

A Core Model for Choreographic Programming

Luís Cruz-Filipe, Fabrizio Montesi

*University of Southern Denmark, Department of Mathematics and Computer Science,
Campusvej 55, 5230 Odense M, Denmark*

Abstract

Choreographic Programming is a paradigm for developing concurrent programs that are deadlock-free by construction, [as a result of](#) programming communications declaratively and then synthesising process implementations automatically. Despite strong interest on choreographies, a foundational model that explains which computations can be performed with the hallmark constructs of choreographies is still missing.

In this work, we introduce Core Choreographies (CC), a model that includes only the core primitives of choreographic programming. Every computable function can be implemented as a choreography in CC, from which we can synthesise a process implementation where independent computations run in parallel. We discuss the design of CC and argue that it constitutes a canonical model for choreographic programming.

Keywords: Choreography, Computability, Process Calculi

1. Introduction

Programming concurrent and distributed systems is hard, because it is challenging to predict how programs executed at the same time in different computers will interact. Empirical studies reveal two important lessons: (i) while programmers have clear intentions about the order in which communication actions should be performed, tools do not adequately support them in translating these wishes to code (Lu et al., 2008); (ii) combining different communication protocols in a single application is a major source of mistakes (Leesatapornwongsa et al., 2016).

The paradigm of Choreographic Programming (Montesi, 2013) was introduced to address these problems. In this paradigm, programmers declaratively write the communications that they wish to take place, as programs called *choreographies*. Choreographies are descriptions of concurrent systems that syntactically disallow writing mismatched I/O actions, inspired by the “Alice and Bob”

Email addresses: 1cf@imada.sdu.dk (Luís Cruz-Filipe), fmontesi@imada.sdu.dk (Fabrizio Montesi)

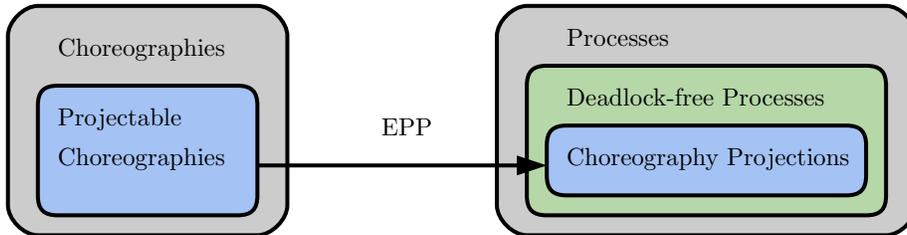


Figure 1: Choreographic Programming

notation of security protocols (Needham and Schroeder, 1978). An EndPoint Projection (EPP) can then be used to synthesise implementations in process models, which faithfully realise the communications given in the choreography and are guaranteed to be deadlock-free by construction even in the presence of arbitrary protocol compositions (Carbone et al., 2012; Carbone and Montesi, 2013).

So far, work on choreographic programming focused on features of practical value – including web services (Carbone et al., 2012), multiparty sessions (Carbone and Montesi, 2013; Chor, 2017), modularity (Montesi and Yoshida, 2013), and runtime adaptation (Dalla Preda et al., 2017). The models proposed all come with differing domain-specific syntaxes, semantics and EPP definitions (e.g., for channel mobility or runtime adaptation), and cannot be considered minimal. Another problem, arguably a consequence of the former, is that choreographic programming is meant for implementation, but we still know little of what can be computed with the code obtained from choreographies (*choreography projections*). The expressivity of the aforementioned models is evaluated just by showing some examples.

In this paper, we propose a canonical model for choreographic programming, called Core Choreographies (CC). CC includes only the core primitives that can be found in most choreography languages, restricted to the minimal requirements to achieve the computational power of Turing machines. In particular, local computation at processes is severely restricted, and therefore nontrivial computations must be implemented by using communications. Therefore, CC is both representative of the paradigm and simple enough to analyse from a theoretical perspective. Our technical development is based on a natural notion of function implementation, and the proof of Turing completeness yields an algorithm for constructing a choreography that implements any given computable function. Since choreographies describe concurrent systems, it is also natural to ask how much parallelism choreographies exhibit. CC helps us in formally defining parallelism in choreographies; we exemplify how to use this notion to reason about the concurrent implementation of functions.

Yet, analysing the expressivity of choreographies is not enough. What we are ultimately interested in is what can be computed with choreography projections, since those are the terms that represent executable code. However, the expres-

sivity of choreographies does not translate directly to expressivity of projections, because EPP is typically an incomplete procedure: it must guarantee deadlock-freedom, which in previous models is obtained by complex requirements, e.g., type systems (Carbone et al., 2012; Carbone and Montesi, 2013). Therefore, only a subset of choreographies (projectable choreographies) can be used to synthesise process implementations. The EPPs of such projectable choreographies form the set of choreography projections, which are deadlock-free processes (see Figure 1).

The main technical contribution of this paper is showing that the set of projectable choreographies in CC is still Turing complete. Therefore, by EPP, the set of corresponding choreography projections is also Turing complete, leading us to a characterisation of a Turing complete and deadlock-free fragment of a process calculus (which follows the same minimal design of CC). Furthermore, the parallel behaviour observed in CC choreographies for function implementations translates directly to parallel execution of the projected processes.

More importantly, the practical consequence of our results is that CC is a simple common setting for the study of foundational questions in choreographies. This makes CC an appropriate foundational model for choreographic programming, akin to λ -calculus for functional programming and π -calculus for mobile processes. As an example of such foundational questions, we describe how the standard communication primitive of label selection can be removed from CC without altering its computational power, yielding a truly minimal choreography language wrt computation called Minimal Choreographies (MC). However, doing so eliminates the clean separation between data and behaviour in message exchanges, which makes the resulting choreography hard to read. [This result suggests that implementations of choreography languages – like \(Montesi, 2013\), \(Dalla Preda et al., 2017\), and \(W3C WS-CDL Working Group, 2004\) – may adopt a simpler intermediate representation of choreographies that does not include label selections, by applying our elimination technique to the input choreography program \(where programmers should be able to use selections, for readability\). A key advantage of having a simpler model without selections is that it bypasses the need for the standard notion of merging \(Carbone et al., 2012\), which is typically one of the most complicated steps in EPP. We formally illustrate this point by showing that our EPP for MC enjoys an elegant definition.](#)

Structure of the paper. § 2 defines Core Choreographies (CC) and its subcalculus of Minimal Choreographies (MC). § 3 introduces Stateful Processes (SP), our target process model, and its sublanguage of Minimal Processes (MP), together with an EndPoint Projection (EPP) from CC to SP (and from MC to MP). We also show that every unprojectable choreography in CC can be amended (transformed into a projectable one) by adding only label selections. We prove that CC and its set of choreography projections are Turing complete in § 4. In § 5 we discuss label selections, and show that they can be encoded by communications; this yields an amendment strategy for MC. In § 6, we show that all the remaining primitives of CC are necessary to achieve Turing completeness,

$$\begin{aligned}
C &::= \eta; C \mid \text{if } p \stackrel{\leq}{=} q \text{ then } C_1 \text{ else } C_2 \mid \text{def } X = C_2 \text{ in } C_1 \mid X \mid \mathbf{0} \\
\eta &::= p.e \rightarrow q \mid p \rightarrow q[l] \quad e ::= \varepsilon \mid c \mid s \cdot c \quad l ::= L \mid R
\end{aligned}$$

Figure 2: Core Choreographies, Syntax.

and discuss the implications of our work for other choreography languages and process calculi – in particular, we identify a Turing complete and deadlock-free fragment of value-passing CCS. Related work and discussion are given in § 7.

Publication history. This is an extended version of the material previously presented in (Cruz-Filipe and Montesi, 2017a). Noteworthy improvements include: detailed proofs of all results (in particular, Theorems 5, 7, 8, and 9); full definitions of our models (namely, the definitions of structural precongruence for CC and SP); full definitions of EPP and the merge operator. Moreover, the encoding of selections into communications (§ 5) and the encoding of CC into Channel Choreographies (§ Appendix A) are new results and have not been published previously.

2. Core Choreographies and Minimal Choreographies

We introduce Core Choreographies (CC), define function implementation and parallel execution of choreographies, and prove some key properties of CC.

2.1. Syntax of CC

The syntax of CC is displayed in Figure 2, where C ranges over choreographies. We use two (infinite) disjoint sets of names: processes (p, q, \dots) and procedures (X, \dots). Processes run in parallel, and each process stores a value – a string of the form $s \cdots s \cdot \varepsilon$ – in a local memory cell. Each process can access its own value, but it cannot read the contents of another process (no data sharing).

