

On Asynchrony and Choreographies

Luís Cruz-Filipe^{a,*}, Fabrizio Montesi^a

^a*Department of Mathematics and Computer Science, University of Southern Denmark*

Abstract

Choreographic Programming is a paradigm for the development of concurrent software, where deadlocks are prevented syntactically and then correct-by-construction implementations in process models are mechanically generated. The formal semantics of choreography languages are typically based on synchronous communications, in order to achieve a simpler theory. However, many real-world systems have asynchronous communications. Previous attempts at enriching choreographies with asynchrony rely on *ad-hoc* constructions, whose adequacy is only argued informally.

In this work, we present a systematic study of asynchronous communications in choreographies. First, we discuss and formalise the properties that an asynchronous semantics for choreographies should have. Then, we explore how out-of-order execution, previously used in choreographies for modelling concurrency, can be used to endow choreographies with an asynchronous semantics that satisfies our properties. Our development yields a pleasant correspondence with FIFO-based asynchronous messaging, modelled in a process calculus. Finally, we investigate the expressivity of choreography languages with respect to asynchrony. Specifically, we find out that choreography languages equipped with process spawning and name mobility primitives are expressive enough to program asynchronous behaviour over a simple synchronous semantics.

Keywords: choreographic programming; asynchrony; concurrent programming

1. Introduction

Choreographic Programming [26] is a paradigm for developing concurrent software, where an “Alice and Bob” notation is used to prevent mismatched I/O actions syntactically. An EndPoint Projection (EPP) can then be used to synthesise correct-by-construction process implementations. Choreographies are used in different settings, including standards [3, 32], languages [6, 20, 28, 31], specification models [4, 5, 23, 29], and design tools [3, 28, 31, 32].

The key to preventing mismatched I/O actions in choreographies is that interactions between two (or more) processes are specified atomically, using terms such as $p.e \rightarrow q; C$ – read “process p sends the evaluation of expression e to process q , and then we proceed as the choreography C ”. Giving a semantics to such terms is relatively easy if we assume that communications are synchronous: we can just reduce $p.e \rightarrow q; C$ to C in a single step (and update q ’s state with the received value, but this is orthogonal to this discussion). For this reason, most research on choreographic programming focused on systems with a synchronous communications semantics.

However, many real-world systems use asynchronous communications. This motivated the introduction of an ad-hoc reduction rule for modelling asynchrony in choreographies (Rule $[C|Async]$ in [5]). As an example, consider the choreography $p.1 \rightarrow q; p.2 \rightarrow r$ (where 1 and 2 are just constants). The special rule would allow for consuming the second communication immediately, thus reducing the choreography to $p.1 \rightarrow q$. In general, roughly, this rule allows a choreography of the form $p.e \rightarrow q; C$ to execute an action in C if this action involves p but not q (sends are non-blocking, receives are blocking). This approach was later

*Corresponding author; tel. +45 6550 2387, fax +45 6550 2373

adopted in other works (see Section 6). Unfortunately, it also comes with a problem: it yields an unintuitive semantics, since a choreography can now reduce to a state that would normally not be reachable in the real world. In our example, specifically, in the real world p would have to send its first message to q before it could proceed to sending its other message to r . This information is lost in the choreography reduction, where it appears that p can just send its messages in any order. In [5], this also translates to a misalignment between the structures of choreographies and their process implementations generated by EPP, since the latter use a standard asynchronous semantics with message buffers; see Section 6 for a detailed discussion of this aspect. Previous work [17] uses intermediate runtime terms in choreographies to represent asynchronous messages in transit, in an attempt at overcoming this problem. However, the adequacy of this approach has never been formally demonstrated.

In this work, we are interested in studying asynchrony for choreographies in a systematic way. Thus, we first analyse the properties that an asynchronous choreography semantics should have. Assuming standard FIFO duplex channels between each pair of processes, messages can be sent without the intended receiver being ready, and sent messages are eventually received. We formulate these properties precisely in a representative choreography language.

Our study leads naturally to the construction of a new choreography model that supports asynchronous communications, by capitalising on the characteristic feature of out-of-order execution found in choreographic programming. We formally establish the adequacy of our asynchronous model, by proving that it respects the formal definitions of our properties. Then, we define an EPP from our new model to an asynchronous process calculus. Thanks to the accurate asynchronous semantics of our choreography model, we prove that the code generated by our EPP and the originating choreography are lockstep operationally equivalent. As a corollary, our generated processes are deadlock-free by construction. Our development also has the pleasant property that programmers do not need to reason about asynchrony: they can just program thinking in the usual terms of synchronous communications, and assume that adopting asynchronous communications will not lead to errors.

Finally, we leverage our new knowledge on asynchronous choreography semantics to study the expressivity of choreography languages with respect to asynchrony. In other words, we want to know when a choreography model is expressive enough to encode asynchronous behaviour using its existing semantics. We find out that choreography languages equipped with process spawning and name mobility enjoy this property, by presenting an endo-encoding that rewrites a choreography to one where all communications are implemented asynchronously. We conclude by linking our two developments, showing that the encoding produces correct asynchronous behaviour according to our formal notion of asynchrony for choreographies.

Contributions. This article presents four main contributions. First, we give an abstract characterisation of asynchronous semantics for choreography languages, which is formalised for a minimal choreography calculus. Secondly, we propose an asynchronous semantics for this minimal language, show that it is an instance of our characterisation, and discuss how it can be applied to other choreography calculi. Thirdly, we prove a lockstep operational correspondence between choreographies and their process implementations, when asynchronous semantics for both systems are considered. Finally, we show that choreography languages equipped with the features of process spawning and name mobility (which our minimal language does not possess) are expressive enough to simulate asynchronous communications as we formalised.

Publication history. This article combines and extends material and results from previous publications. The abstract characterization of asynchrony (Section 3) and the asynchronous semantics for choreographies (Section 4) were originally introduced in [13]. The encoding of asynchronous communications using process spawning was described in [10], while the full extension to the calculi DMC and DCC (Section 5) was first presented in [12]. Theorems 9 and 10, which connect this encoding to the abstract characterization from Section 3, are new, as well as the material in Section 5.4. The discussion on the impact of our results and comparison to related work has also been expanded from the previous presentations.

Structure. We present the representative choreography language MC in which we develop our work, together with its associated process calculus and EPP, in Section 2. In Section 3, we motivate and introduce

$$C ::= p.e \rightarrow q; C \mid p \rightarrow q[\ell]; C \mid \text{if } p.e \text{ then } C_1 \text{ else } C_2 \mid \text{def } X = C_2 \text{ in } C_1 \mid X \mid \mathbf{0}$$

Figure 1: Core Choreographies, syntax.

the properties we would expect of an asynchronous choreography semantics. Section 4 introduces a semantics for MC that satisfies these properties. We also show that we can define an asynchronous variant of the target process calculus and extend the definition of EPP towards it, preserving the precise operational correspondence from the synchronous case. In Section 5, we present a syntactic approach to asynchrony, showing how to encode asynchronous communications in choreography languages that support process spawning and name mobility. We relate our development to other approaches for asynchrony in choreographies in Section 6, before concluding in Section 7 with a discussion on the implications of our work and possible future directions.

2. Core Choreographies and Stateful Processes

We review the choreography model of Core Choreographies (CC) and its target calculus of Stateful Processes (SP), originally introduced in [9].

2.1. Core Choreographies

The language of CC is defined inductively in Figure 1. Processes, ranged over by p, q, r, \dots , are assumed to run concurrently. Processes interact through two types of communication actions, which we collectively denote by η . In value communications $p.e \rightarrow q$, process p evaluates expression e and sends the result to q . The precise syntax of expressions is immaterial for our presentation; in particular, expressions can access values stored in p 's memory. In label selections $p \rightarrow q[\ell]$ process p selects, from several possible behaviours of q , the one identified by label ℓ . We abstract over the set of labels L , requiring only that it be finite and contain at least two elements.

The remaining choreography terms denote conditionals, recursion, and termination. In the conditional $\text{if } p.e \text{ then } C_1 \text{ else } C_2$, process p evaluates expression e to decide whether the choreography should proceed as C_1 or as C_2 . In $\text{def } X = C_2 \text{ in } C_1$, we define variable X to be the choreography term C_2 , which then can be called (as X) inside both C_1 and C_2 . Term $\mathbf{0}$ is the terminated choreography, which we sometimes omit. For a more detailed discussion of these primitives, we refer the reader to [9].¹

The (synchronous) semantics of CC is a reduction semantics that uses a total state function σ to represent the memory state at each process p . Since our development is orthogonal to the details of the memory implementation, we say that $\sigma(p)$ is a representation of the memory state of p (left unspecified) and write $\text{upd}(\sigma, p, v)$ to denote the (uniquely defined) updated memory state of p after receiving a value v . Term $e \Downarrow_p^{\sigma} v$ denotes that locally evaluating expression e at p , with memory state σ , evaluates to v . (This formulation captures the essence of previous memory models for choreographies, cf. [5, 8].)

The transitions are defined over pairs $\langle C, \sigma \rangle$, given by the rules in Figure 2. These rules are mostly standard, and we summarise their intuition. In $\llbracket C \mid \text{Com} \rrbracket$, the state of q is updated with the value received from p (which results from the evaluation of expression e at that process). Label selections are no-ops (rule $\llbracket C \mid \text{Sel} \rrbracket$), but play a role in the implementations as we discuss later. Rule $\llbracket C \mid \text{Cond} \rrbracket$ is as expected, while rule $\llbracket C \mid \text{Ctx} \rrbracket$ allows reductions under recursive definitions. Finally, rule $\llbracket C \mid \text{Struct} \rrbracket$ uses a structural precongruence \preceq , defined in Figure 3, which essentially allows (i) independent communications to be swapped (rule $\llbracket C \mid \text{Eta-Eta} \rrbracket$), (ii) recursive definitions to be unfolded (rule $\llbracket C \mid \text{Unfold} \rrbracket$), and (iii) garbage collection (rule $\llbracket C \mid \text{ProcEnd} \rrbracket$). These rules, taken together, endow the semantics of CC with out-of-order execution: interactions not at the top level may be brought to the top and executed if they do not interfere with other

¹We relaxed the syntax of CC slightly with respect to [9] by leaving the syntax of expressions unspecified, which allows for the simpler conditional $\text{if } p.e \text{ then } C_1 \text{ else } C_2$ in line with typical choreography languages. This minor change simplifies our presentation.

$$\begin{array}{c}
\frac{e \downarrow_{\sigma}^p v}{p.e \rightarrow q; C, \sigma \rightarrow C, \text{upd}(\sigma, q, v)} \text{ [C|Com]} \\
\frac{}{p \rightarrow q[\ell]; C, \sigma \rightarrow C, \sigma} \text{ [C|Sel]} \\
\frac{i = 1 \text{ if } e \downarrow_{\sigma}^p \text{ true} \quad i = 2 \text{ o.w.}}{\text{if } p.e \text{ then } C_1 \text{ else } C_2, \sigma \rightarrow C_i, \sigma} \text{ [C|Cond]} \\
\frac{C_1, \sigma \rightarrow C'_1, \sigma'}{\text{def } X = C_2 \text{ in } C_1, \sigma \rightarrow \text{def } X = C_2 \text{ in } C'_1, \sigma'} \text{ [C|Ctx]} \\
\frac{C_1 \preceq C_2 \quad C_2, \sigma \rightarrow C'_2, \sigma' \quad C'_2 \preceq C'_1}{C_1, \sigma \rightarrow C'_1, \sigma'} \text{ [C|Struct]}
\end{array}$$

Figure 2: Core Choreographies, synchronous semantics.

$$\begin{array}{c}
\frac{\text{pn}(\eta) \cap \text{pn}(\eta') = \emptyset}{\eta; \eta' \equiv \eta'; \eta} \text{ [C|Eta-Eta]} \quad \frac{\text{pn}(C_i) \cap \text{pn}(\eta) = \emptyset}{\text{def } X = C_2 \text{ in } (\eta; C_1) \equiv \eta; \text{def } X = C_2 \text{ in } C_1} \text{ [C|Eta-Rec]} \\
\frac{}{\text{def } X = C_2 \text{ in } C_1[X] \preceq \text{def } X = C_2 \text{ in } C_1[C_2]} \text{ [C|Unfold]} \quad \frac{}{\text{def } X = C \text{ in } \mathbf{0} \preceq \mathbf{0}} \text{ [C|ProcEnd]} \\
\frac{\{p, q\} \cap \text{pn}(\eta) = \emptyset}{\text{if } p.e \text{ then } (\eta; C_1) \text{ else } (\eta; C_2) \equiv \eta; \text{if } p.e \text{ then } C_1 \text{ else } C_2} \text{ [C|Eta-Cond]} \\
\frac{p \neq q}{\text{if } p.e \text{ then } (\text{if } q.e' \text{ then } C_1 \text{ else } C_2) \text{ else } (\text{if } q.e' \text{ then } C'_1 \text{ else } C'_2)} \text{ [C|Cond-Cond]} \\
\equiv \\
\text{if } q.e' \text{ then } (\text{if } p.e \text{ then } C_1 \text{ else } C'_1) \text{ else } (\text{if } p.e \text{ then } C_2 \text{ else } C'_2)
\end{array}$$

Figure 3: Core Choreographies, structural precongruence.

$$\begin{aligned}
B &::= \mathbf{q}!e; B \mid \mathbf{p}?; B \mid \mathbf{q} \oplus \ell; B \mid \mathbf{p}\&\{\ell_i : B_i\}_{i \in I}; B \mid \text{if } e \text{ then } B_1 \text{ else } B_2; B \mid \text{def } X = B_2 \text{ in } B_1 \mid X \mid \mathbf{0} \\
N, M &::= \mathbf{p} \triangleright_{\sigma_p} B \mid (N \mid M) \mid \mathbf{0}
\end{aligned}$$

Figure 4: Stateful Processes, syntax.

interactions that precede them in the choreography’s abstract syntax tree. We write $C \equiv C'$ for $C \preceq C'$ and $C' \preceq C$, and denote the set of process names in a choreography C by $\text{pn}(C)$.

