

The Paths to Choreography Extraction

Luís Cruz-Filipe^(✉), Kim S. Larsen, and Fabrizio Montesi

University of Southern Denmark, Odense, Denmark

{lcf,kslarsen,fmontesi}@imada.sdu.dk

Abstract. Choreographies are global descriptions of interactions among concurrent components, most notably used in the settings of verification and synthesis of correct-by-construction software. They require a top-down approach: programmers first write choreographies, and then use them to verify or synthesize their programs. However, most software does not come with choreographies yet, which prevents their application. To attack this problem, previous work investigated choreography extraction, which automatically constructs a choreography that describes the behavior of a given set of programs or protocol specifications.

We propose a new extraction methodology that improves on the state of the art: we can deal with programs that are equipped with state and internal computation; time complexity is dramatically better; and we capture programs that work by exploiting asynchronous communication.

1 Introduction

Choreographies are global descriptions of interactions among components. They have been used as a basis for different models and tools that aim at tackling the complexity of modern software, where separate units – such as processes, objects, and threads – interact to reach a common goal [3, 25].

Two lines of research are of particular interest. In *choreography specifications*, choreographies specify interaction protocols, e.g., multiparty session types [17]. In *choreographic programming* [20], choreographies are programs that define the behavior of concurrent algorithms [13] and/or distributed systems [5, 6, 14]. The key feature of these works is EndPoint Projection (EPP), a procedure that translates choreographies to correct endpoint behaviors in lower-level models. For choreography specifications, EPP generates the local specifications of each participant; these specifications can then be used for verification, to check whether some programs implement their role in the protocol correctly and will thus interact without problems at runtime [17]. In choreographic programming, instead, EPP generates correct-by-construction implementations in a model of executable code (program synthesis), typically given in terms of a process calculus [6].

Montesi was supported by CRC (Choreographies for Reliable and efficient Communication software), grant DFF-4005-00304 from the Danish Council for Independent Research. Cruz-Filipe and Larsen were supported in part by the Danish Council for Independent Research, Natural Sciences, grant DFF-1323-00247.

EPP implements a top-down development methodology: developers first write choreographies and then use the output mechanically generated by EPP. However, there are scenarios where this methodology is not applicable; for example:

- Analysis or integration of legacy software: either code developed previously, or new code written in a technology without support for choreographies.
- Updates: endpoint programs generated by EPP can later be updated locally (e.g., for configuration or optimizations). Since the original choreography is not automatically updated, rerunning EPP loses these changes.

To attack these issues, previous work investigated a procedure to infer choreographies from arbitrary endpoint descriptions. We call this procedure *choreography extraction*. To the best of our knowledge, the current reference for extracting choreography specifications is [19], where graphical choreographies that represent protocol specifications are extracted from communicating automata [4]. Instead, the state of the art for extraction in choreographic programming is [7], where extraction takes terminating processes typed using a fragment of linear logic as input. We advance both lines of work in several aspects, described below.

1.1 Contributions

Extraction for synchronous systems. We define an extraction procedure that applies directly to both choreography specifications and choreographic programming, by working with representative models. We focus on the more difficult case of choreographic programming, and then show how our approach can be applied to other settings in Sect. 6. First we define an extraction algorithm for processes with synchronous communications (Sect. 4), which showcases the key elements of our construction: building a choreography corresponds to finding paths in a graph that represents the abstract execution of the input processes. Our extraction also helps in debugging: if extraction detects a potential deadlock, we pinpoint it with a special term (**1**). This is the first extraction procedure for choreographic programming that can deal with procedures and infinite behavior [7].

Asynchrony. We extend our development to asynchronous communication (Sect. 5). The key novelty is that we can extract a new class of behaviors where processes progress because of asynchronous communication. The simplest example of this class is a two-way exchange: a network of two processes where each process starts by sending a value to the other, and then consumes the received value. This network is deadlocked under a synchronous semantics, violating the state-of-the-art requirements for extraction [19]. Capturing these behaviors is challenging for two reasons: there is no choreography language capable of representing them; and the extraction algorithms presented so far require the behaviors of processes to be representable also under a synchronous interpretation. We overcome both limitations with a new choreography primitive for multiparty asynchronous exchange and a look-ahead mechanism for asynchronous actions in extraction.

Efficiency. We show that our extraction has exponential worst-case time complexity in both the synchronous and the asynchronous cases (Sects. 4 and 5, respectively), unlike the factorial case of [19], even though we can capture a new class of behaviors. In particular, we need only one phase of exponential complexity, while [19] uses multiple phases applied in sequence. The authors of [19] detail only the complexities of their first two phases: the first has exponential complexity (but in a quantity larger than ours), while the second has factorial complexity in a function exponential in the size of the input. Our better time complexity stems from the design of our process language, which does not allow non-deterministic receives from different channels, and careful algorithm crafting. Despite the restriction, we can still model interesting examples thanks to asynchronous exchange. In Sect. 5, we present a novel formulation of the alternating 2-bit protocol, which is given in [15] and used in [19] as a motivating example. Our formulation is simpler and does not require threads as in [19].

2 Related Work

Choreographic Programming. The state of the art for extraction in choreographic programming is [7], where synchronous processes with finite behavior are typed using the multiplicative-additive fragment of linear logic. Our approach is significantly more expressive, bringing support for recursion and asynchronous communication. Also, the proof theory in [7] requires that there are no cycles in the structure of connections among processes. We do not have this limitation.

Choreography Specifications. To the best of our knowledge, the state of the art for extracting choreography specifications is [19], which captures more behaviors than previous works with similar objectives [18, 22].

Extraction in [19] is more restrictive wrt. to asynchrony, requiring all process traces and choices to be represented in the synchronous transition system of the network. Thus, networks that are safe because of asynchronous communication are not extracted in [19]. Instead, our extraction can deal with programs that use multiparty asynchronous exchange, where multiple processes exchange values by exploiting asynchronous communication. As a consequence, we can extract the alternating 2-bit protocol implemented via asynchronous exchange in Sect. 5, which is deadlocked under a synchronous semantics and thus cannot be extracted in [19]. Our extraction is the first capturing systems that are not correctly approximated by synchronous semantics (cf. [2]). A precise characterization of the class of extractable systems is thus an interesting future direction.

To circumvent the limitation that asynchronous exchange is not supported, choreographies in [19] support local concurrency: processes can have internal threads. This opens up for an alternative formulation of the alternating 2-bit protocol, where the two participants use two threads each. However, these choreographies are harder to read. As an example, compare our choreography for the alternating 2-bit protocol in Sect. 5 to that obtained with the automata in [19] (given in [15], Protocol 7 in Example 2.1). Our formulation is a simple recursive

procedure with two exchanges, whereas the control flow in [15] is rather intricate and uses three different operators (fork, join, and merge) at different places to compose two separate loops. In our opinion, our choreographies follow the principles of structured programming to a greater extent, and are simpler; also because coordination happens only through communication.