Term $\eta; C$ is an interaction between two processes, read “the system may execute η and proceed as C ”. An interaction η is either a value communication – $p.e \rightarrow q$ – or a label selection – $p \rightarrow q[l]$. In $p.e \rightarrow q$, p sends its local evaluation of expression e to q , which stores the received value. Expressions are either the constant ε , the value of the sender (written as c), or an application of the successor operator to c . In $p \rightarrow q[l]$, p communicates label l (either L or R) to q . **Labels are used to communicate decisions about control flow, rather than data (see § 2.3).** In a conditional $\text{if } p \stackrel{\leq}{=} q \text{ then } C_1 \text{ else } C_2$, q sends its value to p , which checks if the received value is equal to its own; the choreography proceeds as C_1 , if that is the case, or as C_2 , otherwise. In value communications, selections and conditionals, the two interacting processes must be different (no self-communications). Definitions and invocations of recursive procedures are standard. The term $\mathbf{0}$, also called *exit point*, is the terminated choreography.

In the remainder, we write $\text{pn}(C)$ for the set of all process names that appear in C .

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

Figure 3: Core Choreographies, Semantics.

2.2. Semantics of CC

The semantics of CC uses reductions of the form $C, \sigma \rightarrow C', \sigma'$. The total state function σ maps each process name to its value. We use v, w, \dots to range over values: $v, w, \dots ::= \varepsilon \mid \mathbf{s} \cdot v$. Values are isomorphic to natural numbers via $\lceil n \rceil = \mathbf{s}^n \cdot \varepsilon$. The reduction relation \rightarrow is defined by the rules given in Figure 3.

Rule C|Com models a value communication $\mathbf{p}.e \rightarrow \mathbf{q}$. In the premise, we write $e[\sigma(\mathbf{p})/c]$ for the result of replacing c with $\sigma(\mathbf{p})$ in e . In the reductum, $\sigma[\mathbf{q} \mapsto v]$ denotes the updated state function σ where \mathbf{q} now maps to v . Rule C|Sel states that selections are no-ops from the choreography's point of view (they do not change the state) – see § 2.3 for a more detailed explanation of why selections are useful. Rule C|Ctx is standard.

Rule C|Struct closes the reduction relation under the *structural precongruence* \preceq , which is the smallest precongruence satisfying the rules in Figure 4. We write $C \equiv C'$ for $C \preceq C'$ and $C' \preceq C$.

Structural precongruence allows non-interfering actions to be executed in any order, modelling concurrent process execution. This is achieved by rules C|Eta-Eta, C|Eta-Cond and C|Cond-Cond, which swap two terms describing actions performed by independent processes. They observe the same conditions as in previous choreography models, where these rules are used for the swapping relation \simeq_C , e.g., in (Carbone and Montesi, 2013). For example, rule C|Eta-Eta allows swapping of two interactions η and η' – $\eta; \eta' \equiv \eta'; \eta$ – whenever the processes that enact them are distinct – $\text{pn}(\eta) \cap \text{pn}(\eta') = \emptyset$.

Rules C|ProcEnd and C|Unfold are standard, and respectively deal with garbage collection of procedure definitions and unfolding of recursive procedures.

In rule C|Unfold, we use the standard notion of context, denoted $C[]$. We borrow its definition (adapted to our language) from (Sangiorgi and Walker, 2001). A context $C[]$ is obtained when the hole \bullet (a new reserved term) replaces one occurrence of $\mathbf{0}$ in a choreography term given by the grammar in Figure 2. We write $C[C']$ for the choreography obtained by replacing the hole \bullet in $C[]$

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

Figure 4: Core Choreographies, Structural precongruence \preceq .

with C' . Therefore, in rule C|Unfold, we are identifying a specific subterm X in the choreography on the left ($C_1[X]$) and then replacing it with the body of the recursive procedure on the right ($C_1[C_2]$).

2.3. Label Selection and Minimal Choreographies

To the reader unfamiliar with choreographies, the role of selection $p \rightarrow q[l]$ – may be unclear at this point. In existing choreography calculi, they are crucial in making choreographies projectable, as we illustrate with an example.

Example 1. Consider the following choreography.

$$C = \text{if } p \stackrel{\leftarrow}{=} q \text{ then } (p.c \rightarrow r; \mathbf{0}) \text{ else } (r.c \rightarrow p; \mathbf{0})$$

Here, p checks whether its value is the same as that of q . If so, p communicates its value to r ; otherwise, it is r that communicates its value to p . Recall that processes are assumed to run independently and share no data. Here, p is the only process that knows which branch of the conditional should be executed. However, r also needs to know this information, since it must behave differently. Intuitively, we have a problem because we are asking r to act differently based on a decision made by another process, p , and there is no propagation of this decision from p to r (either directly or indirectly, through other processes). We can easily fix the example by adding selections:

$$C' = \text{if } p \stackrel{\leftarrow}{=} q \text{ then } (p \rightarrow r[L]; p.c \rightarrow r; \mathbf{0}) \text{ else } (p \rightarrow r[R]; r.c \rightarrow p; \mathbf{0}).$$

Now, p tells r about its choice by sending a different label. This intuition will be formalised in our definition of EndPoint Projection in § 3.3. The choreography C (without label selections) is not projectable, whereas C' is. \square

The example illustrates the role of label selections in choreography languages: they are not meant to carry data, but rather to select one from several possible different behaviours offered by the process receiving the label. Thus, they model interface selections, as in the invocation of a particular method in object-oriented programming, or an operation in service-oriented computing.

One of the main results in this paper is that the same effect can actually be achieved in a language without label selections. However, this hides the way in which control flow is propagated from one process to another. Furthermore, label selections are a typical ingredient of choreography calculi, and it makes sense to include them in a core model. This is also reflected in typical implementations of endpoint languages, where label selections and value communications are implemented as different primitives. See § 7 for a more thorough discussion.

In this work, we also consider the fragment of CC without label selections, which we call Minimal Choreographies (MC). The syntax of MC is the same as that of CC, except that the action $p \rightarrow q[l]$ is disallowed; the semantics of MC is the same as that of CC, disregarding the rules that pertain to terms outside the language.

2.4. Properties

In the remainder of this section we discuss some properties of Core Choreographies. Unless otherwise stated, those properties also hold for Minimal Choreographies.

CC enjoys the usual deadlock-freedom-by-design property of choreographies.

Theorem 1 (Deadlock-freedom by design). *If C is a choreography, then either:*

- $C \preceq \mathbf{0}$ (C has terminated);
- or, for all σ , $C, \sigma \rightarrow C', \sigma'$ for some C' and σ' (C can reduce).

Proof. Assume that $C \not\preceq \mathbf{0}$ (otherwise the result is trivial). The proof is by structural induction on C . If C is not of the form $\text{def } X = C_2 \text{ in } C_1$, then the thesis is a consequence of the semantics, since there is always an applicable reduction rule. If C is of the form $\text{def } X = C_2 \text{ in } C_1$, then also $C_1 \not\preceq \mathbf{0}$, and the thesis follows by induction hypothesis applied to C_1 . \square

Using this theorem, in § 3.3 we prove that the process implementations obtained by projecting choreographies is deadlock-free.

The semantics of CC suggests a natural definition of computation. We write \rightarrow^* for the transitive closure of \rightarrow and $C, \sigma \xrightarrow{*} \mathbf{0}$ for $C, \sigma \xrightarrow{*} \mathbf{0}, \sigma'$ for any σ' .

Definition 1. *A choreography C implements a function $f : \mathbb{N}^n \rightarrow \mathbb{N}$ with input processes p_1, \dots, p_n and output process q if, for all $x_1, \dots, x_n \in \mathbb{N}$ and for every state σ s.t. $\sigma(p_i) = \lceil x_i \rceil$:*

- if $f(\tilde{x})$ is defined, then $C, \sigma \rightarrow^* \mathbf{0}, \sigma'$ where $\sigma'(q) = \lceil f(\tilde{x}) \rceil$;
- if $f(\tilde{x})$ is undefined, then $C, \sigma \not\xrightarrow{*} \mathbf{0}$.

By Theorem 1, in the second case C, σ must reduce infinitely (diverge).

In later sections, we need to characterize choreographies that are equivalent wrt a set of processes. We use the state function σ for this purpose.

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

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

Throughout this paper, we focus on choreographies with only one exit point (a single occurrence of $\mathbf{0}$). When C has a single exit point, we write $C \dot{\;} C'$ for the choreography obtained by replacing $\mathbf{0}$ in C with C' . (Requiring C to have a single exit point makes this construction linear in the sizes of C and C' , and simplifies its theoretical analysis.) This does not add expressivity to CC , but it allows for the usage of macros (as in the examples below). Then, $C \dot{\;} C'$ behaves as a “sequential composition” of C and C' , as induction over C shows.