Example 1. *We introduce the running example we use throughout this article. In the following choreography where a buyer, Alice (a), purchases a product from a seller (s) through her bank (b).*

1. $\mathbf{a.title} \rightarrow \mathbf{s}$;
2. $\mathbf{s.price} \rightarrow \mathbf{a}$;
3. $\mathbf{s.price} \rightarrow \mathbf{b}$;
4. **if** $\mathbf{b.happy}$ **then**
5. $\mathbf{b} \rightarrow \mathbf{s}[ok]$; $\mathbf{b} \rightarrow \mathbf{a}[ok]$;
6. $\mathbf{s.book} \rightarrow \mathbf{a}$
7. **else** $\mathbf{b} \rightarrow \mathbf{s}[ko]$; $\mathbf{b} \rightarrow \mathbf{a}[ko]$

In Line 1, the term $\mathbf{a.title} \rightarrow \mathbf{s}$ denotes an interaction whereby a communicates the title of the book that Alice wishes to buy to s. The seller then sends the price of the book to both a and b. In Line 4, a sends the price she expects to pay to b, which confirms that it is the same amount requested by s (stored internally at b). If so, b notifies both s and a of the successful transaction (Line 5) and s sends the book to a (Line 6). Otherwise, b notifies s and a of the failure (Line 7) and the choreography terminates.

Unsurprisingly, choreographies in CC are always deadlock-free. We use this property later on, to prove that the process code generated from choreographies is also deadlock-free.

Theorem 1 (Deadlock-freedom [9]). *Given a choreography C , either $C \preceq \mathbf{0}$ (termination) or, for every state σ , there exist a choreography C' and a state σ' such that $C, \sigma \rightarrow C', \sigma'$.*

2.2. Stateful Processes and EndPoint Projection

Core Choreographies are meant to be implemented in a minimalistic process calculus, also introduced in [9], called Stateful Processes (SP). We summarize this calculus, noting that we make the same conventions and changes regarding expressions, labels, and states as above.

The syntax of SP is reported in Figure 4. A term $\mathbf{p} \triangleright_{\sigma_p} B$ is a process with name \mathbf{p} , memory state σ_p and behaviour B . Networks, ranged over by N, M , are parallel compositions of processes, with $\mathbf{0}$ being the inactive network.

Behaviours correspond to the local views of choreography actions. The process executing a send term $\mathbf{q}!e; B$ evaluates expression e and sends the result to process \mathbf{q} , proceeding as B . The dual receiving behaviour $\mathbf{p}?; B$ expects a value from process \mathbf{p} , stores it in its memory and proceeds as B . Term $\mathbf{q} \oplus \ell; B$ sends a label ℓ to process \mathbf{q} . Selections are received by the branching term $\mathbf{p}\&\{\ell_i : B_i\}_{i \in I}$, which upon receiving one of the labels ℓ_i proceeds according to B_i . Branching terms must offer at least one branch. The other terms are as in CC.

These intuitions are formalized in the synchronous semantics of SP, which is defined by the rules in Figure 5. Rule [P|Com] models synchronous value communication: a process \mathbf{p} wishing to send a value to \mathbf{q} can synchronise with a receive-from- \mathbf{p} action at \mathbf{q} , and the state of \mathbf{q} is updated accordingly. Rule [P|Sel] is standard selection, with \mathbf{p} selecting one of the branches offered by \mathbf{q} . The remaining rules are standard.

$$\begin{array}{c}
\frac{e \downarrow_{\sigma_p} v}{\mathfrak{p} \triangleright_{\sigma_p} \mathfrak{q}!e; B_1 \mid \mathfrak{q} \triangleright_{\sigma_q} \mathfrak{p}?; B_2 \rightarrow \mathfrak{p} \triangleright_{\sigma_p} B_1 \mid \mathfrak{q} \triangleright_{\text{upd}(\sigma_q, v)} B_2} \text{ [P|Com]} \\
\frac{j \in I}{\mathfrak{p} \triangleright_{\sigma_p} \mathfrak{q} \oplus \ell_j; P \mid \mathfrak{q} \triangleright_{\sigma_q} \mathfrak{p}\&\{\ell_i : Q_i\}_{i \in I} \rightarrow \mathfrak{p} \triangleright_{\sigma_p} P \mid \mathfrak{q} \triangleright_{\sigma_q} Q_j} \text{ [P|Sel]} \\
\frac{i = 1 \text{ if } e \downarrow_{\sigma_p} \text{ true} \quad i = 2 \text{ o.w.}}{\mathfrak{p} \triangleright_{\sigma_p} \text{if } e \text{ then } P_1 \text{ else } P_2 \rightarrow \mathfrak{p} \triangleright_{\sigma_p} P_i} \text{ [P|Cond]} \quad \frac{N \preceq M \quad M \rightarrow M' \quad M' \preceq N'}{N \rightarrow N'} \text{ [P|Struct]} \\
\frac{P \mid N \rightarrow P' \mid N'}{\mathfrak{p} \triangleright_v \text{def } X = Q \text{ in } P \mid N \rightarrow \mathfrak{p} \triangleright_v \text{def } X = Q \text{ in } P' \mid N'} \text{ [P|Ctx]} \quad \frac{N \rightarrow N'}{N \mid M \rightarrow N' \mid M} \text{ [P|Par]}
\end{array}$$

Figure 5: Stateful Processes, synchronous semantics.

$$\begin{array}{c}
\frac{}{\text{def } X = B_2 \text{ in } B_1[X] \preceq \text{def } X = B_2 \text{ in } B_1[B_2]} \text{ [S|Unfold]} \\
\frac{}{\mathfrak{p} \triangleright_v \mathbf{0} \preceq \mathbf{0}} \text{ [S|PZero]} \quad \frac{}{N \mid \mathbf{0} \preceq N} \text{ [S|NZero]} \quad \frac{}{\text{def } X = B \text{ in } \mathbf{0} \preceq \mathbf{0}} \text{ [S|ProcEnd]}
\end{array}$$

Figure 6: Stateful Processes, structural precongurence.

This calculus once again includes a structural precongurence relation \preceq , defined in Figure 6, which allows unfolding of recursive definition and garbage collection (removal of terminated processes).

Networks in SP do not enjoy a counterpart to Theorem 1, as send/receive actions may be unmatched or wrongly ordered. To avoid this happening, we generate networks automatically from choreographies in a faithful way, by EndPoint Projection (EPP).

EPP is defined first at the process level, by means of a partial function $\llbracket C \rrbracket_p$. The rules defining behaviour projection are given in Figure 7. Each choreography term is projected to the local action of the process that we are projecting. For example, a communication term $\mathfrak{p}.e \rightarrow \mathfrak{q}$ is projected into a send action for the sender process \mathfrak{p} , a receive action for the receiver process \mathfrak{q} , or nothing otherwise. The rule for projecting a conditional uses the standard (and 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 (see [9] for details). The definition of merge is given in Figure 8. This aspect is found repeatedly in most choreography models [2, 4, 5, 15].

Definition 1 (EPP from CC to SP). *Given a choreography C and a state σ , the endpoint projection $\llbracket C, \sigma \rrbracket$*

$$\begin{array}{c}
\llbracket \mathfrak{p}.e \rightarrow \mathfrak{q}; C \rrbracket_r = \begin{cases} \mathfrak{q}!e; \llbracket C \rrbracket_r & \text{if } r = \mathfrak{p} \\ \mathfrak{p}?; \llbracket C \rrbracket_r & \text{if } r = \mathfrak{q} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \quad \llbracket \mathfrak{p} \rightarrow \mathfrak{q}[\ell]; C \rrbracket_r = \begin{cases} \mathfrak{q} \oplus \ell; \llbracket C \rrbracket_r & \text{if } r = \mathfrak{p} \\ \mathfrak{p}\&\{\ell : \llbracket C \rrbracket_r\} & \text{if } r = \mathfrak{q} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \text{if } \mathfrak{p}.e \text{ then } C_1 \text{ else } C_2; C \rrbracket_r = \begin{cases} \text{if } e \text{ then } \llbracket C_1 \rrbracket_r \text{ else } \llbracket C_2 \rrbracket_r; \llbracket C \rrbracket_r & \text{if } r = \mathfrak{p} \\ (\llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r); \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \text{def } X = C_2 \text{ in } C_1 \rrbracket_r = \text{def } X = \llbracket C_2 \rrbracket_r \text{ in } \llbracket C_1 \rrbracket_r \quad \llbracket \mathbf{0} \rrbracket_r = \mathbf{0} \quad \llbracket X \rrbracket_r = X
\end{array}$$

Figure 7: Core Choreographies, behaviour projection.

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

Figure 8: Core Choreographies, merge operator in behaviour projection.

is the parallel composition of the EPPs of C wrt all processes in $\text{pn}(C)$:

$$\llbracket C, \sigma \rrbracket = \prod_{p \in \text{pn}(C)} p \triangleright_{\sigma(p)} \llbracket C \rrbracket_p.$$

When $\llbracket C, \sigma \rrbracket = N$ is defined for any σ , we say that C is *projectable* and that N is the projection of C, σ .

Example 2. *The projection of the choreography in Example 1 is the parallel composition of its projections relative to each of the three intervening processes.*

$$\begin{aligned}
a \triangleright_{\sigma(a)} s!title; s?; b\&\{ok : s?, ko : \mathbf{0}\} \\
s \triangleright_{\sigma(s)} a?; a!price; b!price; b\&\{ok : a!book, ko : \mathbf{0}\} \\
b \triangleright_{\sigma(b)} s?; \text{if } happy \text{ then } (s \oplus ok; a \oplus ok) \text{ else } (s \oplus ko; a \oplus ko)
\end{aligned}$$

When these processes execute concurrently, they precisely follow the specification from the original choreography.

Theorem 2 (EPP Theorem [9]). *Let σ be a choreography and σ be a state. If $\llbracket C, \sigma \rrbracket$ is defined, then:*

- (Completeness) *if $C, \sigma \rightarrow C', \sigma'$ for some choreography C' and state σ' , then $\llbracket C, \sigma \rrbracket \rightarrow \succ \llbracket C', \sigma' \rrbracket$;*
- (Soundness) *if $\llbracket C, \sigma \rrbracket \rightarrow N$ for some process network N , then there exist a choreography C' and a state σ' such that $C, \sigma \rightarrow C', \sigma'$ and $\llbracket C', \sigma' \rrbracket \prec N$.*

The *pruning relation* \prec (see [4, 5]) eliminates branches introduced by the merging operator \sqcup when they are not needed anymore to follow the originating choreography. Pruning does not alter reductions, since the eliminated branches are never selected, cf. [4, 15, 23]. Combining Theorem 2 with Theorem 1 we get that the projections of choreographies never deadlock.

Corollary 1 (Deadlock-freedom by construction [9]). *Let N be a process network such that $N = \llbracket C, \sigma \rrbracket$ for some choreography C and state σ . Then either $N \preceq \mathbf{0}$ (termination) or there exists a network N' such that $N \rightarrow N'$.*

Our final definition in this section, also from [9], provides us with a way to compare the behaviour of two choreographies with respect to a given set of processes.

