On Asynchrony and Choreographies

Luís Cruz-Filipe
Department of Mathematics and Computer Science
University of Southern Denmark
Odense, Denmark
lcfilipe@imada.sdu.dk

Fabrizio Montesi
Department of Mathematics and Computer Science
University of Southern Denmark
Odense, Denmark
fmontesi@imada.sdu.dk

Choreographic Programming is a paradigm for the development of concurrent software, where deadlocks are prevented syntactically. However, choreography languages are typically synchronous, whereas 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 formalise the properties that an asynchronous semantics for choreographies should have: messages can be sent without the intended receiver being ready, and all sent messages are eventually received. We explore how out-of-order execution, used in choreographies for modelling concurrency, can be exploited to endow choreographies with an asynchronous semantics. Our approach satisfies the properties we identified. We show how our development yields a pleasant correspondence with FIFO-based asynchronous messaging, modelled in a process calculus, and discuss how it can be adopted in more complex choreography models.

1 Introduction

Choreographic Programming [18] 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 [4, 8, 21]. Choreographies are used in different settings, including standards [2, 23], languages [5, 14, 20, 22], specification models [3, 4, 16], and design tools [2, 20, 22, 23].

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 $\mathsf{p}.e \texttt{ -> } \mathsf{q}; C$ – 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 $\mathsf{p}.e \texttt{ -> } \mathsf{q}; C$ 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 $\lfloor \text{C}|\text{Async} \rceil$ in [4]). As an example, consider the choreography $\mathsf{p}.1 \texttt{ -> } \mathsf{q}; \mathsf{p}.2 \texttt{ -> } \mathsf{r}$ (where 1 and 2 are just constants). The special rule would allow for consuming the second communication immediately, thus reducing the choreography to $\mathsf{p}.1 \texttt{ -> } \mathsf{q}$. In general, roughly, this rule allows a choreography of the form $\mathsf{p}.e \texttt{ -> } \mathsf{q}; C$ 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 adopted in other works (see Section 5). Unfortunately, it also comes with a serious 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.

M. Bartoletti, L. Bocchi, L. Henrio, S. Knight (Eds.):
10th Interaction and Concurrency Experience (ICE 2017)
EPTCS 261, 2017, pp. 76–90, doi:10.4204/EPTCS.261.8

© L. Cruz-Filipe & F. Montesi
This work is licensed under the
Creative Commons Attribution License.

In [4], 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 5 for a detailed discussion of this aspect. Previous work [11] 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 paper, 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 all sent messages are eventually received – and afterwards formulate them 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. We conclude by discussing how our construction can be systematically extended to more complex choreography models.

Contribution. The contribution of this article is threefold. 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. Finally, we prove a lockstep operational correspondence between choreographies and their process implementations, when asynchronous semantics for both systems are considered.

Structure. We present the representative choreography language in which we develop our work, together with its associated process calculus and EPP, in Section 2. In Section 3, we motivate and introduce the properties we would expect of an asynchronous choreography semantics, and introduce a semantics that satisfies these properties. We show that we can define an asynchronous variant of the target process calculus in Section 4, and extend the definition of EPP towards it, preserving the precise operational correspondence from the synchronous case. We relate our development to other approaches for asynchrony in choreographies in Section 5, before concluding in Section 6 with a discussion on the implications of our work and possible future directions.

2 Minimal Choreographies and Stateful Processes

We review the choreography model of Minimal Choreographies (MC) and its target calculus of Stateful Processes (SP), originally introduced in [8].

2.1 Minimal Choreographies

The language of MC is defined inductively in Figure 1. Processes, ranged over by $\mathsf{p},\mathsf{q},\mathsf{r},\ldots$, are assumed

$$C ::= \mathsf{p}.e \texttt{ -> } \mathsf{q}; C \mid \mathsf{if}\, \mathsf{p}.e\, \mathsf{then}\, C_1\, \mathsf{else}\, C_2 \mid \mathsf{def}\, X = C_2\, \mathsf{in}\, C_1 \mid X \mid \mathbf{0}$$

Figure 1: Minimal Choreographies, syntax.

to run concurrently. Processes interact through value communications $\mathsf{p}.e \texttt{ -> } \mathsf{q}$, which we also denote as η. Here, 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.

The remaining choreography terms denote conditionals, recursion, and termination. In the conditional $\mathsf{if}\, \mathsf{p}.e\, \mathsf{then}\, C_1\, \mathsf{else}\, C_2$, process p evaluates expression e to decide whether the choreography should proceed as C_1 or as C_2. In $\mathsf{def}\, X = C_2\, \mathsf{in}\, C_1$, 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 [8].[1]

The (synchronous) semantics of MC 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 $\sigma(\mathsf{p})$ is a representation of the memory state of p (left unspecified) and write $\mathsf{upd}(\sigma,\mathsf{p},v)$ to denote the (uniquely defined) updated memory state of p after receiving a value v. Term $e \downarrow_{\sigma(\mathsf{p})} v$ denotes that locally evaluating expression e at p, with memory state σ, evaluates to v. We assume that expression evaluation is deterministic and always terminates. (This formulation captures the essence of previous memory models for choreographies, cf. [4, 7].)

Transitions are defined over pairs $\langle C, \sigma\rangle$, given by the rules in Figure 2. As usual, we omit the angular brackets in transitions. These rules are mostly standard, and we summarise their intuition. In $\lfloor\text{C}|\text{Com}\rceil$,

$$\frac{e \downarrow_{\sigma(\mathsf{p})} v}{\mathsf{p}.e \texttt{ -> } \mathsf{q}; C, \sigma \to C, \mathsf{upd}(\sigma,\mathsf{q},v)}\ \lfloor\text{C}|\text{Com}\rceil \qquad \frac{C_1,\sigma \to C_1',\sigma'}{\mathsf{def}\, X = C_2\, \mathsf{in}\, C_1,\sigma \to \mathsf{def}\, X = C_2\, \mathsf{in}\, C_1',\sigma'}\ \lfloor\text{C}|\text{Ctx}\rceil$$

$$\frac{e \downarrow_{\sigma(\mathsf{p})} \mathsf{true}}{\mathsf{if}\, \mathsf{p}.e\, \mathsf{then}\, C_1\, \mathsf{else}\, C_2, \sigma \to C_1, \sigma}\ \lfloor\text{C}|\text{Then}\rceil \qquad \frac{e \downarrow_{\sigma(\mathsf{p})} \mathsf{false}}{\mathsf{if}\, \mathsf{p}.e\, \mathsf{then}\, C_1\, \mathsf{else}\, C_2, \sigma \to C_2, \sigma}\ \lfloor\text{C}|\text{Else}\rceil$$

$$\frac{C_1 \preceq C_2 \quad C_2,\sigma \to C_2',\sigma' \quad C_2' \preceq C_1'}{C_1,\sigma \to C_1',\sigma'}\ \lfloor\text{C}|\text{Struct}\rceil$$