Lemma 1. *Let C have one exit point, C' be a choreography, $\sigma, \sigma', \sigma''$ be states.*

1. *If $C, \sigma \rightarrow^* \mathbf{0}, \sigma'$ and $C', \sigma' \rightarrow^* \mathbf{0}, \sigma''$, then $C \dot{\;} C', \sigma \rightarrow^* \mathbf{0}, \sigma''$.*
2. *If $C, \sigma \not\rightarrow^* \mathbf{0}$, then $C \dot{\;} C', \sigma \not\rightarrow^* \mathbf{0}$.*
3. *If $C, \sigma \rightarrow^* \mathbf{0}, \sigma'$ and $C', \sigma' \not\rightarrow^* \mathbf{0}$, then $C \dot{\;} C', \sigma \not\rightarrow^* \mathbf{0}$.*

Proof. Straightforward by structural induction on C . □

Structural precongruence gives $C \dot{\;} C'$ fully parallel behaviour in some cases. Intuitively, C_1 and C_2 run in parallel in $C_1 \dot{\;} C_2$ if their reduction paths to $\mathbf{0}$ can be interleaved in any possible way. Below, we write $C \xrightarrow{\tilde{\sigma}}^* \mathbf{0}$ for $C, \sigma_1 \rightarrow C_2, \sigma_2 \rightarrow \dots \rightarrow \mathbf{0}, \sigma_n$, where $\tilde{\sigma} = \sigma_1, \dots, \sigma_n$, and $\widetilde{\sigma(\mathbf{p})}$ for the sequence $\sigma_1(\mathbf{p}), \dots, \sigma_n(\mathbf{p})$.

Definition 3. *Let \tilde{p} and \tilde{q} be disjoint. Then, $\tilde{\sigma}$ is an interleaving of $\widetilde{\sigma_1}$ and $\widetilde{\sigma_2}$ wrt \tilde{p} and \tilde{q} if $\tilde{\sigma}$ contains two subsequences $\widetilde{\sigma'_1}$ and $\widetilde{\sigma'_2}$ such that:*

- $\widetilde{\sigma'_2} = \tilde{\sigma} \setminus \widetilde{\sigma'_1}$;
- $\widetilde{\sigma'_1(\mathbf{p})} = \widetilde{\sigma_1(\mathbf{p})}$ for all $\mathbf{p} \in \tilde{p}$, and $\widetilde{\sigma'_2(\mathbf{q})} = \widetilde{\sigma_2(\mathbf{q})}$ for all $\mathbf{q} \in \tilde{q}$;
- $\widetilde{\sigma(r)}$ is a constant sequence for all $r \notin \tilde{p} \cup \tilde{q}$.

Definition 4 (Parallel Run). *Let C_1 and C_2 be choreographies such that $\text{pn}(C_1) \cap \text{pn}(C_2) = \emptyset$ and C_1 has only one exit point. We say that C_1 and C_2 run in parallel in $C_1 \dot{\;} C_2$ if: whenever $C_i \xrightarrow{\tilde{\sigma}_i}^* \mathbf{0}$, then $C_1 \dot{\;} C_2 \xrightarrow{\tilde{\sigma}}^* \mathbf{0}$ for every interleaving $\tilde{\sigma}$ of $\tilde{\sigma}_1$ and $\tilde{\sigma}_2$ wrt $\text{pn}(C_1)$ and $\text{pn}(C_2)$.*

Theorem 2 (Parallelisation). *Let C_1 and C_2 be choreographies such that $\text{pn}(C_1) \cap \text{pn}(C_2) = \emptyset$ and C_1 has only one exit point. Then C_1 and C_2 run in parallel in $C_1 \dot{\;} C_2$.*

Proof. The result follows directly by induction over C_1 . \square

Definition 4 and Theorem 2 straightforwardly generalise to an arbitrary number of processes. We provide an example of such parallel behaviour in Theorem 6.

2.5. Examples

We present examples of choreographies in CC, writing them as macros (syntax shortcuts). We use the notation $M(\text{params}) \triangleq C$, where M is the name of the macro, params its parameters, and C its body.

Example 2. The macro $\text{INC}(\mathbf{p}, \mathbf{t})$ increments the value of \mathbf{p} using an auxiliary process \mathbf{t} .

$$\text{INC}(\mathbf{p}, \mathbf{t}) \triangleq \mathbf{p}.c \rightarrow \mathbf{t}; \mathbf{t}.(\mathbf{s} \cdot c) \rightarrow \mathbf{p}; \mathbf{0}$$

Using INC , we write a macro $\text{ADD}(\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{t}_1, \mathbf{t}_2)$ that adds the values of \mathbf{p} and \mathbf{q} and stores the result in \mathbf{p} , using auxiliary processes \mathbf{r} , \mathbf{t}_1 and \mathbf{t}_2 . We follow the intuition as in low-level abstract register machines. First, \mathbf{t}_1 sets the value of \mathbf{r} to zero, and then calls procedure X , which increments the value of \mathbf{p} as many times as the value in \mathbf{q} . In the body of X , \mathbf{r} checks whether its value is the same as \mathbf{q} 's. If so, it informs the other processes that the recursion will terminate (selection of \mathbf{L}); otherwise, it asks them to do another step (selection of \mathbf{R}). In each step, the values of \mathbf{p} and \mathbf{r} are incremented using \mathbf{t}_1 and \mathbf{t}_2 as auxiliary processes. The compositional usage of INC is allowed, as it has exactly one exit point.

$$\begin{aligned} \text{ADD}(\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{t}_1, \mathbf{t}_2) &\triangleq \\ \text{def } X &= \text{if } \mathbf{r} \stackrel{\leftarrow}{=} \mathbf{q} \text{ then } \mathbf{r} \rightarrow \mathbf{p}[\mathbf{L}]; \mathbf{r} \rightarrow \mathbf{q}[\mathbf{L}]; \mathbf{r} \rightarrow \mathbf{t}_1[\mathbf{L}]; \mathbf{r} \rightarrow \mathbf{t}_2[\mathbf{L}]; \mathbf{0} \\ &\quad \text{else } \mathbf{r} \rightarrow \mathbf{p}[\mathbf{R}]; \mathbf{r} \rightarrow \mathbf{q}[\mathbf{R}]; \mathbf{r} \rightarrow \mathbf{t}_1[\mathbf{R}]; \mathbf{r} \rightarrow \mathbf{t}_2[\mathbf{R}]; \text{INC}(\mathbf{p}, \mathbf{t}_1) \text{ ; } \text{INC}(\mathbf{r}, \mathbf{t}_2) \text{ ; } X \\ &\quad \text{in } \mathbf{t}_1.\varepsilon \rightarrow \mathbf{r}; X \end{aligned}$$

By Theorem 2, the calls to $\text{INC}(\mathbf{p}, \mathbf{t}_1)$ and $\text{INC}(\mathbf{r}, \mathbf{t}_2)$ can be executed in parallel. Indeed, applying rule $C\text{Eta-Eta}$ for \preceq repeatedly we can check that:

$$\begin{aligned} &\underbrace{\mathbf{p}.c \rightarrow \mathbf{t}_1; \mathbf{t}_1.(\mathbf{s} \cdot c) \rightarrow \mathbf{p}; \mathbf{r}.c \rightarrow \mathbf{t}_2; \mathbf{t}_2.(\mathbf{s} \cdot c) \rightarrow \mathbf{r}; X}_{\text{expansion of } \text{INC}(\mathbf{p}, \mathbf{t}_1)} \quad \underbrace{\mathbf{r}.c \rightarrow \mathbf{t}_2; \mathbf{t}_2.(\mathbf{s} \cdot c) \rightarrow \mathbf{r}; X}_{\text{expansion of } \text{INC}(\mathbf{r}, \mathbf{t}_2)} \\ &\preceq \underbrace{\mathbf{r}.c \rightarrow \mathbf{t}_2; \mathbf{t}_2.(\mathbf{s} \cdot c) \rightarrow \mathbf{r}; \mathbf{p}.c \rightarrow \mathbf{t}_1; \mathbf{t}_1.(\mathbf{s} \cdot c) \rightarrow \mathbf{p}; X}_{\text{expansion of } \text{INC}(\mathbf{r}, \mathbf{t}_2)} \quad \underbrace{\mathbf{p}.c \rightarrow \mathbf{t}_1; \mathbf{t}_1.(\mathbf{s} \cdot c) \rightarrow \mathbf{p}; X}_{\text{expansion of } \text{INC}(\mathbf{p}, \mathbf{t}_1)} \end{aligned}$$

\square

$$\begin{aligned}
B ::= & \mathbf{q}!\langle e \rangle; B \mid \mathbf{p}?; B \mid \mathbf{q} \oplus l; B \mid \mathbf{p}\&\{l_i : B_i\}_{i \in I} \mid \\
& \mathbf{0} \mid \text{if } \mathbf{c} \stackrel{\leq}{=} \mathbf{q} \text{ then } B_1 \text{ else } B_2 \mid \text{def } X = B_2 \text{ in } B_1 \mid X \\
N, M ::= & \mathbf{p} \triangleright_v B \mid \mathbf{0} \mid N \mid M
\end{aligned}$$

Figure 5: Stateful Processes, syntax.