Definition 2 (Computational equivalence). *Two states σ_1, σ_2 are equivalent wrt a set of process names \tilde{p} , written $\sigma_1 \equiv_{\tilde{p}} \sigma_2$, if $\sigma_1(\mathbf{p}) = \sigma_2(\mathbf{p})$ for every $\mathbf{p} \in \tilde{p}$.*

Two choreographies C_1 and C_2 are equivalent wrt a set of process names \tilde{p} if: whenever $\sigma_1 \equiv_{\tilde{p}} \sigma_2$, if $C_1, \sigma_1 \rightarrow^ \mathbf{0}, \sigma'_1$ then $C_2, \sigma_2 \rightarrow^* \mathbf{0}, \sigma'_2$ with $\sigma'_1 \equiv_{\tilde{p}} \sigma'_2$, and conversely.*

$$\mathcal{C}[\bullet] ::= \bullet; C \mid \eta; \mathcal{C}[\bullet] \mid \text{if } p.e \text{ then } \mathcal{C}_1[\bullet] \text{ else } \mathcal{C}_2[\bullet] \mid \text{def } X = C \text{ in } \mathcal{C}[\bullet] \mid X \mid \mathbf{0}$$

Figure 9: Contexts for asynchrony.

3. Asynchrony in Choreographies

In order to define an asynchronous semantics for CC and SP and prove their equivalence to the synchronous semantics, we need to define precisely what we understand by “asynchronous semantics” and by “equivalence”. Intuitively, communications from p to q are asynchronous if: they occur in two steps (send and receive), the send step does not require q to be ready, and if the send step is executed, then the corresponding receive step is also eventually executed.

We begin by defining asynchrony in CC. We extend the syntax with contexts, which can contain holes denoted as \bullet , formally defined in Figure 9. Structural precongruence is defined for contexts as for choreographies. We do not consider holes to be interactions, so they cannot be swapped with any other action.

Using contexts, we can define a function that identifies the type of the next action in context $\mathcal{C}[\bullet]$ involving a process r : a communication (**comm**), a conditional (**cond**), or dependent on how \bullet is instantiated (\bullet). This function is partial – in particular, it is undefined if $\mathcal{C}[\bullet]$ does not contain either r or \bullet .

Definition 3. *The next action for process r in a context $\mathcal{C}[\bullet]$, denoted $\text{next}_r(\mathcal{C}[\bullet])$, is defined as follows.*

$$\begin{aligned} \text{next}_r(\bullet; C) &= \bullet & \text{next}_r(\mathbf{0}) &= \text{next}_r(X) = \text{undefined} \\ \text{next}_r(\eta; \mathcal{C}[\bullet]) &= \begin{cases} \text{comm} & \text{if } r \in \text{pn}(\eta) \\ \text{next}_r(\mathcal{C}[\bullet]) & \text{otherwise} \end{cases} \\ \text{next}_r(\text{if } p.e \text{ then } \mathcal{C}_1[\bullet] \text{ else } \mathcal{C}_2[\bullet]) &= \begin{cases} \text{cond} & \text{if } p = r \\ \text{next}_r(\mathcal{C}_1[\bullet]) & \text{if } \text{next}_r(\mathcal{C}_1[\bullet]) = \text{next}_r(\mathcal{C}_2[\bullet]) \\ \text{undefined} & \text{otherwise} \end{cases} \\ \text{next}_r(\text{def } X = C \text{ in } \mathcal{C}[\bullet]) &= \text{next}_r((\mathcal{C}[\bullet])[C/X]) \end{aligned}$$

The proviso in the second case of the definition of $\text{next}_r(\text{if } p.e \text{ then } \mathcal{C}_1[\bullet] \text{ else } \mathcal{C}_2[\bullet])$ ensures that the enabled action for r is uniquely defined. Note that recursive definitions are only unfolded once.

This notion is well-defined, since structural precongruence cannot swap actions involving the same process. We denote by $\mathcal{C}[\eta]$ the result of replacing every occurrence of \bullet in $\mathcal{C}[\bullet]$ by η . If η is an interaction and $\mathcal{C}[\eta]$ is projectable, then $\text{next}_r(\mathcal{C}[\eta])$ describes the top action in the EPP of $\llbracket \mathcal{C}[\eta] \rrbracket_r$.

Since we expect asynchronous communications to occur in two steps, it is reasonable to expect CC to be extended with additional actions. By an *extension* of CC, we understand a choreography language that has all the primitives of CC, an arbitrarily larger set of statements η , and an adequately extended function pn . Our definitions of contexts and next extend automatically to these languages.

We can now formalise the intuition given at the beginning of this section.

Definition 4. *A semantics \rightarrow_a for an extension of CC is asynchronous if the following conditions hold for any state σ and context $\mathcal{C}[\bullet]$, where we assume that $\mathbf{0}; C$ stands for C .*

- If $\text{next}_p(\mathcal{C}[\bullet]) = \bullet$, then $\mathcal{C}[p.e \rightarrow q], \sigma \rightarrow_a \mathcal{C}[*_q^v], \sigma$, where $e \downarrow_\sigma^p v$ and $*_q^v$ is a statement that depends on v and q .
- If $\text{next}_p(\mathcal{C}[\bullet]) = \bullet$, then $\mathcal{C}[p \rightarrow q[l]], \sigma \rightarrow_a \mathcal{C}[*_q^l], \sigma$, with $*_q^l$ as above.
- If $\text{next}_q(\mathcal{C}[\bullet]) = \bullet$, then $\mathcal{C}[*_q^v], \sigma \rightarrow_a \mathcal{C}[\mathbf{0}], \text{upd}(\sigma, q, v)$.
- If $\text{next}_q(\mathcal{C}[\bullet]) = \bullet$, then $\mathcal{C}[*_q^l], \sigma \rightarrow_a \mathcal{C}[\mathbf{0}], \sigma$.

$$\eta ::= \dots \mid p.e \xrightarrow{x} \bullet_q \mid \bullet_p \xrightarrow{\hat{v}} q \mid p \xrightarrow{x} \bullet_q[\ell] \mid \bullet_p \xrightarrow{\hat{\ell}} q[\ell] \quad \hat{v} ::= x \mid v \quad \hat{\ell} ::= x \mid \ell$$

Figure 10: Asynchronous Core Choreographies, runtime terms.

The second ingredient we need is a formal definition of equivalence. Our notion is similar to the standard notion of operational correspondence from [19], but stronger, since we require that any additional term introduced by asynchronous reductions can always be consumed.

Definition 5. *An asynchronous semantics \rightarrow_a for an extension of CC is asynchronously equivalent to the semantics \rightarrow of CC if, for all choreographies C and C' and states σ and σ' :*

- if $C, \sigma \rightarrow C', \sigma'$, then $C, \sigma \rightarrow_a^* C', \sigma'$;
- if $C, \sigma \rightarrow_a^* C', \sigma'$, then there exist a choreography C'' and a state σ'' such that $C, \sigma \rightarrow^* C'', \sigma''$ and $C', \sigma' \rightarrow_a^* C'', \sigma''$.

Note that, in the last point of the above definition, the choreography C' might include terms from the extended choreography language.

4. Asynchrony by Design

Now we are ready to present our extension of CC and define its asynchronous semantics.

4.1. Asynchronous Core Choreographies

We extend the syntax of choreographies with the runtime terms in Figure 10. We call the resulting calculus aCC (for asynchronous CC). The key idea is that, at runtime, a communication is expanded into multiple actions. For example, a communication $p.e \rightarrow q$ expands in $p.e \xrightarrow{x} \bullet_q$ – a send action from p – and $\bullet_p \xrightarrow{x} q$ – a receive action by q . The process subscripts at \bullet are immaterial for the semantics, but are used later to define EPP. The tag x is used to specify that the original intention by the programmer was for that message from p to reach that receive action at q . Thus, executing $p.e \xrightarrow{x} \bullet_q$ replaces x in the corresponding receive action with the actual value v computed from e at p , yielding $\bullet_p \xrightarrow{v} q$. Finally, executing $\bullet_p \xrightarrow{v} q$ updates the state of q . The tags in the actions for label selections are not essential for propagating values as above, but they make the semantics uniform. Process names for the new runtime terms are defined as follows, ignoring process names in tags and in \bullet subscripts.

$$\text{pn}(p.e \xrightarrow{x} \bullet_q) = \text{pn}(p \xrightarrow{x} \bullet_q[\ell]) = \{p\} \quad \text{pn}(\bullet_p \xrightarrow{\hat{v}} q) = \text{pn}(\bullet_p \xrightarrow{\hat{\ell}} q[\ell]) = \{q\}$$

The semantics for aCC includes:

- the rules from Figure 11, replacing $[C|\text{Com}]$ and $[C|\text{Sel}]$, and from Figure 12, extending \preceq ;
- rules $[C|\text{Cond}]$ and $[C|\text{Struct}]$ from Figure 2;
- the rules defining \preceq (Figure 2), with η now ranging also over the new runtime terms.

In line with the Barendregt convention, we require the tags introduced by unfolding a value communication or selection to be globally fresh. This allows us to use these tags to maintain the correspondence between the value being sent and that being received.

The key to the asynchronous semantics of aCC lies in the new swaps allowed by \preceq , due to the definition of pn for the new terms.

$$\begin{array}{c}
\frac{e \downarrow_{\sigma}^p v}{p.e \xrightarrow{x} \bullet_q; C, \sigma \rightarrow_a C[v/x], \sigma} \quad [C|Com-S] \qquad \frac{}{\bullet_p \xrightarrow{v} q; C, \sigma \rightarrow_a C, \text{upd}(\sigma, q, v)} \quad [C|Com-R] \\
\frac{}{p \xrightarrow{x} \bullet_q[\ell]; C, \sigma \rightarrow_a G, C[\ell/x], \sigma} \quad [C|Sel-S] \qquad \frac{}{\bullet_p \xrightarrow{\ell} q[\ell]; C, \sigma \rightarrow_a C, \sigma} \quad [C|Sel-R]
\end{array}$$

Figure 11: Asynchronous Core Choreographies, semantics (runtime terms).

$$\frac{}{p.e \rightarrow q \preceq p.e \xrightarrow{x} \bullet_q; \bullet_p \xrightarrow{x} q} \quad [C|Com-Unfold] \qquad \frac{}{p \rightarrow q[\ell] \preceq p \xrightarrow{x} \bullet_q[\ell]; \bullet_p \xrightarrow{x} q[\ell]} \quad [C|Sel-Unfold]$$

Figure 12: Asynchronous Core Choreographies, new rule for structural precongurence.

Example 3. Consider again the choreography from Example 1. We focus on the first three communications. To execute them asynchronously, we must first expand them:

$$\begin{array}{c}
a.\text{title} \rightarrow s; s.\text{price} \rightarrow a; s.\text{price} \rightarrow b \\
\preceq a.\text{title} \xrightarrow{x} \bullet_s; \bullet_a \xrightarrow{x} s; s.\text{price} \xrightarrow{y} \bullet_a; \bullet_s \xrightarrow{y} a; s.\text{price} \xrightarrow{z} \bullet_b; \bullet_s \xrightarrow{z} b
\end{array}$$

The first two communications must be executed in order, since the first action blocks the second one and the second action blocks all subsequent sends. However, the last four actions can be interchanged using \preceq :

$$s.\text{price} \xrightarrow{y} \bullet_a; \bullet_s \xrightarrow{y} a; s.\text{price} \xrightarrow{z} \bullet_b; \bullet_s \xrightarrow{z} b \preceq s.\text{price} \xrightarrow{y} \bullet_a; s.\text{price} \xrightarrow{z} \bullet_b; \bullet_s \xrightarrow{z} b; \bullet_s \xrightarrow{y} a$$

So s can send both its messages at once, while b can receive its message before a .

The semantics of aCC also enjoys deadlock-freedom.

Theorem 3 (Deadlock-freedom). Given a choreography C , either $C \preceq \mathbf{0}$ (termination) or, for every σ , there exist a choreography C' and a state σ' such that $C, \sigma \rightarrow_a C', \sigma'$.

Proof. Straightforward by structural induction on C . \square

Theorem 4. The relation \rightarrow_a is an asynchronous semantics for aCC.

Proof. For the first case, we are only interested in a situation when $\text{next}_p(\mathcal{C}[\bullet]) = \bullet$. A straightforward induction on the definition of next , using the rules in Figure 3, shows that in this case $\mathcal{C}[p.e \rightarrow q] \equiv \mathcal{C}[p.e \xrightarrow{x} \bullet_q; \bullet_p \xrightarrow{x} q] \preceq p.e \xrightarrow{x} \bullet_q; \mathcal{C}[\bullet_p \xrightarrow{x} q]$, as p does not occur in any action encountered during the computation of $\text{next}_p(\mathcal{C}[\bullet])$. The thesis then follows by rule $[C|Com-S]$. The second case is nearly identical. The remaining two cases are similar, arguing over $\text{next}_q(\mathcal{C}[\bullet])$. \square

Theorem 5. The semantics \rightarrow_a is asynchronously equivalent to the semantics of CC.

