

# Connectors meet Choreographies

Anonymous Author(s)

---

## Abstract

---

Choreographic Programming and Exogenous Coordination are two powerful programming paradigms for developing correct concurrent software. Choreographic Programming enables a correct-by-construction synthesis procedure of concurrent process implementations that guarantees deadlock-freedom. Exogenous Coordination enables an elegant compositional approach to the development of inter-process connectors, where complex connectors can be built by assembling simpler ones (like synchronous barriers and asynchronous multi-casts).

We present Cho-Reo-graphies (CR), a new choreography calculus where these two research lines are combined for the first time, merging their benefits. In CR, choreographies are parametric in the connectors through which processes communicate. CR is the first choreography calculus where different communication semantics (determined by the connectors) can be freely mixed in the same program. Since connectors are user-defined, CR also supports many communication semantics that could not be achieved before for choreographies. We develop a decidable static analysis that guarantees that a choreography in CR and its user-defined connectors are compatible, and prove that this compatibility guarantees deadlock-freedom in process implementations synthesised from choreographies.

Digital Object Identifier 10.4230/LIPIcs...

## 1 Introduction

### Background

Programming concurrent software is hard, because it is difficult to reason about how independent processes interact at runtime (e.g., by exchanging messages). Empirical studies reveal that two major reasons for this difficulty are: (i) while programmers have clear intentions about the order in which interactions should take place, existing tools do not adequately support the translation of such intentions to code [36]; (ii) it is easy to cause unintended side-effects when composing multiple interaction protocols together [34].

The challenge of concurrent programming has driven decades of research on new programming models. A particularly fruitful idea was that of providing a native language abstraction for interaction, rather than modelling it as a side-effect. Two research lines in particular were built upon this idea, but following very different directions.

The first research line is that on *Choreographic Programming* [37]. Central to Choreographic Programming is the *choreography*, a programming artefact that specifies a concurrent system in terms of the admissible interactions among its constituent processes, by using an “Alice and Bob” notation that disallows writing mismatched I/O actions (e.g., a send without a corresponding receive). Through *EndPoint Projection* (EPP), implementations in process languages can subsequently be synthesised from a choreography, which faithfully realise the interactions given in the choreography and are guaranteed to be deadlock-free by construction, even in the presence of arbitrary compositions [17, 38].

Previous work studied models for choreographies in settings with different interaction semantics (e.g., synchronous, asynchronous, multi-cast). Regardless of the specifics, common to all existing models of Choreographic Programming is the fact that the nature of interactions is “hardcoded”: each model proposes new syntax and semantics, so results have to be proven from scratch every time and cannot be combined. None of the existing models supports, for instance, mixing synchronous with asynchronous interactions within the same choreography.



© Anonymous Author(s);

licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

This is a serious limitation, which raises the challenge of finding a unifying framework. Furthermore, there is still no indication of how other interesting interaction semantics (e.g., barriers) can be introduced to choreographies.

The second research line is that on *Exogenous Coordination* [2, 3, 33], where interaction protocols are developed separately from the code of the processes that will enact them. Process programs can then be modularly composed with protocols, given as *connectors*, which dictate which interactions take place by accepting/offering messages from/to processes. Connectors offer an elegant way of programming different semantics for interactions, starting from basic instances (e.g., synchronous and asynchronous point-to-point channels) and then composing them to create more sophisticated protocols.

Exogenous Coordination is strong where Choreographic Programming is weak: connectors offer an elegant model for defining different communication semantics, given in terms of arbitrary protocols. However, the opposite also holds: Choreographic Programming is strong where Exogenous Coordination is weak. When we compose processes with connectors, we do not have a global view of the system as in choreographies. Therefore, the programmer has no way to define the global flow of information in the system when multiple connectors are used separately by different parts of the system. Since we cannot know if the entire system is consistent, by considering both processes and their connectors, we may have deadlocks (which are prevented in Choreographic Programming).

### Contribution

Choreographic Programming and Exogenous Coordination have complementary good properties. Can we then hope to combine them and obtain the best of the two worlds? In this paper, we answer this question in the affirmative, developing the first investigation of how connectors and choreographies can be integrated. We believe that our results represent the beginning of an interesting research line on concurrent programming. In the rest of this section, we summarise the achievements reported in this article; we state what we believe this work may lead to in the future in § 6.

We combine the best aspects of Choreographic Programming and Exogenous Coordination by presenting a new calculus of choreographies, called *Cho-Reo-graphies* (CR), in which the interactions among processes specified in a choreography are animated by arbitrary, user-defined, Reo connectors. CR allows for mixing different Reo connectors in the same choreography, making it for the first time possible to write choreographies where different interactions can have different communication semantics. Furthermore, by tapping into the expressivity of Reo connectors, we can endow choreographies with hitherto unexplored communication semantics, such as alternators or barriers. This makes CR more expressive and a generalisation of existing models of Choreographic Programming. Through EPP, choreographies can be compiled into concurrent implementations in a process language. In these implementations, the same Reo connectors as in the original choreography animate interactions among processes. We show that these processes are deadlock-free, provided that the original choreography is compatible with the Reo connectors.

By allowing, for the first time, different communication semantics (i.e., Reo connectors) to be freely mixed in the same choreography, new technical challenges of formalisation and decidability arise. On the formalisation front, the main challenge is to balance expressiveness and comprehensibility (the formal semantics of the calculus should be easy enough to explain and understand). On the decidability front, the main challenges are proving that (1) deadlock-freedom is generally undecidable in our calculus, but (2) we can establish deadlock-freedom for a large subset of the language.

By leveraging existing work, distributed implementations of Reo connectors (e.g., in Scala [39, 40] or in Java [31]) can be automatically generated and deployed on different machines (e.g., for distributed objects and components); our approach is compatible with this existing work. As such, a practical tool based on the work in this paper can in principle be built on top of existing code generators for distributed implementations of Reo connectors.

## Structure

In § 2, we motivate our work by means of an example. In § 3 we give the background on Reo connectors necessary for our development. In § 4 we develop our language integrating choreographies and Reo connectors, and prove the usual results on deadlock-freedom. In § 5 we introduce the target language for EndPoint Projection, show how to compile our choreographies, and prove the operational correspondence between a choreography and its projection. We conclude in § 6 with a discussion on the directions in which this work can be extended. Proofs are given in the Appendix.

## Related work

We already covered the main references to previous work and how it falls short of serving our aim. We briefly recap related work.

In Choreographic Programming, developers write high-level programs that describe the intended communications from a global viewpoint, assuming a fixed communication semantics (e.g., asynchronous point to point). This enables a correct-by-construction synthesis procedure of concurrent process implementations that guarantees deadlock-freedom. Choreographies have been studied in settings with different communication semantics, including synchronous [15, 22, 32], asynchronous [17, 24, 27, 38], one-to-many [18, 20], and many-to-one [16, 35]. In our model, these communication semantics are simply instances of what we can do. But since we tap into the generality of Reo connectors, we can also do more (for example, we illustrate how to use barriers). Our development also extends the line of work on out-of-order execution for choreographies, originally from [17], where non-interfering interactions may proceed concurrently. This style allows for more safe behaviours in the semantics of choreographies (by swapping non-interfering communications), which the programmer gets for free (concurrency is inferred). Having out-of-order execution also simplifies our syntax: as in many other works, we do not need to provide for a parallel operator in choreographies, since most parallel behaviour is already captured by out-of-order execution (cf. [17] for details). As in previous work, our model captures asynchronous behaviour without requiring the programmer to reason about it in choreographies: communications are still specified as atomic interactions, which may be asynchronously reduced at runtime in a safe way (cf. [17, 26, 38]).

In Exogenous Coordination, developers program interaction protocols for communicating processes, called connectors, separately from the internal code of each process. This enables an elegant compositional approach to the development of protocols, where complex protocols can be built by assembling simpler ones (like synchronous barriers and asynchronous multi-casts). Exogenous Coordination has been studied extensively over the last two decades [2, 3, 33]. Examples of models of Exogenous Coordination are the algebras of connectors [11, 12], the algebra of stateless connectors [13], and constraint automata [8]; examples of languages are (interactions in) BIP [9, 10], Ptolemy [14, 41], and Reo [4, 5].

## 2 Motivating Example & Approach

We now present an example to introduce the concept of choreographies and the problem we are interested in studying. This example will be used as running example throughout the whole article to illustrate the different concepts we introduce at each stage.

► **Example 1** (Book sale). Consider a scenario where buyer Alice (**a**) wants to buy a book from seller Carol (**c**), facilitated by a bank (**b**) and a shipper (**s**). First, Alice sends the title of the book to Carol. Carol then replies to Alice with the price of the book. If Alice is happy with the price, she notifies Carol, the bank, and the shipper that the purchase proceeds. Alice subsequently sends the money to the bank (who transfers it to Carol's account), and Carol sends the book to the shipper (who dispatches it to Alice). The choreography for this scenario looks as follows:

1.  $a\langle title \rangle \rightarrow c; c\langle price \rangle \rightarrow a;$
2.  $\text{if } a.\text{happy} \text{ then } ( a \rightarrow c[ok]; a \rightarrow b[ok]; a \rightarrow s[ok];$   
 $\quad a\langle money \rangle \rightarrow b; c\langle book \rangle \rightarrow s )$   
 $\text{else } ( a \rightarrow c[ko]; a \rightarrow b[ko]; a \rightarrow s[ko] )$

In previous choreography models, the nature of the interactions is fixed: depending on the model, these interactions are either all synchronous or all asynchronous. This is not flexible enough, because requirements may be different for each interaction.

► **Example 2** (Book sale). The book sale scenario has the following requirements:

- Because there are no strict timing constraints, it is reasonable for Alice and Carol to communicate asynchronously in line 1.
- Because the same label (*ok* or *ko*) is sent from Alice to Carol, the bank, and the shipper, it makes sense to combine these communications in a multi-cast.
- It is better for Alice to send her money to the bank as late as possible (e.g., because she receives interest on her money). Therefore, Alice does not want to send her money until she knows the others have received her *ok*-label. Thus, the multi-cast of this label should be synchronous (i.e., handshake between Alice, Carol, the bank, and the shipper).
- Alice and Carol may not trust each other: Alice does not want to send her money before Carol has sent her book, and vice versa. To resolve this impasse, Alice and Carol should synchronously (barrier-like) send money and book, so they are sure each of them holds her end of the bargain.

As we show in the next sections, our choreography model equipped with connectors is powerful enough to express all these interactions (and more). The idea is to tag every interaction between a sender  $p$  and a receiver  $q$  with the name of the particular connector through which the interaction transpires. For instance, a value communication from  $p$  to  $q$  through a synchronous channel `sync` is expressed as  $p \rightarrow q \text{ thru sync}$ , while the same communication through an asynchronous channel `async` is expressed as  $p \rightarrow q \text{ thru async}$ . If all interactions in a choreography transpire through the same (type of) connector, as in all existing choreography approaches, tags become redundant and can be omitted.

## 3 Constraint Automata and Reo

We view processes in a concurrent system as black boxes with interfaces consisting of *ports*. For a process to send (receive) a message to (from) another process, it performs an output

action (input action) on one of its own output ports (input ports), but without specifying a receiver (sender). Instead, a separate connector, connected to the ports of the processes, decides how messages flow from senders' output ports to receivers' input ports.

We model connectors using a general existing model of Exogenous Coordination, namely *constraint automata* [8]. The transitions of a constraint automaton model possible synchronous message flows (i.e., interactions) through a connector in a particular state. Constraint automata are parametrised over a language of constraints  $\Phi$  to specify transition labels. In this work, we consider  $\Phi_{P,M}$  to be a language of constraints over two sets  $P$  and  $M$ , of ports and *memory cells* (storage space local to a connector). Constraints are finite sets of formulas of the form  $p_1 \rightarrow p_2$  (port  $p_1$  passes a message to port  $p_2$ ),  $m_1 \rightarrow m_2$  (cell  $m_1$  passes a message to cell  $m_2$ ),  $p \rightarrow m$  (port  $p$  passes a message to cell  $m$ ), and  $m \rightarrow p$  (cell  $m$  passes a message to port  $p$ ).<sup>1</sup> Furthermore, we require that all terms on the right-hand side of formulas in the same constraint in a transition label be distinct, to ensure that every port/cell is assigned a unique message (e.g.,  $p_1 \rightarrow m \wedge p_2 \rightarrow m$  is forbidden, but  $p \rightarrow m_1 \wedge p \rightarrow m_2$  is allowed).

Formally, a constraint automaton is a tuple  $(S, P, M, \rightarrow, s_0, \mu_0)$ , where  $S$  is a finite set of states,  $P$  is a finite set of ports,  $M$  is a finite set of memory cells,  $\rightarrow \subseteq S \times \Phi_{P,M} \times S$  is a transition relation,  $s_0 \in S$  is an initial state, and  $\mu_0$  is an initial memory snapshot. The initial memory snapshot is a function that maps the memory cells in  $M$  to their initial content. We denote  $(s, \phi, s') \in \rightarrow$  as  $s \xrightarrow{\phi} s'$ , for short. A transition  $s \xrightarrow{\phi} s'$  means that, from state  $s$ , a subset of the ports in  $P$  can interact according to  $\phi$ . Constraints in  $\Phi_{P,M}$  also control the evolution of memory snapshots as transitions are made. Instead of formalizing this separately (e.g., in terms of runs and languages [29]), we combine it directly in the semantics of our choreographies, in § 4.

In this work, we restrict ourselves to a subset of constraint automata that satisfy two additional assumptions on their transitions. First, the occurrences of each port in transition labels are either all as sender (it occurs only on the left-hand side of constraints) or all as receiver (it occurs only on the right-hand side of constraints); this constraint is imposed on the whole automaton. For instance, a constraint automaton cannot have two transitions where one is labelled by  $p_1 \rightarrow p_2$  and the other by  $p_2 \rightarrow p_3$ . Secondly, the transition relation of each automaton is deterministic on the ports used, i.e., for any given state  $s$ , if there are two distinct transitions  $s \xrightarrow{\phi_1} s_1$  and  $s \xrightarrow{\phi_2} s_2$ , then the sets of ports used in  $\phi_1$  and  $\phi_2$  must be distinct (but they can overlap, or one be a strict subset of the other). The first assumption simplifies defining the semantics of choreographies in § 4; the second assumption ensures that this semantics is deterministic, in line with previous work on choreographies.

► **Example 3.** Figure 1 shows example constraint automata for useful connectors.