More interestingly than readability, local concurrency makes the complexity of extraction blow up factorially [19]: process threads are represented using non-determinism between different actions in communicating automata. Determining whether the non-deterministic behavior of these automata is extractable takes (super-)factorial time (factorial time in the size of a graph similar to our AES, cf. Definition 2)! Thus, asynchronous exchange supports a more efficient way of capturing an interesting class of behaviors. Nevertheless, we believe that developing efficient extractions of local concurrency may be useful future work.

3 Core Choreographies and Stateful Processes

We review the languages of Core Choreographies (CC) and Stateful Processes (SP), from [11], which respectively model choreographies and endpoint programs. We introduce labels in the reduction semantics for these calculi to formalize the link between choreographies and their process implementations as a bisimilarity.

$$\begin{array}{l}
 C ::= \mathbf{0} \mid \eta; C \mid \text{if } p \stackrel{e}{=} q \text{ then } C_1 \text{ else } C_2 \mid \text{def } X = C_2 \text{ in } C_1 \mid X \\
 \eta ::= p.e \rightarrow q \mid p \rightarrow q[l] \qquad e ::= v \mid * \mid \dots
 \end{array}$$

Fig. 1. Core Choreographies, Syntax.

Core Choreographies (CC). The syntax of CC is given in Fig. 1. A choreography C describes the behavior of a set of processes (p, q, \dots) running concurrently. Each process has an internal memory cell storing a local value (the value of the process). Term $\mathbf{0}$ is the terminated choreography (omitted in examples). Term $\eta; C$ reads “the system executes η and proceeds as C ”. An interaction η is either: a value communication $p.e \rightarrow q$, where process p evaluates e and sends the result to process q , which stores it in its memory cell, replacing its previous value; or a selection $p \rightarrow q[l]$, where p selects l among the branches offered by q . We abstract from the concrete language of expressions e , which models internal computation and is orthogonal to our development, assuming only that: expressions can contain values v and the placeholder $*$, which refers to the value of the process evaluating them; and evaluating expressions always terminates and returns a value. In a conditional $\text{if } p \stackrel{e}{=} q \text{ then } C_1 \text{ else } C_2$, p checks if its value is equal to q ’s to decide whether the system proceeds as C_1 or C_2 . Term $\text{def } X = C_2 \text{ in } C_1$ defines a procedure X with body C_2 , which can be called in C_1 and C_2 by using term X .

$$\begin{array}{c}
\frac{e[\sigma(\mathbf{p})/*] \downarrow v}{\mathbf{p}.e \rightarrow \mathbf{q}; C, \sigma \xrightarrow{\mathbf{p}.v \rightarrow \mathbf{q}} C, \sigma[\mathbf{q} \mapsto v]} \quad [\text{C|Com}] \quad \frac{}{\mathbf{p} \rightarrow \mathbf{q}[l]; C, \sigma \xrightarrow{\mathbf{p} \rightarrow \mathbf{q}[l]} C, \sigma} \quad [\text{C|Sel}] \\
\frac{\sigma(\mathbf{p}) = \sigma(\mathbf{q})}{\text{if } \mathbf{p} \stackrel{\leftarrow}{=} \mathbf{q} \text{ then } C_1 \text{ else } C_2, \sigma \xrightarrow{\mathbf{p} \stackrel{\leftarrow}{=} \mathbf{q}; \text{then}} C_1, \sigma} \quad [\text{C|Then}] \quad \frac{\mathbf{pn}(\eta) \cap \mathbf{pn}(\eta') = \emptyset}{\eta; \eta' \preceq \eta'; \eta} \quad [\text{C|Eta-Eta}]
\end{array}$$

Fig. 2. Core choreographies, semantics and structural precongruence (selected rules).

The semantics of CC is given in terms of labeled reductions $C, \sigma \xrightarrow{\lambda} C', \sigma'$; the main reduction rules are given in Fig. 2. Reductions are also closed under context (procedure definitions) and under a structural precongruence \preceq , allowing procedure calls to be unfolded and non-interfering actions to be executed in any order. The most interesting rule for \preceq is rule [C|Eta-Eta], which swaps communications between disjoint sets of processes (modeling concurrency). The total function σ maps each process name to the value it stores. Labels λ tell us which action has been performed, which helps stating our later results. In rule [C|Com], v is the value obtained by evaluating (\downarrow) the expression e , with $*$ replaced by the value of the sender \mathbf{p} , $\sigma(\mathbf{p})$. In the reductum, σ is updated such that the receiver \mathbf{q} stores v . Rule [C|Sel] does not alter σ : selections model invoking a method/operation available at the receiver. Rules [C|Then] and [C|Else] (omitted) model conditionals in the standard way. Function $\mathbf{pn}(C)$ returns all the process names that appear in C , and $C \equiv C'$ means $C \preceq C'$ and $C' \preceq C$.

Example 1. We define a simple choreography for client authentication. We write $\mathbf{p} \rightarrow \mathbf{c}, \mathbf{s}[l]$ as a shortcut for $\mathbf{p} \rightarrow \mathbf{c}[l]; \mathbf{p} \rightarrow \mathbf{s}[l]$.

$\text{def } X = \left(\mathbf{c}. \text{pwd} \rightarrow \mathbf{a}; \text{if } \mathbf{a} \stackrel{\leftarrow}{=} \mathbf{s} \text{ then } (\mathbf{a} \rightarrow \mathbf{c}, \mathbf{s}[\text{ok}]; \mathbf{s}.t \rightarrow \mathbf{c}) \text{ else } (\mathbf{a} \rightarrow \mathbf{c}, \mathbf{s}[\text{ko}]; X) \right)$ in X

In this choreography, a client process \mathbf{c} sends a password to an authentication process \mathbf{a} , which checks if the password matches that contained in the server-side process \mathbf{s} . If the password is correct, \mathbf{a} notifies \mathbf{c} and \mathbf{s} , and \mathbf{s} sends an authentication token t to \mathbf{c} . Otherwise, \mathbf{a} notifies \mathbf{c} and \mathbf{s} that authentication failed, and a new attempt is made (by recursively invoking X). \square

Stateful Processes. The calculus SP models concurrent/distributed implementations. Thus, unlike in CC, actions are now distributed among processes.