Proof. Proving the first property is straightforward. The only non-trivial case is when the transition $C, \sigma \rightarrow C', \sigma'$ uses rule $[C|Com]$, and in this case in the asynchronous semantics we can unfold the corresponding communication and reduce twice using $[C|Com-S]$ and $[C|Com-R]$.

For the second property, we make two observations.

- (1) The relation \rightarrow_a has the diamond property, namely: if $C, \sigma \rightarrow_a C', \sigma'$ and $C, \sigma \rightarrow_a C'', \sigma''$, then there exist C''' and σ''' such that $C', \sigma' \rightarrow_a C''', \sigma'''$ and $C'', \sigma'' \rightarrow_a C''', \sigma'''$ (see Figure 13, left). This is directly proven by case analysis on the structure of a choreography C with two distinct possible transitions.

- (2) If $C, \sigma \rightarrow_a C', \sigma'$, then there exist C'' and σ'' such that $C, \sigma \rightarrow^* C'', \sigma''$ and $C', \sigma' \rightarrow_a^* C'', \sigma''$. This is similar to the proof for the first property. If the transition uses rule [C|Com-S], then the choreography C'' is obtained by firing rule [C|Com-R], but it might be necessary to execute additional actions in case the receiver is not ready for that transition yet. This is always possible due to Theorem 3.

The property now follows by induction on the length of the reduction chain $C, \sigma \rightarrow_a^* C', \sigma'$. When this length is 0, the result is trivial, and when the length is 1, the result is simply (2) above. Otherwise, we follow the reasoning displayed in Figure 13, right. By assumption, we have that $C, \sigma \rightarrow_a^* C_1, \sigma_1 \rightarrow_a C', \sigma'$ (top of the picture). By induction hypothesis, there exist C'_1 and σ'_1 such that $C, \sigma \rightarrow^* C'_1, \sigma'_1$ and $C_1, \sigma_1 \rightarrow_a^* C'_1, \sigma'_1$ (left triangle). By iterated application of (1) above, there exist C''_1 and σ''_1 such that $C'_1, \sigma'_1 \rightarrow_a^? C''_1, \sigma''_1$ and $C', \sigma' \rightarrow_a^* C''_1, \sigma''_1$ (right square). The reduction $C'_1, \sigma'_1 \rightarrow_a^? C''_1, \sigma''_1$ is optional because it may be the case that the action in the reduction $C_1, \sigma_1 \rightarrow_a C', \sigma'$ is already included in the reduction chain $C_1, \sigma_1 \rightarrow_a^* C'_1, \sigma'_1$. In this case, the thesis follows, taking $C'' = C'_1$ and $\sigma'' = \sigma'_1$. Otherwise, we apply (2) above (lower right triangle) to obtain C'' and σ'' such that $C'_1, \sigma'_1 \rightarrow^* C'', \sigma''$ and $C''_1, \sigma''_1 \rightarrow_a^* C'', \sigma''$, and the thesis follows. \square




Figure 13: The diamond property for \rightarrow_a (left) and the reasoning in the proof of the second property in Theorem 5 (right).

Remark 1. Reasoning on asynchronous behaviour is known to be hard for programmers. That is why our terms for modelling asynchronous communications are runtime terms: they are not intended to be part of the source language used by developers to program. In general, programmers do not need to be aware of the development in this section, thanks to Theorem 5. Instead, they can just write a choreography thinking in terms of synchronous communications as usual (syntax and semantics), and then just assume that adopting an asynchronous communications will not lead to any bad behaviour. In the next subsection, we show that this abstraction carries over to asynchronous process code generated by EPP. Since EPP is automatic, developers do not need to worry about asynchrony in process code either.

4.2. Asynchronous Stateful Processes

In the development above, we characterized asynchrony at an abstract level, and showed that our choreography model satisfies the properties we identified. In this section, we take a more down-to-earth approach: we define directly an asynchronous variant for SP, extend EPP to the runtime terms introduced earlier, and show that the asynchronous semantics for both calculi again satisfy the EPP Theorem (Theorem 2).

Syntactically, we need to make a slight change to our processes: each process is now also equipped with a queue of incoming messages. We name this version of the calculus aSP. So actors in the asynchronous model have the form $\mathbf{p} \triangleright_v^\rho B$, where ρ is a queue of incoming messages. We sometimes omit ρ when it is empty; in particular, this allows us to view SP networks as special cases of aSP networks. A message is a pair $\langle \mathbf{q}, m \rangle$, where \mathbf{q} is the sender process and m is a value, label, or process identifier.

The semantics of aSP consists of rules [P|Cond], [P|Par], [P|Ctx] and [P|Struct] from SP (Figure 5) together with the new rules given in Figure 14. Structural precongruence for aSP is defined exactly as for SP. We write $\rho \cdot \langle \mathbf{q}, m \rangle$ to denote the queue obtained by appending message $\langle \mathbf{q}, m \rangle$ to ρ , and $\langle \mathbf{q}, m \rangle \cdot \rho$ for the queue with $\langle \mathbf{q}, m \rangle$ at the head and ρ as tail. We simulate having one separate FIFO queue for each other

$$\begin{array}{c}
\frac{e \downarrow_{\sigma}^p v \quad \rho'_q = \rho_q \cdot \langle p, v \rangle}{p \triangleright_{\sigma_p}^{\rho_p} q!e; B_p \mid q \triangleright_{\sigma_q}^{\rho_q} B_q \rightarrow p \triangleright_{\sigma_p}^{\rho_p} B_p \mid q \triangleright_{\sigma_q}^{\rho'_q} B_q} \text{ [P|Com-S]} \quad \frac{\rho_q \preceq \langle p, v \rangle \cdot \rho'_q}{q \triangleright_{\sigma_q}^{\rho_q} p?; B \rightarrow q \triangleright_{\text{upd}(\sigma_q, v)}^{\rho'_q} B} \text{ [P|Com-R]} \\
\frac{\rho'_q = \rho_q \cdot \langle p, \ell \rangle}{p \triangleright_{\sigma_p}^{\rho_p} q \oplus \ell; B_p \mid q \triangleright_{\sigma_q}^{\rho_q} B_q \rightarrow p \triangleright_{\sigma_p}^{\rho_p} B_p \mid q \triangleright_{\sigma_q}^{\rho'_q} B_q} \text{ [P|Sel-S]} \quad \frac{j \in I \quad \rho_q \preceq \langle p, v \rangle \cdot \rho'_q}{q \triangleright_{\sigma_q}^{\rho_q} p \& \{ \ell_i : B_i \}_{i \in I} \rightarrow q \triangleright_{\sigma_q}^{\rho'_q} B_j} \text{ [P|Sel-R]}
\end{array}$$

Figure 14: Asynchronous Stateful Processes, semantics (new rules).

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 SP are valid also in aSP.

Theorem 6. *Let N be a process network in SP. If $N \rightarrow N'$ for some network N' , then $N_{\square} \rightarrow^* N'_{\square}$ in aSP, where N_{\square} denotes the asynchronous network obtained by adding an empty queue to each process.*

Proof. Straightforward by case analysis: the only reduction rules in SP that are not present in aSP are those for value communications and label selections, which can be simulated by applying rules [P|Com-S] and [P|Com-R] (respectively, [P|Sel-S] and [P|Sel-R]) in sequence. \square

The converse is not true, so the relation between SP and aSP is not so strong as that between CC and aCC (stated in Theorem 5). This is because of deadlocks: in SP, a communication action can only take place when the sender and receiver are ready to synchronize; in aSP, 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 $p \triangleright_{\sigma_p} q!1 \mid q \triangleright_{\sigma_q} \mathbf{0}$ is deadlocked in SP, but its counterpart in aSP (obtained by adding empty queues at each process) reduces to $q \triangleright_{\sigma_q}^{\langle p, 1 \rangle} \mathbf{0}$. (This network is not equivalent to $\mathbf{0}$, since there is a non-empty queue.)

More interestingly, the network

$$p \triangleright_{\sigma_p} q!e_1; r? \mid q \triangleright_{\sigma_q} r!e_2; p? \mid r \triangleright_{\sigma_r} p!e_3; q? \quad (1)$$

is deadlocked in SP, but reduces to $\mathbf{0}$ in aSP (all queues are eventually emptied). It is possible to extend CC with communication primitives that capture this type of behaviour, as discussed in [7]; we briefly discuss this point in Section 7.

4.3. Asynchronous EndPoint Projection

Defining an EPP from aCC to aSP requires extending the previous definition with clauses for the new runtime terms, which populate the local queues in the projections. Intuitively, when compiling, e.g., $\bullet_p \xrightarrow{q} v$, we add a message from containing v at the top of q 's queue.

There is one problem with this approach: we can write choreographies that use runtime terms in a “wrong” way, for which Theorem 2 no longer holds.

Example 4. *Consider the choreography $C = p.1 \rightarrow q; \bullet_p \xrightarrow{q} 2$. If we naively project it as we described informally, we obtain $p \triangleright_{\sigma(p)}^{\square} q!1 \mid q \triangleright_{\sigma(q)}^{\langle p, 2 \rangle} p?; p?$, where \square is the empty queue, and q will receive 2 before it receives 1.*

To avoid this undesired behaviour, we restrict ourselves to *well-formed* choreographies: those that can arise from executing a choreography that does not contain runtime terms (i.e., a program). Since runtime terms are supposed to be hidden from the programmer anyway, this restriction does not make us lose any generality in practice.

Definition 6 (Well-formedness). *A choreography C in aCC containing runtime terms is well-formed if $\eta_1; \dots; \eta_m; C^{CC} \preceq^- C$, where:*

$$\llbracket \bullet_p \xrightarrow{q} v; C \rrbracket_r = \begin{cases} p?; \llbracket C \rrbracket_r & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \quad \llbracket \bullet_p \xrightarrow{\ell} q[\ell]; C \rrbracket_r = \begin{cases} p\&\{\ell : \llbracket C \rrbracket_r\} & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases}$$

Figure 15: Core Choreographies, asynchronous behaviour projection (new rules).

$$\llbracket \bullet_p \xrightarrow{q} v; C \rrbracket_r = \begin{cases} \langle p, v \rangle \cdot \llbracket C \rrbracket_r & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \quad \llbracket \bullet_p \xrightarrow{\ell} q[\ell]; C \rrbracket_r = \begin{cases} \langle p, \ell \rangle \cdot \llbracket C \rrbracket_r & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases}$$

$$(\text{if } p.e \text{ then } C_1 \text{ else } C_2; C) \rrbracket_r = \llbracket C_1 \rrbracket_r \cdot \llbracket C \rrbracket_r \quad (\eta; C) \rrbracket_r = \llbracket C \rrbracket_r$$

Figure 16: Core Choreographies, projection of messages in transit.

- \leq^- is structural precongurence without rule $\llbracket C \rrbracket \text{Unfold}$;
- each η_i is an instantiated receive action of the form $\bullet_p \xrightarrow{q} v$ or $\bullet_p \xrightarrow{\ell} q[\ell]$;
- C^{CC} is an CC choreography (i.e., a choreography without runtime terms).

Well-formedness is decidable, since the set of choreographies equivalent up to \leq^- 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. Note that the problematic choreography from Example 4 is not well-formed. More generally, we can use well-formedness in the remainder of this section to reason about EPP.

Definition 7 (Asynchronous EPP from aCC to aSP). *Let C be a well-formed aCC choreography and σ be a state. Without loss of generality, we assume that C does not contain $p.e \xrightarrow{x} \bullet_q$ actions.² The EPP of C and σ is defined as*

$$\llbracket C, \sigma \rrbracket = \prod_{p \in \text{pn}(C)} p \triangleright_{\sigma(p)}^{\llbracket C \rrbracket_p} \llbracket C \rrbracket_p$$

where $\llbracket C \rrbracket_p$ is defined as in Figure 7 with the extra rule in Figure 15 and $\llbracket C \rrbracket_p$ is defined by the rules in Figure 16.

In the last case of the definition of $\llbracket C \rrbracket_p$ (bottom-right), η ranges 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. We could alternatively define projection only for well-formed choreographies in the “canonical form” implicit in the definition of well-formedness.

With this definition, we can state and prove an asynchronous variant of Theorem 2.