*Sync* models a synchronous channel. Indefinitely, this connector lets two processes synchronously send and receive a message through ports  $p_1$  and  $p_2$ . *Async1* models an asynchronous channel with a 1-capacity buffer (using a memory cell  $m$ ). Indefinitely, first, this connector lets a process send a message through port  $p_1$ ; subsequently, it lets a process receive the message through port  $p_2$ . *Async2* models an asynchronous channel with a 2-capacity buffer. *SyncMulti2* and *SyncMulti3* model synchronous multi-cast connectors for two and three receivers. *Async1Multi2* models an asynchronous multi-cast connector for two receivers. *Barrier* models a barrier send/receive connector. ◀

Choreography programmers can model connectors by explicitly defining constraint automata. Alternatively, programmers can model connectors by *composing* constraint automata

<sup>1</sup> We can formalise  $\Phi$  as an equational theory as in [8, 29], but this is unnecessary for our development.

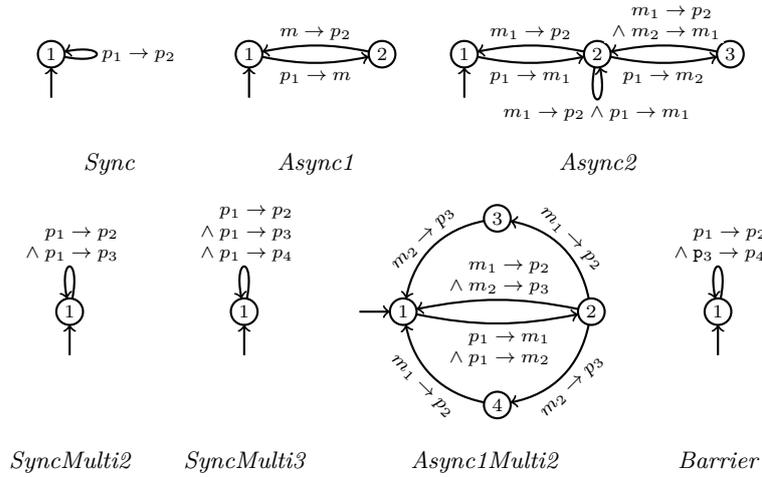


Figure 1 Example constraint automata.

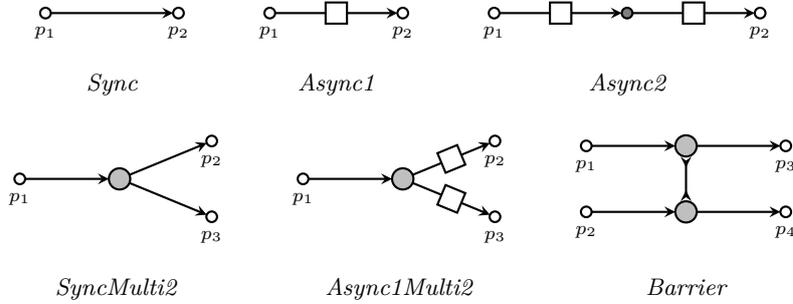


Figure 2 Example Reo connectors for the constraint automata in Figure 1.

from basic primitive constraint automata, using a synchronous product operator [8]. A significant advantage of this latter approach is that there exists an intuitive and user-friendly graphical syntax for constraint automaton product expressions. In this graphical syntax, called *Reo* [4, 5], programmers draw constraint automaton product expressions as data-flow graphs between (ports of) processes. As an example, Figure 2 shows Reo connectors for the constraint automata in Figure 1. Reo conveniently hides from choreography programmers the intimidating act of explicitly composing constraint automata, without sacrificing generality: Reo is complete for constraint automata (under the instantiation of  $\Phi$  considered in this paper), meaning that every constraint automaton can be expressed as a Reo connector [6, 7]. Moreover, tooling exists to animate flows of messages through Reo connectors (<http://reo.project.cwi.nl>).

We use (constraint) automata and (Reo) connectors interchangeably: at the semantics level, we use automata, while at the syntax level, we use connectors.

## 4 Cho-Reo-graphies

### 4.1 Overview

We now present our choreography calculus that combines choreographies with Reo connectors, called *Cho-Reo-graphies* (CR). A choreography describes the behaviour of a set of processes.

$$\begin{aligned}
C &::= \tilde{\eta} \text{ thru } \gamma; C \mid \text{if } p.e \text{ then } C_1 \text{ else } C_2 \mid \text{def } X = C_2 \text{ in } C_1 \mid X \mid \mathbf{0} \\
\eta &::= p\langle e \rangle \rightarrow q.x \mid p \rightarrow q[\ell] \mid \boxed{q.x?v} \mid \boxed{q[\ell]}
\end{aligned}$$

■ **Figure 3** Cho-Reo-grammies, syntax. The boxed terms are runtime terms, necessary for defining the semantics, but which are not meant to be used by programmers.

For simplicity, we assume that values are untyped; treating value types is straightforward and analogous to [15, 17, 23].

The syntax of CR is displayed in Figure 3. We use  $C, C', C_1, \dots$  to range over choreographies;  $p, q, \dots$  to range over processes;  $\eta, \eta', \eta_1, \dots$  to range over interactions (and  $\tilde{\eta}, \tilde{\eta}', \tilde{\eta}_1, \dots$  to range over sets of them);  $\gamma, \gamma', \gamma_1, \dots$  to range over connector names;  $X$  to range over procedure names;  $e, e', \dots$  to range over (side-effect free) expressions;  $v, v', \dots$  to range over values; and  $\ell, \ell', \dots$  to range over selection labels. Each process  $p$  has associated with it a finite set of local variables  $\text{var}_p = \{x_1, \dots, x_n\}$ , and expressions are assumed to be an inductively defined set including  $\text{var}_p$  as base cases.

We comment briefly on label selections. It is standard practice in Choreographic Programming to distinguish value communications, which are used to exchange data, from label selections, which are used to propagate decisions regarding control flow. In the choreography in Example 1, Alice uses label selections (*ok* or *ko*) to communicate her choice of whether to buy the book or not to Carol and the bank. Although this can be encoded using value communications [22], it is useful to distinguish them, as they are usually treated differently in implementations. Also, as discussed in § 5, label selections are instrumental in generating process implementations automatically. The syntax of labels is unspecified.

We write  $p\langle e \rangle \rightarrow \{q_1.x_1, \dots, q_n.x_n\}$  (value multicast) and  $p \rightarrow \{q_1, \dots, q_n\}[\ell]$  (label multicast) to abbreviate  $\{p\langle e \rangle \rightarrow q_1.x_1, \dots, p\langle e \rangle \rightarrow q_n.x_n\}$  and  $\{p \rightarrow q_1[\ell], \dots, p \rightarrow q_n[\ell]\}$ , respectively. Also, we write  $p\langle x \rangle \rightarrow q$  for  $p\langle x \rangle \rightarrow q.x$  (i.e., if sender  $p$  sends the value in local variable  $x$ , and if receiver  $q$  stores the received value in a local variable with the same name, we omit the name). If  $\tilde{\eta}$  is a singleton, we omit curly braces.

► **Example 4** (Book sale). The choreography for our running example (cf. Example 1) can be written as follows in CR.

1.  $a\langle \text{title} \rangle \rightarrow c \text{ thru } a2c;$
2.  $c\langle \text{price} \rangle \rightarrow a \text{ thru } c2a;$
3.  $\text{if } a.\text{happy} \text{ then } ( a \rightarrow \{c, b, s\}[\text{ok}] \text{ thru } a2cbs;$
4.  $\quad \{a\langle \text{money} \rangle \rightarrow b, c\langle \text{book} \rangle \rightarrow s\} \text{ thru } ac2bs; \mathbf{0}$
5.  $\quad \text{else } ( a \rightarrow \{c, b, s\}[\text{ko}] \text{ thru } a2cbs; \mathbf{0}$

◀

The semantics of most terms is standard. In  $\text{if } p.e \text{ then } C_1 \text{ else } C_2$ , process  $p$  evaluates expression  $e$ ; if this results in *true*, the choreography proceeds as  $C_1$ , and otherwise, as  $C_2$ . In  $\text{def } X = C_2 \text{ in } C_1$ , procedure  $X$  is defined as  $C_2$ ; it can then be invoked as  $X$  from both  $C_1$  and  $C_2$ .  $\mathbf{0}$  indicates successful termination.

The interesting new bit is the semantics of a term  $\tilde{\eta} \text{ thru } \gamma$ . Informally,  $\tilde{\eta} \text{ thru } \gamma$  specifies that all communications in  $\tilde{\eta}$  occur through connector  $\gamma$ . More precisely, in  $p\langle e \rangle \rightarrow q.x \text{ thru } \gamma$ , process  $p$  (the *sender*) evaluates expression  $e$  and offers the resulting value to connector  $\gamma$ ; the connector eventually accepts and offers it to process  $q$  (the *receiver*), who stores it in its local variable  $x$ . The behaviour of  $p \rightarrow q[\ell] \text{ thru } \gamma$  is similar, except that  $p$  offers a selected

label instead of a normal value. Label selections do not change the state of the receiving process; their role in synthesizing process implementations is discussed in § 5.

The boxed terminals in Figure 3 are *runtime terms*, meant to be used only in defining the formal semantics and not by programmers. They arise because connectors may have a multi-step semantics (i.e., they do not necessarily synchronise sends with receives). In particular,  $q.x?v$  is obtained when a communication  $p\langle e \rangle \rightarrow q.x$  is partially executed, and  $p$  has already sent its value, but  $q$  has not yet received it;  $q[\ell]$  arises similarly – see Example 7 below.

For the semantics of  $\tilde{\eta} \text{ thru } \gamma$  to be well-defined,  $\tilde{\eta}$  must satisfy two conditions. First, all interactions in  $\tilde{\eta}$  must have *distinct receivers*: if  $p_1\langle e_1 \rangle \rightarrow q_1.x_1 \in \tilde{\eta}$  and  $p_2\langle e_2 \rangle \rightarrow q_2.x_2 \in \tilde{\eta}$ , then  $q_1 \neq q_2$ . This ensures that the value received by a receiver is uniquely defined. Second, all sends must be *consistent*: if a process  $p$  is involved in multiple interactions in the same set, then they are either all communications of the same expression or all selections of the same label.

► **Example 5** (Book sale). In the context of our running example, the following interaction sets are allowed (✓) or disallowed (✗) by our conditions for distinct receivers and consistent sends.

- ✓  $\{a\langle money \rangle \rightarrow b, c\langle book \rangle \rightarrow s\}$  – Alice sends money to the bank, while Carol sends a book to the shipper (distinct receivers; consistent sends).
- ✗  $\{a\langle money \rangle \rightarrow b, c\langle money \rangle \rightarrow b\}$  – Both Alice and Carol send money to the bank, but the bank can receive only from one sender at a time (receivers are not distinct).
- ✗  $\{c\langle price \rangle \rightarrow a, c\langle book \rangle \rightarrow s\}$  – Carol sends both the price to Alice and the book to the shipper, but Carol can send only one value at a time (sends are not consistent).
- ✓  $\{a \rightarrow b[ok], a \rightarrow c[ok], a \rightarrow s[ok]\}$  – Alice sends label *ok* to Carol, the bank, and the shipper (distinct receivers; consistent sends). ◀

## 4.2 Formal semantics

The semantics of CR is a reduction semantics parametrised over a *connector mapping*: a function  $\mathcal{G}$  from connector names to automata. Intuitively,  $\mathcal{G}(\gamma)$  denotes the automaton that models connector  $\gamma$  used in the choreography; the set  $P$  of ports in each  $\mathcal{G}(\gamma)$  is simply a one-to-one mapping (i.e., a renaming of) the set of processes that use the connector.<sup>2</sup>

► **Example 6** (Book sale). The connector names that occur in the choreography in Example 4 are *a2c*, *c2a*, *a2cbs*, and *ac2bs*. The requirements in Example 2 subsequently give rise to the following connector mapping:

$$\left\{ \begin{array}{l} a2c \mapsto Async1[a/p_1, c/p_2], \\ c2a \mapsto Async1[c/p_1, a/p_2], \\ a2cbs \mapsto SyncMulti3[a/p_1, b/p_2, c/p_3, s/p_4], \\ ac2bs \mapsto Barrier[a/p_1, b/p_2, c/p_3, s/p_4] \end{array} \right\}$$

where  $Async1[a/p_1, c/p_2]$  denotes automaton *Async1* in Figure 1, with  $a$  substituted for  $p_1$ , and  $c$  for  $p_2$  (and likewise in the other mappings). Under this connector mapping, thus, Alice and Carol communicate via asynchronous channels (*a2c* and *c2a*) to exchange title and price;

<sup>2</sup> This means that each process can interact at most through one port in each automaton.

Alice, Carol, the bank, and the shipper communicate via synchronous multi-cast (**a2cbs**) to exchange *ok* or *ko*, and via barrier sends/receives (**ac2bs**) to exchange money and book. ◀

► **Remark.** The book sale scenario illustrates an important design decision, namely the separation between *intention* and *realisation*: a choreography defines *what* interactions are intended (e.g., communications of the money from Alice to the Bank and the book from Carol to the shipper), while the connectors define *how* these communications are realised (e.g., synchronously or asynchronously). As a result, every interaction has to be expressed in two places, serving two complementary purposes: as “specifications” in the choreography and as “implementations” in the connectors (automaton transitions). As usual, implementations should respect specifications; we address this in § 4.4. ◀

The reduction relation for CR under a given  $\mathcal{G}$  is denoted as  $\rightsquigarrow_{\mathcal{G}}$ ; it ranges over triples  $C, \sigma, \mathcal{A}$ , where  $C$  is a choreography,  $\sigma$  is a choreography state function (mapping each process to a mapping of its variables to values, i.e.,  $\sigma(p.x)$  is the value stored at variable  $x$  in process  $p$ ), and  $\mathcal{A}$  is an automaton state function (mapping each connector name  $\gamma$  in the domain of  $\mathcal{G}$  to a pair  $\langle s, \mu \rangle$  of the state and memory snapshot of the automaton  $\mathcal{G}(\gamma)$ ). Before introducing the formal rule for communications, we give an example that discusses the intuition.

► **Example 7 (Book sale).** Returning to our running example, suppose that Alice wants to buy the book titled Foo.