The syntax of SP is given in Fig. 3. Networks N are parallel compositions of processes $\mathbf{p} \triangleright B$, read “process \mathbf{p} has behavior B ”. An output term $\mathbf{q}!(e); B$ sends the result of evaluating e to \mathbf{q} , and then proceeds as B . Outputs are meant to

$$\begin{array}{l}
B ::= \mathbf{q}!(e); B \mid \mathbf{p}?: B \mid \mathbf{q} \oplus l; B \mid \mathbf{p} \& \{l_i : B_i\}_{i \in I} \mid N ::= \mathbf{p} \triangleright B \mid \mathbf{0} \mid N \mid N \\
\mid \mathbf{0} \mid \text{if } * \stackrel{\leftarrow}{=} \mathbf{q} \text{ then } B_1 \text{ else } B_2 \mid \text{def } X = B_2 \text{ in } B_1 \mid X
\end{array}$$

Fig. 3. Stateful processes, syntax.

$$\begin{array}{c}
 \frac{e[\sigma(\mathbf{p})/*] \downarrow v}{\mathbf{p} \triangleright \mathbf{q}!(e); B_1 \mid \mathbf{q} \triangleright \mathbf{p}^?; B_2, \sigma \xrightarrow{\mathbf{p}.v \rightarrow \mathbf{q}} \mathbf{p} \triangleright B_1 \mid \mathbf{q} \triangleright B_2, \sigma[\mathbf{q} \mapsto v]} \text{ [S|Com]} \\
 \\
 \frac{j \in I}{\mathbf{p} \triangleright \mathbf{q} \oplus l_j; B \mid \mathbf{q} \triangleright \mathbf{p}\&\{l_i : B_i\}_{i \in I}, \sigma \xrightarrow{\mathbf{p} \rightarrow \mathbf{q}[l]} \mathbf{p} \triangleright B \mid \mathbf{q} \triangleright B_j, \sigma} \text{ [S|Sel]} \\
 \\
 \frac{e[\sigma(\mathbf{q})/*] \downarrow \sigma(\mathbf{p})}{\mathbf{p} \triangleright \text{if } * \stackrel{\leq}{=} \mathbf{q} \text{ then } B_1 \text{ else } B_2 \mid \mathbf{q} \triangleright \mathbf{p}!(e); B', \sigma \xrightarrow{\mathbf{p} \stackrel{\leq}{=} \mathbf{q}; \text{then}} \mathbf{p} \triangleright B_1 \mid \mathbf{q} \triangleright B', \sigma} \text{ [S|Then]}
 \end{array}$$

Fig. 4. Stateful processes, semantics (selected rules).

synchronize with input terms at the target process, i.e., $\mathbf{p}^?; B$, which receives a value from \mathbf{p} to be stored locally and then proceeds as B . Term $\mathbf{q} \oplus l; B$ sends the selection of the branch labeled l to \mathbf{q} . Branches are offered by the receiver with term $\mathbf{p}\&\{l_i : B_i\}_{i \in I}$, which offers a choice among the labels l_i to \mathbf{p} . When one of these labels is selected, the respective behavior B_i is run. Term $\text{if } * \stackrel{\leq}{=} \mathbf{q} \text{ then } B_1 \text{ else } B_2$ communicates with process \mathbf{q} to check whether it stores the same value as the process running this behavior, in order to choose between the continuations B_1 and B_2 . Terms $\text{def } X = B_2 \text{ in } B_1$ and X are procedure definition and call, respectively.

The semantics of SP is given by labeled reductions $N, \sigma \xrightarrow{\lambda} N', \sigma'$, with labels λ as in CC.¹ Figure 4 shows the key rules (see the appendix for the complete set). Two processes can synchronize when they refer to each other. In rule [S|Com], an output at \mathbf{p} directed at \mathbf{q} synchronizes with the dual input action at \mathbf{q} – intention to receive from \mathbf{p} ; in the reductum, \mathbf{q} 's value is updated. The reduction receives the same label as the equivalent communication term in CC. The other rules shown are similar. The omitted rules are standard, and close the semantics under parallel composition, structural precongruence, and procedure definitions.

Example 2. The following network implements the choreography in Example 1.

$$\begin{array}{l}
 \mathbf{c} \triangleright \text{def } X = \mathbf{a}!(\mathbf{p}\mathbf{w}\mathbf{d}); \mathbf{a}\&\{\mathbf{ok} : \mathbf{s}^?, \mathbf{ko} : X\} \text{ in } X \\
 \mid \mathbf{a} \triangleright \text{def } X = \mathbf{c}^?; \text{if } * \stackrel{\leq}{=} \mathbf{s} \text{ then } (\mathbf{c} \oplus \mathbf{ok}; \mathbf{s} \oplus \mathbf{ok}) \text{ else } (\mathbf{c} \oplus \mathbf{ko}; \mathbf{s} \oplus \mathbf{ko}; X) \text{ in } X \\
 \mid \mathbf{s} \triangleright \text{def } X = \mathbf{a}!(*) ; \mathbf{a}\&\{\mathbf{ok} : \mathbf{c}!(t), \mathbf{ko} : X\} \text{ in } X
 \end{array}$$

EndPoint Projection (EPP). As shown in [11], there exists a partial function $\llbracket \cdot \rrbracket : \text{CC} \rightarrow \text{SP}$, called EndPoint Projection (EPP), that produces correct implementations of choreographies. EPP produces a parallel composition of processes, one for each process name in the original choreography: $\llbracket C \rrbracket = \prod_{\mathbf{p} \in \text{pn}(C)} \mathbf{p} \triangleright \llbracket C \rrbracket_{\mathbf{p}}$. The rules for computing $\llbracket C \rrbracket$ project the local action performed by the process of interest. For example, $\llbracket \mathbf{p}.e \rightarrow \mathbf{q} \rrbracket_{\mathbf{p}} = \mathbf{q}!(e)$ and $\llbracket \mathbf{p}.e \rightarrow \mathbf{q} \rrbracket_{\mathbf{q}} = \mathbf{p}^?$.

¹ Deviating from [11], we model process values using σ as for CC, for simplicity.

The network presented in Example 2 is exactly the EPP of the choreography in Example 1. Observe that the projection of the conditional in the original choreography for the processes c and s is a branching that supports all the possible choices made by process a in its projected conditional. Producing these branching terms is possible only if, whenever there is a conditional at a process (a in our example), all other processes receive a label that tells them which branch such a process has chosen. (In case the behaviors of the other processes are the same in both cases, producing branching terms is not necessary.) When this cannot be done for a choreography C , the EPP for C is undefined, and we say that C is unprojectable. Conversely, C is *projectable* if $\llbracket C \rrbracket$ is defined.

In the remainder, we relate choreographies to network implementations via a strong labeled reduction bisimilarity \sim . Bisimilarity is defined as usual [24]: it is the union of all bisimulation relations \mathcal{R} , which in our case relate choreographies to networks. A relation \mathcal{R} is one such bisimulation if whenever $C\mathcal{R}N$ we have that, for all σ : (i) $C, \sigma \xrightarrow{\lambda} C', \sigma'$ implies $N, \sigma \xrightarrow{\lambda} N', \sigma'$ for some N' with $C'\mathcal{R}N'$; (ii) $N, \sigma \xrightarrow{\lambda} N', \sigma'$ implies $C, \sigma \xrightarrow{\lambda} C', \sigma'$ for some C' with $C'\mathcal{R}N'$.