Theorem 7 (Asynchronous EPP Theorem). *If C is a projectable and well-formed CC choreography, then, for all states σ :*

- (Completeness) if $C, \sigma \rightarrow C', \sigma'$ for some choreography C' and state σ' , then $\llbracket C, \sigma \rrbracket \rightarrow \succ \llbracket C', \sigma' \rrbracket$;
- (Soundness) if $\llbracket C, \sigma \rrbracket \rightarrow N$ for some network N , then there exist a choreography C' and a state σ' such that $C, \sigma \rightarrow C', \sigma'$ and $\llbracket C', \sigma' \rrbracket \prec N$.

²By well-formedness, we can always rewrite C to an equivalent choreography satisfying this condition. Such a choreography is also guaranteed not to contain $\bullet_p \xrightarrow{x} q$ or $p \xrightarrow{x} q$ actions.

As a consequence, Corollary 1 applies also to the asynchronous case: the processes projected from CC into aSP are deadlock-free, even when they contain runtime terms.

As usual, the hypotheses in Theorem 7 are not necessary, and this result also holds for some choreographies that are not well-formed. For example, the network in (1) is the projection of

$$p.e_1 \xrightarrow{x} \bullet_q; q.e_2 \xrightarrow{y} \bullet_r; r.e_3 \xrightarrow{z} \bullet_p; \bullet_p \xrightarrow{r} x; \bullet_q \xrightarrow{p} y; \bullet_r \xrightarrow{q} z;$$

with an adequately defined state σ , but this choreography is not structurally precongruent to any choreography obtained from execution of an CC choreography.

5. Asynchrony by Encoding

We now show an alternative implementation of asynchrony, by encoding communications in a synchronous semantics in another extension of CC. We take inspiration from how asynchrony is modelled in foundational process models, specifically the π -calculus [25]. The key idea there is to use processes to represent messages in transit, allowing the sender to proceed immediately after having sent a message without having to synchronise with the receiver [30]. In an asynchronous system, there is no bound to the number of messages that could be transiting in the network; this means that CC is not powerful enough for our purposes, because it can only capture a finite number of processes. For this reason, we extend CC with two standard notions, borrowed from process calculi and previous choreography models: process spawning – the ability to create new processes at runtime – and name mobility – the ability to send process references, or names. We call this extended language Dynamic Core Choreographies (DCC).

For simplicity, we first present our development for the computational fragments of CC and DCC, called Minimal Choreographies (MC) and Dynamic Minimal Choreographies (DMC), respectively. The relationships between the different choreography calculi discussed in this section are summarized in Figure 17. Strict inclusions between calculi are denoted as \hookrightarrow ; in particular, MC is a strict sublanguage of DMC and CC is a strict sublanguage of DCC.

The dotted arrow (1) in Figure 17 is the cornerstone of the development in this section: every choreography in MC can be encoded in an asynchronous implementation in DMC, by using auxiliary processes to represent messages in transit. Since DMC extends MC with new primitives, it makes sense to extend this encoding to the whole language of DMC (2). Due to the syntactic minimality of MC, our results imply that asynchronous communications can be modelled in choreographies using well-known notions, i.e., process spawning and name mobility (studied, e.g., in [5, 11]), without the need for *ad-hoc* constructions.

The fact that our encoding can be extended from MC to DMC is evidence that our approach is robust, and the simplicity of DMC makes it a convenient foundational calculus to use in future developments of choreographies. However, one of the expected advantages of using a foundational theory such as DMC for capturing asynchrony is indeed that we can reuse existing formal techniques based on standard primitives for choreographies. (This is a common scenario in π -calculus, where many techniques apply to its sublanguages [30].) An example of such reuses is the extension of our results to CC. An important property of CC, shown in [9], is that selections can be encoded as value communications preserving projectability, yielding a map from CC to MC – the dashed arrow (a) in Figure 17. What happens if we consider DCC, which extends DMC with label selections? Ideally, this calculus should *both* have an asynchronous interpretation through the techniques introduced in this section, and still possess the property that selections are encodable in the simpler language DMC. This is indeed the case. We extend our encoding to yield an interpretation of asynchronous selections, yielding arrows (3) and (4) in Figure 17. Encodability of selections in DMC follows immediately from the fact that CC is a sublanguage of DCC, giving us (b) for free.

5.1. Dynamic Core Choreographies

We begin by describing the mapping (1) in Figure 17. This mapping $\{\cdot\} : MC \rightarrow DMC$ transforms every communication action $p.e \rightarrow q \in C \in MC$ into a send/receive pair in $\{\{C\}\} \in DMC$, with the properties that: p can continue executing without waiting for q to receive its message (and even send further messages to




Figure 17: Choreography calculi and encodings.

$$\begin{aligned}
C &::= \dots \mid \text{def } X(\tilde{p}) = C_2 \text{ in } C_1 \mid X\langle\tilde{p}\rangle \\
\eta &::= \dots \mid \mathbf{p} \text{ start } \mathbf{q}^T \quad e ::= \mathbf{p} \mid \dots
\end{aligned}$$

Figure 18: Dynamic Core Choreographies, syntax (new terms).

\mathbf{q}); and messages from \mathbf{p} to \mathbf{q} are delivered in the same order as they were originally sent. These properties essentially reflect the requirements concerning value communications in Definition 4.

In MC, there is a bound on the number of values that can be stored at any given time by the system: since each process can hold a single value, the maximum number of values the system can know is equal to the number of processes in the choreography, which is fixed. However, in an asynchronous setting, the number of values that need to be stored is unbounded: a process \mathbf{p} may loop forever sending values to \mathbf{q} , and \mathbf{q} may wait an arbitrary long time before receiving any of them. Therefore, we need to extend MC with the capability to generate new processes. As discussed in [11], this requires enriching the language with two additional abilities: parameters to recursive procedures (in order to be able to use a potentially unbounded number of processes at the same time) and a new action to communicate process names.

Since the extensions from MC to DMC and from CC to DCC are similar, and we later extend our results to CC/DCC (mapping (4) in Figure 17, we first describe the full calculus DCC, which extends CC with process spawning and name mobility.

Formally, the differences between the syntaxes of CC and DCC are highlighted in Figure 18: procedure definitions and calls now have parameters; there is a new term for generating processes; and the expressions sent by processes can also be process names. The possibility of communicating a process name ($\mathbf{p}.\mathbf{q} \rightarrow r$) ensures name mobility. We use the abbreviation $\mathbf{p} : r \leftrightarrow \mathbf{q}$ as shorthand for $\mathbf{p}.\mathbf{q} \rightarrow r$; $\mathbf{p}.r \rightarrow \mathbf{q}$.

The semantics for DCC includes an additional ingredient, borrowed from [11]: a graph of connections G , keeping track of which pairs of processes are allowed to communicate. This graph is directed, and an edge from \mathbf{p} to \mathbf{q} in G (written $\mathbf{p} \xrightarrow{G} \mathbf{q}$) means that \mathbf{p} knows the name of \mathbf{q} . In order for an actual message to flow between \mathbf{p} and \mathbf{q} , both processes need to know each other, which we write as $\mathbf{p} \xleftrightarrow{G} \mathbf{q}$.³ The reduction relation now has the form $G, C, \sigma \rightarrow G', C', \sigma'$, where G and G' are the connection graphs before and after executing C , respectively. The complete rules are given in Figure 19, with \preceq defined as in CC.⁴ In rule $\lfloor C \rfloor \text{Start}$, the fresh process \mathbf{q} is assigned a default value \perp .

The proof of Theorem 1 can be generalised to DCC, but this requires an extra ingredient: a simple type system (which we do not detail, as it is a subsystem of that presented in [11]). This type system checks that all processes that attempt at communicating are connected in the communication graph, e.g., by being properly introduced using name mobility. Furthermore, we can define a target process calculus for

³In some process calculi, the weaker condition $\mathbf{p} \xrightarrow{G} \mathbf{q}$ is typically sufficient for \mathbf{p} to send a message to \mathbf{q} . Our condition is equivalent to that found in the standard model of Multiparty Session Types [22]. This choice is orthogonal to our development.

⁴The function $\text{pn}(\eta)$ is extended to communications of process names by the clause $\text{pn}(\mathbf{p}.r \rightarrow \mathbf{q}) = \{\mathbf{p}, \mathbf{q}, r\}$.

$$\begin{array}{c}
\frac{p \xleftarrow{G} q \quad \downarrow_{e[\sigma(p)/c]}^v}{G, p.e \rightarrow q; C, \sigma \rightarrow G, C, \text{upd}(\sigma, q, v)} \text{ [C|Com]} \quad \frac{p \xleftarrow{G} q}{G, p \rightarrow q[\ell]; C, \sigma \rightarrow G, C, \sigma} \text{ [C|Sel]} \\
\frac{}{G, p \text{ start } q^T; C, \sigma \rightarrow \text{upd}(p, q), C, \text{upd}(\sigma, q, \perp)} \text{ [C|Start]} \\
\frac{p \xleftarrow{G} q \quad p \xrightarrow{G} r}{G, p.r \rightarrow q; C, \sigma \rightarrow G \cup \{q \rightarrow r\}, C, \sigma} \text{ [C|Intro]} \quad \frac{i = 1 \text{ if } \sigma(p) = \sigma(q), \quad i = 2 \text{ otherwise}}{G, \text{if } p.e \text{ then } C_1 \text{ else } C_2, \sigma \rightarrow G, C_i, \sigma} \text{ [C|Cond]} \\
\frac{C_1 \preceq C_2 \quad G, C_2, \sigma \rightarrow G', C'_2, \sigma' \quad C'_2 \preceq C'_1}{G, C_1, \sigma \rightarrow G', C'_1, \sigma'} \text{ [C|Struct]}
\end{array}$$

Figure 19: Dynamic Core Choreographies, semantics.

DCC and an EndPoint Projection that will automatically synthesise correct-by-construction deadlock-free implementations of (projectable) choreographies, using techniques from [11]. Although these constructions are not technically challenging, we omit them at this point, since they are immaterial for our results. In Section 5.4 we briefly explain how the EndPoint Projection in [11] applies to DCC, and how the results related to it can be combined with our development.

The fragment of DCC that does not contain label selections is called Dynamic Minimal Choreographies (DMC). Amendment and label selection elimination, described in [9], hold for DMC and DCC just as for MC and CC, so that, for any DMC choreography C , we have the following results.

Lemma 1 (Amendment Lemma). *There is a function $\text{Amend} : CC \rightarrow CC$ such that, for every choreography C :*

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

Projectability: for every state σ , $\llbracket \text{Amend}(C), \sigma \rrbracket$ is defined;

Correspondence: for every other choreography C' and states σ and σ' , $C, \sigma \rightarrow^* C', \sigma'$ iff $\text{Amend}(C), \sigma \rightarrow^* \text{Amend}(C'), \sigma'$.

Eliminating selections can be done for any CC choreography; however, the interesting application (which also suffices for our development) is when it is applied to choreographies obtained by amendment.

Theorem 8 (Selection elimination). *There is a function $\llbracket \cdot \rrbracket : CC \rightarrow MC$ such that, for every choreography $C \in MC$, $\llbracket \text{Amend}(C) \rrbracket$ is projectable.*

The function $\llbracket \cdot \rrbracket$ essentially replaces all label selections by a pair value communication/conditional. Intuitively, instead of p sending a label to q from a possible set of labels, it sends this label as a value, and q tests the value received in order to decide on its behaviour. To adapt the construction in [9] to the variant of MC we consider in this work, the translation must also change the type of the value stored at each process p from T_p to $T_p \times L$, where L is the set of possible labels; the state update function $\text{upd}(t, h)$ changes either component of the pair depending on the type of the value communicated.⁵ The following lemmas state that selection elimination does not change the semantics of the original choreography in an essential way. In these lemmas, we write σ^+ for a state obtained from σ by adding some label to the value stored at each process.