Let  $\mathcal{G}$  denote the connector mapping in Example 6, let  $C'$  denote lines 2–3 in Example 4, let  $\sigma_0$  denote the initial choreography state function such that  $\sigma_0(\mathbf{a}.title) = \text{"foo"}$ , let  $\sigma'_0 = \sigma_0[\mathbf{c}.title \mapsto \text{"foo"}]$ , and let  $\mathcal{A}_0 = \{\mathbf{a2c} \mapsto \langle 1, \{m \mapsto \perp\} \rangle\} \cup \mathcal{A}_0^{\text{rest}}$  denote the initial automaton state function, where:

$$\mathcal{A}_0^{\text{rest}} = \{\mathbf{c2a} \mapsto \langle 1, \{m \mapsto \perp\} \rangle, \mathbf{a2cbs} \mapsto \langle 1, \emptyset \rangle, \mathbf{ac2bs} \mapsto \langle 1, \emptyset \rangle\}$$

Initially, thus, all connectors are in their initial state (state 1). Furthermore, connectors **a2c** and **c2a** have an empty memory cell ( $m \mapsto \perp$ ); connectors **a2cbs** and **ac2bs** have no memory cells (memory snapshot  $\emptyset$ ). By rule [C|Com] (presented after this example), the choreography in Example 4 reduces under  $\mathcal{G}$  as follows:

$$\begin{aligned} & \mathbf{a}\langle title \rangle \rightarrow \mathbf{c} \text{ thru } \mathbf{a2c}; C', \quad \sigma_0, \quad \{\mathbf{a2c} \mapsto \langle 1, \{m \mapsto \perp\} \rangle\} \cup \mathcal{A}_0^{\text{rest}} \\ \rightsquigarrow_{\mathcal{G}} & \boxed{\text{In state 1, according to } \mathcal{G}, \text{ connector } \mathbf{a2c} \text{ only has a transition that allows Alice to send (asynchronously) to Carol. By performing such a send, Alice enables the choreography to make a reduction, in which the first half of the communication completes; the (asynchronous) receive remains. In the same step, } \mathbf{a2c} \text{ moves to state 2, and the value sent by Alice is stored in } \mathbf{a2c}\text{'s internal memory cell (} title \text{ evaluates to "foo", based on } \sigma_0\text{).}} \\ & \mathbf{c}.title? \text{"foo"} \text{ thru } \mathbf{a2c}; C', \quad \sigma_0, \quad \{\mathbf{a2c} \mapsto \langle 2, \{m \mapsto \text{"foo"} \} \rangle\} \cup \mathcal{A}_0^{\text{rest}} \\ \rightsquigarrow_{\mathcal{G}} & \boxed{\text{In state 2, according to } \mathcal{G}, \text{ connector } \mathbf{a2c} \text{ only has a transition that allows Carol to receive (asynchronously) from Alice. By performing such a receive, Carol enables the choreography to make a reduction in which the whole communication completes. In the same step, } \mathbf{a2c} \text{ moves to state 1 (the internal memory cell is not cleared).}} \\ & \emptyset \text{ thru } \mathbf{a2c}; C', \quad \sigma'_0, \quad \{\mathbf{a2c} \mapsto \langle 1, \{m \mapsto \text{"foo"} \} \rangle\} \cup \mathcal{A}_0^{\text{rest}} \end{aligned}$$

These intuitions are captured in the rule for communications [C|Com], which is the key rule defining  $\rightsquigarrow_{\mathcal{G}}$ . We discuss this rule in detail.

$$\frac{\mathcal{A}(\gamma) = \langle s, \mu \rangle \quad \tilde{\eta}, \sigma, \mu \xrightarrow{\phi} \tilde{\eta}', \sigma', \mu' \quad s \xrightarrow{\phi}_{\gamma} s'}{\tilde{\eta} \text{ thru } \gamma; C, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} \tilde{\eta}' \text{ thru } \gamma; C, \sigma', \mathcal{A}[\gamma \mapsto \langle s', \mu' \rangle]} \text{ [C|Com]}$$

$$\begin{array}{c}
 \frac{e \downarrow_p^\sigma v}{\{\mathbf{p}\langle e \rangle \rightarrow \mathbf{q}.x\}, \sigma, \mu \xrightarrow{\mathbf{p} \rightarrow \mathbf{q}} \emptyset, \sigma[\mathbf{q}.x \mapsto v], \mu} [\text{C|SyncVal}] \\
 \frac{e \downarrow_p^\sigma v}{\{\mathbf{p}\langle e \rangle \rightarrow \mathbf{q}.x\}, \sigma, \mu \xrightarrow{\mathbf{p} \rightarrow m} \mathbf{q}.x?v, \sigma, \mu[m \mapsto v]} [\text{C|SendVal}] \\
 \frac{\mu(m) = v}{\{\mathbf{q}.x?v\}, \sigma, \mu \xrightarrow{m \rightarrow \mathbf{q}} \emptyset, \sigma[\mathbf{q}.x \mapsto v], \mu} [\text{C|RecvVal}] \quad \frac{\mu(m) = \ell}{\{\mathbf{q}[\ell]\}, \sigma, \mu \xrightarrow{m \rightarrow \mathbf{q}} \emptyset, \sigma, \mu} [\text{C|RecvSel}] \\
 \frac{}{\{\mathbf{p} \rightarrow \mathbf{q}[\ell]\}, \sigma, \mu \xrightarrow{\mathbf{p} \rightarrow \mathbf{q}} \emptyset, \sigma, \mu} [\text{C|SyncSel}] \quad \frac{}{\{\mathbf{p} \rightarrow \mathbf{q}[\ell]\}, \sigma, \mu \xrightarrow{\mathbf{p} \rightarrow m} \{\mathbf{q}[\ell]\}, \sigma, \mu[m \mapsto \ell]} [\text{C|SendSel}] \\
 \frac{\mu(m_1) = v}{\emptyset, \sigma, \mu \xrightarrow{m_1 \rightarrow m_2} \emptyset, \sigma, \mu[m_2 \mapsto v]} [\text{C|Mem}] \quad \frac{\tilde{\eta}_1, \sigma, \mu \xrightarrow{\phi} \tilde{\eta}_2, \sigma', \mu'}{(\tilde{\eta}_1 \uplus \tilde{\eta}'), \sigma, \mu \xrightarrow{\phi} (\tilde{\eta}_2 \uplus \tilde{\eta}'), \sigma', \mu'} [\text{C|Mon}] \\
 \frac{\tilde{\eta}_1, \sigma, \mu \xrightarrow{\phi_1} \tilde{\eta}'_1, \sigma', \mu' \quad \tilde{\eta}_2, \sigma', \mu' \xrightarrow{\phi_2} \tilde{\eta}'_2, \sigma'', \mu'' \quad (\dagger)}{(\tilde{\eta}_1 \uplus \tilde{\eta}_2), \sigma, \mu \xrightarrow{\phi_1 \cup \phi_2} (\tilde{\eta}'_1 \uplus \tilde{\eta}'_2), \sigma'', \mu''} [\text{C|Join}]
 \end{array}$$

■ **Figure 4** Semantics of individual communications. The side condition  $(\dagger)$  in rule  $[\text{C|Join}]$  reads: if a memory cell  $m$  occurs in both  $\phi_1$  and  $\phi_2$ , then it is not both written to in  $\phi_1$  and read from in  $\phi_2$ .

Rule  $[\text{C|Com}]$  allows some of the communications in  $\tilde{\eta}$  to reduce only if the state of the automaton corresponding to connector  $\gamma$  allows it. The rule reads: under  $\mathcal{G}$ , triple  $\tilde{\eta} \text{ thru } \gamma; C, \sigma, \mathcal{A}$  can reduce if automaton  $\mathcal{G}(\gamma)$  can fire a transition out of its current state that is compatible with the interactions specified in  $\tilde{\eta}$ .

More formally, the first premise of this rule retrieves the current state  $s$  and memory snapshot  $\mu$  of the connector  $\gamma$  controlling the communication. In the second premise, the labelled reduction  $\tilde{\eta}, \sigma, \mu \xrightarrow{\phi} \tilde{\eta}', \sigma', \mu'$  (defined below) states that reducing  $\tilde{\eta}$  to  $\tilde{\eta}'$  transforms the state of processes  $\sigma$  into  $\sigma'$  and the memory snapshot of the connector  $\mu$  into  $\mu'$ . The label  $\phi$  represents the actions executed in this reduction. The third premise checks that these actions are allowed by the automaton, by checking that  $\phi$  labels an outgoing transition of  $s$ .

The rules defining labelled reductions  $\tilde{\eta}, \sigma, \mu \xrightarrow{\phi} \tilde{\eta}', \sigma', \mu'$  are given in Figure 4. They are obtained by considering the different possible cases for terms  $\tilde{\eta}$  and matching them to appropriate constraints  $\phi$ . Let  $e \downarrow_p^\sigma v$  denote that expression  $e$  evaluates to value  $v$  under  $\sigma$  (i.e., after substituting every free variable  $x$  in  $e$  by  $\sigma(\mathbf{p}.x)$ ). In rule  $[\text{C|SyncVal}]$ , an entire communication  $\mathbf{p}\langle e \rangle \rightarrow \mathbf{q}.x$  is executed in one step. Accordingly, the label  $\mathbf{p} \rightarrow \mathbf{q}$  denotes that the automaton should support a synchronous communication between  $\mathbf{p}$  and  $\mathbf{q}$ . The state of the receiver  $\mathbf{q}$  is updated with the value sent by  $\mathbf{p}$ . Rule  $[\text{C|SendVal}]$  applies in the case where the message from  $\mathbf{p}$  should be stored in a memory cell  $m$  of the automaton (label  $\mathbf{p} \rightarrow m$ ). This is used for asynchronous communications, where the message will be received later on by the receiver. In the reductum, we keep a runtime term signalling that the receiver is still waiting to receive the message ( $\mathbf{q}.x?v$ ). This kind of runtime terms is handled by rule  $[\text{C|RecvVal}]$ , whose label specifies that  $\mathbf{q}$  should receive the message stored in some memory cell  $m$ . The premise of this rule checks that the value  $\mathbf{q}$  is expecting to receive ( $v$ , defined in the choreography term) is the one stored in  $m$ .<sup>3</sup> Rules  $[\text{C|SyncSel}]$ ,  $[\text{C|SendSel}]$

<sup>3</sup> In other words, the automaton delivers the messages as specified in the choreography.

and [C|RecvSel] deal with the case of label selection in a similar way. Rule [C|Mem] covers internal transitions in the automaton that only modify memory. Finally, rules [C|Mon] and [C|Join] extend this notion to transitions labelled by non-singleton sets. Rule [C|Mon] states that some communications in  $\tilde{\eta}$  may not be executed at all (i.e., they are postponed until a later reduction). Rule [C|Join] allows executing several communications at the same time; note that the labels  $\phi_1$  and  $\phi_2$  may share constraints (e.g., with multi-cast), while  $\tilde{\eta}_1$  and  $\tilde{\eta}_2$  must be disjoint ( $\uplus$  is the disjoint union operator).

The syntactic assumptions on choreography terms (page 8), namely distinct receivers (all terms on the right-hand sides of interactions are distinct) and consistent sends (if a process sends to several processes, then it always sends the same value or label), ensure that the sequentialisation in rule [C|Join] is of no consequence (building the final set  $\phi$  in any order always yields the same final states  $\sigma'$  and  $\mu'$  in rule [C|Com]), except if the same memory cell is both written to and read from. In that case, the read must precede the write; this is guaranteed by side condition ( $\dagger$ ), which ensures that concurrent accesses to a memory cell are done in a consistent way.

► **Example 8** (Book sale). We formally derive the reductions in Example 7, using rule [C|Com] and Figure 4. Let  $\mathcal{G}$ ,  $\sigma_0$ ,  $\sigma'_0$ , and  $\mathcal{A}_0$  be defined as in Example 7.

For the first reduction, first,  $\mathcal{A}_0(\text{a2c}) = \langle 1, \{m \mapsto \perp\} \rangle$  (rule [C|Com], first premise). Next, automaton  $\mathcal{G}(\text{a2c})$  has one transition out of state 1, namely  $1 \xrightarrow{\text{a} \rightarrow m} \text{a2c} 2$  (rule [C|Com], third premise). Finally, using Figure 4, we need to derive  $\text{a}\langle \text{title} \rangle \rightarrow \text{c}, \sigma_0, \{m \mapsto \perp\} \xrightarrow{\text{a} \rightarrow m} \dots$  (rule [C|Com], second premise); this follows from rule [C|SendVal]. Thus, we derive:

$$\begin{array}{c} \text{title} \downarrow_a^{\sigma_0} \text{"foo"} \\ \hline \text{a}\langle \text{title} \rangle \rightarrow \text{c}, \sigma_0, \{m \mapsto \perp\} \xrightarrow{\text{a} \rightarrow m} \text{c.title?"foo", } \sigma_0, \{m \mapsto \text{"foo"}\} \quad 1 \xrightarrow{\text{a} \rightarrow m} \text{a2c} 2 \\ \hline \mathcal{A}_0(\text{a2c}) = \langle 1, \{m \mapsto \perp\} \rangle \\ \hline \text{a}\langle \text{title} \rangle \rightarrow \text{c thru a2c}; C', \sigma_0, \{ \text{a2c} \mapsto \langle 1, \{m \mapsto \perp\} \rangle \} \cup \mathcal{A}_0^{\text{rest}} \\ \rightsquigarrow_{\mathcal{G}} \text{c.title?"foo" thru a2c}; C', \sigma_0, \{ \text{a2c} \mapsto \langle 2, \{m \mapsto \text{"foo"} \rangle \} \} \cup \mathcal{A}_0^{\text{rest}} \end{array}$$

Henceforth, let  $\mathcal{A}'_0 = \{ \text{a2c} \mapsto \langle 2, \{m \mapsto \text{"foo"} \rangle \} \} \cup \mathcal{A}_0^{\text{rest}}$ .