Theorem 1 (adapted from [11]). *If C is projectable, then $C \sim \llbracket C \rrbracket$.*

4 Extraction from SP

The finite case. We first investigate *finite SP*, the fragment of SP without recursive definitions, which we use to discuss the intuition behind our extraction.

Definition 1. *We define a rewriting relation \rightsquigarrow on the language of CC extended with terms $\llbracket N \rrbracket$, where N is a network in finite SP, as the transitive closure of:*

$$\begin{array}{c} \frac{N \equiv \mathbf{0}}{\llbracket N \rrbracket \rightsquigarrow \mathbf{0}} \quad \frac{N \equiv \mathbf{p} \triangleright \mathbf{q}! \langle e \rangle; N_p \mid \mathbf{q} \triangleright \mathbf{p}^?; N_q \mid N'}{\llbracket N \rrbracket \rightsquigarrow \mathbf{p}.e \rightarrow \mathbf{q}; \llbracket N_p \mid N_q \mid N' \rrbracket} \\ \frac{N \equiv \mathbf{p} \triangleright \mathbf{q} \oplus l_k; N_p \mid \mathbf{q} \triangleright \mathbf{p} \& \{l_1 : N_{q_1}, \dots, l_n : N_{q_n}\} \mid N'}{\llbracket N \rrbracket \rightsquigarrow \mathbf{p} \rightarrow \mathbf{q}[l_k]; \llbracket N_p \mid N_{q_k} \mid N' \rrbracket} \\ \frac{N \equiv \mathbf{p} \triangleright \text{if } * \stackrel{\leq}{=} \mathbf{q} \text{ then } N_{p_1} \text{ else } N_{p_2} \mid \mathbf{q} \triangleright \mathbf{p}! \langle e \rangle; N_q \mid N'}{\llbracket N \rrbracket \rightsquigarrow \text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } \llbracket N_{p_1} \mid N_q \mid N' \rrbracket \text{ else } \llbracket N_{p_2} \mid N_q \mid N' \rrbracket} \quad \frac{\text{no other rule applies}}{\llbracket N \rrbracket \rightsquigarrow \mathbf{1}} \end{array}$$

A network N in finite SP extracts to a choreography C if $\llbracket N \rrbracket \rightsquigarrow C$.

The last rule guarantees that every network is extractable. Extraction uses structural precongruence (namely, commutativity and associativity of parallel composition) to find matching actions. For finite SP, this is not a problem (the set of networks equivalent to a given one is finite), but it makes extraction nondeterministic, e.g., the network $\mathbf{p} \triangleright \mathbf{q}! \langle e \rangle \mid \mathbf{q} \triangleright \mathbf{p}^? \mid \mathbf{r} \triangleright \mathbf{s}! \langle e' \rangle \mid \mathbf{s} \triangleright \mathbf{r}^?$ extracts both to $\mathbf{p}.e \rightarrow \mathbf{q}; \mathbf{r}.e' \rightarrow \mathbf{s}$ and $\mathbf{r}.e' \rightarrow \mathbf{s}; \mathbf{p}.e \rightarrow \mathbf{q}$. These choreographies are equivalent by Rule [C|Eta-Eta] (Fig. 2). This holds in general, as stated below.

Lemma 1. *If $\llbracket N \rrbracket \rightsquigarrow C_1$ and $\llbracket N \rrbracket \rightsquigarrow C_2$, then $C_1 \equiv C_2$.*

$$\begin{array}{c}
 \frac{}{\rho \triangleright q!(e); B_1 \mid q \triangleright p?; B_2 \xrightarrow{p.e \rightarrow q} \rho \triangleright B_1 \mid q \triangleright B_2} \text{[S|Com]} \\
 \hline
 \frac{}{\rho \triangleright \text{if } * \stackrel{\leftarrow}{=} q \text{ then } B_1 \text{ else } B_2 \mid q \triangleright p!(e); B' \xrightarrow{p \stackrel{\leftarrow}{=} q; \text{then}} \rho \triangleright B_1 \mid q \triangleright B'} \text{[S|Then]}
 \end{array}$$

Fig. 5. Stateful Processes, Abstract Semantics (selected rules).

There is one important design option to consider: what to do with actions that cannot be matched, i.e., processes that will deadlock. There are two alternatives: restrict extraction to lock-free networks (networks where all processes eventually progress, in the sense of [8]); or extract stuck processes to a new choreography term $\mathbf{1}$, with the same semantics as $\mathbf{0}$. We choose the latter option for debugging reasons. Specifically, practical applications of extraction may annotate $\mathbf{1}$ with the code of the deadlocked processes, giving the programmer a chance to see exactly where the system is unsafe, and attempt at fixing it manually. Better yet: since the code to unlock deadlocked processes in process calculi can be efficiently synthesized [8], our method may be integrated with the technique in [8] to suggest an automatic system repair.

Remark 1. If $([N]) \rightsquigarrow C$ and C does not contain $\mathbf{1}$, then N is lock-free. However, even if C contains $\mathbf{1}$, N may still be lock-free: the code causing the deadlock may be dead code in a conditional branch that is never chosen during execution.

Extraction is sound: it yields a choreography that is bisimilar to the original network. Also, for finite SP, it behaves as an inverse of EPP.

Theorem 2. *Let N be in finite SP. If $([N]) \rightsquigarrow C$, then $C \sim N$. Furthermore, if $N = [[C']]$ for some C' , then $([N]) \rightsquigarrow C'$.*

As we show later, the second part of this theorem does not hold in the presence of recursive definitions.

We now restate extraction in terms of a particular graph, which is the hallmark of our development: when we add recursion to SP, we can no longer define extraction as a set of rewriting rules. We first introduce a new abstract semantics for networks, $N \xrightarrow{\alpha} N'$, defined as in Fig. 4 except for the rules for value communication and conditionals, which are replaced by those in Fig. 5 (we omit the obvious rule [S|Else]). In particular, conditionals are now nondeterministic. Labels α are like λ but may now contain expressions (see the new rule [S|Com]); in all other rules, λ is replaced by α . We write $N \xrightarrow{\hat{\alpha}^*} N'$ for $N \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} N'$.

Definition 2. *Given a network N , the Abstract Execution Space (AES) of N is the directed graph obtained by considering all possible abstract reduction paths from N . Its vertices are all the networks N' such that $N \xrightarrow{\hat{\alpha}^*} N'$, and there is an edge between two vertices N_1 and N_2 labeled α if $N_1 \xrightarrow{\alpha} N_2$.*

A Symbolic Execution Graph (SEG) for N is a subgraph of its AES that contains N and such that each vertex $N' \not\stackrel{\leftarrow}{=} \mathbf{0}$ has either one outgoing edge labeled by an η or two outgoing edges labeled $p \stackrel{\leftarrow}{=} q : \text{then}$ and $p \stackrel{\leftarrow}{=} q : \text{else}$.