Lemma 2. *Choreographies C and $\llbracket \text{Amend}(C) \rrbracket^+$ are equivalent wrt $\text{pn}(C)$.*

⁵In the context of the original work [9] there were only two possible labels, making the encoding slightly simpler.

$$\begin{aligned}
\{\{p.e \rightarrow q; C\}\}_M &= p.e \rightarrow pq^M; p \text{ start } pq^{M+}; p : pq^M \leftrightarrow pq^{M+}; \\
&\quad pq^M.q \rightarrow pq^{M+}; pq^M.pq^{M+} \rightarrow q; pq^M.c \rightarrow q; \\
&\quad \{\{C\}\}_{M[(p,q) \mapsto M(p,q)+1]} \\
\{\{ \text{if } p.e \text{ then } C_1 \text{ else } C_2 \}\}_M &= \text{if } p.e \text{ then } \{\{C_1\}\}_M \text{ else } \{\{C_2\}\}_M \\
\{\{0\}\}_M &= 0 \\
\{\{X\}\}_M &= X\langle \bar{M} \rangle \\
\{\{ \text{def } X = C_2 \text{ in } C_1 \}\}_M &= \text{def } X\langle \bar{M}_0 \rangle = \{\{C_2\}\}_{M_0} \text{ in } \{\{C_1\}\}_M
\end{aligned}$$

Figure 20: Encoding MC in DMC

Lemma 3. *Let C and C' be MC choreographies and σ and σ' be states. If $C, \sigma \rightarrow C', \sigma'$, then $(\text{Amend}(C)), \sigma^+ \rightarrow^* (\text{Amend}(C')), \sigma'^+$ for some σ^+ and σ'^+ . Furthermore, the latter reduction consists of only one step except for the case when the former uses rule $\lfloor C \rfloor \text{Cond}$.*

Conversely, if $(\text{Amend}(C)), \sigma^+ \rightarrow C', \sigma'$, then there exist a choreography C'' and a state σ'' such that $C, \sigma \rightarrow C'', \sigma''$ and $C', \sigma' \rightarrow^ (\text{Amend}(C'')), \sigma''^+$ for some σ''^+ . Furthermore, the latter reduction is non-empty only in the case when the former uses rule $\lfloor C \rfloor \text{Cond}$.*

The proofs of all these results are straightforward adaptations of the corresponding proofs in [9]: all proofs either by induction on the structure of the choreography or by case analysis on its possible transitions, and the new cases that arise due to the presence of new constructors are trivially dealt with, since both Amend and (\cdot) treat these constructors homomorphically.

5.2. Asynchronously encoding MC in DMC

We focus now on the mapping (1) from Figure 17, as this is the key ingredient to establish the remaining connections in that figure. Let C be a choreography in MC. In order to encode C in DMC, we use a function $M_C : \mathcal{P}^2 \rightarrow \mathbb{N}$, where $\mathcal{P} = \text{pn}(C)$ is the set of process names in C . We omit the subscript C when the choreography is clear. The countable set of auxiliary processes $\{pq^i \mid p, q \in \mathcal{P}, i \in \mathbb{N}\}$ is used to store messages in transit, with pq^i holding the i th message from p to q .

First, we setup initial channels for communications between all processes occurring in C .

$$\{\{C\}\} = \{p \text{ start } pq^0; p : q \leftrightarrow pq^0\}_{p,q \in \mathcal{P}, p \neq q}; \{\{C\}\}_{M_0}$$

Here, $M_0(p, q) = 0$ for all p and q . For simplicity, we write pq^M for $pq^{M(p,q)}$ and pq^{M+} for $pq^{M(p,q)+1}$. The definition of $\{\{C\}\}_M$ is given in Figure 20. We write \bar{M} for $\{pq^M \mid p, q \in \mathcal{P}, p \neq q\}$, where we assume that the order of the values of M is fixed. In recursive definitions, we reset M to M_0 ; note that the parameter declarations act as binders, so these process names are still fresh. We can use α -renaming on $\{\{C\}\}$ to make all bound names distinct.

In order to encode $p.e \rightarrow q$, p uses the auxiliary process pq^M to store the value it wants to send to q . Then, p creates a fresh process (to use in the next communication) and sends its name to pq^M . Afterwards, p is free to proceed with execution. In turn, pq^M communicates q 's name to the new process, which now is ready to receive the next message from p . Finally, pq^M waits for q to be ready to receive both the value being communicated and the name of the process that will store the next value.

The encoding from C to $\{\{C\}\}$ does not constitute an asynchronous semantics for MC, since it expands communication actions syntactically, rather than semantically. We support our claim that $\{\{C\}\}$ is an asynchronous encoding of C by showing that, nevertheless, the two properties in Definition 5 hold. In these lemmas, we abuse notation and, given a state σ , we also write σ to denote an unspecified extension of σ to the auxiliary processes introduced by the encoding.⁶

⁶This imprecision significantly simplifies the statements and proofs, which do not depend on the precise values assigned to those spurious states.

Theorem 9. *Let C be an MC choreography and σ be a state, and assume that $C, \sigma \rightarrow C', \sigma'$ for some C' and σ' . For each labeling function M , let G_M be the graph such that $\mathbf{p} \xleftrightarrow{G} \mathbf{q}$, $\mathbf{p} \xleftrightarrow{G} \mathbf{pq}^M$ and $\mathbf{pq}^M \xleftrightarrow{G} \mathbf{q}$ for all $\mathbf{p}, \mathbf{q} \in \mathbf{pn}(C)$. Then for every labeling function M there exists M' such that $G_M, \{\!\{C}\!\}_M, \sigma \rightarrow^* G', \{\!\{C'\}\!\}_{M'}, \sigma'$ for some graph G' containing $G_{M'}$.*

Proof. By induction on the derivation of the transition $C, \sigma \rightarrow C', \sigma'$.

- If the transition is justified directly by rule $\lfloor \text{C|Com} \rfloor$, then C is of the form $\mathbf{p}.e \rightarrow \mathbf{q}; C'$, $e \downarrow_{\sigma}^{\mathbf{p}} v$ and $\sigma' = \text{upd}(\sigma, \mathbf{q}, v)$. In this case,

$$\begin{aligned} \{\!\{C}\!\}_M &= \{\!\{\mathbf{p}.e \rightarrow \mathbf{q}; C\}\!\}_M = \mathbf{p} \text{ start } \mathbf{pq}^{M+}; \mathbf{p} : \mathbf{pq}^M \leftarrow \mathbf{pq}^{M+}; \\ &\quad \mathbf{pq}^M . \mathbf{q} \rightarrow \mathbf{pq}^{M+}; \mathbf{pq}^M . \mathbf{pq}^{M+} \rightarrow \mathbf{q}; \mathbf{pq}^M . c \rightarrow \mathbf{q}; \\ &\quad \{\!\{C'\}\!\}_{M[(\mathbf{p}, \mathbf{q}) \mapsto M(\mathbf{p}, \mathbf{q})+1]} \end{aligned}$$

and therefore $G_M, \{\!\{C}\!\}_M, \sigma$ reduces in six steps to $G', \{\!\{C'\}\!\}_{M[(\mathbf{p}, \mathbf{q}) \mapsto M(\mathbf{p}, \mathbf{q})+1]}, \text{upd}(\sigma, \mathbf{p}, v)$, where $G' = G_M \cup \{(\mathbf{p}, \mathbf{pq}^{M+}), (\mathbf{pq}^M, \mathbf{pq}^{M+}), (\mathbf{q}, \mathbf{pq}^{M+})\}$.

- If the transition is justified directly by rule $\lfloor \text{C|Cond} \rfloor$ and $e \downarrow_{\sigma}^{\mathbf{p}} \text{true}$, then C is of the form $\text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2$, $\sigma' = \sigma$, $C' = C_1$, and $\{\!\{C}\!\}_M = \text{if } \mathbf{p}.e \text{ then } \{\!\{C_1}\!\}_M \text{ else } \{\!\{C_2}\!\}_M$. In this case, $G_M, \{\!\{C}\!\}_M, \sigma \rightarrow G_M, \{\!\{C_1}\!\}_M, \sigma$. The case where $e \downarrow_{\sigma}^{\mathbf{p}} \text{false}$ is analogous.

- If the transition is justified by rule $\lfloor \text{C|Ctx} \rfloor$, then C is of the form $\text{def } X = C_2 \text{ in } C_1$, C' is of the form $\text{def } X = C_2 \text{ in } C'_1$ and $C_1, \sigma \rightarrow C'_1, \sigma'$. Furthermore, $\{\!\{C}\!\}_M = \text{def } X(\overline{M}_0) = \{\!\{C_2}\!\}_{M_0} \text{ in } \{\!\{C_1}\!\}_M$. By induction hypothesis, $G_M, \{\!\{C_1}\!\}_M, \sigma \rightarrow^* G', \{\!\{C'_1}\!\}_{M'}, \sigma'$ for some G' and M' satisfying the thesis, from which we deduce that $G_M, \{\!\{\text{def } X = C_2 \text{ in } C_1\}\!\}_M, \sigma \rightarrow G', \text{def } X(\overline{M}_0) = \{\!\{C_2}\!\}_{M_0} \text{ in } \{\!\{C'_1}\!\}_{M'}, \sigma'$, and the thesis follows by observing that $\text{def } X(\overline{M}_0) = \{\!\{C_2}\!\}_{M_0} \text{ in } \{\!\{C'_1}\!\}_{M'} = \{\!\{\text{def } X(\overline{M}_0) = C_2 \text{ in } C'_1\}\!\}_{M'}$.

- If the transition is justified by rule $\lfloor \text{C|Struct} \rfloor$, then there exist C_0 and C'_0 such that $C \preceq C_0$, $C'_0 \preceq C'$ and $C_0, \sigma \rightarrow C'_0, \sigma'$.

By induction hypothesis, $G_M, \{\!\{C_0}\!\}_M, \sigma \rightarrow^* G', \{\!\{C'_0}\!\}_{M'}, \sigma'$ for some G' and M' in the conditions of the thesis. By induction, it is straightforward to prove that $C_1 \preceq C_2$ implies $\{\!\{C_1}\!\}_{M_1} \preceq C_2 M_1$ for any M_1 , from which it follows that $\{\!\{C}\!\}_M \preceq \{\!\{C_0}\!\}_M$ and $\{\!\{C'_0}\!\}_{M'} \preceq \{\!\{C'\}\!\}_{M'}$. By applying rule $\lfloor \text{C|Struct} \rfloor$ to the first and last transitions in the sequence $G_M, \{\!\{C_0}\!\}_M, \sigma \rightarrow^* G', \{\!\{C'_0}\!\}_{M'}, \sigma'$, we conclude that $G_M, \{\!\{C}\!\}_M, \sigma \rightarrow^* G', \{\!\{C'\}\!\}_{M'}, \sigma'$.

□

Since the semantics of DMC is confluent, Theorem 9 states that messages sent from \mathbf{p} to \mathbf{q} are eventually received by \mathbf{q} in every terminating computation.

Theorem 10. *Let C be an MC choreography, G be a connection graph, M be a labeling function and σ be a state, and assume that $G, \{\!\{C}\!\}_M, \sigma \rightarrow^* G', C', \sigma'$ for some graph G' , choreography C' and state σ' . Then there exist a graph G'' , a choreography C'' and a state σ'' such that $C, \sigma \rightarrow^* C'', \sigma''$, and $G', C', \sigma' \rightarrow^* G'', \{\!\{C''}\!\}_{M}, \sigma''$.*

Sketch. This proof is very similar to the proof of Theorem 5, except that in the second property the relevant case is when $G, \{\!\{C}\!\}_M, \sigma$ reduces by application of rule $\lfloor \text{C|Start} \rfloor$. Then C must be of the form $\mathbf{p}.e \rightarrow \mathbf{q}; C_0$, and we can model the reduction $C, \sigma \rightarrow C_0, \sigma_0$ by consuming the five additional communication actions obtained by encoding $\mathbf{p}.e \rightarrow \mathbf{q}$. □

Theorem 10 states that the encoding does not add any additional behaviour to the original choreography, aside from expanding communications into several actions.

Example 5. We again return to the setting of Example 1, and partially show the result of applying this transformation to the first three communication actions. We only include the initializations of the channels that are used in this fragment; the numbers indicated refer to the line numbers in the original choreography.

```

a start as0; a : as0 <-> s;
s start sa0; s : sa0 <-> a;
s start sb0; s : sb0 <-> b;
1. a.title -> as0; a start as1; a : as1 <-> as0;
   as0.as1 -> s; as0.s -> as1; as0.c -> s;
2. s.price -> sa0; s start sa1; s : sa1 <-> sa0;
   sa0.sa1 -> a; sa0.a -> sa1; sa0.c -> a;
3. s.price -> sb0; s start sb1; s : sb1 <-> sb0;
   sb0.sb1 -> b; sb0.b -> sb1; sb0.c -> b;
...

```

The first three lines initialize three channels: from a to s ; from s to a ; and from s to b . Then one message is passed in each of these channels, as dictated by the encoding. All communications are asynchronous in the sense explained above, as in each case the main sender process sends its message to a dedicated intermediary (as^0 , sa^0 or sb^0 , respectively), who will eventually deliver it to the recipient. Moreover, causal dependencies are kept: in Step 2, s can only send its message to sa^0 after receiving the message sent by a in Step 1. However, in Step 3 s can send its message to sb^0 without waiting for a to receive the previous message, as the action $s.price \rightarrow sa^0$ can swap with the three actions immediately preceding it. (This is similar to the behaviour previously observed in Example 3.)