For the second reduction, similarly, we derive:

$$\begin{array}{c} \{m \mapsto \text{"foo"}\}(m) = \text{"foo"} \\ \hline \text{c.title?"foo", } \sigma_0, \{m \mapsto \text{"foo"}\} \xrightarrow{m \rightarrow c} \emptyset, \sigma'_0, \{m \mapsto \text{"foo"}\} \quad 2 \xrightarrow{m \rightarrow c} \text{a2c} 1 \\ \hline \mathcal{A}'_0(\text{a2c}) = \langle 2, \{m \mapsto \text{"foo"}\} \rangle \\ \hline \text{c.title?"foo" thru a2c}; C', \sigma_0, \{ \text{a2c} \mapsto \langle 2, \{m \mapsto \text{"foo"} \rangle \} \} \cup \mathcal{A}_0^{\text{rest}} \\ \rightsquigarrow_{\mathcal{G}} \emptyset \text{ thru a2c}; C', \sigma'_0, \{ \text{a2c} \mapsto \langle 1, \{m \mapsto \text{"foo"} \rangle \} \} \cup \mathcal{A}_0^{\text{rest}} \end{array}$$

► **Example 9** (Book sale). We formally derive the reduction of line 4 in Example 4. Let  $\mathcal{G}$  denote the connector mapping in Example 6, let  $\sigma$  denote a choreography state function such that  $\sigma(\text{a.money}) = \text{\$10}$  and  $\sigma(\text{c.book}) = \text{foo.pdf}$ , let  $\sigma' = \sigma[\text{b.money} \mapsto \text{\$10}]$ , let  $\sigma'' = \sigma'[\text{s.book} \mapsto \text{foo.pdf}]$ , and let  $\mathcal{A} = \mathcal{A}^{\text{rest}} \cup \{ \text{ac2bs} \mapsto \{1, \emptyset\} \}$  denote an automaton state function, for some  $\mathcal{A}^{\text{rest}}$ . We derive:

$$\begin{array}{c} \text{money} \downarrow_a^{\sigma} \text{\$10} \quad \text{book} \downarrow_c^{\sigma} \text{foo.pdf} \\ \hline \{ \text{a}\langle \text{money} \rangle \rightarrow \text{b} \}, \sigma, \emptyset \quad \{ \text{c}\langle \text{book} \rangle \rightarrow \text{s} \}, \sigma', \emptyset \\ \hline \text{a} \xrightarrow{\text{a} \rightarrow \text{b}} \emptyset, \sigma', \emptyset \quad \text{c} \xrightarrow{\text{c} \rightarrow \text{s}} \emptyset, \sigma'', \emptyset \\ \hline \{ \text{a}\langle \text{money} \rangle \rightarrow \text{b}, \text{c}\langle \text{book} \rangle \rightarrow \text{s} \}, \sigma, \emptyset \\ \hline \mathcal{A}(\text{ac2bs}) = \langle 1, \emptyset \rangle \\ \hline \text{a} \xrightarrow{\text{a} \rightarrow \text{b} \wedge \text{c} \rightarrow \text{s}} \emptyset, \sigma'', \emptyset \quad 1 \xrightarrow{\text{a} \rightarrow \text{b} \wedge \text{c} \rightarrow \text{s}} \text{ac2bs} 1 \\ \hline \{ \text{a}\langle \text{money} \rangle \rightarrow \text{b}, \text{c}\langle \text{book} \rangle \rightarrow \text{s} \} \text{ thru ac2bs}; \mathbf{0}, \sigma, \mathcal{A} \\ \rightsquigarrow_{\mathcal{G}} \emptyset \text{ thru ac2bs}; \mathbf{0}, \sigma'', \mathcal{A} \end{array}$$

$$\begin{array}{c}
 \frac{i = 1 \text{ if } e \downarrow_p^\sigma \text{ true, } i = 2 \text{ otherwise}}{\text{if } p.e \text{ then } C_1 \text{ else } C_2, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} C_i, \sigma, \mathcal{A}} \quad [\text{C|Cond}] \\
 \\
 \frac{C_1 \preceq C_2 \quad C_2, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} C'_2, \sigma', \mathcal{A}' \quad C'_2 \preceq C'_1}{C_1, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} C'_1, \sigma', \mathcal{A}'} \quad [\text{C|Struct}] \\
 \\
 \frac{C_1, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} C'_1, \sigma', \mathcal{A}'}{\text{def } X = C_2 \text{ in } C_1, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} \text{def } X = C_2 \text{ in } C'_1, \sigma', \mathcal{A}'} \quad [\text{C|Ctx}]
 \end{array}$$

■ **Figure 5** Cho-Reo-graphy, semantics.

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

■ **Figure 6** Cho-Reo-graphy, structural precongrence.

Rule  $[\text{C|Com}]$  is the only rule in the semantics that can cause a choreography to get stuck; we discuss this in more detail shortly.

The remaining rules defining  $\rightsquigarrow_{\mathcal{G}}$  are standard from other choreography calculi (e.g., [22]), and they are given in Figure 5. Rule  $[\text{C|Struct}]$  uses a structural precongrence that allows for actions to be swapped if they do not interfere. Its definition, given in Figure 6 (where  $\text{pn}(\tilde{\eta})$  denotes the process names that occur in  $\tilde{\eta}$ ), contains two interesting new rules:  $[\text{C|Eta-Split}]$  and  $[\text{C|EtaEnd}]$ . Rule  $[\text{C|Eta-Split}]$  allows interactions through the same connector to be joined in one  $\tilde{\eta}$  or split among several ones, which is necessary for correct interaction with rule  $[\text{C|Eta-Eta}]$ . (The example below illustrates this interaction; see also [21] for a similar discussion: this rule is needed whenever one  $\tilde{\eta}$  can specify several communications.) Rule  $[\text{C|EtaEnd}]$  removes completed interactions from the head of a choreography.

► **Example 10.** Consider the following choreography, where for simplicity we abstract from the actual values being communicated.

$$C \equiv \{p \rightarrow q, r \rightarrow s\} \text{ thru } \gamma; \{p \rightarrow q, t \rightarrow v\} \text{ thru } \gamma'$$

and assume that both  $\mathcal{G}(\gamma)$  and  $\mathcal{G}(\gamma')$  allow the interactions between the processes they connect to occur independently and in any order. Then it is actually possible that the communication between  $\mathfrak{t}$  and  $\mathfrak{v}$  is the first one to take place. In order for our choreography language to allow this behaviour, we need to use both  $[C|E\eta\text{-Split}]$  and  $[C|E\eta\text{-Eta}]$  to exchange actions, as follows.

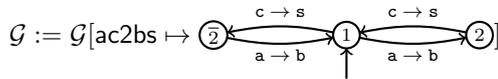
$$\begin{aligned} C &\equiv \{p \rightarrow q, r \rightarrow s\} \text{ thru } \gamma; \{p \rightarrow q, \mathfrak{t} \rightarrow \mathfrak{v}\} \text{ thru } \gamma' \\ &\equiv \{p \rightarrow q, r \rightarrow s\} \text{ thru } \gamma; \mathfrak{t} \rightarrow \mathfrak{v} \text{ thru } \gamma'; p \rightarrow q \text{ thru } \gamma' && \text{by } [C|E\eta\text{-Split}] \\ &\equiv \mathfrak{t} \rightarrow \mathfrak{v} \text{ thru } \gamma'; \{p \rightarrow q, r \rightarrow s\} \text{ thru } \gamma; p \rightarrow q \text{ thru } \gamma' && \text{by } [C|E\eta\text{-Eta}] \end{aligned}$$

◀

### 4.3 Flexibility

An immediate advantage of CR is that different communication semantics can freely be mixed in the same choreography. A second advantage is that CR enables programmers to change the semantics of a choreography modularly, by altering the behaviour of the connectors through which the processes interact with each other, *without* the need to change the choreography itself. The following example further motivates this feature and illustrates its use.

► **Example 11** (Book sale). The original book sale scenario (Examples 1 and 2) requires Alice and Carol to send money and book to the bank and the shipper synchronously, as they initially do not trust each other. Now, suppose Alice and Carol establish mutual trust after successfully completing a number of book sales, such that their communications with the bank and the shipper no longer need to occur synchronously. Instead of redeveloping the choreography from scratch, we need to redefine only the connector mapping  $\mathcal{G}$  in Example 6, as follows:



Thus, we updated the mapping for  $\text{ac2bs}$ ; for all other connectors, the mapping remains the same as in Example 6. The new automaton for  $\text{ac2bs}$  allows either a communication between Alice and the bank, asynchronously followed by a communication between Carol and the shipper (via state 2), or the same two communications in the reverse order (via state  $\bar{2}$ ).<sup>4</sup>

Redefining the connector mapping for  $\text{ac2bs}$  is the only change we need to make: the choreography itself is *exactly* the same as in Example 4. This means that also the first reductions remain *exactly* the same as in Examples 7 and 8. By contrast, the reduction in Example 9 is no longer valid, as it relies on the semantics of  $\text{ac2bs}$ . To show the difference formally, let  $\sigma$ ,  $\sigma'$ ,  $\sigma''$ , and  $\mathcal{A}$  be defined as in Example 9. Let also  $\mathcal{A}' = \mathcal{A}[\text{ac2bs} \mapsto \langle 2, \emptyset \rangle]$  and  $\mathcal{A}'' = \mathcal{A}$ . We derive:

$$\begin{array}{c} \frac{\frac{\frac{\text{money} \downarrow_a^\sigma \$10}{\{a \langle \text{money} \rangle \rightarrow b\}, \sigma, \emptyset \xrightarrow{a \rightarrow b} \emptyset, \sigma', \emptyset} [C|SyncVal]}{\{a \langle \text{money} \rangle \rightarrow b, c \langle \text{book} \rangle \rightarrow s\}, \sigma, \emptyset} [C|Mon]}{\mathcal{A}(\text{ac2bs}) = \langle 1, \emptyset \rangle \xrightarrow{a \rightarrow b} \{c \langle \text{book} \rangle \rightarrow s\}, \sigma', \emptyset} \quad 1 \xrightarrow{a \rightarrow b}_{\text{ac2bs}} 2} \\ \xrightarrow{\mathcal{G}} \{a \langle \text{money} \rangle \rightarrow b, c \langle \text{book} \rangle \rightarrow s\} \text{ thru } \text{ac2bs}; \mathbf{0}, \sigma, \mathcal{A} \\ \{c \langle \text{book} \rangle \rightarrow s\} \text{ thru } \text{ac2bs}; \mathbf{0}, \sigma', \mathcal{A}' \end{array}$$

<sup>4</sup> The communications between Alice and the bank, and between Carol and the shipper, are synchronous in this automaton. We can easily make those communications asynchronous as well, but we skip this here to save space (the automaton gets larger).

## XX:14 Connectors meet Choreographies

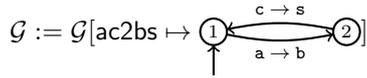
Next, we derive:

$$\frac{A(\text{ac2bs}) = \langle 2, \emptyset \rangle \quad \frac{\text{book} \downarrow_c^\sigma \text{foo.pdf}}{\{c\langle \text{book} \rangle \rightarrow s\}, \sigma', \emptyset \xrightarrow{c \rightarrow s} \emptyset, \sigma'', \emptyset} \text{[C|SyncVal]} \quad 2 \xrightarrow{c \rightarrow s} \text{ac2bs } 1}{\frac{\{c\langle \text{book} \rangle \rightarrow s\} \text{ thru ac2bs; } \mathbf{0}, \sigma', \mathcal{A}'}{\emptyset \text{ thru ac2bs; } \mathbf{0}, \sigma'', \mathcal{A}''}}{\rightsquigarrow_{\mathcal{G}}}$$

Thus, as intended, our reduction rules let us derive two separate reductions with one communication each (first Alice and the bank, then Carol and the shipper) instead of one reduction with two communications (Example 9). Similarly, we can derive two separate reductions whereby Carol and the shipper communicate first, followed by Alice and the bank. ◀

► **Example 12.** The previous example works also “in the opposite direction”, from a trusting Alice and Carol (using connector mapping  $\mathcal{G}$  in Example 11) to cautious ones (using connector mapping  $\mathcal{G}$  in Example 6).

Our choreography is also compatible with the case where we have a trusting Alice and a cautious Carol, who only sends the book after receiving payment. A connector mapping that implements this behaviour is the following.



This connector mapping is still compatible with the choreography in Example 4. The symmetric case where Carol is trusting and Alice is cautious is similar. ◀

► **Example 13.** Note that the changes to connector mappings in Examples 11 and 12 would still be possible if the programmer had written, e.g.,

$$\dots; \{a\langle \text{money} \rangle \rightarrow b\} \text{ thru ac2bs; } \{c\langle \text{book} \rangle \rightarrow s\} \text{ thru ac2bs; } \mathbf{0}$$

instead of the choreography in Example 4. Indeed, these two choreographies are equivalent due to the congruence rule [C|Eta-Split], and thus the sets of connectors that are compatible with each of them are the same.

This might be surprising at first, but it fits with the view of choreographies as global specifications of independent processes. Specifically in this case, no choreography can impose a causal dependency between the two communications  $a\langle \text{money} \rangle \rightarrow b$  and  $c\langle \text{book} \rangle \rightarrow s$  unless it includes an additional communication in the middle involving a process that can observe both. The lack of causal dependencies in this example thus leaves the connector for  $\text{ac2bs}$  free to decide the order in which the interactions are performed. ◀

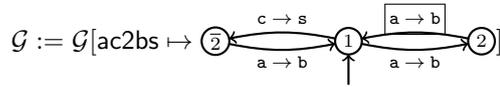
### 4.4 Deadlock-freedom

Rule [C|Com] is the only reduction rule in the semantics that can cause a choreography to get stuck: in choreography  $\tilde{\eta} \text{ thru } \gamma; C$ , there can be incompatibilities between the communications allowed by connector  $\gamma$  and the intended communications in  $\tilde{\eta}$ , causing none of the communications in  $\tilde{\eta}$  to be permitted by  $\gamma$ . In this case, we say that  $\gamma$  does not *respect* the choreography.<sup>5</sup>

<sup>5</sup> A choreography expresses the intentions of the programmer. Although she may instantiate connectors however she likes, we assume they do not violate her intentions.

In more detail, the first premise in rule  $[C|Com]$  is always true (assuming the connector mapping is defined for all connector names in  $\tilde{\eta} \text{ thru } \gamma; C$ ). This gives us unique bindings for  $s$  and  $\mu$ . The third premise in  $[C|Com]$  is also always true (assuming every state of a connector has at least one outgoing transition; this can trivially be checked). For every outgoing transition of  $s$ , this gives us bindings for  $\phi$  and  $s'$ . Now, the choreography gets stuck if for each of those bindings, the second premise in  $[C|Com]$  is false. This can happen in two cases: either  $\tilde{\eta}, \sigma, \mu \xrightarrow{\varphi} \tilde{\eta}', \sigma', \mu'$  can be derived (using the rules in Figure 4) and  $\phi \neq \varphi$  for every derivation, or  $\tilde{\eta}, \sigma, \mu \xrightarrow{\varphi} \tilde{\eta}', \sigma', \mu'$  cannot be derived at all. The former happens if every  $\varphi$  contains different processes than  $\phi$  (Example 14 below), or the same processes but in different send/receive pairs (Example 15 below); the latter happens if  $\tilde{\eta}$  contains only asynchronous receives for which rules  $[C|RecvVal]$  and  $[C|RecvSel]$  in Figure 4 are inapplicable (Example 16 below). We now exemplify these cases.

► **Example 14** (Book sale). Suppose we mistakenly redefine the connector mapping  $\mathcal{G}$  in Example 6 as follows (cf. Example 11; i.e., the boxed label is wrong):



Thus, connector  $ac2bs$  initially allows a communication either between Alice and the bank, or between Carol and the shipper. In the latter case,  $ac2bs$  subsequently allows a communication between Alice and the bank, as in Example 11. But in the former case,  $ac2bs$  subsequently allows a second communication between Alice and the bank (instead of between Carol and the shipper).