Intuitively, the AES of N represents all possible evolutions of N (each evolution is a path in this graph). A SEG fixes the order of execution of actions, but still abstracts from the state (and thus considers both branches of conditionals). For networks in finite SP, these graphs are finite.

Given a network N , there is a one-to-one correspondence between SEGs for N and choreographies C such that $([N]) \rightsquigarrow C$. Indeed, given a SEG we can extract a choreography as follows. We start from the initial vertex, labeled N . If there is an outgoing edge with label η to N' , we add η to the choreography and continue from N' . If there are two outgoing edges with labels $\mathbf{p} \stackrel{\leq}{=} \mathbf{q} : \text{then}$ and $\mathbf{p} \stackrel{\leq}{=} \mathbf{q} : \text{else}$ to N_1 and N_2 , respectively, we extract a conditional whose branches are the choreographies extracted by continuing exploration from N_1 and N_2 , respectively. When we reach a leaf, we extract $\mathbf{0}$ or $\mathbf{1}$, according to whether its label is equivalent to $\mathbf{0}$ or not. Conversely, we can build a SEG from a particular rewriting of $([N])$ by following the choreography actions one at a time.

Treating recursive definitions. We now extend extraction to networks with recursive definitions, using SEGs. We need to be careful with the definition of the AES, since including all possible (abstract) executions now may make it infinite (due to recursion unfolding), and thus extraction may not terminate. To avoid this, we only allow recursive definitions to be unfolded (once) if they occur at the head of a process involved in a reduction. With this restriction, we can define the AES and SEGs for a network as in the finite case. These graphs may now contain cycles: a network may evolve into the same term after a few reductions.

Example 3. Consider the following network.

$$\begin{aligned} & \mathbf{p} \triangleright \text{def } X = \mathbf{q}! \langle * \rangle; \mathbf{q} \& \{ \mathbf{L} : \mathbf{q}! \langle * \rangle; X, \mathbf{R} : \mathbf{0} \} \text{ in } \mathbf{q}! \langle * \rangle; X \\ & | \mathbf{q} \triangleright \text{def } Y = \mathbf{p} ?; \mathbf{p} ?; \text{if } * \stackrel{\leq}{=} \mathbf{r} \text{ then } \mathbf{p} \oplus \mathbf{L}; Y \text{ else } \mathbf{p} \oplus \mathbf{R}; \mathbf{0} \text{ in } Y \mid \mathbf{r} \triangleright \text{def } Z = \mathbf{q}! \langle * \rangle; Z \text{ in } Z \end{aligned}$$

This network generates the AES in Fig. 6, which is also its SEG. □

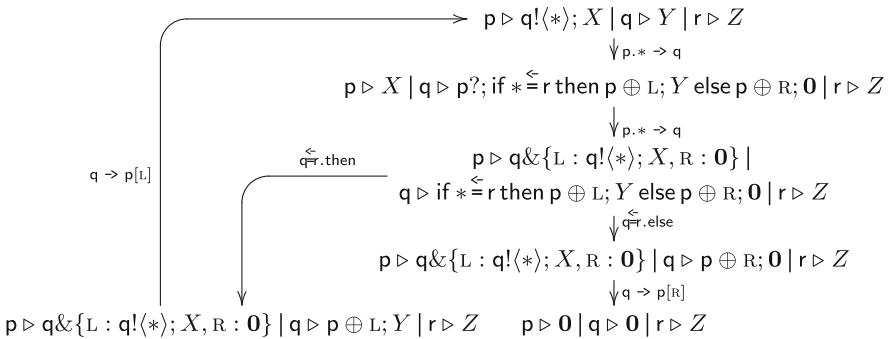


Fig. 6. The AES and SEG for the network in Example 3.

The key insight is that the definitions of recursive procedures are extracted from the loops in the SEG, rather than from the recursive definitions in the source network. This construction typically yields mutually recursive definitions, motivating a small change to CC that does not add expressivity: we replace the constructor $\text{def } X = C_2 \text{ in } C_1$ by top-level procedure definitions, in the style of [12]. A choreography now becomes a pair $\langle \mathcal{D}, C \rangle$, where $\mathcal{D} = \{X_i = C_i\}$ and all procedure calls in either C or the C_i are to some X_i defined in \mathcal{D} .

Definition 3. *The choreography extracted from a SEG is defined as follows. We annotate each node that has more than one incoming edge with a unique procedure identifier. Then, for every node annotated with an identifier, say X , we replace each of its incoming edges with an edge to a new leaf node that contains a special term X (so now the node annotated with X has no incoming edges). This eliminates all loops in the SEG, allowing us to reuse the extraction procedure for the non-recursive case to extract the desired pair $\langle \mathcal{D}, C \rangle$. We get C by extraction starting from the initial network. Then, for each node that we annotated with an X , we extract a choreographic procedure X in \mathcal{D} that has as body the choreography extracted from the graph that starts from that annotated node. Any new leaf node containing a special term X is extracted as a procedure call X .*

Example 4. Consider the SEG in Fig. 3. To extract a choreography, we annotate the topmost node with a procedure identifier X and replace the incoming edge to that node with an edge to a new leaf X . We thus extract X to be

$$p.*q; p.* \rightarrow q; \text{if } q \stackrel{\leftarrow}{=} r \text{ then } q \rightarrow p[L]; X \text{ else } q \rightarrow p[R]; \mathbf{1}$$

and the extracted choreography itself is simply X . The body of X is not projectable (the branches for r are not mergeable, cf. [11]), but it faithfully describes the behavior of the original network. \square

The procedure in Definition 3 always terminates, but sometimes it yields choreographies that starve some processes. As an example, the network

$$\begin{aligned} p \triangleright \text{def } X = q!(*); X \text{ in } X \quad | \quad q \triangleright \text{def } Y = p?; Y \text{ in } Y \\ | \quad r \triangleright \text{def } Z = s!(*); Z \text{ in } Z \quad | \quad s \triangleright \text{def } W = r?; W \text{ in } W \end{aligned} \quad (1)$$


has two SEGs, which extract to the choreographies $\text{def } X = p.* \rightarrow q; X \text{ in } X$ and $\text{def } X = r.* \rightarrow s; X \text{ in } X$, none of which captures all the behaviors of N .

To avoid this problem, we change the definitions of AES and SEGs slightly. We annotate all procedure calls in networks with either \circ or \bullet . The node in the AES corresponding to the initial network has all procedure calls annotated with \circ . There is an edge from N to N' with label α if $N \xrightarrow{\alpha} N'$ and the procedure calls in N' are annotated as follows.

- If executing α does not require unfolding procedure calls, then all calls in N' are annotated as in N .
- If executing α requires unfolding procedure calls, then we annotate all the calls in N' introduced by these unfoldings with \bullet . If N' now has *all* procedure calls annotated with \bullet , we change all annotations to \circ .

