projection (New Rules).

$$\begin{array}{l}
\llbracket \bullet_{\mathbf{p}} \xrightarrow{\hat{v}} \mathbf{q}.f; C \rrbracket_r = \begin{cases} \langle \mathbf{p}, v \rangle \cdot \llbracket C \rrbracket_r & \text{if } r = \mathbf{q} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \quad \llbracket \bullet_{\mathbf{p}} \xrightarrow{\hat{l}} \mathbf{q}[l]; C \rrbracket_r = \begin{cases} \langle \mathbf{p}, l \rangle \cdot \llbracket C \rrbracket_r & \text{if } r = \mathbf{q} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \bullet_{\mathbf{p}}.r \xrightarrow{\hat{r}} \mathbf{q}; C \rrbracket_s = \begin{cases} \langle \mathbf{p}, r \rangle \cdot \llbracket C \rrbracket_s & \text{if } s = \mathbf{q} \\ \llbracket C \rrbracket_s & \text{otherwise} \end{cases} \\
(\text{if } \mathbf{p} \leq \mathbf{q} \text{ then } C_1 \text{ else } C_2; C)_r = \llbracket C_1 \rrbracket_r \cdot \llbracket C_2 \rrbracket_r \quad \llbracket \eta; C \rrbracket_r = \llbracket I; C \rrbracket_r = \llbracket C \rrbracket_r
\end{array}$$

**Fig. 22.** Asynchronous PC, State Projection.