We briefly illustrate Theorems 9 and 10 in this setting. From Theorem 9 we obtain that there is a reduction path of the encoded choreography that passes through the encodings of each reductum of the original choreography. Furthermore, since the original choreography terminates, we also know that, e.g., the action $a.title \rightarrow as^0$ is eventually followed by a communication of $title$ from as^0 to s . Theorem 10 implies that if, e.g., $a.title \rightarrow as^0$ is executed, then it must be “part” of an action in the original choreography (in this case, $a.title \rightarrow s$), and furthermore it is possible to find an execution path that will execute the remaining actions generated from that one (the remaining five actions in Step 1).

5.3. Asynchrony in DMC and DCC

The calculus DMC is in itself synchronous, just like MC. We now show that we can extend $\{\cdot\}$ to the full language of DMC – arrow (2) in Figure 17 – thereby obtaining a systematic way to write asynchronous communications in DMC. By further marking which communications we want to treat as synchronous (so that they are untouched by $\{\cdot\}$) we obtain a calculus in which we can have both synchronous and asynchronous communication, compiled in itself. This is similar (albeit dual) to the situation in asynchronous π -calculus, where we can also encode synchronous communication without extending the language.

The main challenge is dealing with M , as the source choreography can now include process spawning. This means that the domain of M can be dynamically extended throughout the computation of $\{\{C\}\}_M$, which renders our parameter-passing in recursive calls invalid (since the number of parameters in the procedures generated by our encoding is fixed). However, since each procedure $X(p_1, \dots, p_n)$ in DMC can only use (by convention) the processes p_1, \dots, p_n in its body, we can restrict the additional parameters introduced by the encoding to the $n(n-1)$ auxiliary processes currently assigned by M to communications between the p_i s. For example, $\{\{\text{def } X(p, q) = C_2 \text{ in } C_1\}\}_M$ becomes $\text{def } X(p, q, pq^0, qp^0) = \{\{C_2\}\}_{M_0} \text{ in } \{\{C_1\}\}_M$. We assume the convention that these additional arguments are given in lexicographic ordering (where processes are ordered by their position in the original definition or call).

With this in mind, we can easily define the new cases for $\{\{C\}\}_M$.

$$\begin{aligned}
\{\{p \text{ start } q; C\}\}_M &= p \text{ start } q; p \text{ start } pq^0; q \text{ start } qp^0; \\
& p : q \leftrightarrow pq^0; q : p \leftrightarrow qp^0; \\
& \{\{C\}\}_{M[(p,q) \mapsto 0, (q,p) \mapsto 0]} \\
\{\{p.q \rightarrow r; C\}\}_M &= p \text{ start } qr^0; p.qr^0 \rightarrow pq^M; p.qr^0 \rightarrow pr^M; \\
& p \text{ start } pq^{M+}; p : pq^M \leftrightarrow pq^{M+}; \\
& p \text{ start } pr^{M+}; p : pr^M \leftrightarrow pr^{M+}; \\
& pq^M.q \rightarrow pq^{M+}; pr^M.r \rightarrow pr^{M+}; \\
& pq^M.pq^{M+} \rightarrow q; pq^M.qr^0 \rightarrow q; \\
& pr^M.pr^{M+} \rightarrow r; pr^M.qr^0 \rightarrow r; \\
& \{\{C\}\}_{M[(p,q) \mapsto M(p,q)+1, (p,r) \mapsto M(p,r)+1, (q,r) \mapsto 0]}
\end{aligned}$$

In $\{\{p \text{ start } q; C\}\}_M$, we simply create the asynchronous communication channels between p and q – the only step where these process will need to synchronize – and extend M in the continuation. The encoding of $p.q \rightarrow r$ is better understood by reading it as a composition: first, p creates the new asynchronous communication channel from q to r , then uses its own channels to send this name to these processes. Note that the auxiliary channels do not communicate, so this encoding will introduce asymmetries in the graph of communications.

It is easy to verify the new cases of the proofs of Theorems 9 and 10, establishing that they still hold for this extended encoding.

Finally, we extend this encoding to the whole language of DCC – yielding arrow (4) in Figure 17 – by adding the clause

$$\begin{aligned}
\{\{p \rightarrow q[\ell]; C\}\}_M &= p \rightarrow pq^M[\ell]; p \text{ start } pq^{M+}; \\
& p : pq^M \leftrightarrow pq^{M+}; pq^M.q \rightarrow pq^{M+}; \\
& pq^M.pq^{M+} \rightarrow q; pq^M \rightarrow q[\ell]; \\
& \{\{C\}\}_{M[(p,q) \mapsto M(p,q)+1]}
\end{aligned}$$

to the definition of $\{\{C\}\}_M$. Restricting this encoding to the language of CC yields arrow (3) in Figure 17.

Example 6. *Again in the setting of Example 1, we now assume that communications with b are supposed to be asynchronous, while the remaining ones are synchronous. The corresponding implementation would*

thus leave the communications in Lines 1, 2 and 6 unaffected, and only transform the remaining ones.

```

a start ab0; a : ab0 <-> b;
b start ba0; b : ba0 <-> a;
b start bs0; b : bs0 <-> s;
s start sb0; s : sb0 <-> b;
1. a.title -> s
2. s.price -> a
3. s.price -> sb0; s start sb1; s : sb1 <-> sb0;
   sb0.sb1 -> b; sb0.b -> sb1; sb0.c -> b;
4. a.c -> ab0; a start ab1; a : ab1 <-> ab0;
   ab0.ab1 -> b; ab0.b -> ab1;
   if b.happy then
5.     s.price -> sb0; s start sb1; s : sb1 <-> sb0;
       sb0.sb1 -> b; sb0.b -> sb1; sb0.c -> b;
...

```

5.4. Projectability

In general, DCC choreographies may deadlock because of attempted communications between processes that are not joined in the connection graph. We have not discussed this aspect yet, as it is orthogonal to the results in this section: Theorems 9 and 10 guarantee that a choreography deadlocks if and only if the choreography obtained by encoding its communications asynchronously deadlocks.

In order to regain properties such as Theorem 1 for DCC and DMC, we need to have some syntactic criteria to identify deadlock-free choreographies. A standard way to do this in choreography models is by means of a type system. We omit a formal discussion of a type system for DCC, since it can be constructed by restricting the type system for the choreography model of [11] to the constructors present in this language. From the results in that work, this type system enjoys type inference and decidability of type checking. Typable choreographies are deadlock-free. Furthermore, we can define an EndPoint Projection to a process calculus of Dynamic Stateful Processes (DSP) such that the correspondence in Theorem 2 holds. As a consequence, projections of typable choreographies are again deadlock-free by construction.

More interestingly, our encoding for asynchronous communications is designed to ensure that typability and projectability are preserved, i.e., if C is typable/projectable, then $\{\{C\}\}$ is also typable/projectable. As a result, the correspondences described in Theorems 9 and 10 for the different choreography models also hold at the process level.

5.5. Overview of correspondences

We now summarize the operational correspondences between the main calculi in this article (see Figure 21). Similar relationships exist between the restricted calculi without label selections, as they are all subcalculi of the languages included here.

The solid arrows in Figure 21 refer to the one-to-one correspondence described by the EPP theorems for each of the calculi (stated in Theorem 2 for CC and SP, in Theorem 7 for aCC and aSP, and discussed informally in the previous section, for DCC and DSP).

The dashed arrows between choreography languages express the expansion simulation obtained by either endowing CC with asynchronous semantics (arrow from CC to aCC, formalized as Theorem 5, first property) or by encoding communications asynchronously (arrow from CC to DCC, Theorem 9). The arrow from aCC to DCC can be obtained by extending the encoding from CC to DCC to runtime terms in aCC, using ideas similar to those in Definition 7: messages in transit need to be stored at the right auxiliary processes, and




Figure 21: The diagram on the right summarizes the simulation relationships between the different choreography and process calculi: one-to-one (solid), one-to-many (dashed) or many-to-one (dotted). The correspondences marked (*) also hold for arbitrary networks, as explained in the main text. For reference, the diagram on the left indicates the choreography/process language each term lives in.

some assumptions need to be made regarding the connection graphs. We do not formalize this translation, as it is a tedious but relatively straightforward extension of our development.

The correspondence between SP and aSP was stated explicitly as Theorem 6. In order to obtain similar results for DSP, we need to define translations from SP/aSP networks into DSP networks. Again, this is a simple extension of our constructions, since we can use the EPP to infer how these translations should be defined. The results in this paper imply that these translations give expansion simulations for networks in the fragments of SP and aSP obtained by projecting choreographies in CC or aCC, but the correspondences can be shown to hold for the entire language in both cases.

The last correspondences are the dotted arrows, which state many-to-one correspondences: if, for example, a choreography in aCC makes a reduction, then the reductum may need to make some more reductions in order to reach a reductum of the original choreography in CC (property (2) in the proof of Theorem 5). The corresponding relationship between DCC and CC is also stated as part of the proof of Theorem 10. The dotted arrow from DCC to aCC is similar, and states that the asynchronous encoding is more fine-grained than the asynchronous semantics: this is not surprising, since a communication action in CC becomes two actions in aCC, but six actions in DCC.

These correspondences do not propagate automatically to the process level, since the EPP theorems only guarantee that similar correspondences hold for projections of choreographies (see the discussion after Theorem 6: a deadlocked SP choreography may reduce under asynchronous semantics). Therefore, we can only state them for networks that are obtained as choreography projections. The exception is the correspondence between aSP and DSP: since both these calculi are inherently asynchronous, it is still the case that encoding a network from aSP into DSP by making its communications pass through intermediate auxiliary processes does not add new possible reductions.

6. Related Work

Asynchronous semantics for choreographies. The first work that introduced an asynchronous semantics to choreographies is [5], as described in the introduction. The same approach was later adapted to compositional choreographies [27] and multiparty session types [22]. The models presented in these works all suffer from the shortcoming discussed in the introduction: choreographies can reduce to states that would normally not be reachable in the real world.

We can now give a formal example of the consequences of this discrepancy.

Example 7. Let C be the choreography from the introduction: $C \triangleq p.1 \rightarrow q; p.2 \rightarrow r$. If we adopted the

asynchronous semantics from [5], then the reduction

$$C, \sigma \rightarrow \mathbf{p}.1 \rightarrow \mathbf{q}, \text{upd}(\sigma, r, 2)$$

would be possible for any σ . This shows that Theorem 7 (the completeness direction) does not hold in this semantics, since $\llbracket C, \sigma \rrbracket$ cannot reduce to any N such that $\llbracket \mathbf{p}.1 \rightarrow \mathbf{q}, \text{upd}(\sigma, r, 2) \rrbracket \preceq N$ (we have to consume the first output by \mathbf{p} first).

Therefore, if we want to formalise the correspondence between a choreography and its EPP in this setting, we need to consider multiple steps. In this example, specifically, we can observe that

$$\mathbf{p}.1 \rightarrow \mathbf{q}, \text{upd}(\sigma, r, 2) \rightarrow \mathbf{0}, \text{upd}(\text{upd}(\sigma, r, 2), \mathbf{q}, 1)$$

and that $\llbracket C, \sigma \rrbracket \rightarrow^* \llbracket \mathbf{0}, \text{upd}(\text{upd}(\sigma, r, 2), \mathbf{q}, 1) \rrbracket$.

In general, the EPP Theorem with this kind of asynchronous semantics is weaker and more complicated. We report it here by adapting the formulation from [26, Chapter 2] (the full version of [5]) to our notation:

Theorem 11. *If C is a projectable choreography, then, for all σ :*

- (Completeness) *if $C, \sigma \rightarrow C', \sigma'$, then $C', \sigma' \rightarrow^* C'', \sigma''$ for some C'', σ'' and $\llbracket C, \sigma \rrbracket \rightarrow^* \llbracket C'', \sigma'' \rrbracket$;*
- (Soundness) *if $\llbracket C, \sigma \rrbracket \rightarrow^* N$, then $N \rightarrow^* N'$ for some N' , and $C, \sigma \rightarrow^* C', \sigma'$ for some σ' and C' such that $\llbracket C', \sigma' \rrbracket \preceq N'$.*

Compared to our Theorem 7 above, Theorem 11 is missing the step-by-step correspondence between choreographies and their projections, and therefore does not say much about the intermediate steps.

In [17], multiparty session types are equipped with runtime terms that represent messages in transit in asynchronous communications, similarly to our approach. Differently from our model, the semantics in [17] uses a labelled transition system (instead of out-of-order execution) to identify when a communication can be executed asynchronously. The difference between these two systems seems to be mostly a matter of presentation, but their expressive powers cannot be directly compared because a counterpart to Theorem 5 is not provided for the model of [17].