The first derivation in Example 11 is still valid, but the second derivation is not: rule  $[C|SyncVal]$  is still applied to derive  $\{c\langle book \rangle \rightarrow s\}, \sigma', \emptyset \xrightarrow{c \rightarrow s} \emptyset, \sigma'', \emptyset$  to fulfill the second premise of rule  $[C|Com]$ , but  $ac2bs$  has no transition in state 2 labelled with  $c \rightarrow s$ . As there are no other derivations to fulfill the second premise of rule  $[C|Com]$ , the choreography gets stuck. ◀

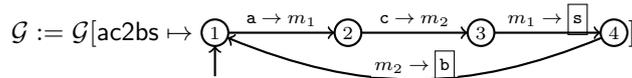
► **Example 15** (Book sale). Suppose we mistakenly redefine connector mapping  $\mathcal{G}$  in Example 6 as follows (i.e., the boxed process names are wrong/swapped):

$$\mathcal{G} := \mathcal{G}[ac2bs \mapsto Barrier[a/p_1, \boxed{s}/p_2, c/p_3, \boxed{b}/p_4]]$$

Formally, automaton  $\mathcal{G}(ac2bs)$  has the following transition:  $1 \xrightarrow{a \rightarrow s \wedge c \rightarrow b} ac2bs \ 1$ . Thus, connector  $ac2bs$  allows communications between Alice and the shipper (instead of the bank), and between Carol and the bank (instead of the shipper).

The derivation in Example 9 is no longer valid: rules  $[C|SyncVal]$  and  $[C|Join]$  are still applied to derive  $\{a\langle money \rangle \rightarrow b, c\langle book \rangle \rightarrow s\}, \sigma, \emptyset \xrightarrow{a \rightarrow b \wedge c \rightarrow s} \emptyset, \sigma'', \emptyset$  to fulfill the second premise of rule  $[C|Com]$ , but  $ac2bs$  has no transition labelled with  $a \rightarrow b \wedge c \rightarrow s$ . As there are no other derivations to fulfill the second premise of rule  $[C|Com]$ , the choreography gets stuck. ◀

► **Example 16** (Book sale). Suppose we mistakenly redefine the connector mapping  $\mathcal{G}$  in Example 6 as follows (i.e., the boxed process names are wrong/swapped):



Thus, connector  $ac2bs$  allows an asynchronous send by Alice, followed by an asynchronous send by Carol, followed by an asynchronous receive by the shipper, followed by an asynchronous

## XX:16 Connectors meet Choreographies

receive by the bank. However, the shipper receives the value sent by Alice (instead of Carol), while the bank receives the value sent by Carol (instead of Alice)

Let  $\sigma$ ,  $\sigma'$ ,  $\sigma''$ , and  $\mathcal{A}$  be defined as in Example 9. Furthermore, let  $\mu = \{m_1 \mapsto \perp, m_2 \mapsto \perp\}$ , let  $\mu' = \{m_1 \mapsto \$10, m_2 \mapsto \perp\}$ , and let  $\mu'' = \{m_1 \mapsto \$10, m_2 \mapsto \text{foo.pdf}\}$ . The following reductions can be derived using rule [C|Com]:

$$\begin{aligned} & \{a(\text{money}) \rightarrow b, c(\text{book}) \rightarrow s\} \text{ thru } \text{ac2bs}; \mathbf{0}, \quad \sigma, \quad \mathcal{A}[\text{ac2bs} \mapsto \langle 1, \mu \rangle] \\ \rightsquigarrow_{\mathcal{G}} & \quad \{b.\text{money}?\$10, c(\text{book}) \rightarrow s\} \text{ thru } \text{ac2bs}; \mathbf{0}, \quad \sigma, \quad \mathcal{A}[\text{ac2bs} \mapsto \langle 2, \mu' \rangle] \\ \rightsquigarrow_{\mathcal{G}} & \quad \{b.\text{money}?\$10, s.\text{book}? \text{foo.pdf}\} \text{ thru } \text{ac2bs}; \mathbf{0}, \quad \sigma, \quad \mathcal{A}[\text{ac2bs} \mapsto \langle 3, \mu'' \rangle] \end{aligned}$$

At this point, the choreography gets stuck: there are no derivations to fulfill the second premise of [C|Com]. To see this, note that only rule [C|RecvVal] may be applicable (together with rule [C|Join]), but  $\mu''(m_2) = \text{foo.pdf}$ , whereas the choreography states  $b.\text{money}?\$10$ . In other words, the choreography expects the bank to receive  $\$10$ , but connector  $\text{ac2bs}$  allows the bank only to receive  $\text{foo.pdf}$ , out of memory cell  $m_2$ . ◀

None of these examples can be constructed in existing choreography models: in existing models, all (a)synchronous channels are guaranteed to respect their choreographies, because the choreography syntax is carefully tuned to the *fixed* communication semantics of these channels. In CR, we have no fixed communication semantics: the fact that connectors in CR may not respect their choreography is, thus, a consequence of the added expressiveness and flexibility CR provides.

We proceed with a more formal account.

► **Definition 17.** Connector mapping  $\mathcal{G}$  in automaton state function  $\mathcal{A}$  *respects* choreography  $C$  if: for every  $\sigma, \tilde{\eta}, \gamma, \sigma'$  and  $\mathcal{A}'$ , if  $C, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}}^* \tilde{\eta} \text{ thru } \gamma; C', \sigma', \mathcal{A}'$ , then there exist  $\sigma''$  and  $\mathcal{A}''$  such that  $\tilde{\eta} \text{ thru } \gamma; C', \sigma', \mathcal{A}' \rightsquigarrow_{\mathcal{G}}^* C', \sigma'', \mathcal{A}''$ .

Connector mapping  $\mathcal{G}$  respects choreography  $C$  if  $\mathcal{G}$  respects  $C$  in initial automaton state function  $\mathcal{A}_0$  (which assigns each automaton to its initial state and memory snapshot, as specified in  $\mathcal{G}$ ).

By definition, respectfulness is equivalent to deadlock-freedom.

► **Definition 18.**  $C, \sigma, \mathcal{A}_0$  is *deadlock-free* for every  $\sigma$  iff  $\mathcal{G}$  respects  $C$ .

We can show respectfulness/deadlock-freedom to be undecidable using a classical recursion-theoretic argument.

► **Theorem 19** (Undecidability of Deadlock-Freedom). *In general, it is undecidable whether a connector mapping  $\mathcal{G}$  respects a choreography  $C$ .*

► **Remark.** Undecidability of deadlock-freedom arises exactly because of the new ways in which a choreography and connectors can affect each other, which did not exist in previous work. Specifically, deadlock occurs if a connector's current state has no transitions for the interactions in the choreography's current  $\tilde{\eta}$ . In previous choreography models, this can never happen, since the choreography syntax matches the hard-wired communication semantics *by definition*. Violation of respectfulness is, thus, a unique byproduct of allowing custom communication semantics, through connectors. Concretely, the proof of Theorem 19 relies on the existence of a communication action that does not respect  $\mathcal{G}$ . Such an action does not exist in previous models. ◀

However, we can approximate respectfulness by a decidable one, called *compatibility*, essentially by abstracting away from data. The key point is that a conditional satisfies compatibility only if both its branches satisfy compatibility.

$$\begin{array}{c}
\mathcal{A}(\gamma) = \langle s, \mu \rangle \left\{ \Gamma \vdash_{\mathcal{A}[\gamma \rightarrow \langle s', \mu' \rangle]} \tilde{\eta}' \text{ thru } \gamma; C' \mid \begin{array}{l} C \equiv \tilde{\eta} \text{ thru } \gamma; C' \\ \text{and } \tilde{\eta}, \mu \xrightarrow{\phi} \tilde{\eta}', \mu' \\ \text{and } s \xrightarrow{\phi} s' \end{array} \right\} \quad (\dagger) \\
\hline
\Gamma \vdash_{\mathcal{A}} C \quad [\text{CC|Com}] \\
\\
\frac{\Gamma \vdash_{\mathcal{A}} C}{\Gamma \vdash_{\mathcal{A}} \emptyset \text{ thru } \gamma; C} [\text{CC|Done}] \quad \frac{}{\Gamma \vdash_{\mathcal{A}} \mathbf{0}} [\text{CC|Nil}] \quad \frac{\Gamma \vdash_{\mathcal{A}} C_1 \quad \Gamma \vdash_{\mathcal{A}} C_2}{\Gamma \vdash_{\mathcal{A}} \text{if } p.e \text{ then } C_1 \text{ else } C_2} [\text{CC|Cond}] \\
\\
\frac{\Gamma, (X : \mathcal{A}_X) \vdash_{\mathcal{A}} C_1 \quad \Gamma, (X : \mathcal{A}_X) \vdash_{\mathcal{A}} C_2}{\Gamma \vdash_{\mathcal{A}} \text{def } X = C_2 \text{ in } C_1} [\text{CC|Def}] \quad \frac{(X : \mathcal{A}) \in \Gamma}{\Gamma \vdash_{\mathcal{A}} X} [\text{CC|Call}]
\end{array}$$

■ **Figure 7** Cho-Reo-graphy, compatibility relation. The side condition  $(\dagger)$  reads: the set of judgments on the left is nonempty. By abuse of notation, we use set builder notation in rule  $[\text{CC|Com}]$  to indicate that all judgments in the set must be true.

► **Definition 20.** Let  $C$  be a choreography,  $\mathcal{G}$  be a connector mapping, and  $\mathcal{A}$  be an automaton state function. We say that  $C$  and  $\mathcal{G}$  are *compatible* by automaton state function  $\mathcal{A}$  if  $\vdash_{\mathcal{A}}^{\mathcal{G}} C$ , where the relation  $\vdash$  is defined by the rules in Figure 7. We say  $C$  and  $\mathcal{G}$  are compatible, written  $\vdash^{\mathcal{G}} C$ , if  $\vdash_{\mathcal{A}_0}^{\mathcal{G}} C$  with  $\mathcal{A}_0$  as in Definition 17.

Relation  $\vdash$  uses a context  $\Gamma$ , defined inductively as  $\Gamma ::= (X : \mathcal{A}), \Gamma \mid \cdot$ , and an abstraction of the labelled reductions for communications from Figure 4,  $\tilde{\eta}, \mu \xrightarrow{\phi} \tilde{\eta}', \mu'$ . The latter models a symbolic execution of communications; it is defined as in Figure 4, with two differences: (i)  $\sigma$  is removed from the domain of the reduction and (ii) in rule  $[\text{C|SendVal}]$ ,  $v$  is a fresh token value.

Intuitively,  $X : \mathcal{A} \in \Gamma$  indicates that procedure  $X$  can be called only whenever the automata have current states  $\mathcal{A}$ ; this is encoded in rules  $[\text{CC|Def}]$  and  $[\text{CC|Call}]$  (in the former rule, a unique automaton state function  $\mathcal{A}_X$  is stipulated; in the latter rule, it is checked against the current automaton state function  $\mathcal{A}$ ). Together with the fact that we allow actions to be swapped in rule  $[\text{CC|Com}]$ , but not recursive calls to be unfolded, this means that the recursive structures of the choreography and the automata in the connector mapping must be similar (i.e., the loops in the automata must match the recursions in the choreography). Furthermore, in order for these rules to ensure respectfulness, the transition relation in the automaton also needs to be confluent (cf. Theorem 24).

► **Remark.** Compatibility can become more robust/modular by disregarding connectors not occurring in procedure bodies in  $[\text{CC|Def}]$ . We chose our current formulation for simplicity. ◀

We now revisit our previous examples to demonstrate cases where the compatibility relation constitutes a precise approximation of respectfulness.

► **Example 21 (Book sale).** We illustrate how compatibility works in the context of our running example by revisiting choreography  $\{\mathbf{a}\langle \text{money} \rangle \rightarrow \mathbf{b}, \mathbf{c}\langle \text{book} \rangle \rightarrow \mathbf{s}\} \text{ thru } \text{ac2bs}; \mathbf{0}$  with five different connector mappings from previous examples.

✓ Let  $\mathcal{G}$  and  $\mathcal{A}$  be defined as in Example 9. Using Figure 7, we derive:

$$\frac{\frac{\frac{}{\cdot \vdash_{\mathcal{A}[\text{ac2bs} \rightarrow \langle 1, \emptyset \rangle]} \mathbf{0}}{\cdot \vdash_{\mathcal{A}[\text{ac2bs} \rightarrow \langle 1, \emptyset \rangle]} \emptyset \text{ thru } \text{ac2bs}; \mathbf{0}} [\text{CC|Done}]}{\cdot \vdash_{\mathcal{A}} \{\mathbf{a}\langle \text{money} \rangle \rightarrow \mathbf{b}, \mathbf{c}\langle \text{book} \rangle \rightarrow \mathbf{s}\} \text{ thru } \text{ac2bs}; \mathbf{0}} [\text{CC|Com}]} [\text{CC|Nil}]$$



$\mu'' = \{m_1 \mapsto \diamond, m_2 \mapsto \heartsuit\}$ , let  $\mathcal{A}' = \mathcal{A}[\text{ac2bs} \mapsto \langle 2, \mu' \rangle]$ , and let  $\mathcal{A}'' = \mathcal{A}[\text{ac2bs} \mapsto \langle 2, \mu'' \rangle]$ . Using Figure 7, we attempt:

$$\begin{array}{c}
 \mathcal{A}''(\text{ac2bs}) \\
 = \langle 3, \mu'' \rangle \quad \boxed{\phantom{\text{derivation}}} \\
 \hline
 \mathcal{A}'(\text{ac2bs}) \quad \cdot \vdash_{\mathcal{A}''}^{\mathcal{G}} \{b.\text{money}?\diamond, s.\text{book}?\heartsuit\} \\
 = \langle 2, \mu' \rangle \quad \text{thru ac2bs; } \mathbf{0} \\
 \hline
 \mathcal{A}(\text{ac2bs}) \quad \cdot \vdash_{\mathcal{A}'}^{\mathcal{G}} \{b.\text{money}?\diamond, c\langle \text{book} \rangle \rightarrow s\} \\
 = \langle 1, \mu \rangle \quad \text{thru ac2bs; } \mathbf{0} \\
 \hline
 \cdot \vdash_{\mathcal{A}}^{\mathcal{G}} \{a\langle \text{money} \rangle \rightarrow b, c\langle \text{book} \rangle \rightarrow s\} \text{ thru ac2bs; } \mathbf{0} \\
 \hline
 \end{array}
 \begin{array}{l}
 \text{[CC|Com]} \\
 \text{[CC|Com]} \\
 \text{[CC|Com]}
 \end{array}$$

This attempted derivation fails, because the intended subderivation inside box (the receive of  $\diamond$ , followed by the receive of  $\heartsuit$ ) is invalid: under our current  $\mathcal{G}$ , connector  $\text{ac2bs}$  in state 3 has no transition labelled with  $m_1 \rightarrow b$ .