Figure 2: Minimal Choreographies, synchronous semantics.

the state of q is updated with the value received from p (which results from the evaluation of expression e at that process). Rules $\lfloor\text{C}|\text{Then}\rceil$ and $\lfloor\text{C}|\text{Else}\rceil$ are as expected, while rule $\lfloor\text{C}|\text{Ctx}\rceil$ allows reductions under recursive definitions. Finally, rule $\lfloor\text{C}|\text{Struct}\rceil$ uses a structural precongruence $\preceq$, defined in Figure 3, which essentially allows (i) independent communications to be swapped (rule $\lfloor\text{C}|\text{Eta-Eta}\rceil$), (ii) recursive definitions to be unfolded (rule $\lfloor\text{C}|\text{Unfold}\rceil$), and (iii) garbage collection (rule $\lfloor\text{C}|\text{ProcEnd}\rceil$). The

[1] We relaxed the syntax of MC slightly with respect to [8] by leaving the syntax of expressions unspecified, which allows for the simpler conditional $\mathsf{if}\, \mathsf{p}.e\, \mathsf{then}\, C_1\, \mathsf{else}\, C_2$ in line with typical choreography languages. This minor change simplifies our presentation.

$$\frac{\mathsf{pn}(\eta)\cap\mathsf{pn}(\eta')=\emptyset}{\eta;\eta'\equiv\eta';\eta}\ \lfloor\text{C}|\text{Eta-Eta}\rceil \qquad \frac{\mathsf{pn}(C_i)\cap\mathsf{pn}(\eta)=\emptyset}{\mathsf{def}\,X=C_2\,\mathsf{in}\,(\eta;C_1)\equiv\eta;\mathsf{def}\,X=C_2\,\mathsf{in}\,C_1}\ \lfloor\text{C}|\text{Eta-Rec}\rceil$$

$$\frac{}{\mathsf{def}\,X=C_2\,\mathsf{in}\,C_1\preceq\mathsf{def}\,X=C_2\,\mathsf{in}\,C_1[C_2/X]}\ \lfloor\text{C}|\text{Unfold}\rceil \qquad \frac{}{\mathsf{def}\,X=C\,\mathsf{in}\,\mathbf{0}\preceq\mathbf{0}}\ \lfloor\text{C}|\text{ProcEnd}\rceil$$

$$\frac{\{\mathsf{p},\mathsf{q}\}\cap\mathsf{pn}(\eta)=\emptyset}{\mathsf{if}\,\mathsf{p}.e\,\mathsf{then}\,(\eta;C_1)\,\mathsf{else}\,(\eta;C_2)\equiv\eta;\mathsf{if}\,\mathsf{p}.e\,\mathsf{then}\,C_1\,\mathsf{else}\,C_2}\ \lfloor\text{C}|\text{Eta-Cond}\rceil$$

$$\frac{\mathsf{p}\neq\mathsf{q}}{\begin{array}{c}\mathsf{if}\,\mathsf{p}.e\,\mathsf{then}\,(\mathsf{if}\,\mathsf{q}.e'\,\mathsf{then}\,C_1\,\mathsf{else}\,C_2)\,\mathsf{else}\,(\mathsf{if}\,\mathsf{q}.e'\,\mathsf{then}\,C_1'\,\mathsf{else}\,C_2')\\ \equiv\\ \mathsf{if}\,\mathsf{q}.e'\,\mathsf{then}\,(\mathsf{if}\,\mathsf{p}.e\,\mathsf{then}\,C_1\,\mathsf{else}\,C_1')\,\mathsf{else}\,(\mathsf{if}\,\mathsf{p}.e\,\mathsf{then}\,C_2\,\mathsf{else}\,C_2')\end{array}}\ \lfloor\text{C}|\text{Cond-Cond}\rceil$$

Figure 3: Minimal Choreographies, structural precongruence.

remaining rules are additional rules allowing communications to swap with other constructs, required for achieving (i). These rules, taken together, endow the semantics of MC 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 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 $\mathsf{pn}(C)$.

Unsurprisingly, choreographies in MC 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)**.** *Given a choreography C, either $C\preceq\mathbf{0}$ (termination) or, for every σ, there exist C' and σ' such that $C,\sigma\to C',\sigma'$.*

2.2 Stateful Processes and EndPoint Projection

Minimal Choreographies are meant to be implemented in a minimalistic process calculus, also introduced in [8], 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 $\mathsf{p}\triangleright_{\sigma_\mathsf{p}} B$ is a process with name p, memory state σ_p

$$B ::= \mathsf{q}!e;B \mid \mathsf{p}?;B \mid \mathsf{if}\,e\,\mathsf{then}\,B_1\,\mathsf{else}\,B_2;B \mid \mathsf{def}\,X=B_2\,\mathsf{in}\,B_1 \mid X \mid \mathbf{0} \qquad N,M ::= \mathsf{p}\triangleright_{\sigma_p} B \mid (N \mid M) \mid \mathbf{0}$$

Figure 4: Stateful Processes, syntax.

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 $\mathsf{q}!e;B$ evaluates expression e and sends the result to process q, proceeding as B. The dual receiving behaviour $\mathsf{p}?;B$ expects a value from process p, stores it in its memory and proceeds as B. The other terms are as in MC.