3. Stateful Processes, Minimal Processes and EndPoint Projections

We present Stateful Processes (SP), our target process model, and define an EndPoint Projection (EPP) that synthesises process implementations from choreographies in CC. By restricting SP adequately, we obtain a target process calculus for MC, which we call Minimal Processes (MP), with a corresponding, simpler, EPP.

3.1. Syntax of SP and MP

The syntax of SP is given in Figure 5.

Networks (N, M) are either the inactive network $\mathbf{0}$ or parallel compositions of processes $\mathbf{p} \triangleright_v B$, where \mathbf{p} is the name of the process, v its stored value, and B its behaviour.

We comment on behaviours. Expressions and labels are as in CC. A send term $\mathbf{q}!\langle e \rangle; B$ sends the evaluation of expression e to \mathbf{q} , proceeding as B . Term $\mathbf{p}?; B$, the dual receiving action, stores the value received from \mathbf{p} in the process executing the behaviour, proceeding as B . A selection term $\mathbf{q} \oplus l; B$ sends l to \mathbf{q} . Dually, a branching term $\mathbf{p}\&\{l_i : B_i\}_{i \in I}$ receives one of the labels l_i and proceeds as B_i . A process offers either: a single branch (labeled L or R); or two branches (with distinct labels). In a conditional $\text{if } \mathbf{c} \stackrel{\leq}{=} \mathbf{q} \text{ then } B_1 \text{ else } B_2$, the process receives a value from process \mathbf{q} and compares it with its own value to choose the continuation B_1 or B_2 . The other terms (definition/invoke of recursive procedures, termination) are standard.

The syntax of MP is obtained by disallowing the terms $\mathbf{q} \oplus l; B$ and $\mathbf{p}\&\{l_i : B_i\}_{i \in I}$.

3.2. Semantics

The reduction rules for SP are mostly standard, from process calculi, and are included in Figure 6. The key difference from CC is that execution is now distributed over processes. Rule S|Com follows the standard communication rule in process calculi. A process \mathbf{p} executing a send action towards a process \mathbf{q} can synchronise with a receive-from- \mathbf{p} action at \mathbf{q} ; in the reduct, \mathbf{q} 's value is updated with the value sent by \mathbf{p} , obtained by replacing the placeholder \mathbf{c} in e with the value of \mathbf{p} . Rule S|Sel is selection from session types (Honda et al., 1998), with the sender selecting one of the branches offered by the receiver. In rule S|Cond, \mathbf{p} (executing the conditional) acts as a receiver for the value sent by the process whose value it wants to read (\mathbf{q}). [All other rules are standard, and](#)

$$\begin{array}{c}
\frac{u = e[v/c]}{\mathfrak{p} \triangleright_v \mathfrak{q}!(e); B_1 \mid \mathfrak{q} \triangleright_w \mathfrak{p}^?; B_2 \rightarrow \mathfrak{p} \triangleright_v B_1 \mid \mathfrak{q} \triangleright_u B_2} \text{S|Com} \\
\frac{j \in I}{\mathfrak{p} \triangleright_v \mathfrak{q} \oplus l_j; B \mid \mathfrak{q} \triangleright_w \mathfrak{p} \&\{l_i : B_i\}_{i \in I} \rightarrow \mathfrak{p} \triangleright_v B \mid \mathfrak{q} \triangleright_w B_j} \text{S|Sel} \\
\frac{i = 1 \text{ if } v = e[w/c], \quad i = 2 \text{ otherwise}}{\mathfrak{p} \triangleright_v \text{ if } c \stackrel{\leq}{=} \mathfrak{q} \text{ then } B_1 \text{ else } B_2 \mid \mathfrak{q} \triangleright_w \mathfrak{p}!(e); B' \rightarrow \mathfrak{p} \triangleright_v B_i \mid \mathfrak{q} \triangleright_w B'} \text{S|Cond} \\
\frac{\mathfrak{p} \triangleright_v B_1 \rightarrow \mathfrak{p} \triangleright_v B'_1}{\mathfrak{p} \triangleright_v \text{ def } X = B_2 \text{ in } B_1 \rightarrow \mathfrak{p} \triangleright_v \text{ def } X = B_2 \text{ in } B'_1} \text{S|Ctx} \\
\frac{N \rightarrow N'}{N \mid M \rightarrow N' \mid M} \text{S|Par} \quad \frac{N \preceq M \quad M \rightarrow M' \quad M' \preceq N'}{N \rightarrow N'} \text{S|Struct}
\end{array}$$

Figure 6: Stateful Processes, Semantics.

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

Figure 7: Stateful Processes, Structural Precongruence.

use a structural precongruence that includes associativity and commutativity of parallel composition, together with the rules in Figure 7, supporting recursion unfolding and garbage collection of terminated processes and unused definitions.

As for CC, we can define function implementation in SP.

Definition 5 (Function implementation in SP). *A network N implements a function $f : \mathbb{N}^n \rightarrow \mathbb{N}$ with input processes $\mathfrak{p}_1, \dots, \mathfrak{p}_n$ and output process \mathfrak{q} if $N \preceq (\prod_{i \in [1, n]} \mathfrak{p}_i \triangleright_{v_i} B_i) \mid \mathfrak{q} \triangleright_w B' \mid N'$ and, for all $x_1, \dots, x_n \in \mathbb{N}$:*

- if $f(\tilde{x})$ is defined, then $N(\tilde{x}) \rightarrow^* \mathfrak{q} \triangleright_{\Gamma_{f(\tilde{x})} \cap} \mathbf{0}$;
- if $f(\tilde{x})$ is not defined, then $N(\tilde{x}) \not\rightarrow^* \mathbf{0}$.

where $N(\tilde{x})$ is a shorthand for $N[\widetilde{\Gamma_{x_i} \cap} / v_i]$, the network obtained by replacing in N the values of the input processes with the arguments of the function.

3.3. EndPoint Projection

We now define an EndPoint Projection (EPP) from CC to SP.

$$\begin{aligned}
\llbracket p.e \rightarrow q; C \rrbracket_r &= \begin{cases} q!(e); \llbracket C \rrbracket_r & \text{if } r = p \\ p?; \llbracket C \rrbracket_r & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{o.w.} \end{cases} \\
\llbracket p \rightarrow q[l]; C \rrbracket_r &= \begin{cases} q \oplus l; \llbracket C \rrbracket_r & \text{if } r = p \\ p \& \{l : \llbracket C \rrbracket_r\} & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{o.w.} \end{cases} \\
\llbracket \text{if } p \stackrel{\leq}{=} q \text{ then } C_1 \text{ else } C_2 \rrbracket_r &= \begin{cases} \text{if } c \stackrel{\leq}{=} q \text{ then } \llbracket C_1 \rrbracket_r \text{ else } \llbracket C_2 \rrbracket_r & \text{if } r = p \\ p!(c); (\llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r) & \text{if } r = q \\ \llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r & \text{o.w.} \end{cases} \\
\llbracket \text{def } X^{\tilde{p}} = C_2 \text{ in } C_1 \rrbracket_r &= \begin{cases} \text{def } X = \llbracket C_2 \rrbracket_r \text{ in } \llbracket C_1 \rrbracket_r & \text{if } r \in \tilde{p} \\ \llbracket C_1 \rrbracket_r & \text{o.w.} \end{cases} \\
\llbracket \mathbf{0} \rrbracket_r = \mathbf{0} \quad \llbracket X^{\tilde{p}} \rrbracket_r &= \begin{cases} X & \text{if } r \in \tilde{p} \\ \mathbf{0} & \text{o.w.} \end{cases}
\end{aligned}$$

Figure 8: Core Choreographies, Behaviour Projection.

We first discuss the rules for projecting the behaviour of a single process p , a partial function $\llbracket C \rrbracket_p$ defined by the rules in Figure 8. Selections are projected similarly to communications, and $\mathbf{0}$ is projected to $\mathbf{0}$. All rules follow the intuition of projecting, for each choreography term, the local action performed by the process that we are projecting. For example, for a communication term $p.e \rightarrow q$, we project a send action for the sender p , a receive action for the receiver q , or just the continuation otherwise. The rule for selection is similar.

The rules for projecting recursive definitions and recursive calls are defined assuming that procedure names have been annotated with the process names appearing inside the body of the procedure, in order to avoid projecting unnecessary procedure code – see (Carbone et al., 2012). This is an easy preprocessing of the choreography, which we assume is performed before using projection. The preprocessing recursively visits the choreography and leaves it untouched, except when it meets a recursive definition term $\text{def } X = C_2 \text{ in } C_1$. In that case, it annotates the choreography – $\text{def } X^{\tilde{p}} = C_2\{X^{\tilde{p}}/X\} \text{ in } C_1\{X^{\tilde{p}}/X\}$ where $\{\tilde{p}\} = \text{pn}(C_2)$ – and then proceeds recursively. In other words, the definition of X and all its invocations in C_2 and C_1 are annotated with \tilde{p} .