Thus, the choreography and the connector mapping are incompatible. In fact, in this case, the choreography deadlocks under  $\mathcal{G}$ . ◀

The restriction that a choreography and the automata in its connector mapping must have similar recursive structures (for them to be judged compatible), implies there exist connector mappings that respect their choreographies, but that cannot be shown to do so by means of the compatibility relation – which is unavoidable in view of our undecidability result. We illustrate this by some examples.

► **Example 22.** Let  $C$  be the simple choreography:

$$\text{def } X = p \rightarrow q \text{ thru } \gamma; p \rightarrow q \text{ thru } \gamma; r \rightarrow s \text{ thru } \gamma; X \text{ in } X$$

and  $\mathcal{G}(\gamma)$  a connector that allows communications from  $p$  to  $q$  to occur simultaneously with communications from  $r$  to  $s$  (e.g., *Barrier* in Figure 1). Then  $C$  is deadlock-free, since structural precongruence allows the second communication from  $p$  to  $q$  to be “delayed” and the communication from  $r$  to  $s$  to be “pushed forward”:

$$\begin{array}{ll}
 p \rightarrow q \text{ thru } \gamma; p \rightarrow q \text{ thru } \gamma; r \rightarrow s \text{ thru } \gamma & \\
 \equiv p \rightarrow q \text{ thru } \gamma; \{p \rightarrow q, r \rightarrow s\} \text{ thru } \gamma & \text{by [C|Eta-Split]} \\
 \equiv p \rightarrow q \text{ thru } \gamma; r \rightarrow s \text{ thru } \gamma; p \rightarrow q \text{ thru } \gamma & \text{by [C|Eta-Split]}
 \end{array}$$

However,  $\not\vdash^{\mathcal{G}} C$ , since the second communication from  $p$  to  $q$  in the body of  $X$  cannot be consumed without unfolding the definition of  $X$ .

In this example, the recursive structure of  $X$  (two communications from  $p$  to  $q$  and one from  $r$  to  $s$ ) differs from the recursive structure of  $\mathcal{G}(\gamma)$  (one communication between each pair of processes). ◀

► **Example 23.** Consider now the choreography  $C$  defined as

$$\text{def } X = p \rightarrow q \text{ thru } \gamma; r \rightarrow s \text{ thru } \gamma; X \text{ in } X$$

where  $\mathcal{G}(\gamma)$  only allows communications from  $p$  to  $q$ . Again  $C$  is deadlock-free, since structural precongruence allows the communications from  $r$  to  $s$  to be indefinitely postponed. However,  $\not\vdash^{\mathcal{G}} C$ . In general, compatibility ensures that the choreography is not only deadlock-free, but also that there is a correspondence between the recursive structure of the choreography and the recursive structure of the connectors: the connector must allow all interactions in the body of a definition to be executed before calling other procedures. ◀

► **Theorem 24** (Preservation of Compatibility). *Let  $C$  and  $C'$  be choreographies,  $\mathcal{G}$  be a connector mapping, such that the transition relation in each automaton  $\mathcal{G}(\gamma)$  is confluent,  $\sigma$  and  $\sigma'$  be choreography states, and  $\mathcal{A}$  and  $\mathcal{A}'$  be automata state functions. If  $\vdash_{\mathcal{A}}^{\mathcal{G}} C$  and  $C, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} C', \sigma', \mathcal{A}'$ , then  $\vdash_{\mathcal{A}'}^{\mathcal{G}} C'$ .*

The hypothesis that the transition relations of automata are confluent is required to make sure that unfolding cannot add unwanted reductions.

► **Corollary 25** (Soundness of Compatibility). *Under the assumptions of Theorem 24, if  $\vdash_{\mathcal{A}}^{\mathcal{G}} C$ , then  $\mathcal{G}$  in  $\mathcal{A}$  respects  $C$ .*

Furthermore, compatibility is decidable.

► **Theorem 26** (Decidability of Compatibility). *There is an algorithm that, given  $C$ ,  $\mathcal{G}$  and  $\mathcal{A}$ , returns YES if  $\vdash_{\mathcal{A}}^{\mathcal{G}} C$  and NO if  $\not\vdash_{\mathcal{A}}^{\mathcal{G}} C$ .*

**Proof.** A simple finiteness argument suffices for establishing decidability of compatibility, since the number of automaton states is finite, the number of applicable rules at each step is finite, and all rules have a finite number of premises, and the size of the choreographies in the premises is always smaller than the size of the choreographies in the conclusions. Therefore, by non-deterministically guessing the types of all procedures, we can decide whether  $\vdash_{\mathcal{A}}^{\mathcal{G}} C$  or not. ◀

In the Appendix, we give a more intelligent proof that constructs the types for the recursive definitions.

► **Theorem 27** (Progress). *Let  $C$  be a choreography,  $\mathcal{G}$  be a connector mapping,  $\sigma$  be a choreography state and  $\mathcal{A}$  be an automaton state function such that  $\vdash_{\mathcal{A}}^{\mathcal{G}} C$ . Then, either  $C \preceq \mathbf{0}$  ( $C$  has terminated) or there exist  $C'$ ,  $\sigma'$  and  $\mathcal{A}'$  such that  $C, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} C', \sigma', \mathcal{A}'$ .*

**Proof.** If  $C \not\preceq \mathbf{0}$ , then  $C$  is of the form  $\tilde{\eta} \text{ thru } \gamma; C'$  or **if p.e then**  $C_1$  **else**  $C_2$ , eventually inside some recursive definitions. In the latter case,  $C$  can always reduce; in the former case, compatibility guarantees that  $C$  can reduce. ◀

Given the previous theorem, if  $C \not\preceq \mathbf{0}$ , then by Theorem 24 we also have that  $\vdash_{\mathcal{A}'}^{\mathcal{G}} C'$  whenever  $C, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} C', \sigma', \mathcal{A}'$ . By induction, we therefore can state a stronger result: every choreography starting from a compatible automaton state function always keeps reducing unless it reaches  $\mathbf{0}$ . Thus, choreographies starting from an automaton state function they are compatible with can never deadlock.

► **Theorem 28** (Deadlock-Freedom by Design). *Let  $C$  be a choreography,  $\sigma$  be a choreography state function, and  $\mathcal{A}$  be an automaton state function. If  $\vdash_{\mathcal{A}}^{\mathcal{G}} C$  and  $C, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}}^* C', \sigma', \mathcal{A}'$ , then either  $C' \preceq \mathbf{0}$  or there exist  $C''$ ,  $\sigma''$  and  $\mathcal{A}''$  such that  $C', \sigma', \mathcal{A}' \rightsquigarrow_{\mathcal{G}} C'', \sigma'', \mathcal{A}''$ .*

## 5 Connected Processes

The results presented so far show that choreographies can be combined with connectors, but do not indicate yet how we can use CR to obtain executable implementations of concurrent systems. The missing link is determining how, given a choreography, we can synthesise terms that represents executable concurrent processes that communicate through connectors. We address this aspect in this section, first by defining a process calculus based on standard I/O actions and then by presenting a translation (compilation procedure) from CR to this calculus.

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

■ **Figure 8** Connected Processes, Syntax.

## 5.1 Syntax and semantics

We define Connected Processes (CP), the process calculus that we use to represent concrete implementations of choreographies. The syntax of CP is given in Figure 8. A network  $N$  is a parallel composition of processes. A process is written  $\text{p} \triangleright_{\rho} B$ , where  $\text{p}$  is its identifier,  $\rho$  its current state (mapping variable names to values), and  $B$  its behaviour. Behaviours correspond to local views of choreography interactions. Procedure definitions and calls, conditionals, and termination ( $\mathbf{0}$ ) follow the same ideas as in choreographies. Communication actions implement the local behaviour of each process in a choreography interaction: sending a value through an output port ( $\text{ol}\langle e \rangle$ ); receiving a value through an input port ( $\text{i}?x$ ); selecting a label through an output port ( $\text{o} \oplus \ell$ ); and offering a choice on some labels through an input port ( $\text{i}\&\{\ell_i : B_i\}_{i \in I}$ ).

The key difference with respect to choreographies is that communications now refer to actual ports, instead of to connectors (i.e., we have no “**thru**  $\gamma$ ” for communications in the process calculus). This reflects the principle that processes should not know how they are connected [5, 29, 30].

The semantics of CP is parameterised on connectors represented as a set of automata  $\mathcal{C}$  that do not share any ports. Differently from the automata used in choreographies, the automata in  $\mathcal{C}$  use the names of the actual ports to which they are connected (and which are also used by the processes). Reductions in CP have the form  $N, A \rightsquigarrow_{\mathcal{C}} N', A'$ , where  $A$  maps each automaton in  $\mathcal{C}$  to a pair  $\langle s, \mu \rangle$  of its state  $s$  and memory snapshot  $\mu$ . The key reduction rule of CP is the one for communications (the remaining ones are standard; see the Appendix):

$$\frac{a \in \mathcal{C} \quad A(a) = \langle s, \mu \rangle \quad N, \mu \xrightarrow{\phi} N', \mu' \quad s \xrightarrow{\phi}_a s'}{N, A \rightsquigarrow_{\mathcal{C}} N', A[a \mapsto \langle s', \mu' \rangle]} \text{ [CP|Com]}$$

This rule is reminiscent of rule [C|Com] for choreographies. In particular, it uses a similar auxiliary reduction relation on pairs of networks and memory snapshots (stated in premise  $N, \mu \xrightarrow{\phi} N', \mu'$ ), whose main rules are given in Figure 9 (see the Appendix for the others).

## 5.2 EndPoint Projection (EPP)

The EPP of a choreography  $C$  from CR into CP follows the usual construction, but with an additional ingredient: we need to add port names associated with communication actions. This is visible in the rules for projecting the individual behaviour of each process (Figure 10), notably in the rule for projecting communications.

► **Remark.** In Figure 10,  $\text{o}$  and  $\text{i}$  denote variables that range over concrete ports. Thus, a process  $\text{p}$  has output ports  $\text{o}_{\gamma_1}, \text{o}_{\gamma_2}, \dots$ , and input ports  $\text{i}_{\gamma_1}, \text{i}_{\gamma_2}, \dots$ , where  $\text{o}$  and  $\text{i}$  actually stand for  $\text{o}^{\text{p}}$  (“output port at  $\text{p}$ ”) and  $\text{i}^{\text{p}}$  (“input port at  $\text{p}$ ”), while connector  $\gamma_1$  knows output ports  $\text{o}_{\gamma_1}^{\text{p}_1}, \text{o}_{\gamma_1}^{\text{p}_2}, \dots$ , and similarly for input ports. ◀

The rule for projecting conditionals uses the standard partial merging operator  $\sqcup$ , where  $B \sqcup B'$  is isomorphic to  $B$  and  $B'$  up to branching with different labels (see [22] for details).

$$\begin{array}{c}
 \frac{e \downarrow^{\rho_p} v}{\mathbf{p} \triangleright_{\rho_p} \mathbf{o}!\langle e \rangle; B_p \mid \mathbf{q} \triangleright_{\rho_q} \mathbf{i}?x; B_q, \mu \xrightarrow{\mathbf{p}.\mathbf{o} \rightarrow \mathbf{q}.\mathbf{i}} \mathbf{p} \triangleright_{\rho_p} B_p \mid \mathbf{q} \triangleright_{\rho_q[x \mapsto v]} B_q, \mu} \text{[CP|SyncVal]} \\
 \frac{e \downarrow^{\rho} v}{\mathbf{p} \triangleright_{\rho} \mathbf{o}!\langle e \rangle; B, \mu \xrightarrow{\mathbf{p}.\mathbf{o} \rightarrow m} \mathbf{p} \triangleright_{\rho} B, \mu[m \mapsto v]} \text{[CP|SendVal]} \\
 \frac{\mu(m) = v}{\mathbf{q} \triangleright_{\rho} \mathbf{i}?x; B, \mu \xrightarrow{m \rightarrow \mathbf{q}.\mathbf{i}} \mathbf{q} \triangleright_{\rho[x \mapsto v]} B, \mu} \text{[CP|RecvVal]} \quad \frac{\mu(m_1) = v}{\mathbf{0}, \mu \xrightarrow{m_1 \rightarrow m_2} \mathbf{0}, \mu[m_2 \mapsto v]} \text{[CP|Mem]} \\
 \frac{N_1, \mu \xrightarrow{\phi_1} N'_1, \mu' \quad N_2, \mu' \xrightarrow{\phi_2} N'_2, \mu'' \quad (\dagger)}{(N_1 \mid N_2), \mu \xrightarrow{\phi_1 \cup \phi_2} (N'_1 \mid N'_2), \sigma'', \mu''} \text{[CP|Join]}
 \end{array}$$

■ **Figure 9** Semantics of communications (process level, main rules). Side condition  $(\dagger)$  in  $\text{[CP|Join]}$  is the same as in  $\text{[C|Join]}$ .

$$\begin{aligned}
 \llbracket \tilde{\eta} \text{ thru } \gamma; C \rrbracket_{\mathbf{r}} &= \begin{cases} \mathbf{o}_{\gamma}!\langle e \rangle; \llbracket C \rrbracket_{\mathbf{r}} & \text{if } \mathbf{r} = \mathbf{p} \text{ and } \mathbf{p}\langle e \rangle \rightarrow \mathbf{q}.x \in \tilde{\eta} \\ \mathbf{i}_{\gamma}?x; \llbracket C \rrbracket_{\mathbf{r}} & \text{if } \mathbf{r} = \mathbf{q} \text{ and } (\mathbf{p}\langle e \rangle \rightarrow \mathbf{q}.x \in \tilde{\eta} \text{ or } \mathbf{q}.x?v \in \tilde{\eta}) \\ \mathbf{o}_{\gamma} \oplus \ell; \llbracket C \rrbracket_{\mathbf{r}} & \text{if } \mathbf{r} = \mathbf{p} \text{ and } \mathbf{p} \rightarrow \mathbf{q}[\ell] \in \tilde{\eta} \\ \mathbf{i}_{\gamma} \& \{ \ell : \llbracket C \rrbracket_{\mathbf{r}} \} & \text{if } \mathbf{r} = \mathbf{q} \text{ and } (\mathbf{p} \rightarrow \mathbf{q}[\ell] \in \tilde{\eta} \text{ or } \mathbf{q}[\ell] \in \tilde{\eta}) \end{cases} \\
 \llbracket \text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2 \rrbracket_{\mathbf{r}} &= \begin{cases} \text{if } e \text{ then } \llbracket C_1 \rrbracket_{\mathbf{r}} \text{ else } \llbracket C_2 \rrbracket_{\mathbf{r}} & \text{if } \mathbf{r} = \mathbf{p} \\ \llbracket C_1 \rrbracket_{\mathbf{r}} \sqcup \llbracket C_2 \rrbracket_{\mathbf{r}} & \mathbf{r} \neq \mathbf{p} \end{cases} \\
 \llbracket \text{def } X = C_2 \text{ in } C_1 \rrbracket_{\mathbf{r}} = \text{def } X = \llbracket C_2 \rrbracket_{\mathbf{r}} \text{ in } \llbracket C_1 \rrbracket_{\mathbf{r}} & \quad \llbracket \mathbf{0} \rrbracket_{\mathbf{r}} = \mathbf{0} \quad \llbracket X \rrbracket_{\mathbf{r}} = X
 \end{aligned}$$

■ **Figure 10** Cho-Reo-graphics, Behaviour Projection.