These intuitions are formalized in the synchronous semantics of SP, which is defined by the rules in Figure 5. Rule $\lfloor$P|Com$\rceil$ models synchronous value communication: a process p wishing to send a value

$$\frac{e\downarrow_{\sigma_\mathsf{p}} v}{\mathsf{p}\triangleright_{\sigma_\mathsf{p}} \mathsf{q}!e;B_1 \mid \mathsf{q}\triangleright_{\sigma_\mathsf{q}} \mathsf{p}?;B_2 \to \mathsf{p}\triangleright_{\sigma_\mathsf{p}} B_1 \mid \mathsf{q}\triangleright_{\mathsf{upd}(\sigma_\mathsf{q},v)} B_2}\ \lfloor\mathrm{P|Com}\rceil$$

$$\frac{e\downarrow_{\sigma_\mathsf{p}} \mathsf{true}}{\mathsf{p}\triangleright_{\sigma_\mathsf{p}} \mathsf{if}\, e\,\mathsf{then}\, P_1\,\mathsf{else}\, P_2 \to \mathsf{p}\triangleright_{\sigma_\mathsf{p}} P_1}\ \lfloor\mathrm{P|Then}\rceil \qquad \frac{e\downarrow_{\sigma_\mathsf{p}} \mathsf{false}}{\mathsf{p}\triangleright_{\sigma_\mathsf{p}} \mathsf{if}\, e\,\mathsf{then}\, P_1\,\mathsf{else}\, P_2 \to \mathsf{p}\triangleright_{\sigma_\mathsf{p}} P_2}\ \lfloor\mathrm{P|Else}\rceil$$

$$\frac{\mathsf{p}\triangleright_{\sigma_\mathsf{p}} P \mid N \to \mathsf{p}\triangleright_{\sigma_\mathsf{p}} P' \mid N'}{\mathsf{p}\triangleright_{\sigma_\mathsf{p}} \mathsf{def}\, X = Q\,\mathsf{in}\, P \mid N \to \mathsf{p}\triangleright_{\sigma_\mathsf{p}} \mathsf{def}\, X = Q\,\mathsf{in}\, P' \mid N'}\ \lfloor\mathrm{P|Ctx}\rceil$$

$$\frac{N \preceq M \quad M \to M' \quad M' \preceq N'}{N \to N'}\ \lfloor\mathrm{P|Struct}\rceil \qquad \frac{N \to N'}{N \mid M \to N' \mid M}\ \lfloor\mathrm{P|Par}\rceil$$

Figure 5: Stateful Processes, synchronous semantics.

to q can synchronise with a receive-from-p action at q, and the state of q is updated accordingly. The remaining rules are standard. This calculus once again includes a structural precongruence relation $\preceq$, defined in Figure 6, which allows unfolding of recursive definition and garbage collection (removal of terminated processes).

$$\frac{}{\mathsf{def}\, X = B_2\,\mathsf{in}\, B_1 \preceq \mathsf{def}\, X = B_2\,\mathsf{in}\, B_1[B_2/X]}\ \lfloor\mathrm{S|Unfold}\rceil$$

$$\frac{}{\mathsf{p}\triangleright_{\sigma_\mathsf{p}} \mathbf{0} \preceq \mathbf{0}}\ \lfloor\mathrm{S|PZero}\rceil \qquad \frac{}{N \mid \mathbf{0} \preceq N}\ \lfloor\mathrm{S|NZero}\rceil \qquad \frac{}{\mathsf{def}\, X = B\,\mathsf{in}\,\mathbf{0} \preceq \mathbf{0}}\ \lfloor\mathrm{S|ProcEnd}\rceil$$

Figure 6: Stateful Processes, structural precongruence.

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 focus on networks that are generated 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 $[\![C]\!]_\mathsf{p}$. The rules defining behaviour projection are given in Figure 7. Each choreography term is projected to the local action of

$$[\![\mathsf{p}.e\ \texttt{->}\ \mathsf{q};C]\!]_\mathsf{r} = \begin{cases} \mathsf{q}!e;[\![C]\!]_\mathsf{r} & \text{if } \mathsf{r}=\mathsf{p} \\ \mathsf{p}?;[\![C]\!]_\mathsf{r} & \text{if } \mathsf{r}=\mathsf{q} \\ [\![C]\!]_\mathsf{r} & \text{otherwise} \end{cases}$$

$$[\![\mathsf{if}\ \mathsf{p}.e\,\mathsf{then}\, C_1\,\mathsf{else}\, C_2;C]\!]_\mathsf{r} = \begin{cases} \mathsf{if}\, e\,\mathsf{then}\,[\![C_1]\!]_\mathsf{r}\,\mathsf{else}\,[\![C_2]\!]_\mathsf{r};[\![C]\!]_\mathsf{r} & \text{if } \mathsf{r}=\mathsf{p} \\ [\![C_1]\!]_\mathsf{r};[\![C]\!]_\mathsf{r} & \text{if } \mathsf{r}\neq\mathsf{p} \text{ and } [\![C_1]\!]_\mathsf{r} = [\![C_2]\!]_\mathsf{r} \end{cases}$$

$$[\![\mathsf{def}\, X = C_2\,\mathsf{in}\, C_1]\!]_\mathsf{r} = \mathsf{def}\, X = [\![C_2]\!]_\mathsf{r}\,\mathsf{in}\,[\![C_1]\!]_\mathsf{r} \qquad [\![\mathbf{0}]\!]_\mathsf{r} = \mathbf{0} \qquad [\![X]\!]_\mathsf{r} = X$$

Figure 7: Minimal Choreographies, behaviour projection.

the process that we are projecting. For example, a communication term $\mathsf{p}.e \ \texttt{->}\ \mathsf{q}$ is projected into a send action for the sender process p, a receive action for the receiver process q, or nothing otherwise. The rule for projecting a conditional requires that the behaviours of processes not aware of the choice are independent of that choice (see [8] for details).

Definition 1 (EPP from MC to SP)**.** *Given a choreography C and a state σ, the endpoint projection $[\![C,\sigma]\!]$ is the parallel composition of the EPPs of C wrt all processes in $\mathsf{pn}(C)$:*

$$[\![C,\sigma]\!] = \prod_{\mathsf{p}\in\mathsf{pn}(C)} \mathsf{p} \triangleright_{\sigma(\mathsf{p})} [\![C]\!]_{\mathsf{p}}\,.$$

Given any two states σ and σ', $[\![C,\sigma]\!]$ is defined iff $[\![C,\sigma']\!]$ is defined. If this is the case, we say that C is *projectable* and that N is the projection of C,σ.

Theorem 2 (EPP Theorem)**.** *If C is projectable, then, for any σ:*

- *(Completeness) if $C,\sigma \to C',\sigma'$, then $[\![C,\sigma]\!] \to [\![C',\sigma']\!]$;*
- *(Soundness) if $[\![C,\sigma]\!] \to N$, then $C,\sigma \to C',\sigma'$ for some C' and σ' such that $[\![C',\sigma']\!] \preceq N$.*

Combining Theorem 2 with Theorem 1 we get that the projections of choreographies never deadlock.

Corollary 1 (Deadlock-freedom by construction)**.** *Let $N = [\![C,\sigma]\!]$ for some C and σ. Then either $N \preceq \mathbf{0}$ (termination) or there exists N' such that $N \to N'$.*

3 Asynchronous Choreographies

In order to define an asynchronous semantics for MC 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 MC. We define the syntactic category of contexts, which can contain holes denoted as $\bullet$, formally defined in Figure 8. Structural precongruence is defined for contexts

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

Figure 8: Contexts for asynchrony.