The rule for projecting a conditional is more involved, using the partial merging operator \sqcup to merge the possible behaviours of a process that does not know which branch will be chosen. The formal definition is found in Figure 9. Merging is a homomorphic binary operator; for all terms but branchings it requires isomorphism, $q!(e); B \sqcup q!(e); B' = q!(e); (B \sqcup B')$. The only case where

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

Figure 9: Core Choreographies, Merge Operator in Behaviour Projection.

branching terms can have unmergeable continuations is when they are guarded by distinct labels, in which case merge returns a larger branching including all options (merging branches with the same label).

Merging explains the role of selections in CC, common in choreography models (Coppo et al., 2016; Carbone et al., 2012; Carbone and Montesi, 2013; Honda et al., 2008; Dalla Preda et al., 2017; Qiu et al., 2007). Recall the choreographies from Example 1. In choreography C , the behaviour of r cannot be projected because we cannot merge its different behaviours in the two branches of the conditional (a send with a receive). Choreography C' is projectable, and the behaviour of r is $\llbracket C \rrbracket_r = \mathbf{p} \&\{\mathbf{L} : \mathbf{p}?; \mathbf{0}, \mathbf{R} : \mathbf{p}!\langle c \rangle; \mathbf{0}\}$.

Definition 6 (EPP from CC to SP). *Given a choreography C and a state σ , the endpoint projection of C and σ is the parallel composition of the projections of the processes in C .*

$$\llbracket C, \sigma \rrbracket = \prod_{\mathbf{p} \triangleright_{\sigma}(\mathbf{p})} \llbracket C \rrbracket_{\mathbf{p}}$$

The EPP from MC to MP is defined by restricting the EPP from CC to SP to the relevant cases. Since choreographies in MC do not have selections, process projections of choreographies in MC never have branchings. This means that, in the case of MC, the merging operator \sqcup used in EPP is exactly syntactic equality (since the only nontrivial case was that of branchings). Consequently, we can replace the rule for projecting conditionals with the following, simpler, one.

$$\llbracket \text{if } p \stackrel{\leq}{=} q \text{ then } C_1 \text{ else } C_2 \rrbracket_r = \begin{cases} \text{if } c \stackrel{\leq}{=} q \text{ then } \llbracket C_1 \rrbracket_r \text{ else } \llbracket C_2 \rrbracket_r & \text{if } r = p \\ p!(c); \llbracket C_1 \rrbracket_r & \text{if } r = q \text{ and } \llbracket C_1 \rrbracket_r = \llbracket C_2 \rrbracket_r \\ \llbracket C_1 \rrbracket_r & \text{if } r \notin \{p, q\} \text{ and } \llbracket C_1 \rrbracket_r = \llbracket C_2 \rrbracket_r \end{cases}$$

Since the state function σ is total, $\llbracket C, \sigma \rrbracket$ is defined for some σ iff $\llbracket C, \sigma' \rrbracket$ is defined for all other σ' . In this case, we say that C is *projectable*.

Example 3. Recall the definition of $\text{INC}(p, t)$.

$$\text{INC}(p, t) \triangleq p.c \rightarrow t; t.(s \cdot c) \rightarrow p; \mathbf{0}$$

Given any σ , the EPP of $\text{INC}(p, t)$ is:

$$\llbracket \text{INC}(p, t), \sigma \rrbracket = p \triangleright_{\sigma(p)} t!(c); t?; \mathbf{0} \mid t \triangleright_{\sigma(t)} p?; p!(s \cdot c); \mathbf{0}$$

In order to characterize the correspondence between choreographies and their projections, we introduce a *pruning relation* between networks, inspired by (Carbone et al., 2012; Carbone and Montesi, 2013).

Definition 7. The pruning relation \sqsupseteq on networks is defined as $\mathbf{0} \sqsupseteq \mathbf{0}$ and $p \triangleright_v B \mid N \sqsupseteq p \triangleright_v B' \mid N'$ if $B \sqcup B' = B$ and $N \sqsupseteq N'$.

Intuitively, $N \sqsupseteq N'$ if both N and N' are parallel compositions of the same processes, and each process in N offers at least the same behaviours as in N' , with potentially extra options in its branching terms.

Using pruning, we can formalise the operational correspondence guaranteed by EPP.

Theorem 3 (Operational Correspondence (CC \leftrightarrow SP)). *Let C be a projectable choreography. Then, for all σ :*

Completeness: *If $C, \sigma \rightarrow C', \sigma'$, then $\llbracket C, \sigma \rrbracket \rightarrow \sqsupseteq \llbracket C', \sigma' \rrbracket$;*

Soundness: *If $\llbracket C, \sigma \rrbracket \rightarrow N$, then there exist C', σ' and N' such that $C, \sigma \rightarrow C', \sigma'$ and $N \preceq N' \sqsupseteq \llbracket C', \sigma' \rrbracket$.*

Proof. The proof is an adaptation of the proofs of similar results in (Carbone et al., 2012; Carbone and Montesi, 2013); we sketch a few typical cases, also to illustrate where pruning plays a role.

Completeness: by induction on the derivation of the reduction $C, \sigma \rightarrow C', \sigma'$.

- Consider the case of rule C|Com. Then $C = p.e \rightarrow q; C', \sigma' = \sigma[q \mapsto e[\sigma(p)/c]]$ and

$$\begin{aligned} \llbracket C, \sigma \rrbracket &= p \triangleright_{\sigma(p)} q!(e); B_p \mid q \triangleright_{\sigma(q)} p?; B_q \mid N'' \\ &\rightarrow p \triangleright_{\sigma(p)} B_p \mid q \triangleright_{e[\sigma(p)/v]} B_q \mid N'' \end{aligned}$$

If $B_p, B_q \not\preceq \mathbf{0}$, then the result of the reduction is precisely $\llbracket C', \sigma' \rrbracket$. If one or both of these behaviours are $\mathbf{0}$, then \preceq needs to be applied to remove the corresponding processes from the network in order to obtain $\llbracket C', \sigma' \rrbracket$.

- Consider the case of rule C|Cond with $i = 1$ (i.e. $\sigma(\mathbf{p}) = \sigma(\mathbf{q})$). Then $C = \text{if } \mathbf{p} \stackrel{\leftarrow}{=} \mathbf{q} \text{ then } C_1 \text{ else } C_2, C' = C_1, \sigma' = \sigma$ and

$$\begin{aligned} \llbracket C, \sigma \rrbracket &= \mathbf{p} \triangleright_{\sigma(\mathbf{p})} \text{if } \mathbf{c} \stackrel{\leftarrow}{=} \mathbf{q} \text{ then } B_1 \text{ else } B_2 \mid \mathbf{q} \triangleright_{\sigma(\mathbf{q})} \mathbf{p}! \langle \mathbf{c} \rangle; B_{\mathbf{q}} \mid N'' \\ &\rightarrow \mathbf{p} \triangleright_{\sigma(\mathbf{p})} B_1 \mid \mathbf{q} \triangleright_{\sigma(\mathbf{q})} B_{\mathbf{q}} \mid N'' \end{aligned}$$

Assume first for simplicity that $B_1, B_{\mathbf{q}} \not\stackrel{\neq}{=} \mathbf{0}$. For each process $r \neq \mathbf{p}$, the network obtained above is $\llbracket C_1, \sigma \rrbracket_r \sqcup \llbracket C_2, \sigma \rrbracket_r$, which in general does not coincide with $\llbracket C_1, \sigma \rrbracket_r$. However, since merge is trivially associative and idempotent, it immediately follows that $\mathbf{p} \triangleright_{\sigma(\mathbf{p})} B_1 \mid \mathbf{q} \triangleright_{\sigma(\mathbf{q})} B_{\mathbf{q}} \mid N'' \sqsupseteq \llbracket C_1, \sigma \rrbracket$.

In the case that $B_1, B_{\mathbf{q}} \stackrel{\neq}{=} \mathbf{0}$, structural precongurence also needs to be applied to remove the respective processes from the resulting network.

- Finally consider the case when C|Struct applies, i.e. $C \preceq C_1, C_1, \sigma \rightarrow C'_1, \sigma'$ and $C'_1 \preceq C'$. The thesis then follows by observing that rules C|Eta-Eta, C|Eta-Cond, C|Eta-Rec and C|Cond-Cond do not change the projected choreographies, while the behaviour of C|Unfold is reproducible by S|Unfold. Therefore $\llbracket C, \sigma \rrbracket \preceq \llbracket C_1, \sigma \rrbracket$ and the induction hypothesis applies. It remains to be shown that $\llbracket C'_1, \sigma' \rrbracket \preceq \llbracket C', \sigma' \rrbracket$, which also requires considering rule C|ProcEnd – whose behaviour again can be directly replicated by S|ProcEnd.

Soundness: the proof is now by induction on the structure of C . We detail one representative case.