A more recent development that is nearer to ours is the model of Applied Choreographies [18], where out-of-order execution (modelled as structural equivalences) is used to swap independent partial choreographic actions, which contain terms that are similar to our runtime asynchronous terms. However, the asynchronous terms in Applied Choreographies are not enough to ensure that the semantics is sound; therefore, these choreographies also need to include queues to store messages in transit. This makes their development substantially more complicated than the one we propose. Furthermore, the work in [18] focuses on implementation models, and does not provide a formal definition of asynchrony for choreographies as in this paper. There is also no correspondence result relating Applied Choreographies to a standard synchronous choreography semantics, although we conjecture that it is possible to map a fragment of that language into an asynchronous extension of CC in the sense of our definition.

Our approach relies on out-of-order execution for choreographic interactions, which was first introduced in [5] to capture parallel execution. We extended it to support the swapping of interactions supported by asynchronous message passing. Out-of-order execution does not introduce any burden on the programmer, since only non-interfering interactions can be swapped and all swappings are thus safe by design (more precisely, swaps correspond to the parallel semantics of typical process calculi [26]); instead, they just model different safe ways of executing the same choreography.

Encoding asynchrony in choreographies. To the best of our knowledge, this is the first work that presenting an interpretation of asynchronous communications in choreographies based solely on the expressive power of primitives for the creation of processes and their connections, via name mobility.

Our development in Section 5 recalls the development of the asynchronous π -calculus [21] ($A\pi$ for short, using the terminology from [30]). $A\pi$ has a synchronous semantics, in the sense that two processes can communicate when they are both ready to, respectively, perform compatible input and output actions.

However, an output action can have no sequential continuation, but can instead only be composed in parallel with other behaviour. Thus, the interpretation of communications in $A\pi$ is asynchronous, since outputs can be seen as messages in transit over a network. The synchronisation between (the process holding) a message in transit and the intended receiver models then the extraction of the message from the medium by the receiver. Differently from our work, $A\pi$ is obtained from the standard π -calculus by *restricting* the syntax of processes such that all communications necessarily conform to this asynchronous interpretation. It is then shown that $A\pi$ is expressive enough to encode the synchronous communications from standard π -calculus, by using acknowledgement messages. DMC and DCC exhibit the dual behaviour: communications are naturally synchronous, but we can always encode them to be asynchronous by passing them through intermediary processes.

Process spawning and name mobility are the key additions to DCC and DMC, from CC and MC, that yield the expressive power to represent asynchronous communications. Process spawning in choreographies has been studied also in the works [4, 5, 27], but in a different form where processes have to synchronise over a shared channel to proceed. Name mobility in choreographies was introduced in [5], but for channel rather than process names. Our process spawning and name mobility primitives are simplifications of those presented in [11], which makes all results from that work applicable to DCC (and thus DMC).

7. Conclusions

This article discusses asynchrony in choreographic programming. Within the context of the minimal choreography language CC, we first identified abstractly the properties that an asynchronous semantics should have. Then, we extended CC with runtime terms that store messages in transit in order to endow the resulting calculus aCC with a semantics that enjoys these properties (Theorem 4) and is equivalent to the original semantics for CC in a precisely defined sense (Theorem 5). We showed how the EndPoint projection for CC can be extended to aCC, yielding terms in the same target process calculus SP, and that this extension still enjoys a correspondence theorem (Theorem 7) with respect to an asynchronous semantics for SP that uses queues of incoming messages at each process. In comparison with the traditional approach to asynchrony in choreographies [5], the correspondence guaranteed by the EPP theorem is more precise in our case, and the proof of this theorem is significantly simpler.

Our construction was deliberately made in a modular way: it changes only the rules for communications, leaving the rest untouched, so that it can be easily applied to more expressive languages. In particular, our semantics can be orthogonally applied to the models in [4, 5, 7, 8, 16, 24, 27], since it only changes the behaviour of communication actions. Furthermore, other types of communication (e.g. name mobility [5, 11]) can be treated in a similar way.

A second approach to asynchrony relies on encoding it syntactically, rather than semantically, by relying on auxiliary processes to store messages in transit. We showed that CC is not expressive enough for this purpose, given that the number of messages in transit is in general unbounded. Therefore, we extended CC with a primitive for spawning new processes at runtime. Extending CC with only this primitive would yield a minimal calculus for our purposes, but in order to obtain a realisable model we also argued that we need a primitive for communicating process names along the network. The resulting calculus DCC is essentially a subcalculus of the choreography model from [11]. We showed that we could simulate asynchronous communications by having senders spawn a new process for each message, storing the message until the receiver is ready to consume it, and thus unblocking the sender for further actions. We showed that this encoding also does not fundamentally change the behaviour of the original choreography (Theorems 9 and 10). Furthermore, it can be extended to process spawning and name mobility to an encoding of the whole language of DCC in itself.

As before, this construction can be applied in any calculus that contains the primitives in DCC, such as the language of Procedural Choreographies [11].

The extension of CC to DCC consists of two primitives: one for spawning new processes, such that we can have an unbounded number of processes at runtime, and another for name mobility, such that the spawned processes can be discovered. Intuitively, this is a minimal extension, in the sense that we strictly

need both primitives to simulate asynchronous behaviour in CC. In particular, spawning new processes is necessary for concurrency reasons. To understand why, consider the choreography snippet $p.e \rightarrow q; p.e' \rightarrow q$. In an asynchronous setting, there is no order restriction between p sending e' and q receiving e . If we relied on a single process (acting as a channel) to relay messages from p to q , then this process would need the ability to offer both those capabilities after receiving message e from p . This requires changing the semantics of both choreographies and processes to allow for swapping communication actions performed by the same process, as in [14]. Such a change, however, is much more intrusive than adding a new primitive (process spawn), because it alters fundamentally the semantics of communications themselves. As a result, this alternative development would not be modular and would not be readily applicable to other existing choreography models.

Instead of out-of-order execution, some choreography models also include explicit parallel composition, e.g., $C \mid C'$ [4, 29]. Most behaviours of $C \mid C'$ are captured by out-of-order execution, for example $p.e \rightarrow q \mid r.e' \rightarrow s$ is equivalent to $p.e \rightarrow q; r.e' \rightarrow s$ in CC (due to rule [C|Eta-Eta] in Figure 2) – see [5] for a deeper discussion. Generalising our construction to supporting also an explicit parallel operator is straightforward, since structural precongruence can be extended homomorphically without surprises.

A more interesting extension is multicom [7], a primitive that considerably extends the expressivity of choreographies by allowing for general criss-cross communications. A prime example is the asynchronous exchange $\{p.e \rightarrow q, q.e' \rightarrow p\}$, which allows two processes to exchange message without waiting for each other (this is the building block, e.g., of the alternating 2-bit protocol given in [7]). Multicom crucially depend on asynchrony, making them an interesting case study in our context. Our development applies immediately to multicom, since we can just apply the expansion rule [C|Com-Unfold] to all communications in a multicom simultaneously.

Acknowledgements. This work was partially supported by CRC (Choreographies for Reliable and efficient Communication software), grant no. DFF-4005-00304 from the Danish Council for Independent Research, the Open Data Framework project at the University of Southern Denmark, and by the Independent Research Fund Denmark, Natural Sciences, grant DFF-7014-00041.

References

- [1] Elvira Albert and Ivan Lanese, editors. *Formal Techniques for Distributed Objects, Components, and Systems – 36th IFIP WG 6.1 International Conference, FORTE 2016*, volume 9688 of *LNCS*. Springer, 2016.
- [2] Lorenzo Bettini, Mario Coppo, Loris D’Antoni, Marco De Luca, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. Global progress in dynamically interleaved multiparty sessions. In *CONCUR*, volume 5201 of *LNCS*, pages 418–433. Springer, 2008.
- [3] Business Process Model and Notation. <http://www.omg.org/spec/BPMN/2.0/>.
- [4] Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.*, 34(2):8, 2012.
- [5] Marco Carbone and Fabrizio Montesi. Deadlock-freedom-by-design: multiparty asynchronous global programming. In *POPL*, pages 263–274. ACM, 2013.
- [6] Chor. Programming Language. <http://www.chor-lang.org/>.
- [7] Luís Cruz-Filipe, Kim S. Larsen, and Fabrizio Montesi. The paths to choreography extraction. In Javier Esparza and Andrzej S. Murawski, editors, *FoSSaCS*, volume 10203 of *LNCS*, pages 424–440. Springer, 2017.
- [8] Luís Cruz-Filipe and Fabrizio Montesi. Choreographies in practice. In Albert and Lanese [1], pages 114–123.
- [9] Luís Cruz-Filipe and Fabrizio Montesi. A core model for choreographic programming. In Olga Kouchnarenko and Ramtin Khosravi, editors, *FACS*, volume 10231 of *LNCS*, pages 17–35. Springer, 2017.
- [10] Luís Cruz-Filipe and Fabrizio Montesi. Encoding asynchrony in choreographies. In D. Shin and M. Lencastre, editors, *Proceedings of SAC 2017*, pages 1175–1177, 2017.
- [11] Luís Cruz-Filipe and Fabrizio Montesi. Procedural choreographic programming. In Ahmed Bouajjani and Alexandra Silva, editors, *Proceedings of FORTE 2017*, volume 10321 of *LNCS*, pages 92–107. Springer, 2017.
- [12] Luís Cruz-Filipe and Fabrizio Montesi. That’s enough: Asynchrony with standard choreography primitives. In Carlos Caleiro, Francisco Dionísio, Paula Gouveia, Paulo Mateus, and João Rasga, editors, *Logic and Computation: Essays in Honour of Amílcar Sernadas*, volume 33 of *Tributes*, pages 125–142. College Publications, 2017.
- [13] Luís Cruz-Filipe and Fabrizio Montesi. On asynchrony and choreographies. In *Proceedings of ICE 2017*, accepted for publication.
- [14] Luís Cruz-Filipe, Fabrizio Montesi, and Marco Peressotti. Communications in choreographies, revisited. In *Proceedings of SAC 2018*, 2018. Accepted for publication.
- [15] Mila Dalla Preda, Maurizio Gabbriellini, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro. Deadlock freedom by construction for distributed adaptive applications. *CoRR*, abs/1407.0970, 2014.

- [16] Mila Dalla Preda, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro. Dynamic choreographies. In *COORDINATION*, LNCS, pages 67–82. Springer, 2015.
- [17] Pierre-Malo Deniérou and Nobuko Yoshida. Multiparty compatibility in communicating automata: Characterisation and synthesis of global session types. In *ICALP (2)*, pages 174–186, 2013.
- [18] Maurizio Gabbrielli, Saverio Giallorenzo, and Fabrizio Montesi. Applied choreographies. *CoRR*, abs/1510.03637, 2015.
- [19] Daniele Gorla. Towards a unified approach to encodability and separation results for process calculi. *Inf. Comput.*, 208(9):1031–1053, 2010.
- [20] Kohei Honda, Aybek Mukhamedov, Gary Brown, Tzu-Chun Chen, and Nobuko Yoshida. Scribbling interactions with a formal foundation. In Raja Natarajan and Adegboyega K. Ojo, editors, *Proceedings of ICDCIT 2011*, volume 6536 of *LNCS*, pages 55–75. Springer, 2011.
- [21] Kohei Honda and Mario Tokoro. An object calculus for asynchronous communication. In Pierre America, editor, *Proceedings of ECOOP'91*, volume 512 of *LNCS*, pages 133–147. Springer, 1991.
- [22] Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty Asynchronous Session Types. *J. ACM*, 63(1):9, 2016.
- [23] Ivan Lanese, Claudio Guidi, Fabrizio Montesi, and Gianluigi Zavattaro. Bridging the gap between interaction- and process-oriented choreographies. In *SEFM*, pages 323–332, 2008.
- [24] Hugo A. López, Flemming Nielson, and Hanne Riis Nielson. Enforcing availability in failure-aware communicating systems. In Albert and Lanese [1], pages 195–211.
- [25] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, I and II. *Information and Computation*, 100(1):1–40,41–77, September 1992.
- [26] Fabrizio Montesi. *Choreographic Programming*. Ph.D. Thesis, IT University of Copenhagen, 2013. http://fabriziomontesi.com/files/choreographic_programming.pdf.
- [27] Fabrizio Montesi and Nobuko Yoshida. Compositional choreographies. In *CONCUR*, volume 8052 of *LNCS*, pages 425–439. Springer, 2013.
- [28] PI4SOA. <http://www.pi4soa.org>, 2008.
- [29] Zongyan Qiu, Xiangpeng Zhao, Chao Cai, and Hongli Yang. Towards the theoretical foundation of choreography. In Carey L. Williamson, Mary Ellen Zurko, Peter F. Patel-Schneider, and Prashant J. Shenoy, editors, *Proceedings of WWW 2007*. ACM, 2007.
- [30] Davide Sangiorgi and David Walker. *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
- [31] Savara. JBoss Community. <http://www.jboss.org/savara/>.
- [32] 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.