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 $\mathscr{C}[\bullet]$ 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 $\mathscr{C}[\bullet]$ does not contain either r or $\bullet$.

Definition 2. *The* next *action for* r *in a context* $\mathscr{C}[\bullet]$*, denoted* $\mathsf{next}_\mathsf{r}(\mathscr{C}[\bullet])$*, is defined as follows.*

$$\mathsf{next}_\mathsf{r}(\bullet;C) = \bullet \qquad \mathsf{next}_\mathsf{r}(\mathbf{0}) = \mathsf{next}_\mathsf{r}(X) = \textit{undefined}$$

$$\mathsf{next}_\mathsf{r}(\eta;\mathscr{C}[\bullet]) = \begin{cases} \mathsf{comm} & \textit{if } \mathsf{r} \in \mathsf{pn}(\eta) \\ \mathsf{next}_\mathsf{r}(\mathscr{C}[\bullet]) & \textit{otherwise} \end{cases}$$

$$\mathsf{next}_\mathsf{r}(\mathsf{if}\ \mathsf{p}.e\ \mathsf{then}\ \mathscr{C}_1[\bullet]\ \mathsf{else}\ \mathscr{C}_2[\bullet]) = \begin{cases} \mathsf{cond} & \textit{if } \mathsf{p} = \mathsf{r} \\ \mathsf{next}_\mathsf{r}(\mathscr{C}_1[\bullet]) & \textit{if } \mathsf{next}_\mathsf{r}(\mathscr{C}_1[\bullet]) = \mathsf{next}_\mathsf{r}(\mathscr{C}_2[\bullet]) \\ \textit{undefined} & \textit{otherwise} \end{cases}$$

$$\mathsf{next}_\mathsf{r}(\mathsf{def}\ X = C\ \mathsf{in}\ \mathscr{C}[\bullet]) = \mathsf{next}_\mathsf{r}((\mathscr{C}[\bullet])[C/X])$$

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

This notion is well-defined, since structural precongruence cannot swap actions involving the same process. We denote by $\mathscr{C}[\eta]$ the result of replacing every occurrence of $\bullet$ in $\mathscr{C}[\bullet]$ by η.

Since we expect asynchronous communications to occur in two steps, it is reasonable to expect MC to be extended with additional actions. By an *extension* of MC, we understand a choreography language that has all the primitives of MC, an arbitrarily larger set of interaction 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 3. *A semantics* $\to_a$ *for an extension of MC is* asynchronous *if the following conditions hold for any state* σ*. (We assume that* $\mathbf{0};C$ *stands for* C*.)*

- *If* $\mathsf{next}_\mathsf{p}(\mathscr{C}[\bullet]) = \bullet$*, then* $\mathscr{C}[\mathsf{p}.e \texttt{ -> } \mathsf{q}],\sigma \to_a \mathscr{C}[*^v_\mathsf{q}],\sigma$*, where* $e \downarrow_{\sigma(\mathsf{p})} v$ *and* $*^v_\mathsf{q}$ *is a statement that depends on* v *and* q.
- *If* $\mathsf{next}_\mathsf{q}(\mathscr{C}[\bullet]) = \bullet$*, then* $\mathscr{C}[*^v_\mathsf{q}],\sigma \to_a \mathscr{C}[\mathbf{0}],\mathsf{upd}(\sigma,\mathsf{q},v)$.

The second ingredient we need is a formal definition of equivalence. Our notion is similar to the standard notion of operational correspondence from [13], but stronger, since we require that any additional term introduced by asynchronous reductions can always be consumed.

Definition 4. *An asynchronous semantics* $\to_a$ *for an extension of MC is asynchronously equivalent to the semantics* $\to$ *of MC if:*

- *if* $C,\sigma \to C',\sigma'$*, then* $C,\sigma \to_a^* C',\sigma'$*;*
- *if* $C,\sigma \to_a^* C',\sigma'$*, then there exist* C'' *and* σ'' *such that* $C,\sigma \to^* C'',\sigma''$ *and* $C',\sigma' \to_a^* C'',\sigma''$*.*

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

Now we are ready to present our extension of MC and define its asynchronous semantics. We extend the syntax of choreographies with the runtime terms in Figure 9. We call the resulting calculus aMC (for asynchronous MC). The key idea is that, at runtime, a communication is expanded into multiple actions.

$$\eta ::= \ldots \mid \mathsf{p}.e \texttt{ -> } x \mid \hat{v} \texttt{ -> } \mathsf{q} \qquad \hat{v} ::= x \mid v$$

Figure 9: Asynchronous Minimal Choreographies, runtime terms.

For example, a communication $\mathsf{p}.e \texttt{ -> } \mathsf{q}$ expands in $\mathsf{p}.e \texttt{ -> } x$ – a send action from p – and $x \texttt{ -> } \mathsf{q}$ – a receive

action by q. The variable 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 $\mathsf{p}.e \texttt{ -> } x$ replaces x in the corresponding receive action with the actual value v computed from e at p, yielding $v \texttt{ -> } \mathsf{q}$. Finally, executing $v \texttt{ -> } \mathsf{q}$ updates the state of q. Process names for the new runtime terms are defined in the obvious way.

The semantics for aMC includes:

- the rules from Figure 10, replacing $\lfloor$C|Com$\rceil$, and from Figure 11, extending $\preceq$;
- rules $\lfloor$C|Then$\rceil$, $\lfloor$C|Else$\rceil$ and $\lfloor$C|Struct$\rceil$ from Figure 2;
- the rules defining $\preceq$ (Figure 2), with η now ranging also over the new runtime terms.

$$\frac{e \downarrow_{\sigma(\mathsf{p})} v}{\mathsf{p}.e \texttt{ -> } x; C, \sigma \to_a C[v/x], \sigma} \lfloor \text{C|Com-S} \rceil \qquad \frac{}{v \texttt{ -> } \mathsf{q}; C, \sigma \to_a C, \mathsf{upd}(\sigma, \mathsf{q}, v)} \lfloor \text{C|Com-R} \rceil$$

Figure 10: Asynchronous Minimal Choreographies, semantics (runtime terms).

$$\frac{}{\mathsf{p}.e \texttt{ -> } \mathsf{q} \preceq \mathsf{p}.e \texttt{ -> } x; x \texttt{ -> } \mathsf{q}} \lfloor \text{C|Com-Unfold} \rceil$$

Figure 11: Asynchronous Minimal Choreographies, new rule for structural precongruence.