- Suppose that $C = \mathbf{p}.e \rightarrow \mathbf{q}; C'$. Then $\llbracket C, \sigma \rrbracket = \mathbf{p} \triangleright_{\sigma(\mathbf{p})} \mathbf{q}! \langle e \rangle; B_{\mathbf{p}} \mid \mathbf{q} \triangleright_{\sigma(\mathbf{q})} \mathbf{p}?: B_{\mathbf{q}} \mid N''$, and there are two cases.

If $\llbracket C, \sigma \rrbracket \rightarrow N$ involves the communication between \mathbf{p} and \mathbf{q} , then the choreography can mimic this action by consuming $\mathbf{p}.e \rightarrow \mathbf{q}$ and evolving to C' . Any applications of structural precongurence to $\llbracket C, \sigma \rrbracket$ must involve only rules S|Unfold and S|ProcEnd, which can be directly applied also to C (and may require additional applications of the same rules to other processes in N , namely those that occur in the definition of the relevant variable X). Likewise, applications of structural precongurence to N either involve the same rules (and the same argument applies to $\llbracket C', \sigma' \rrbracket$) or garbage collection (which is transparent, since it can never be applied to $\llbracket C', \sigma' \rrbracket$). Finally, additional garbage collection rules may be needed to obtain $\llbracket C', \sigma' \rrbracket$. In any case, $N \preceq \llbracket C', \sigma' \rrbracket$.

If $\llbracket C, \sigma \rrbracket \rightarrow N$ does not involve \mathbf{p} and \mathbf{q} , then we observe that necessarily N has the form $\mathbf{p} \triangleright_{\sigma(\mathbf{p})} \mathbf{q}! \langle e \rangle; B_{\mathbf{p}} \mid \mathbf{q} \triangleright_{\sigma(\mathbf{q})} \mathbf{p}?: B_{\mathbf{q}} \mid N''$, where $\llbracket C', \sigma' \rrbracket \rightarrow \mathbf{p} \triangleright_{\sigma(\mathbf{p})} B_{\mathbf{p}} \mid \mathbf{q} \triangleright_{\sigma(\mathbf{q})} B_{\mathbf{q}} \mid N'' = N^-$. Then we can apply the induction hypothesis to C' , and find C'', σ' such that $N^- \preceq N' \sqsupseteq \llbracket C'', \sigma' \rrbracket$ for some N' . Since the transition from C' to C'' does not involve \mathbf{p} or \mathbf{q} , a simple induction argument shows that

$p.e \rightarrow q; C', \sigma \rightarrow p.e \rightarrow q; C'', \sigma'$, and from the definition of \preceq and \sqsupseteq it follows that $N \preceq N''' \sqsupseteq \llbracket p.e \rightarrow q; C'', \sigma' \rrbracket$ where N''' is obtained from N' by prepending the original communication actions between p and q at the head of these processes' behaviours.

□

As a consequence of Theorems 1 and 3, choreography projections never deadlock.

Theorem 4 (Deadlock-freedom by construction). *Let $N = \llbracket C, \sigma \rrbracket$ for some C and σ . Then, either $N \preceq \mathbf{0}$ (N has terminated), or $N \rightarrow N'$ for some N' (N can reduce).*

Proof. If $N \preceq \mathbf{0}$ then the theorem clearly holds. Otherwise, the thesis follows from Theorems 1 and 3. □

3.4. Choreography Amendment

An important property of CC is that all unprojectable choreographies can be made projectable by adding some selections. We annotate recursion variables as for EPP, assuming that $\text{pn}(X^{\bar{p}}) = \{\bar{p}\}$.

Definition 8 (Amendment). *Given C in CC, the transformation $\text{Amend}(C)$ repeatedly applies the following procedure until no longer possible, starting from the innermost subterms in C . For each conditional subterm $\text{if } p \stackrel{\leq}{=} q \text{ then } C_1 \text{ else } C_2$ in C , let $\tilde{r} \subseteq (\text{pn}(C_1) \cup \text{pn}(C_2))$ be the largest set such that $\llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r$ is undefined for all $r \in \tilde{r}$; then $\text{if } p \stackrel{\leq}{=} q \text{ then } C_1 \text{ else } C_2$ in C is replaced with:*

$$\text{if } (p \stackrel{\leq}{=} q) \text{ then } (p \rightarrow r_1[L]; \dots ; p \rightarrow r_n[L]; C_1) \text{ else } (p \rightarrow r_1[R]; \dots ; p \rightarrow r_n[R]; C_2)$$

Lemma 2 (Amendment Lemma). *For every choreography C :*

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

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

Correspondence: for all $\sigma, C, \sigma \rightarrow^* C', \sigma'$ iff $\text{Amend}(C), \sigma \rightarrow^* \text{Amend}(C'), \sigma'$.

Proof. This procedure clearly terminates, since each iteration processes one conditional that is never changed again. By construction, the resulting choreography is projectable: if $\llbracket C_1, \sigma \rrbracket_r$ and $\llbracket C_2, \sigma \rrbracket_r$ are not mergeable, then the corresponding terms in $\llbracket \text{Amend}(\text{if } p \stackrel{\leq}{=} q \text{ then } C_1 \text{ else } C_2), \sigma \rrbracket_r$ are both branching terms, making them mergeable. Correspondence is immediate, since the only new actions in $\text{Amend}(C')$ are label selections, which do not change the state when executed. □

Example 4. *Applying Amend to the choreography C in Example 1 yields the choreography C' in the same example.* □

Example 5. *Thanks to merging, amendment can also recognise some situations where additional selections are not needed. For example, in the choreography $C = \text{if } p \stackrel{\leftarrow}{=} q \text{ then } (p.(s \cdot c) \rightarrow r; \mathbf{0}) \text{ else } (p.(c) \rightarrow r; \mathbf{0})$, r does not need to know the choice made by p , as it always performs the same input action. Here, C is projectable and $\text{Amend}(C) = C$. \square*

Note that amending a choreography from MC returns a choreography in CC. In § 5 we show that we can prove an amendment lemma for MC, but this will require much more work.

4. Turing Completeness of MC and Its Consequences

We now move to our main result: the set of choreography projections of CC (the processes synthesised by EPP) is not only deadlock-free, but also capable of computing all partial recursive functions, as defined by Kleene (Kleene, 1952), and hence Turing complete. To this aim, the design and properties of CC give us a considerable pay off. First, by Theorem 3, the problem reduces to establishing that a projectable fragment of CC is Turing complete. Second, by Lemma 2, this simpler problem is reduced to establishing that MC is Turing complete, since any choreography in MC can be amended to one in CC that is projectable and computes the same values. We also exploit the concurrent semantics of CC and Theorem 2 to parallelise independent sub-computations (Theorem 6). By projecting our choreographies via EPP, we obtain corresponding function implementations in the process calculus SP.

Establishing that CC is Turing complete is long, but not difficult. Our proof is in line with other traditional proofs of computational completeness (Cutland, 1980; Kleene, 1952; Turing, 1937), where data and programs are distinct. This differs from other proofs of similar results for, e.g., π -calculus (Sangiorgi and Walker, 2001) and λ -calculus (Barendregt, 1984), which encode data as particular programs. The advantages are: our proof can be used to build choreographies that compute particular functions; and we can parallelise independent sub-computations in functions (Theorem 6).

4.1. Partial Recursive Functions

Our definition of the class of partial recursive functions \mathcal{R} is slightly simplified, but equivalent to, that in (Kleene, 1952), where it is also shown that \mathcal{R} is the class of functions computable by a Turing machine. \mathcal{R} is defined inductively as follows.

Unary zero: $Z \in \mathcal{R}$, where $Z : \mathbb{N} \rightarrow \mathbb{N}$ is s.t. $Z(x) = 0$ for all $x \in \mathbb{N}$.

Unary successor: $S \in \mathcal{R}$, where $S : \mathbb{N} \rightarrow \mathbb{N}$ is s.t. $S(x) = x + 1$ for all $x \in \mathbb{N}$.

Projections: If $n \geq 1$ and $1 \leq m \leq n$, then $P_m^n \in \mathcal{R}$, where $P_m^n : \mathbb{N}^n \rightarrow \mathbb{N}$ is s.t. $P_m^n(x_1, \dots, x_n) = x_m$ for all $x_1, \dots, x_n \in \mathbb{N}$.

Composition: if $f, g_i \in \mathcal{R}$ for $1 \leq i \leq k$, with each $g_i : \mathbb{N}^n \rightarrow \mathbb{N}$ and $f : \mathbb{N}^k \rightarrow \mathbb{N}$, then $h = C(f, \tilde{g}) \in \mathcal{R}$, where $h : \mathbb{N}^n \rightarrow \mathbb{N}$ is defined by composition from f and g_1, \dots, g_k as: $h(\tilde{x}) = f(g_1(\tilde{x}), \dots, g_k(\tilde{x}))$.