We now define the projection of  $C$  given a state  $\sigma$ . As usual, this is the parallel composition of the projections of all processes in  $C$ .

$$\llbracket C, \sigma \rrbracket = \prod_{\mathbf{p} \in \text{pn}(C)} \mathbf{p} \triangleright_{\rho_p} \llbracket C \rrbracket_{\mathbf{p}} \quad \text{where } \rho_p(x) = \sigma(\mathbf{p}, x) \text{ for each variable } x \text{ at } \mathbf{p}$$

We say that  $C$  is projectable when  $\llbracket C, \sigma \rrbracket$  is defined for some  $\sigma$ . This is equivalent to saying that  $\llbracket C, \sigma \rrbracket$  is defined for all  $\sigma$ .

► **Example 29** (Book sale). Continuing with our running example, the choreography presented in Example 4 is projectable, and yields the following network of connected processes each state  $\sigma$ .

$$\begin{aligned}
 & \mathbf{a} \triangleright_{\rho_a} \mathbf{o}_{a2c}!\langle \text{title} \rangle; \mathbf{i}_{c2a}?price; \text{if } (\text{happy}) \text{ then } (\mathbf{o}_{a2cbs} \oplus ok; \mathbf{o}_{ac2bs}!\langle \text{money} \rangle; \mathbf{0}) \\
 & \quad \quad \quad \text{else } (\mathbf{o}_{a2cbs} \oplus ko; \mathbf{0}) \\
 & \mid \mathbf{b} \triangleright_{\rho_b} \mathbf{i}_{a2cbs} \& \{ ok : \mathbf{i}_{ac2bs}?money; \mathbf{0}; ko : \mathbf{0} \} \\
 & \mid \mathbf{c} \triangleright_{\rho_c} \mathbf{i}_{a2c}?title; \mathbf{o}_{c2a}!\langle \text{price} \rangle; \mathbf{i}_{a2cbs} \& \{ ok : \mathbf{o}_{ac2bs}!\langle \text{book} \rangle; \mathbf{0}; ko : \mathbf{0} \} \\
 & \mid \mathbf{s} \triangleright_{\rho_s} \mathbf{i}_{a2cbs} \& \{ ok : \mathbf{i}_{ac2bs}?book; \mathbf{0}; ko : \mathbf{0} \}
 \end{aligned}$$

► **Example 30.** It is also worthwhile to note that the following choreographies are *not*

congruent, and that they have *different* EPPs:

$$\begin{aligned} C_1 &= \mathbf{p} \rightarrow \{\mathbf{q}, \mathbf{r}\}[\ell] \text{ thru } \gamma; \mathbf{0} \\ C_2 &= \mathbf{p} \rightarrow \mathbf{q}[\ell] \text{ thru } \gamma; \mathbf{p} \rightarrow \mathbf{r}[\ell] \text{ thru } \gamma; \mathbf{0} \end{aligned}$$

Choreography  $C_1$  is syntactic sugar for  $\{\mathbf{p} \rightarrow \mathbf{q}[\ell], \mathbf{p} \rightarrow \mathbf{r}[\ell]\} \text{ thru } \gamma; \mathbf{0}$ . The EPP to  $\mathbf{p}$ , thus, consists of *one* send; connector  $\gamma$  must subsequently ensure that label  $\ell$  is replicated to, and received by, both  $\mathbf{q}$  and  $\mathbf{r}$  (i.e., formally,  $\gamma$  must be compatible with  $C_1$ ).<sup>6</sup>

Choreography  $C_2$ , in contrast, is not congruent to  $\{\mathbf{p} \rightarrow \mathbf{q}[\ell], \mathbf{p} \rightarrow \mathbf{r}[\ell]\} \text{ thru } \gamma; \mathbf{0}$ . Specifically, we cannot use rule [C|Eta-Split] to merge the two interactions in  $C_2$ , because its disjointness premise does not hold ( $\mathbf{p}$  occurs in both interactions). Accordingly, the EPP of choreography  $C_2$  on  $\mathbf{p}$  consists of *two* sends. ◀

To state the operational correspondence between a choreography and its projection, we need to map the process names used as ports in a connector mapping  $\mathcal{G}$  to the actual port names used in networks. We define  $\llbracket \mathcal{G} \rrbracket$  to be the set of all automata in the codomain of  $\mathcal{G}$ , where each output port  $\mathbf{p}$  in automaton  $\mathcal{G}(\gamma)$  becomes  $\mathbf{p.o}_\gamma$  (and likewise for input ports). We define a similar function for automaton state function  $\mathcal{A}$ .

► **Theorem 31** (Operational Correspondence). *Let  $C$  be a projectable choreography. Then, for all  $\sigma, \mathcal{G}$ , and  $\mathcal{A}$ :*

**Completeness:** *If  $C, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} C', \sigma', \mathcal{A}'$ , then  $\llbracket C, \sigma \rrbracket, \llbracket \mathcal{A} \rrbracket \rightsquigarrow_{\llbracket \mathcal{G} \rrbracket} \llbracket C', \sigma' \rrbracket, \llbracket \mathcal{A}' \rrbracket$ ;*

**Soundness:** *If  $\llbracket C, \sigma \rrbracket, \llbracket \mathcal{A} \rrbracket \rightsquigarrow_{\llbracket \mathcal{G} \rrbracket} N, \mathcal{A}'$ , then  $C, \sigma, \mathcal{A} \rightsquigarrow_{\mathcal{G}} C', \sigma', \mathcal{A}'$  for some  $\sigma'$  and  $\mathcal{A}'$  with  $\llbracket C', \sigma' \rrbracket \prec N$  and  $\llbracket \mathcal{A}' \rrbracket = \mathcal{A}'$ .*

In the soundness result, the *pruning relation*  $\prec$  [15, 17] states that the processes in the network  $N'$  may offer more branches than those present in  $\llbracket C', \sigma' \rrbracket$ . Such branches can never be selected [15, 32, 25].

In particular, if  $\vdash^{\mathcal{G}} C$ , then  $\llbracket C, \sigma \rrbracket$  is guaranteed to be deadlock-free when executed with all automata in  $\llbracket \mathcal{G} \rrbracket$  in their initial states.

► **Example 32** (Book sale). For any connector mapping  $\mathcal{G}$ , the process network in Example 29 operates under  $\llbracket \mathcal{G} \rrbracket$  exactly as the choreography in Example 4 under  $\mathcal{G}$ . In particular, if  $\mathcal{G}$  respects the original choreography, then this implementation never deadlocks under  $\llbracket \mathcal{G} \rrbracket$ . ◀

## 6 Conclusions

Choreographic approaches to concurrent programming have been heavily investigated [1, 28], but these approaches typically adopt some fixed (and restrictive) definition (like point-to-point synchronous) for the semantics of communications. CR is the first model that allows choreographies to be modularly integrated with what runs “under the hood” of communications, allowing for user-defined communication semantics specified as connectors. Thanks to the notion of compatibility (Definition 20), CR inherits the good properties of both Choreographic Programming and Exogenous Coordination. Thus, we have significantly extended the applicability of choreographies. Not only can we capture new kinds of behaviours in choreographies (like barriers, cf. Example 6, and alternators, cf. Example 12), but

<sup>6</sup> This behaviour was the motivation for introducing multicasts as abbreviations: the notation in  $C_1$  better conveys how communications really happen; EPP follows this intuition.

we can even use choreographies to describe systems that integrate different parts with different communication semantics *and* check whether such integration will lead to deadlocks. This is essential in many real-world scenarios, where different components with different communication semantics are usually combined (e.g., some microservices in a distributed system may asynchronously exchange data to be used later in a synchronous multiparty transition, similarly to our example).

This work lays the foundations for applying the combined power of choreographies and connectors to the challenge of concurrent programming, in that CR contains all the necessary foundations to obtain a concrete implementation. The results in § 5 specify how to use CR to obtain code in a process model supported by connectors. Thus, a natural next step will be to implement CR by combining implementations of processes generated from choreographies [19, 25] with distributed implementations of Reo connectors [31, 39, 40]. The main pieces exist; the main challenge lies in their effective composition, and CR is the first essential step towards this objective.

CR also provides a very explicit direction for future developments of this new combined research line: allowing for more kinds of connectors would make the model immediately more expressive. By relaxing the requirements we imposed on the automata in CR (see page 5), we can introduce non-deterministic communication semantics to choreographies, to cater to applications that require lossy channels and safe communication races. Likewise, a more fine-grained semantics that splits communications into two independent send and receive actions (similar to [24]) would enrich the class of behaviours that are captured.

We have followed the traditional approach of viewing choreographies as precise specifications of the intended interactions. However, it would be reasonable to allow choreographies to underspecify communications, such that the underlying connectors were allowed to exchange messages also to participants not defined in the choreography. For example, the semantics for the choreography term  $p\langle e \rangle \rightarrow q.x \text{ thru } \gamma$  can allow  $\gamma$  to send the message from  $p$  to  $q$  via an intermediate process  $r$  that may perform additional actions (like logging the message, or sharing it through another connector). This generalisation can, in particular, provide a novel way for studying how choreographies can be applied to open-ended systems, where the processes projected from *multiple* choreographies execute in parallel and share connectors.

---

**References**

---

- 1 D. Ancona, V. Bono, M. Bravetti, J. Campos, G. Castagna, P. Deniérou, S. J. Gay, N. Gessbert, E. Giachino, R. Hu, E. B. Johnsen, F. Martins, V. Mascardi, F. Montesi, R. Neykova, N. Ng, L. Padovani, V. T. Vasconcelos, and N. Yoshida. Behavioral types in programming languages. *Found. Trends Progr. Lang.*, 3(2-3):95–230, 2016.
- 2 F. Arbab. The IWIM model for coordination of concurrent activities. In *COORDINATION*, pages 34–56. Springer, 1996.
- 3 F. Arbab. What do you mean, coordination. *Bulletin of the Dutch Association for Theoretical Computer Science, NVTI*, 1122:1–18, 1998.
- 4 F. Arbab. Reo: a channel-based coordination model for component composition. *MSCS*, 14(3):329–366, 2004.
- 5 F. Arbab. Puff, the magic protocol. In *Formal Modeling: Actors, Open Systems, Biological Systems*, volume 7000 of *LNCS*, pages 169–206. Springer, 2011.
- 6 F. Arbab, C. Baier, F. S. de Boer, J. J. M. M. Rutten, and M. Sirjani. Synthesis of reo circuits for implementation of component-connector automata specifications. In *COORDINATION*, volume 3454 of *LNCS*, pages 236–251. Springer, 2005.
- 7 C. Baier, J. Klein, and S. Klüppelholz. Synthesis of reo connectors for strategies and controllers. *Fundam. Inform.*, 130(1):1–20, 2014.
- 8 C. Baier, M. Sirjani, F. Arbab, and J. Rutten. Modeling component connectors in Reo by constraint automata. *Sci. Comput. Program.*, 61(2):75–113, 2006.
- 9 A. Basu, S. Bensalem, M. Bozga, J. Combaz, M. Jaber, T. Nguyen, and J. Sifakis. Rigorous component-based system design using the BIP framework. *IEEE Software*, 28(3):41–48, 2011.
- 10 A. Basu, M. Bozga, and J. Sifakis. Modeling heterogeneous real-time components in BIP. In *SEFM*, pages 3–12. IEEE Computer Society, 2006.
- 11 S. Bliudze and J. Sifakis. The algebra of connectors—structuring interaction in BIP. *IEEE Transactions on Computers*, 57(10):1315–1330, 2008.
- 12 S. Bliudze and J. Sifakis. Causal semantics for the algebra of connectors. *Formal methods in system design*, 36(2):167–194, 2010.
- 13 R. Bruni, I. Lanese, and U. Montanari. A basic algebra of stateless connectors. *Theoretical Computer Science*, 366(1-2):98–120, 2006.
- 14 J. T. Buck, S. Ha, E. A. Lee, and D. G. Messerschmitt. Ptolemy: A framework for simulating and prototyping heterogeneous systems. *Int. Journal in Computer Simulation*, 4(2), 1994.
- 15 M. Carbone, K. Honda, and N. Yoshida. Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.*, 34(2):8, 2012.
- 16 M. Carbone, S. Lindley, F. Montesi, C. Schürmann, and P. Wadler. Coherence generalises duality: A logical explanation of multiparty session types. In *CONCUR*, volume 59 of *LIPICs*, pages 33:1–33:15. Schloss Dagstuhl, 2016.
- 17 M. Carbone and F. Montesi. Deadlock-freedom-by-design: multiparty asynchronous global programming. In *POPL*, pages 263–274. ACM, 2013.
- 18 M. Carbone, F. Montesi, C. Schürmann, and N. Yoshida. Multiparty session types as coherence proofs. *Acta Informatica*, pages 1–27, 2016.
- 19 Chor. Programming Language. <http://www.chor-lang.org/>.
- 20 M. Coppo, M. Dezani-Ciancaglini, N. Yoshida, and L. Padovani. Global progress for dynamically interleaved multiparty sessions. *Math. Struct. Comp. Sci.*, 26(2):238–302, 2016.
- 21 L. Cruz-Filipe, K. S. Larsen, and F. Montesi. The paths to choreography extraction. In *FoSSaCS*, volume 10203 of *LNCS*, pages 424–440. Springer, 2017.
- 22 L. Cruz-Filipe and F. Montesi. A core model for choreographic programming. In *FACS 2016*, volume 10231 of *LNCS*, pages 17–35. Springer, 2017.