By the Barendregt convention, the variables introduced by unfolding a value communication are globally fresh. (We assume that their scope is global.) This allows us to use these variables to maintain the correspondence between the value being sent and that being received.

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

Example 1. *Let* $C \triangleq \mathsf{p}.e \texttt{ -> } \mathsf{q}; \mathsf{p}.e' \texttt{ -> } \mathsf{r}$. *To execute* C *asynchronously, we must expand it:*

$$C \preceq \mathsf{p}.e \texttt{ -> } x;\ x \texttt{ -> } \mathsf{q};\ \mathsf{p}.e' \texttt{ -> } y;\ y \texttt{ -> } \mathsf{r}$$

Using $\preceq$*, we can swap the second term to the end:*

$$C \preceq \mathsf{p}.e \texttt{ -> } x;\ \mathsf{p}.e' \texttt{ -> } y;\ y \texttt{ -> } \mathsf{r};\ x \texttt{ -> } \mathsf{q}$$

So p *can send both messages immediately, and* r *can receive its message before* q.

The semantics of aMC also enjoys deadlock-freedom.

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

Theorem 4. *The relation* $\to_a$ *is an asynchronous semantics for aMC.*

Proof (Sketch). By induction on the definition of next. For the first point, observe that the only process name in $\mathsf{p}.e \texttt{ -> } x$ is p, and therefore we can expand $\mathsf{p}.e \texttt{ -> } \mathsf{q}$ and use swapping ($\preceq$) to rewrite $\mathscr{C}[\mathsf{p}.e \texttt{ -> } \mathsf{q}]$ as $\mathsf{p}.e \texttt{ -> } x; \mathscr{C}[x \texttt{ -> } \mathsf{q}]$, if $\mathsf{next}_{\mathsf{p}}(\mathscr{C}[\bullet]) = \bullet$. The second point follows from a similar observation for $v \texttt{ -> } \mathsf{q}$. $\square$

Theorem 5. *The semantics $\to_a$ is asynchronously equivalent to the semantics of MC.*

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

For the second property, we make two observations.

(1) The relation $\to_a$ has the diamond property, namely: if $C,\sigma \to_a C',\sigma'$ and $C,\sigma \to_a C'',\sigma''$, then there exist C''' and σ''' such that $C',\sigma' \to_a C''',\sigma'''$ and $C'',\sigma'' \to_a C''',\sigma'''$ (see Figure 12, left). This is directly proven by case analysis on the structure of a choreography C with two distinct possible transitions.

(2) If $C,\sigma \to_a C',\sigma'$, then there exist C'' and σ'' such that $C,\sigma \to^* C'',\sigma''$ and $C',\sigma' \to_a^* C'',\sigma''$. This is similar to the proof for the first property. If the transition uses rule $\lfloor$C|Com-S$\rceil$, then the choreography C'' is obtained by firing rule $\lfloor$C|Com-R$\rceil$, 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 $C,\sigma \to_a^* C',\sigma'$. 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 12, right. By assumption, we have that $C,\sigma \to_a^* C_1,\sigma_1 \to_a C',\sigma'$ (top of the picture). By induction hypothesis, there exist C_1' and σ_1' such that $C,\sigma \to^* C_1',\sigma_1'$ and $C_1,\sigma_1 \to_a^* C_1',\sigma_1'$ (left triangle). By iterated application of (1) above, there exist C_1'' and σ_1'' such that $C_1',\sigma_1' \to_a^? C_1'',\sigma_1''$ and $C',\sigma' \to_a^* C_1'',\sigma_1''$ (right square). The reduction $C_1',\sigma_1' \to_a^? C_1'',\sigma_1''$ is optional because it may be the case that the action in the reduction $C_1,\sigma_1 \to_a C',\sigma'$ is already included in the reduction chain $C_1,\sigma_1 \to_a^* C_1',\sigma_1'$. 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 $C_1',\sigma_1' \to^* C'',\sigma''$ and $C_1'',\sigma_1'' \to_a^* C'',\sigma''$, and the thesis follows. □

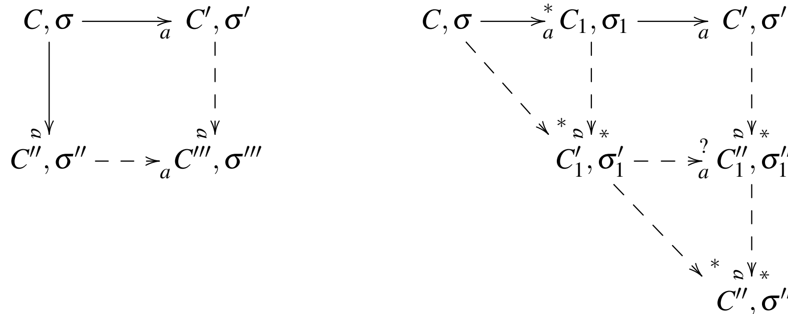

Figure 12: The diamond property for $\to_a$ (left) and the reasoning in the proof of the second property in Theorem 5 (right).

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 section,*

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 Asynchronous Processes

In the previous section, 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).

4.1 Asynchronous Stateful Processes

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 $\mathsf{p} \triangleright_{\sigma}^{\rho} B$, 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 $\langle \mathsf{q}, m \rangle$, where q is the sender process and m is a value, label, or process identifier.

The semantics of aSP consists of rules $\lfloor$P|Then$\rceil$, $\lfloor$P|Else$\rceil$, $\lfloor$P|Par$\rceil$, $\lfloor$P|Ctx$\rceil$ and $\lfloor$P|Struct$\rceil$ from SP (Figure 5) together with the new rules given in Figure 13. Structural precongruence for aSP is defined exactly as for SP. We write $\rho \cdot \langle \mathsf{q}, m \rangle$ to denote the queue obtained by appending message $\langle \mathsf{q}, m \rangle$ to ρ,

$$\frac{e \downarrow_{\sigma_\mathsf{p}} v \qquad \rho'_\mathsf{q} = \rho_\mathsf{q} \cdot \langle \mathsf{p}, v \rangle}{\mathsf{p} \triangleright_{\sigma_\mathsf{p}}^{\rho_\mathsf{p}} \mathsf{q}!e; B_\mathsf{p} \mid \mathsf{q} \triangleright_{\sigma_\mathsf{q}}^{\rho_\mathsf{q}} B_\mathsf{q} \to \mathsf{p} \triangleright_{\sigma_\mathsf{p}}^{\rho_\mathsf{p}} B_\mathsf{p} \mid \mathsf{q} \triangleright_{\sigma_\mathsf{q}}^{\rho'_\mathsf{q}} B_\mathsf{q}} \lfloor \text{P|Com-S} \rceil \qquad \frac{\rho_\mathsf{q} \preceq \langle \mathsf{p}, v \rangle \cdot \rho'_\mathsf{q}}{\mathsf{q} \triangleright_{\sigma_\mathsf{q}}^{\rho_\mathsf{q}} \mathsf{p}?; B \to \mathsf{q} \triangleright_{\mathsf{upd}(\sigma_\mathsf{q}, v)}^{\rho'_\mathsf{q}} B} \lfloor \text{P|Com-R} \rceil$$