We then require loops in a SEG to contain a node where every procedure call is annotated with \circ . This ensures that every procedure call is unfolded at least once before returning to the same node. This holds even if $p \triangleright X$ unfolds to a behavior that calls different procedures, but not X : in order to return to the same node, the newly invoked procedures themselves need to be unfolded.

Example 5. The annotated AES for the network (1) is:



This AES now has the following two SEGs:

$$\begin{array}{cc}
 p \triangleright X^\circ | q \triangleright Y^\circ | r \triangleright Z^\circ | s \triangleright W^\circ & p \triangleright X^\circ | q \triangleright Y^\circ | r \triangleright Z^\circ | s \triangleright W^\circ \\
 \left. \begin{array}{l} r.* \to s \\ p.* \to q \end{array} \right\} & \left. \begin{array}{l} p.* \to q \\ r.* \to s \end{array} \right\} \\
 p \triangleright X^\bullet | q \triangleright Y^\bullet | r \triangleright Z^\circ | s \triangleright W^\circ & p \triangleright X^\circ | q \triangleright Y^\circ | r \triangleright Z^\bullet | s \triangleright W^\bullet
 \end{array}$$

Observe that the self-loops are discarded because they do not go through a node with all \circ annotations. From these SEGs, we can extract two definitions for X :

$$\text{def } X = p.* \to q; r.* \to s; X \text{ in } X \quad \text{and} \quad \text{def } X = r.* \to s; p.* \to q; X \text{ in } X$$

Both of these definitions correctly capture all behaviors of the network. \square

A similar situation may occur if there are processes with finite behavior (no procedure calls): the network

$$p \triangleright \text{def } X = q!(*); X \text{ in } X \mid q \triangleright \text{def } Y = p?; Y \text{ in } Y \mid r \triangleright s!(*) \mid s \triangleright r?$$

can be extracted to the choreography X , with $X = p.* \to q; X$, where r and s never communicate. Hence, we require that if a node in a SEG has more than one incoming edge (it is a “loop” node) and contains processes with finite behavior, then these processes must be deadlocked (being finite, this is trivially verifiable). This ensures that if finite processes are able to reduce, they cannot be in a loop.

Definition 4. A SEG for a network N is valid if all its loops:

- pass through a node where all recursive calls are marked with \circ ;
- start in a node where all processes with finite behavior are deadlocked.

A network N extracts to a choreography C if C can be constructed (as in Definition 3) from a valid SEG for N .

Validity implies, however, that there are some non-deadlocked networks that are not extractable, such as

$$p \triangleright \text{def } X = q!(*); X \text{ in } X \mid q \triangleright \text{def } Y = p?; Y \text{ in } Y \mid r \triangleright \text{def } Z = p?; Z \text{ in } Z$$

for which there is no valid SEG. This is to be expected, since deadlock-freedom is undecidable in SP. We can generalize this observation as a necessary condition for extraction to be defined, in the following theorem.

Theorem 3. *If the AES for a network N does not contain nodes from which a process is always deadlocked, then N is extractable.*

Lemma 1 and the first part of Theorem 2 still hold for extraction in SP with recursion, but the second part of Theorem 2 does not: in general, the projection of a choreography is extracted to a choreography with different procedures, since extraction ignores the actual definitions in the source network.

Theorem 4. *If C is a choreography extracted from a network N , then $N \sim C$.*

We conclude this section with some complexity theoretical considerations.

Lemma 2. *The annotated AES for a network of size n has at most $e^{\frac{2n}{e}}$ vertices.*

Theorem 5. *Extraction from a network of size n terminates in time $O(ne^{\frac{2n}{e}})$.*

As discussed earlier, this time complexity is a dramatic improvement over earlier, comparable work. However, in practice, we may be able to perform even better. Algorithmically, all the required work stems from traversals of the AES, so any reduction in its (explored) size will lead to proportional runtime improvements. Thus, instead of first computing the entire AES and then a valid SEG, we can compute the relevant parts of the AES lazily as we need them, so parts of the AES that are never explored while computing a valid SEG are never generated.

5 Asynchrony

We now discuss an asynchronous semantics for SP, with which we can express new safe behaviors. Most notably, SP can now express asynchronous exchange (Example 6). We also show a novel choreography primitive that successfully captures this pattern, which cannot be described in previous works on choreographic programming, and extend our algorithm to extract it from networks.

Asynchronous SP. Asynchronous communication can be added to SP using standard techniques for process calculi. In the semantics of networks, we add a FIFO queue for each pair of processes. Communications now synchronize with these queues: send actions append a message in the queue of the receiver, and receive actions remove the first message from the queue of the receiver (see [12] for a formalization in an extension of SP).

Example 6. The network $p \triangleright q!(*); q? \mid q \triangleright p!(*); p?$ exemplifies the pattern of asynchronous exchange. This network is deadlocked in synchronous SP, but runs without errors in asynchronous SP: both p and q can send their respective values, becoming ready to receive each other's messages. This behavior is not representable in any previous work on choreographies (including CC from Sect. 3), since all choreographies presented so far can only describe processes that are not deadlocked under a synchronous semantics (see [12] for a formal argument). \square

The multicom. The situation in Example 6 is prototypical of programs that are safe only in an asynchronous setting: a group of processes wants to send messages to a group of receivers, with circular dependencies among communications.

We deal with this situation by means of a new choreography action, which we call a *multicom*. Syntactically, a multicom is a list of communication actions with distinct receivers, which we write $(\tilde{\eta})$. In the unary case, we obtain the usual communications and selections; by removing these from the syntax of CC and adding the multicom, we obtain a more expressive calculus with fewer primitives. The semantics of multicom is given by the following rule, which generalizes (and replaces) both $\llbracket \text{C|Com} \rrbracket$ and $\llbracket \text{C|Sel} \rrbracket$.

$$\frac{I = \{i \mid \mathbf{p}_i.e_i \rightarrow \mathbf{q}_i \in \tilde{\eta}\} \quad v_i = e_i[\sigma(\mathbf{p}_i)/*]}{(\tilde{\eta}); C, \sigma \xrightarrow{(\tilde{\eta})[e_i/v_i]_{i \in I}} C, \sigma[\mathbf{q}_i \mapsto v_i]_{i \in I}} \llbracket \text{C|MCom} \rrbracket$$

Structural precongruence rules for the multicom are motivated by its intuitive semantics: actions inside a multicom can be permuted as long as the senders differ, and sequential multicoms can be merged as long as they do not share receivers and there are no sequential constraints between them (i.e., none of the receivers in the first multicom is a sender in the second one).