- 23 L. Cruz-Filipe and F. Montesi. Procedural choreographic programming. In *FORTE 2017*, volume 10321 of *LNCS*, pages 92–107. Springer, 2017.
- 24 L. Cruz-Filipe and F. Montesi. On asynchrony and choreographies. In *ICE 2017*, EPTCS. EPTCS, accepted for publication.
- 25 M. Dalla Preda, M. Gabbrielli, S. Giallorenzo, I. Lanese, and J. Mauro. Dynamic choreographies: Theory and implementation. *LMCS*, 13(2), 2017.
- 26 P.-M. Deniélou and N. Yoshida. Multipart compatibility in communicating automata: Characterisation and synthesis of global session types. In *ICALP (II)*, volume 7966 of *LNCS*, pages 174–186. Springer, 2013.
- 27 S. Giallorenzo. *Real-World Choreographies*. PhD thesis, University of Bologna, Italy, 2016.
- 28 H. Hüttel, I. Lanese, V. T. Vasconcelos, L. Caires, M. Carbone, P. Deniélou, D. Mostrous, L. Padovani, A. Ravara, E. Tuosto, H. T. Vieira, and G. Zavattaro. Foundations of session types and behavioural contracts. *ACM Comput. Surv.*, 49(1):3:1–3:36, 2016.
- 29 S.-S. Jongmans. *Automata-Theoretic Protocol Programming*. PhD thesis, Leiden University, 2016.
- 30 S.-S. Jongmans and F. Arbab. Prdk: Protocol programming with automata. In *TACAS*, pages 547–552. Springer, 2016.
- 31 S.-S. T. Jongmans, F. Santini, and F. Arbab. Partially distributed coordination with Reo and constraint automata. *Service Oriented Computing and Applications*, 9(3-4):311–339, 2015.
- 32 I. Lanese, C. Guidi, F. Montesi, and G. Zavattaro. Bridging the gap between interaction- and process-oriented choreographies. In *SEFM*, pages 323–332. IEEE, 2008.
- 33 K.-K. Lau, P. V. Elizondo, and Z. Wang. Exogenous connectors for software components. In *CBSE*, pages 90–106. Springer, 2005.
- 34 T. Leesatapornwongsa, J. F. Lukman, S. Lu, and H. S. Gunawi. TaxDC: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In *ASPLOS*, pages 517–530. ACM, 2016.
- 35 H. A. López, E. R. B. Marques, F. Martins, N. Ng, C. Santos, V. T. Vasconcelos, and N. Yoshida. Protocol-based verification of message-passing parallel programs. In *OOPSLA*, pages 280–298. ACM, 2015.
- 36 S. Lu, S. Park, E. Seo, and Y. Zhou. Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. In *ASPLOS*, pages 329–339. ACM, 2008.
- 37 F. Montesi. *Choreographic Programming*. Ph.D. thesis, IT University of Copenhagen, 2013.
- 38 F. Montesi and N. Yoshida. Compositional choreographies. In *CONCUR*, volume 8052 of *LNCS*, pages 425–439. Springer, 2013.
- 39 J. Proença. *Synchronous coordination of distributed components*. PhD thesis, Leiden University, 2011.
- 40 J. Proença, D. Clarke, E. De Vink, and F. Arbab. Dreams: a framework for distributed synchronous coordination. In *SAC*, pages 1510–1515. ACM, 2012.
- 41 C. Ptolemaeus, editor. *System Design, Modeling, and Simulation using Ptolemy II*. Ptolemy.org, 2014.

## Additional Proofs

**Theorem 19.** Let  $\eta$  be a communication action that does not respect  $\mathcal{A}$ , and assume that connector  $\gamma$  has synchronous links  $\mathbf{p}2\mathbf{q}$  and  $\mathbf{q}2\mathbf{p}$ , from  $\mathbf{p}$  to  $\mathbf{q}$  and conversely (e.g., *Sync* in Figure 1). Synchronous links are always enabled and do not change  $\mathcal{A}$ .

Let  $f$  be a total function implemented at  $\mathbf{p}$  and consider the choreography

$$\begin{aligned} C \equiv \text{def } X = & \mathbf{q}\langle y \rangle \rightarrow \mathbf{p}.x \text{ thru } \gamma; \\ & \text{if } (\mathbf{p}.f(x) = 0) \text{ then } \mathbf{p}\langle x+1 \rangle \rightarrow \mathbf{q}.y \text{ thru } \gamma; X \\ & \text{else } \eta \text{ thru } \gamma \\ \text{in } & \mathbf{p}\langle 0 \rangle \rightarrow \mathbf{q}.y \text{ thru } \gamma; X \end{aligned}$$

In this choreography,  $\mathbf{q}$  sequentially sends the natural numbers to  $\mathbf{p}$ , which applies  $f$  to its input and proceeds if the result is 0. If  $\mathbf{q}$  sends a value where  $f$  is not 0, the choreography attempts to perform  $\eta$  and deadlocks. Then  $C$  respects  $\mathcal{A}$  iff  $f$  is constantly equal to 0, which by Rice's theorem is not decidable.  $\blacktriangleleft$

**Theorem 24.** Straightforward by case analysis on the reduction from  $C, \sigma, \mathcal{A}$  to  $C', \sigma', \mathcal{A}'$ , using the fact that the automata are confluent (to make sure unfolding cannot add unwanted reductions), and therefore compatibility is preserved by structural precongruence.  $\blacktriangleleft$

**Corollary 25.** If  $C, \sigma \mathcal{A} \rightsquigarrow_G^* \eta; C', \sigma', \mathcal{A}'$ , then, by induction on the length of this sequence of reductions, we use Theorem 24 to show that  $\eta; C', \sigma', \mathcal{A}' \rightsquigarrow_G C'', \sigma'', \mathcal{A}''$  for some  $\mathcal{A}''$  and  $\sigma''$ .  $\blacktriangleleft$

**Theorem 26, alternative proof.** We assume that every procedure defined in  $C$  is called at least once outside of its body.

The idea behind our algorithm is to construct a derivation for  $\vdash_{\mathcal{A}}^G C$  by applying the rules in Figure 7 bottom-up. When we meet a term of the form  $\text{def } X = C_2 \text{ in } C_1$ , we focus on  $C_1$  first, and leave  $\mathcal{A}_X$  (see rule [CC|Def]) unspecified. We instantiate  $\mathcal{A}_X$  later, when we meet  $X$  for the first time inside of  $C_1$ . More precisely:

1. Initialize a list  $\mathcal{L} = [\cdot \vdash_{\mathcal{A}}^G C]$ .
2. While  $\mathcal{L}$  is not empty:
  - a. Remove the first pending judgement  $\Gamma \vdash_{\mathcal{A}}^G C$  from  $\mathcal{L}$ .
  - b. If  $C$  is  $\mathbf{0}$ , proceed to the next iteration.
  - c. If  $C$  is of the form  $\text{if } \mathbf{p}.e \text{ then } C_1 \text{ else } C_2$ , then add  $\Gamma \vdash_{\mathcal{A}}^G C_1$  and  $\Gamma \vdash_{\mathcal{A}}^G C_2$  at the beginning of  $\mathcal{L}$ .
  - d. If  $C$  is of the form  $\text{def } X = C_2 \text{ in } C_1$ , then add  $\Gamma, (X : \mathcal{A}_X) \vdash_{\mathcal{A}}^G C_1$  and  $\Gamma, (X : \mathcal{A}_X) \vdash_{\mathcal{A}}^G C_2$ , in this order, at the beginning of  $\mathcal{L}$ . Here,  $\mathcal{A}_X$  is a unique variable representing an unknown state function.
  - e. If  $C$  is of the form  $X$ , there are two cases. If  $\Gamma$  contains  $(X : \mathcal{A}_X)$  with  $\mathcal{A}_X$  instantiated, check whether  $\mathcal{A}_X = \mathcal{A}$ ; if so, proceed to the next iteration, otherwise return NO. If  $\Gamma$  contains  $(X : \mathcal{A}_X)$  with  $\mathcal{A}_X$  uninstantiated, replace all occurrences of  $\mathcal{A}_X$  in  $\mathcal{L}$  by  $\mathcal{A}$  and proceed to the next iteration.<sup>7</sup>
  - f. Otherwise,  $C$  is of the form  $\tilde{\eta} \text{ thru } \gamma; C'$ . Consider all possible ways of rewriting  $C$  as  $\tilde{\eta}' \text{ thru } \gamma; C'$  by swapping independent actions, without unfolding recursive definitions. Let  $\mathcal{A}(\gamma) = \langle s, \mu \rangle$ . For each such  $\tilde{\eta}'$ , check whether  $\tilde{\eta}', \mu \xrightarrow{\phi} \tilde{\eta}'', \mu'$  for some  $\phi$ , and in

<sup>7</sup> Note that  $\Gamma$  must contain  $(X : \mathcal{A}_X)$ , otherwise the initial choreography is not well-formed.

$$\begin{array}{c}
 \frac{e \downarrow^{\rho_p} v}{\mathbf{p} \triangleright_{\rho_p} \mathbf{o}!(e); B_p \mid \mathbf{q} \triangleright_{\rho_q} \mathbf{i}?x; B_q, \mu \xrightarrow{\mathbf{p.o} \rightarrow \mathbf{q.i}} \mathbf{p} \triangleright_{\rho_p} B_p \mid \mathbf{q} \triangleright_{\rho_q[x \mapsto v]} B_q, \mu} \text{[CP|SyncVal]} \\
 \\
 \frac{e \downarrow^{\rho} v}{\mathbf{p} \triangleright_{\rho} \mathbf{o}!(e); B, \mu \xrightarrow{\mathbf{p.o} \rightarrow m} \mathbf{p} \triangleright_{\rho} B, \mu[m \mapsto v]} \text{[CP|SendVal]} \\
 \\
 \frac{\mu(m) = v}{\mathbf{q} \triangleright_{\rho} \mathbf{i}?x; B, \mu \xrightarrow{m \rightarrow \mathbf{q.i}} \mathbf{q} \triangleright_{\rho[x \mapsto v]} B, \mu} \text{[CP|RecvVal]} \\
 \\
 \frac{j \in I}{\mathbf{p} \triangleright_{\rho_p} \mathbf{o} \oplus \ell_j; B \mid \mathbf{q} \triangleright_{\rho_q} \mathbf{i}\&\{\ell_i : B_i\}_{i \in I}, \mu \xrightarrow{\mathbf{p.o} \rightarrow \mathbf{q.i}} \mathbf{p} \triangleright_{\rho_p} B \mid \mathbf{q} \triangleright_{\rho_q} B_j, \mu} \text{[CP|SyncSel]} \\
 \\
 \frac{}{\mathbf{p} \triangleright_{\rho} \mathbf{o} \oplus \ell; B, \mu \xrightarrow{\mathbf{p.o} \rightarrow m} \mathbf{p} \triangleright_{\rho} B, \mu} \text{[CP|SendSel]} \\
 \\
 \frac{\mu(m) = \ell_j \quad j \in I}{\mathbf{q} \triangleright_{\rho} \mathbf{i}\&\{\ell_i : B_i\}_{i \in I}, \mu \xrightarrow{m \rightarrow \mathbf{q.i}} \mathbf{q} \triangleright_{\rho} B_j, \mu} \text{[CP|RecvSel]} \\
 \\
 \frac{N_1, \mu \xrightarrow{\phi_1} N'_1, \mu' \quad N_2, \mu' \xrightarrow{\phi_2} N'_2, \mu''}{(N_1 \mid N_2), \mu \xrightarrow{\phi_1 \cup \phi_2} (N'_1 \mid N'_2), \sigma'', \mu''} \text{[CP|Join]}
 \end{array}$$

■ **Figure 11** Semantics of communications (process level).

the affirmative case compute  $s'$  such that  $s \xrightarrow{\phi}_{\gamma} s'$  and add  $\Gamma \vdash_{\mathcal{A}[\gamma \rightarrow \langle s', \mu' \rangle]} \tilde{\eta}'' \text{ thru } \gamma; C'$  at the beginning of  $\mathcal{L}$ . If no such transitions exist, return NO.

3. Return YES.

Termination of this algorithm is straight-forward: the sum of the sizes of all the choreographies in  $\mathcal{L}$  strictly decreases at each iteration, and each step terminates in finite time. (The size of a choreography is the number of nodes in its abstract syntax tree, except that  $\mathbf{p}(e) \rightarrow \mathbf{q}.x$  and  $\mathbf{p} \rightarrow \mathbf{q}[\ell]$  count as 2, while  $\mathbf{q}.x?v$  and  $\mathbf{q}[\ell]$  count as 1.) Soundness is immediate by observing that the judgements stored in  $\mathcal{L}$  are exactly those that are necessary to construct a derivation of  $\vdash_{\mathcal{A}}^{\mathcal{G}} C$ , since at each stage there is only one rule that can be applied to build such a derivation, and this rule is determined by the structure of  $C$ . If the algorithm returns YES, then a valid derivation for  $\vdash_{\mathcal{A}}^{\mathcal{G}} C$  can be built. If the algorithm returns NO because of a mismatch between the state of the automata and a communication action (Step 2.f), then clearly  $\not\vdash_{\mathcal{A}}^{\mathcal{G}} C$ . If the algorithm returns NO because of an incompatibility between the state assigned to a procedure name  $X$  in  $\Gamma$  and the state in the current judgement (Step 2.e), then this failure means that we constructed two judgements involving  $X$  with different automaton state functions, which also implies that  $\not\vdash_{\mathcal{A}}^{\mathcal{G}} C$ . ◀

## Complete semantics of Connected Processes

Figure 11 gives the full definition of the auxiliary reduction from rule [CP|Com]. Figure 12 gives the semantics of CP. Rule [CP|Struct] uses the structural precongruence relation  $\preceq$  (including associativity and commutativity of  $\mid$ ), defined in the standard way [22].

$$\begin{array}{c}
\frac{i = 1 \text{ if } e \downarrow^\rho \text{ true, } i = 2 \text{ otherwise}}{\mathfrak{p} \triangleright_\rho \text{ if } e \text{ then } B_1 \text{ else } B_2, A \rightsquigarrow_{\mathcal{C}} \mathfrak{p} \triangleright_\rho B_i, A} \text{ [CP|Cond]} \\
\frac{\mathfrak{p} \triangleright_v B_1 \mid N, A \rightsquigarrow_{\mathcal{C}} \mathfrak{p} \triangleright_v B'_1 \mid N', A'}{\mathfrak{p} \triangleright_v \text{ def } X = B_2 \text{ in } B_1 \mid N, A \rightsquigarrow_{\mathcal{C}} \mathfrak{p} \triangleright_v \text{ def } X = B_2 \text{ in } B'_1 \mid N', A'} \text{ [CP|Ctx]} \\
\frac{N, A \rightsquigarrow_{\mathcal{C}} N', A'}{N \mid M, A \rightsquigarrow_{\mathcal{C}} N' \mid M, A'} \text{ [CP|Par]} \quad \frac{N \preceq M \quad M, A \rightsquigarrow_{\mathcal{C}} M', A' \quad M' \preceq N'}{N, A \rightsquigarrow_{\mathcal{C}} N', A'} \text{ [CP|Struct]}
\end{array}$$

■ **Figure 12** Connected Processes, Semantics.