Figure 13: Asynchronous Stateful Processes, semantics (new rules).

and $\langle \mathsf{q}, m \rangle \cdot \rho$ for the queue with $\langle \mathsf{q}, m \rangle$ at the head and ρ as tail. We simulate having one separate FIFO queue for each other process by allowing incoming messages from different senders to be exchanged, which we represent using the congruence $\rho \preceq \rho'$ defined by the rule $\langle \mathsf{p}, m \rangle \cdot \langle \mathsf{q}, m' \rangle \preceq \langle \mathsf{q}, m' \rangle \cdot \langle \mathsf{p}, m \rangle$ if $\mathsf{p} \neq \mathsf{q}$.

All behaviours of SP are valid also in aSP.

Theorem 6. *Let N be an SP network. If $N \to N'$ (in SP), then $N_{[]} \to^* N'_{[]}$ (in aSP), where $N_{[]}$ denotes the asynchronous network obtained by adding an empty queue to each process.*

Proof. Straightforward by case analysis: the only reduction rule in SP that is not present in aSP is that for communications, which can be simulated by applying rules $\lfloor$P|Com-S$\rceil$ and $\lfloor$P|Com-R$\rceil$ in sequence. □

The converse is not true, so the relation between SP and aSP is not so strong as that between MC and aMC (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 $\mathsf{p} \triangleright_{\sigma_\mathsf{p}} \mathsf{q}!1 \mid \mathsf{q} \triangleright_{\sigma_\mathsf{q}} \mathbf{0}$ is deadlocked in SP, but its counterpart in aSP (obtained by adding empty queues at each process) reduces to $\mathsf{q} \triangleright_{\sigma_\mathsf{q}}^{\langle \mathsf{p}, 1 \rangle} \mathbf{0}$. (This network is not equivalent to $\mathbf{0}$, since there is a non-empty queue.)

More interestingly, the network

$$\mathsf{p} \triangleright_{\sigma_\mathsf{p}} \mathsf{q}!e_1; \mathsf{r}? \mid \mathsf{q} \triangleright_{\sigma_\mathsf{q}} \mathsf{r}!e_2; \mathsf{p}? \mid \mathsf{r} \triangleright_{\sigma_\mathsf{r}} \mathsf{p}!e_3; \mathsf{q}? \tag{1}$$

is deadlocked in SP, but reduces to **0** in aSP (all queues are eventually emptied). It is possible to extend MC with communication primitives that capture this type of behaviour, as discussed in [6]; we briefly discuss this point in Section 6.

4.2 Asynchronous EndPoint Projection

Defining an EPP from aMC 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., v -> q, we add a message containing v at the top of q's queue.

There are two problems with this approach. The first is a syntactic issue: each message in the queues of aSP processes must include the name of its sender, but that information is not present in the runtime term v -> q. This is a minor issue that can be dealt with by annotating send and receive actions with the name of the process sending them – writing, e.g., $\langle \mathsf{p}, x \rangle$ -> q and $\langle \mathsf{p}, v \rangle$ -> q. Changing the syntax of runtime terms and rule $\lfloor$C$|$Com-Unfold$\rceil$ in this way is trivial, and we assume this annotated terms in the remainder of this session.[2]

The second problem arises because we can write choreographies that use runtime terms in a "wrong" way, for which Theorem 2 no longer holds.

Example 2. *Consider the choreography $C = \mathsf{p}.1$ -> q; $\langle \mathsf{p}, 2 \rangle$ -> q. If we naively project it as described informally, we obtain* $\mathsf{p} \triangleright^{[]}_{\sigma(\mathsf{p})} \mathsf{q}!1 \mid \mathsf{q} \triangleright^{\langle \mathsf{p},2 \rangle}_{\sigma(\mathsf{q})} \mathsf{p}?; \mathsf{p}?$, *where* $[]$ *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 5 (Well-formedness)**.** *A choreography C in aMC containing runtime terms is* well-formed *if* $\eta_1; \ldots; \eta_n; C^{MC} \preceq^- C$*, where:*

- $\preceq^-$ *is structural precongruence without rule* $\lfloor C|Unfold \rceil$*;*
- *each* η_i *is an instantiated receive action of the form* $\langle \mathsf{p}_i, v_i \rangle$ -> q_i*;*
- C^{MC} *is an MC choreography (i.e., a choreography without runtime terms).*

Well-formedness is decidable, since the set of choreographies equivalent up to $\preceq^-$ is decidable. More efficiently, one can check that C is well-formed by swapping all runtime actions to the beginning and folding all paired send/receive terms. Furthermore, choreography execution preserves well-formedness; in particular, the problematic choreography from Example 2 is not well-formed. More generally, we can use well-formedness in the remainder of this section to reason about EPP.

[2]We could have defined the runtime terms for aMC annotated from the start. However, we felt that it would be unnatural to include them in a choreographic presentation, as they are only needed for projecting runtime terms – a feature that is a technicality required for the proof of Theorem 7 below, since the programmer should not write runtime terms anyway.