$$\frac{\text{pn}(\eta_1) \cap \text{pn}(\eta_2) = \emptyset}{(\dots, \eta_1, \eta_2, \dots) \equiv (\dots, \eta_2, \eta_1, \dots)} \llbracket \text{C|MCom-Perm} \rrbracket$$

$$\frac{\text{rcv}(\eta) \cap \text{rcv}(\nu) = \emptyset \quad \text{rcv}(\tilde{\eta}) \cap \text{snd}(\tilde{\nu}) = \emptyset}{(\tilde{\eta}); (\tilde{\nu}) \equiv (\tilde{\eta}, \tilde{\nu})} \llbracket \text{C|MCom-MCom} \rrbracket$$

From these rules we can derive all instances of $\llbracket \text{C|Eta-Eta} \rrbracket$, e.g.:

$$\mathbf{p}.* \rightarrow \mathbf{q}; \mathbf{r}.* \rightarrow \mathbf{s} \equiv \left(\begin{array}{l} \mathbf{p}.* \rightarrow \mathbf{q} \\ \mathbf{r}.* \rightarrow \mathbf{s} \end{array} \right) \equiv \left(\begin{array}{l} \mathbf{r}.* \rightarrow \mathbf{s} \\ \mathbf{p}.* \rightarrow \mathbf{q} \end{array} \right) \equiv \mathbf{r}.* \rightarrow \mathbf{s}; \mathbf{p}.* \rightarrow \mathbf{q}$$

The problematic program in Example 6 can now be written as $\left(\begin{array}{l} \mathbf{p}.* \rightarrow \mathbf{q} \\ \mathbf{q}.* \rightarrow \mathbf{p} \end{array} \right)$.

Structural precongruence rules for multicom also allow us to define a normal form for choreographies, where no multicom can be split in smaller multicoms.

Extraction. In order to extract choreographies containing multicoms, we alter the definition of the AES for a process network by allowing multicoms as labels for the edges. These can be computed using the following iterative algorithm.

1. For a process \mathbf{p} with behavior $\mathbf{q}!\langle e \rangle; B$ (or $\mathbf{q} \oplus l; B$), set $\text{actions} = \emptyset$ and $\text{waiting} = \{\mathbf{p}.e \rightarrow \mathbf{q}\}$ (resp. $\text{waiting} = \{\mathbf{p} \rightarrow \mathbf{q}[l]\}$).
2. While $\text{waiting} \neq \emptyset$:
 - (a) Move an action η from waiting to actions . Assume η is of the form $\mathbf{r}.e \rightarrow \mathbf{s}$ (the case for label selection is similar).

- (b) If the behavior of s is of the form $a_1; \dots; a_k; r?; B$ where each a_i is either the sending of a value or a label selection, then: for each a_i , if the corresponding choreography action is not in **actions**, add it to **waiting**.

3. Return **actions**.

This algorithm may fail (the behavior of s in step 2(b) is not of the required form), in which case the action initially chosen cannot be unblocked by a multicom.

Example 7. Consider the network from Example 6. Starting with action $q!\langle * \rangle$ at process p , we initialize **actions** = \emptyset and **waiting** = $\{p.* \rightarrow q\}$. We pick the action $p.* \rightarrow q$ from **waiting** and move it to **actions**. The behavior of q is $p!\langle * \rangle; p?$, which is of the form described in step 2(b); the choreography action corresponding to $p!\langle * \rangle$ is $q.* \rightarrow p$, so we add this action to **waiting**, obtaining **actions** = $\{p.* \rightarrow q\}$ and **waiting** = $\{q.* \rightarrow p\}$. Now we consider the action $q.* \rightarrow p$, which we move from **waiting** to **action**, and look at p 's behavior, which is $q!\langle * \rangle; q?$. The choreography action corresponding to $q!\langle * \rangle$ is $p.* \rightarrow q$, which is already in **actions**, so we do not change **waiting**. The set **waiting** is now empty, and the algorithm terminates, returning $\begin{pmatrix} p.* \rightarrow q \\ q.* \rightarrow p \end{pmatrix}$. We would obtain the equivalent $\begin{pmatrix} q.* \rightarrow p \\ p.* \rightarrow q \end{pmatrix}$ by starting with the send action at q . \square

Example 8. As a more sophisticated example, we show how our new choreographies with multicom can model the alternating 2-bit protocol. Here, Alice alternates between sending a 0 and a 1 to Bob; in turn, Bob sends an acknowledgment for every bit he receives, and Alice waits for the acknowledgment before sending another copy of the same bit. Since we are in an asynchronous semantics, we only consider the time when the messages arrive. With this in mind, we can write this protocol as the following network.

$$\begin{aligned} a \triangleright \text{def } X &= (b?; b!\langle 0 \rangle; b?; b!\langle 1 \rangle; X) \text{ in } (b!\langle 0 \rangle; b!\langle 1 \rangle; X) \\ | \quad b \triangleright \text{def } Y &= (a?; a!\langle \text{ack}_0 \rangle; a?; a!\langle \text{ack}_1 \rangle; Y) \text{ in } Y \end{aligned}$$

This implementation imposes exactly the dependencies dictated by the protocol. For example, Alice can receive Bob's acknowledgment to the first 0 before or after Bob receives the first 1. This network extracts to the choreography

$$a.0 \rightarrow b; X \quad \text{where} \quad X = \begin{pmatrix} a.1 \rightarrow b \\ b.\text{ack}_0 \rightarrow a \end{pmatrix}; \begin{pmatrix} a.0 \rightarrow b \\ b.\text{ack}_1 \rightarrow a \end{pmatrix}; X$$

which is a simple and elegant representation of the alternating 2-bit protocol. \square

Extraction for asynchronous SP is still sound, but behavioral equivalence is now an expansion [1, 24], as each communication now takes two steps in asynchronous SP. Its complexity is also no larger than for the synchronous case. The algorithm computing the multicom takes linear time in the size of the multicom produced. Via a one-time preprocessing of the network, we can assume direct references from communication terms in one process to the process it directs its communication at, and from there to the current state of that process. Other

than the above, all constant steps in the algorithm can be seen as an extension of the multicom. Since adding a communication to a multicom removes a potential node in the AES (as we are combining communications), the worst-case time complexity is no worse than in the synchronous case. In practice, this complexity actually gets better when larger multicomms are created, since building these is a much cheaper local operation than exploring graphs that would be larger in terms of nodes as well as edges without the multicomms.

6 Extensions and Applications

We discuss some straightforward modifications of our extraction to cover other scenarios occurring in the literature.

More expressive communications and processes. In real-world contexts, the values stored and communicated by processes are typed, and the receiver process can also specify how to treat incoming messages [12]. This means that communication actions now have the form $p.e \rightarrow q.f$, where f is the function consuming the received message, and systems may deadlock because of typing errors. Our construction applies without changes to this scenario.

Some works allow processes to store several values, used via variables [5, 6]. Again, dealing with this situation does not require any changes to our algorithm.

Local conditionals. Many choreography models allow for a local conditional construct, i.e., if $p.e$ then C_1 else C_2 [6, 14, 21]. Dealing with this construct is simple: the if and then transitions now can occur whenever a process has a conditional as top action, since they no longer require synchronization with other processes.