Primitive recursion: if $f, g \in \mathcal{R}$, with $f : \mathbb{N}^n \rightarrow \mathbb{N}$ and $g : \mathbb{N}^{n+2} \rightarrow \mathbb{N}$, then $h = R(f, g) \in \mathcal{R}$, where $h : \mathbb{N}^{n+1} \rightarrow \mathbb{N}$ is defined by primitive recursion from f and g as: $h(0, \tilde{x}) = f(\tilde{x})$ and $h(x_0 + 1, \tilde{x}) = g(x_0, h(x_0, \tilde{x}), \tilde{x})$.

Minimization: If $f \in \mathcal{R}$, with $f : \mathbb{N}^{n+1} \rightarrow \mathbb{N}$, then $h = M(f) \in \mathcal{R}$, where $h : \mathbb{N}^n \rightarrow \mathbb{N}$ is defined by minimization from f as: $h(\tilde{x}) = y$ iff (1) $f(\tilde{x}, y) = 0$ and (2) $f(\tilde{x}, z)$ is defined and different from 0 for all $z < y$.

Example 6 (Addition and Subtraction). *We show that $\text{add} \in \mathcal{R}$, where $\text{add} : \mathbb{N}^2 \rightarrow \mathbb{N}$ adds its two arguments. Since $\text{add}(0, y) = y$ and $\text{add}(x + 1, y) = \text{add}(x, y) + 1$, we can define add by recursion as $\text{add} = R(P_1^1, C(S, P_2^3))$. Indeed, the function $y \mapsto \text{add}(0, y)$ is simply P_1^1 , whereas $1 + \text{add}(x, y)$ is $h(x, \text{add}(x, y), y)$ where $h(x, y, z) = y + 1$, which is the composition of the successor function with P_2^3 .*

From addition, we can define subtraction by minimization, since $\text{sub}(x, y) = x - y$ is the smallest z such that $y + z = x$ (subtraction is not defined if $y > x$). We use an auxiliary function $\text{eq}(x, y)$ that returns 0 if $x = y$ and a non-zero value otherwise, which is known to be partial recursive. Then we can define subtraction as $\text{sub} = M(C(\text{eq}, C(\text{add}, P_2^3, P_3^3), P_1^3))$. Indeed, composing add with P_2^3 and P_3^3 produces $(x, y, z) \mapsto y + z$, and the outer composition yields $(x, y, z) \mapsto \text{eq}(y + z, x)$. This function evaluates to 0 precisely when $z = y - x$, and applying minimization computes this value from x and y .

4.2. Encoding Partial Recursive Functions in MC and CC

All functions in \mathcal{R} can be implemented in CC, in the sense of Definition 1. Since selections can be inferred by amendment, we develop our encoding in MC and discuss projectability later.

Given $f : \mathbb{N}^n \rightarrow \mathbb{N}$, we denote its implementation in MC by $\{\{f\}\}^{\tilde{p} \rightarrow \mathfrak{q}}$, where \tilde{p} and \mathfrak{q} are parameters. All choreographies we build have a single exit point, and we combine them using the sequential composition operator $\mathfrak{;}$ from § 2.

We use auxiliary processes (r_0, r_1, \dots) for intermediate computation, and annotate the encoding with the index ℓ of the first free auxiliary process name $(\{\{f\}\}_\ell^{\tilde{p} \rightarrow \mathfrak{q}})$. To alleviate the notation, the encoding assigns mnemonic names to these processes and their correspondence to the actual process names is formalised in the text using $\pi(f)$ for the number of auxiliary processes needed for encoding $f : \mathbb{N}^n \rightarrow \mathbb{N}$, defined by

$$\begin{aligned} \pi(S) = \pi(Z) = \pi(P_m^n) &= 0 & \pi(R(f, g)) &= \pi(f) + \pi(g) + 3 \\ \pi(C(f, g_1, \dots, g_k)) &= \pi(f) + \sum_{i=1}^k \pi(g_i) + k & \pi(M(f)) &= \pi(f) + 3 \end{aligned}$$

For simplicity, we write \tilde{p} for p_1, \dots, p_n (when n is known) and $\{A_i\}_{i=1}^n$ for $A_1 \mathfrak{;} \dots \mathfrak{;} A_n$.

The encoding of the base cases is straightforward.

$$\{\!|Z|\!\}_{\ell}^{\mathbf{p} \mapsto \mathbf{q}} = \mathbf{p}.\varepsilon \rightarrow \mathbf{q} \quad \{\!|S|\!\}_{\ell}^{\mathbf{p} \mapsto \mathbf{q}} = \mathbf{p}.\langle \mathbf{s} \cdot \mathbf{c} \rangle \rightarrow \mathbf{q} \quad \{\!|P_m^n|\!\}_{\ell}^{\bar{\mathbf{p}} \mapsto \mathbf{q}} = \mathbf{p}_m.\mathbf{c} \rightarrow \mathbf{q}$$

Composition is also simple. Let $h = C(f, g_1, \dots, g_k) : \mathbb{N}^n \rightarrow \mathbb{N}$. Then:

$$\{\!|h|\!\}_{\ell}^{\bar{\mathbf{p}} \mapsto \mathbf{q}} = \left\{ \{\!|g_i|\!\}_{\ell_i}^{\bar{\mathbf{p}} \mapsto \mathbf{r}'_i} \right\}_{i=1}^k \ ; \ \{\!|f|\!\}_{\ell_{k+1}}^{\mathbf{r}'_1, \dots, \mathbf{r}'_k \mapsto \mathbf{q}}$$

where $\mathbf{r}'_i = \mathbf{r}_{\ell+i-1}$, $\ell_1 = \ell + k$ and $\ell_{i+1} = \ell_i + \pi(g_i)$. Each auxiliary process \mathbf{r}'_i connects the output of g_i to the corresponding input of f . Choreographies obtained inductively use these process names as parameters; name clashes are prevented by increasing ℓ . By definition of $\ ; \$ $\{\!|g_{i+1}|\!\}$ is substituted for the (unique) exit point of $\{\!|g_i|\!\}$, and $\{\!|f|\!\}$ is substituted for the exit point of $\{\!|g_k|\!\}$. The resulting choreography also has only one exit point (that of $\{\!|f|\!\}$). Below we discuss how to modify this construction slightly so that the g_i s are computed in parallel.

For the recursion operator, we need to use recursive procedures. Let $h = R(f, g) : \mathbb{N}^{n+1} \rightarrow \mathbb{N}$. Then, using the macro `INC` from Example 2 for brevity:

$$\begin{aligned} \{\!|h|\!\}_{\ell}^{\mathbf{p}_0, \dots, \mathbf{p}_n \mapsto \mathbf{q}} = & \text{def } T = \text{if } (r_c \stackrel{\leq}{=} \mathbf{p}_0) \text{ then } (\mathbf{q}'.\mathbf{c} \rightarrow \mathbf{q}; \mathbf{0}) \\ & \text{else } \{\!|g|\!\}_{\ell_g}^{\mathbf{r}_c, \mathbf{q}', \mathbf{p}_1, \dots, \mathbf{p}_n \mapsto \mathbf{r}_t} \ ; \ \mathbf{r}_t.\mathbf{c} \rightarrow \mathbf{q}'; \text{INC}(\mathbf{r}_c, \mathbf{r}_t) \ ; \ T \\ & \text{in } \{\!|f|\!\}_{\ell_f}^{\mathbf{p}_1, \dots, \mathbf{p}_n \mapsto \mathbf{q}'} \ ; \ \mathbf{r}_t.\varepsilon \rightarrow \mathbf{r}_c; T \end{aligned}$$

where $\mathbf{q}' = \mathbf{r}_{\ell}$, $\mathbf{r}_c = \mathbf{r}_{\ell+1}$, $\mathbf{r}_t = \mathbf{r}_{\ell+2}$, $\ell_f = \ell + 3$ and $\ell_g = \ell_f + \pi(f)$. Process \mathbf{r}_c is a counter, \mathbf{q}' stores intermediate results, and \mathbf{r}_t is temporary storage; T checks the value of \mathbf{r}_c and either outputs the result or recurs. Note that $\{\!|h|\!\}$ has only one exit point (after the communication from \mathbf{r} to \mathbf{q}), as the exit points of $\{\!|f|\!\}$ and $\{\!|g|\!\}$ are replaced by code ending with calls to T .