Definition 6 (Asynchronous EPP from aMC to aSP)**.** *Let C be a well-formed aMC choreography and σ be a state. Without loss of generality, we assume that C does not contain $\mathsf{p}.e$ -> x actions.*[3] *The EPP of C and σ is defined as*

$$[\![C,\sigma]\!] = \prod_{\mathsf{p}\in\mathsf{pn}(C)} \mathsf{p} \triangleright^{(\!|C|\!)_\mathsf{p}}_{\sigma(\mathsf{p})} [\![C]\!]_\mathsf{p}$$

where $[\![C]\!]_\mathsf{p}$ is defined as in Figure 7 with the extra rule in Figure 14 and $(\!|C|\!)_\mathsf{p}$ is defined by the rules in Figure 15.

$$[\![\langle \mathsf{p}, v\rangle \text{ -> } \mathsf{q}; C]\!]_\mathsf{r} = \begin{cases} \mathsf{p}?;[\![C]\!]_\mathsf{r} & \text{if } \mathsf{r}=\mathsf{q} \\ [\![C]\!]_\mathsf{r} & \text{otherwise} \end{cases}$$

Figure 14: Minimal Choreographies, asynchronous behaviour projection (new rule).

$$(\!|\langle \mathsf{p}, v\rangle \text{ -> } \mathsf{q}; C|\!)_\mathsf{r} = \begin{cases} \langle \mathsf{p}, v\rangle \cdot (\!|C|\!)_\mathsf{r} & \text{if } \mathsf{r}=\mathsf{q} \\ (\!|C|\!)_\mathsf{r} & \text{otherwise} \end{cases} \qquad \begin{array}{l} (\!|\mathsf{if}\ \mathsf{p}.e\ \mathsf{then}\ C_1\ \mathsf{else}\ C_2; C|\!)_\mathsf{r} = (\!|C_1|\!)_\mathsf{r} \cdot (\!|C|\!)_\mathsf{r} \\ (\!|\eta; C|\!)_\mathsf{r} = (\!|C|\!)_\mathsf{r} \end{array}$$

Figure 15: Minimal Choreographies, projection of messages in transit.

In the last case of the definition of $(\!|C|\!)_\mathsf{p}$ (bottom-right), η ranges over all cases that are not covered previously. The rule for the conditional may seem a bit surprising: projectability of choreographies implies 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 an asynchronous variant of Theorem 2.

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

- *(Completeness) if $C,\sigma \to C',\sigma'$, then $[\![C,\sigma]\!] \to [\![C',\sigma']\!]$;*
- *(Soundness) if $[\![C,\sigma]\!] \to N$, then $C,\sigma \to C',\sigma'$ for some C' and σ' such that $[\![C',\sigma']\!] \preceq N$.*

As a consequence, Corollary 1 applies also to the asynchronous case: the processes projected from MC 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) can be written as the projection of a choreography with runtime terms, but this choreography is not structurally precongruent to any choreography obtained from execution of an MC choreography.

5 Related Work

The first work introducing an asynchronous semantics to choreographies is [4], as described in the introduction. The same approach was later adapted to compositional choreographies [19] and multiparty

[3]By well-formedness, we can always rewrite C to an equivalent choreography satisfying this condition; such a choreography is also guaranteed not to contain x -> q actions.

session types [15]. 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 3. *Let C be the choreography from the introduction: $C \triangleq \mathsf{p}.1 \mathrel{\texttt{->}} \mathsf{q}; \mathsf{p}.2 \mathrel{\texttt{->}} \mathsf{r}$. If we adopted the asynchronous semantics from [4], then the reduction*

$$C, \sigma \to \mathsf{p}.1 \mathrel{\texttt{->}} \mathsf{q}, \mathsf{upd}(\sigma, \mathsf{r}, 2)$$

would be possible for any σ. This shows that Theorem 7 (the completeness direction) does not hold in this semantics, since $[\![C,\sigma]\!]$ cannot reduce to any N such that $[\![\mathsf{p}.1 \mathrel{\texttt{->}} \mathsf{q}, \mathsf{upd}(\sigma,\mathsf{r},2)]\!] \preceq N$ (we have to consume the first output by 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

$$\mathsf{p}.1 \mathrel{\texttt{->}} \mathsf{q}, \mathsf{upd}(\sigma, \mathsf{r}, 2) \to \mathbf{0}, \mathsf{upd}(\mathsf{upd}(\sigma, \mathsf{r}, 2), \mathsf{q}, 1)$$

and that $[\![C,\sigma]\!] \to^* [\![\mathbf{0}, \mathsf{upd}(\mathsf{upd}(\sigma,\mathsf{r},2),\mathsf{q},1)]\!]$.

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 [18, Chapter 2] (the full version of [4]) to our notation:

Theorem 8. *If C is a projectable choreography, then, for all σ:*

- *(Completeness) if $C,\sigma \to C',\sigma'$, then $C',\sigma' \to^* C'',\sigma''$ for some C'',σ'' and $[\![C,\sigma]\!] \to^* [\![C'',\sigma'']\!]$;*
- *(Soundness) if $[\![C,\sigma]\!] \to^* N$, then $N \to^* N'$ for some N', and $C,\sigma \to^* C',\sigma'$ for some σ' and C' such that $[\![C',\sigma']\!] \preceq N'$.*

Compared to our Theorem 7 above, Theorem 8 is missing the step-by-step correspondence between choreographies and their projections, and therefore does not say much about the intermediate steps.

In [11], 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 [11] 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 are difficult to compare formally because a counterpart to Theorem 5 is not provided for the model of [11].

A more recent development that is nearer to ours is the model of Applied Choreographies [12], 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 [12] 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 MC in the sense of our definition.

Our approach relies on out-of-order execution for choreographic interactions, which was first introduced in [4] 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 [18]); instead, they just model different safe ways of executing the same choreography.

6 Conclusions and Future Work

We presented a definition of asynchrony in the minimal choreography language MC and showed that we can define an asynchronous semantics for an extension of MC (aMC) that respects this definition. This 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. We discuss a few relevant cases.

The first case regards label selection, usually written $\mathsf{p} \texttt{->} \mathsf{q}[\ell]$, which is a widespread primitive in choreographies [3]. In a selection, the sender informs the receiver of a local choice, and the receiver can use this information to change its behaviour. Selections are mostly relevant for extending the domain of EPP, which is orthogonal to this work. Our approach applies directly also to selections, by adding asynchronous runtime terms that are similar to the ones for value communications but carry selection labels instead of values. Name mobility [4, 9] can also be treated in a similar way.

Instead of out-of-order execution, some choreography models also include explicit parallel composition, e.g., $C \mid C'$ [3, 21]. Most behaviours of $C \mid C'$ are captured by out-of-order execution, for example $\mathsf{p}.e \texttt{->} \mathsf{q} \mid \mathsf{r}.e' \texttt{->} \mathsf{s}$ is equivalent to $\mathsf{p}.e \texttt{->} \mathsf{q};\mathsf{r}.e' \texttt{->} \mathsf{s}$ in MC (due to rule $\lfloor$C|Eta-Eta$\rceil$ in Figure 2) – see [4] for a deeper discussion. Generalising our construction to supporting also an explicit parallel operator is straightforward (structural precongruence is extended homomorphically without surprises).