Choreography Specifications. So far, we have considered choreographies that describe concrete implementations, i.e., processes are equipped with storage and local computational capabilities. However, choreographies have also been advocated for the specification of communication protocols. Most notably, multiparty session types use choreographies to define types used in the verification of process calculi [17]. While there are multiple variants of multiparty session types, the one so far most used in practice is almost identical to a simplification of SP. In this variant, each pair of participants has a dedicated channel, and communication actions refer directly to the intended sender/recipient as in SP (see, e.g., the theory of [6, 9, 10, 21] and the practical implementations in [16, 20, 23]). To obtain multiparty session types from SP (and CC), we just need to: remove the capability of storing values at processes; replace message values with constants (representing types, which could also be extended to subtyping in the straightforward way); and make conditionals nondeterministic (since in types we abstract from the precise values and expression used by the evaluator). These modifications do not require any significant change to our approach, since our AES already abstracts from data and thus our treatment of the conditional is already nondeterministic. For reference, we can simply treat the standard construct for an internal choice at a process $p - C_1 \oplus_p C_2 -$ as syntactic sugar for a local conditional like if $p.\text{coinflip}$ then C_1 else C_2 .

References

1. Arun-Kumar, S., Hennessy, M.: An efficiency preorder for processes. *Acta Inf.* **29**(8), 737–760 (1992)
2. Basu, S., Bultan, T.: Choreography conformance via synchronizability. In: WWW, pp. 795–804 (2011)
3. Business Process Model and Notation. <http://www.omg.org/spec/BPMN/2.0/>
4. Brand, D., Zafropulo, P.: On communicating finite-state machines. *J. ACM* **30**(2), 323–342 (1983)
5. Carbone, M., Honda, K., Yoshida, N.: Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.* **34**(2), 8 (2012)
6. Carbone, M., Montesi, F.: Deadlock-freedom-by-design: multiparty asynchronous global programming. In: Giacobazzi, R., Cousot, R. (eds.) POPL, pp. 263–274. ACM (2013)
7. Carbone, M., Montesi, F., Schürmann, C.: Choreographies, logically. In: Baldan, P., Gorla, D. (eds.) CONCUR 2014. LNCS, vol. 8704, pp. 47–62. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-44584-6_5](https://doi.org/10.1007/978-3-662-44584-6_5)
8. Carbone, M., Dardha, O., Montesi, F.: Progress as compositional lock-freedom. In: Kühn, E., Pugliese, R. (eds.) COORDINATION 2014. LNCS, vol. 8459, pp. 49–64. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-43376-8_4](https://doi.org/10.1007/978-3-662-43376-8_4)
9. Carbone, M., Lindley, S., Montesi, F., Schürmann, C., Wadler, P.: Coherence generalises duality: A logical explanation of multiparty session types. In: Desharnais, J., Jagadeesan, R. (eds.) 27th International Conference on Concurrency Theory, CONCUR 2016, August 23–26, 2016, Québec City, Canada, vol. 59 of LIPIcs, pp. 33:1–33:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2016)
10. Coppo, M., Dezani-Ciancaglini, M., Yoshida, N., Padovani, L.: Global progress for dynamically interleaved multiparty sessions. *Math. Struct. Comput. Sci.* **26**(2), 238–302 (2016)
11. Cruz-Filipe, L., Montesi, F.: A core model for choreographic programming. In: FACS 2016. LNCS. Springer (accepted for publication)
12. Cruz-Filipe, L., Montesi, F.: Choreographies, divided and conquered. CoRR, abs/1602.03729 (2016). Submitted for publication
13. Cruz-Filipe, L., Montesi, F.: Choreographies in practice. In: Albert, E., Lanese, I. (eds.) FORTE 2016. LNCS, vol. 9688, pp. 114–123. Springer, Cham (2016). doi:[10.1007/978-3-319-39570-8_8](https://doi.org/10.1007/978-3-319-39570-8_8)
14. Preda, M., Gabbrielli, M., Giallorenzo, S., Lanese, I., Mauro, J.: Dynamic choreographies. In: Holvoet, T., Viroli, M. (eds.) COORDINATION 2015. LNCS, vol. 9037, pp. 67–82. Springer, Cham (2015). doi:[10.1007/978-3-319-19282-6_5](https://doi.org/10.1007/978-3-319-19282-6_5)
15. Deniérou, P.-M., Yoshida, N.: Multiparty session types meet communicating automata. In: Seidl, H. (ed.) ESOP 2012. LNCS, vol. 7211, pp. 194–213. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-28869-2_10](https://doi.org/10.1007/978-3-642-28869-2_10)
16. Honda, K., Mukhamedov, A., Brown, G., Chen, T.-C., Yoshida, N.: Scribbling interactions with a formal foundation. In: Natarajan, R., Ojo, A. (eds.) ICD-CIT 2011. LNCS, vol. 6536, pp. 55–75. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-19056-8_4](https://doi.org/10.1007/978-3-642-19056-8_4)
17. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. *J. ACM* **63**(1), 9 (2016)
18. Lange, J., Tuosto, E.: Synthesising choreographies from local session types. In: Koutny, M., Ulidowski, I. (eds.) CONCUR 2012. LNCS, vol. 7454, pp. 225–239. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-32940-1_17](https://doi.org/10.1007/978-3-642-32940-1_17)

19. Lange, J., Tuosto, E., Yoshida, N.: From communicating machines to graphical choreographies. In: Rajamani, S.K., Walker, D. (eds.) Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, pp. 221–232. ACM, 15–17 January 2015
20. Montesi, F.: Choreographic Programming. Ph.D. thesis, IT University of Copenhagen (2013). http://fabriziomontesi.com/files/choreographic_programming.pdf
21. Montesi, F., Yoshida, N.: Compositional choreographies. In: D’Argenio, P.R., Mellgratti, H. (eds.) CONCUR 2013. LNCS, vol. 8052, pp. 425–439. Springer, Heidelberg (2013). doi:10.1007/978-3-642-40184-8_30
22. Mostrous, D., Yoshida, N., Honda, K.: Global principal typing in partially commutative asynchronous sessions. In: Castagna, G. (ed.) ESOP 2009. LNCS, vol. 5502, pp. 316–332. Springer, Heidelberg (2009). doi:10.1007/978-3-642-00590-9_23
23. Ng, N., Yoshida, N.: Pabble: Parameterised scribble for parallel programming. In: 22nd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP 2014, Torino, Italy, pp. 707–714. IEEE Computer Society, 12–14 February 2014
24. Sangiorgi, D., Walker, D.: The π -calculus: A Theory of Mobile Processes. Cambridge University Press, Cambridge (2001)
25. W3C WS-CDL Working Group. Web services choreography description language version 1.0 (2004). <http://www.w3.org/TR/2004/WD-ws-cdl-10-20040427/>