A more interesting extension is multicom [6], a primitive that considerably extends the expressivity of choreographies by allowing for general criss-cross communications. A prime example is the asynchronous exchange $\{\mathsf{p}.e \texttt{->} \mathsf{q}, \mathsf{q}.e' \texttt{->} \mathsf{p}\}$, 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 [6]). Multicoms crucially depend on asynchrony, making them an interesting case study in our context. Our development applies immediately to multicoms, since we can just apply the expansion rule $\lfloor$C|Com-Unfold$\rceil$ to all communications in a multicom simultaneously.

Likewise, we conjecture that our development is applicable to the models in [10, 17].

Acknowledgements. This work was supported by CRC (Choreographies for Reliable and efficient Communication software), grant no. DFF–4005-00304 from the Danish Council for Independent Research, by grant DFF-1323-00247 from the Danish Council for Independent Research, Natural Sciences, and by the Open Data Framework project at the University of Southern Denmark.

References

[1] Elvira Albert & Ivan Lanese, editors (2016): *Formal Techniques for Distributed Objects, Components, and Systems – 36th IFIP WG 6.1 International Conference, FORTE 2016*. *LNCS* 9688, Springer, doi:10.1007/978-3-319-39570-8.

[2] *Business Process Model and Notation*. `http://www.omg.org/spec/BPMN/2.0/`.

[3] Marco Carbone, Kohei Honda & Nobuko Yoshida (2012): *Structured Communication-Centered Programming for Web Services*. *ACM Trans. Program. Lang. Syst.* 34(2), pp. 8:1–8:78, doi:10.1145/2220365.2220367.

[4] Marco Carbone & Fabrizio Montesi (2013): *Deadlock-freedom-by-design: multiparty asynchronous global programming*. In Roberto Giacobazzi & Radhia Cousot, editors: *POPL*, ACM, pp. 263–274. Available at `http://doi.acm.org/10.1145/2429069.2429101`.

[5] Chor: *Programming Language*. `http://www.chor-lang.org/`.

[6] Luís Cruz-Filipe, Kim S. Larsen & Fabrizio Montesi (2017): *The Paths to Choreography Extraction*. In Javier Esparza & Andrzej S. Murawski, editors: *FoSSaCS*, *LNCS* 10203, Springer, pp. 424–440, doi:10.1007/978-3-662-54458-7_25.

[7] Luís Cruz-Filipe & Fabrizio Montesi (2016): *Choreographies in Practice*. In Albert & Lanese [1], pp. 114–123, doi:10.1007/978-3-319-39570-8_8.

[8] Luís Cruz-Filipe & Fabrizio Montesi (2017): *A Core Model for Choreographic Programming*. In Olga Kouchnarenko & Ramtin Khosravi, editors: *FACS*, *LNCS* 10231, Springer, pp. 17–35, doi:10.1007/978-3-319-57666-4_3.

[9] Luís Cruz-Filipe & Fabrizio Montesi (2017): *Procedural Choreographic Programming*. In Ahmed Bouajjani & Alexandra Silva, editors: *FORTE 2017*, *LNCS* 10321, Springer, pp. 92–107, doi:10.1007/978-3-319-60225-7_7.

[10] Mila Dalla Preda, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese & Jacopo Mauro (2015): *Dynamic Choreographies – Safe Runtime Updates of Distributed Applications*. In Tom Holvoet & Mirko Viroli, editors: *COORDINATION*, *LNCS* 9037, Springer, pp. 67–82, doi:10.1007/978-3-319-19282-6_5.

[11] Pierre-Malo Deniélou & Nobuko Yoshida (2013): *Multiparty Compatibility in Communicating Automata: Characterisation and Synthesis of Global Session Types*. In Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska & David Peleg, editors: *ICALP (2)*, *LNCS* 7966, Springer, pp. 174–186, doi:10.1007/978-3-642-39212-2_18.

[12] Maurizio Gabbrielli, Saverio Giallorenzo & Fabrizio Montesi (2015): *Applied Choreographies*. *CoRR* abs/1510.03637. Available at `http://arxiv.org/abs/1510.03637`.

[13] Daniele Gorla (2010): *Towards a unified approach to encodability and separation results for process calculi*. *Inf. Comput.* 208(9), pp. 1031–1053, doi:10.1016/j.ic.2010.05.002.

[14] Kohei Honda, A. Mukhamedov, G. Brown, T.-C. Chen & Nobuko Yoshida (2011): *Scribbling Interactions with a Formal Foundation*. In Raja Natarajan & Adegboyega K. Ojo, editors: *ICDCIT*, *LNCS* 6536, Springer, pp. 55–75, doi:10.1007/978-3-642-19056-8_4.

[15] Kohei Honda, Nobuko Yoshida & Marco Carbone (2016): *Multiparty Asynchronous Session Types*. *J. ACM* 63(1), pp. 9:1–9:67, doi:10.1145/2827695.

[16] Ivan Lanese, Claudio Guidi, Fabrizio Montesi & Gianluigi Zavattaro (2008): *Bridging the Gap between Interaction- and Process-Oriented Choreographies*. In Antonio Cerone & Stefan Gruner, editors: *SEFM*, IEEE, pp. 323–332, doi:10.1109/SEFM.2008.11.

[17] Hugo A. López, Flemming Nielson & Hanne Riis Nielson (2016): *Enforcing Availability in Failure-Aware Communicating Systems*. In Albert & Lanese [1], pp. 195–211, doi:10.1007/978-3-319-39570-8_13.

[18] Fabrizio Montesi (2013): *Choreographic Programming*. Ph.D. Thesis, IT University of Copenhagen. http://fabriziomontesi.com/files/choreographic_programming.pdf .

[19] Fabrizio Montesi & Nobuko Yoshida (2013): *Compositional Choreographies*. In Pedro R. D'Argenio & Hernán C. Melgratti, editors: *CONCUR*, *LNCS* 8052, Springer, pp. 425–439, doi:10.1007/978-3-642-40184-8_30.

[20] (2008): *PI4SOA*. `http://www.pi4soa.org`.

[21] Zongyan Qiu, Xiangpeng Zhao, Chao Cai & Hongli Yang (2007): *Towards the theoretical foundation of choreography*. In Carey L. Williamson, Mary Ellen Zurko, Peter F. Patel-Schneider & Prashant J. Shenoy, editors: *WWW*, ACM, pp. 973–982, doi:10.1145/1242572.1242704.

[22] Savara: *JBoss Community*. `http://www.jboss.org/savara/`.

[23] W3C WS-CDL Working Group (2004): *Web Services Choreography Description Language Version 1.0*. `http://www.w3.org/TR/2004/WD-ws-cdl-10-20040427/`.