The strategy for minimization is similar, but simpler. Let $h = M(f) : \mathbb{N}^n \rightarrow \mathbb{N}$. Again we use a counter \mathbf{r}_c and compute successive values of f , stored in \mathbf{q}' , until a zero is found. This procedure may loop forever, either because $f(\tilde{x}, x_{n+1})$ is never 0 or because one of the evaluations itself never terminates.

$$\begin{aligned} \{\!|h|\!\}_{\ell}^{\mathbf{p}_1, \dots, \mathbf{p}_{n+1} \mapsto \mathbf{q}} = & \text{def } T = \{\!|f|\!\}_{\ell_f}^{\mathbf{p}_1, \dots, \mathbf{p}_n, \mathbf{r}_c \mapsto \mathbf{q}'} \ ; \ \mathbf{r}_c.\varepsilon \rightarrow \mathbf{r}_z; \\ & \text{if } (\mathbf{r}_z \stackrel{\leq}{=} \mathbf{q}') \text{ then } (\mathbf{r}_c.\mathbf{c} \rightarrow \mathbf{q}; \mathbf{0}) \text{ else } (\text{INC}(\mathbf{r}_c, \mathbf{r}_z) \ ; \ T) \\ & \text{in } \mathbf{r}_z.\varepsilon \rightarrow \mathbf{r}_c; T \end{aligned}$$

where $\mathbf{q}' = \mathbf{r}_{\ell}$, $\mathbf{r}_c = \mathbf{r}_{\ell+1}$, $\mathbf{r}_z = \mathbf{r}_{\ell+2}$, $\ell_f = \ell + 3$ and $\ell_g = \ell_f + \pi(f)$. In this case, the whole if-then-else is inserted at the exit point of $\{\!|f|\!\}$; the only exit point of this choreography is again after communicating the result to \mathbf{q} .

Definition 9 (Encoding). *Let $f \in \mathcal{R}$. The encoding of f in MC is $\{\!|f|\!\}_{\bar{\mathbf{p}} \mapsto \mathbf{q}} = \{\!|f|\!\}_{\mathbf{0}}^{\bar{\mathbf{p}} \mapsto \mathbf{q}}$.*

Example 7. We illustrate this construction by showing the encoding of the add and sub functions given in Example 6. Recall that $\text{add} = R(P_1^1, C(S, P_2^3))$. Expanding $\{\{\text{add}\}\}^{\text{p}_x, \text{p}_y \mapsto \text{q}}$ we obtain:

$$\begin{aligned} \{\{\text{add}\}\}_0^{\text{p}_x, \text{p}_y \mapsto \text{q}} = \\ \text{def } T = & \text{if } (r_1 \stackrel{\leq}{=} \text{p}_x) \text{ then } (r_0.c \rightarrow \text{q}; \mathbf{0}) \\ & \text{else } \underbrace{r_0.c \rightarrow r_3}_{\{\{P_2^3\}\}_4^{r_1, r_0, \text{p}_y \mapsto r_3}} ; \underbrace{r_3.(s \cdot c) \rightarrow r_2}_{\{\{S\}\}_4^{r_3 \mapsto r_2}} ; r_2.c \rightarrow r_0 ; \underbrace{r_1.c \rightarrow r_2; r_2.(s \cdot c) \rightarrow r_1}_{\text{INC}(r_1, r_2)} ; T \\ \text{in } & \underbrace{\text{p}_y.c \rightarrow r_0; r_2.\varepsilon \rightarrow r_1; T}_{\{\{P_1^1\}\}_3^{\text{p}_y \mapsto r_0}} \end{aligned}$$

The first two actions in the else branch correspond to $\{\{C(S, P_2^3)\}\}_3^{r_1, r_0, \text{p}_y \mapsto r_2}$.

For subtraction, we first show how to implement equality directly in MC, without resorting to its proof of membership in \mathcal{R} . This choreography is not the simplest possible because we want it to have only one exit point; its construction illustrates how any choreography can be transformed to have this property.

$$\begin{aligned} \text{EQ}(\text{p}_x, \text{p}_y, \text{q}, r) \triangleq & \text{def } T = (r.c \rightarrow \text{q}; \mathbf{0}) \text{ in} \\ & \text{if } (\text{p}_x \stackrel{\leq}{=} \text{p}_y) \text{ then } (\text{p}_x.\varepsilon \rightarrow r; T) \text{ else } (\text{p}_x.(s \cdot c) \rightarrow r; T) \end{aligned}$$

Recall now that $\text{sub} = M(C(\text{eq}, C(\text{add}, P_2^3, P_3^3), P_1^3))$. Unfolding the encoding of minimization and composition, we obtain that $\{\{\text{sub}\}\}_0^{\text{p}_x, \text{p}_y \mapsto \text{q}}$ is

$$\begin{aligned} \text{def } T = & \{\{P_2^3\}\}_7^{\text{p}_x, \text{p}_y, r_1 \mapsto r_5} \circ \{\{P_2^3\}\}_7^{\text{p}_x, \text{p}_y, r_1 \mapsto r_6} \circ \{\{\text{add}\}\}_7^{r_5, r_6 \mapsto r_3} \circ \\ & \{\{P_1^3\}\}_{11}^{\text{p}_x, \text{p}_y, r_1 \mapsto r_4} \circ \text{EQ}(r_3, r_4, r_0, r_{11}) \circ r_1.\varepsilon \rightarrow r_2; \\ & \text{if } r_2 \stackrel{\leq}{=} r_0 \text{ then } (r_1.c \rightarrow \text{q}; \mathbf{0}) \text{ else } (\text{INC}(r_1, r_2) \circ T) \\ \text{in } & r_2.\varepsilon \rightarrow r_1; T \end{aligned}$$

The first line in the definition of T is $\{\{C(\text{add}, P_2^3, P_3^3)\}\}_5^{\text{p}_x, \text{p}_y, r_1 \mapsto r_3}$; the first five processes composed therewithin are

$$\{\{C(\text{eq}, C(\text{add}, P_2^3, P_3^3), P_1^3)\}\}_3^{\text{p}_x, \text{p}_y, r_1 \mapsto r_0}.$$

Fully unfolding the base cases, we obtain

$$\begin{aligned}
\{\text{sub}\}_0^{\mathbf{p}_x, \mathbf{p}_y \mapsto \mathbf{q}} &= \text{def } T = \mathbf{p}_y \cdot \mathbf{c} \rightarrow \mathbf{r}_5; \mathbf{r}_1 \cdot \mathbf{c} \rightarrow \mathbf{r}_6; \\
&\text{def } R = \text{if } (\mathbf{r}_8 \stackrel{\leftarrow}{=} \mathbf{r}_5) \\
&\quad \text{then } \mathbf{r}_7 \cdot \mathbf{c} \rightarrow \mathbf{r}_3; \mathbf{p}_x \cdot \mathbf{c} \rightarrow \mathbf{r}_4; \\
&\quad \text{def } S = \mathbf{r}_{11} \cdot \mathbf{c} \rightarrow \mathbf{r}_0; \mathbf{r}_1 \cdot \varepsilon \rightarrow \mathbf{r}_2; \\
&\quad \quad \text{if } (\mathbf{r}_2 \stackrel{\leftarrow}{=} \mathbf{r}_0) \text{ then } (\mathbf{r}_1 \cdot \mathbf{c} \rightarrow \mathbf{q}; \mathbf{0}) \text{ else } (\mathbf{r}_1 \cdot \mathbf{c} \rightarrow \mathbf{r}_2; \mathbf{r}_2 \cdot (\mathbf{s} \cdot \mathbf{c}) \rightarrow \mathbf{r}_1; T) \\
&\quad \quad \text{in if } (\mathbf{r}_3 \stackrel{\leftarrow}{=} \mathbf{r}_4) \text{ then } (\mathbf{r}_3 \cdot \varepsilon \rightarrow \mathbf{r}_{11}; S) \text{ else } (\mathbf{r}_3 \cdot (\mathbf{s} \cdot \mathbf{c}) \rightarrow \mathbf{r}_{11}; S) \\
&\quad \text{else } \mathbf{r}_7 \cdot \mathbf{c} \rightarrow \mathbf{r}_{10}; \mathbf{r}_{10} \cdot (\mathbf{s} \cdot \mathbf{c}) \rightarrow \mathbf{r}_9; \\
&\quad \quad \mathbf{r}_9 \cdot \mathbf{c} \rightarrow \mathbf{r}_7; \mathbf{r}_8 \cdot \mathbf{c} \rightarrow \mathbf{r}_9; \mathbf{r}_9 \cdot (\mathbf{s} \cdot \mathbf{c}) \rightarrow \mathbf{r}_8; R \\
&\quad \text{in } \mathbf{r}_6 \cdot \mathbf{c} \rightarrow \mathbf{r}_7; \mathbf{r}_9 \cdot \varepsilon \rightarrow \mathbf{r}_8; R \\
&\quad \text{in } \mathbf{r}_2 \cdot \varepsilon \rightarrow \mathbf{r}_1; T
\end{aligned}$$

Due to the way sequential composition works, the structure of the definition of `sub` is no longer clear in this fully unfolded encoding.

4.3. Soundness and Main Results

By induction we now show that the construction presented above is sound. In the proof, we use partial specifications of states. For example, $C, \{ \mathbf{p} \mapsto v \} \rightarrow C', \{ \mathbf{q} \mapsto w \}$ denotes that execution of C from *any* state where \mathbf{p} contains value v will yield C' in *some* state where \mathbf{q} contains value w .

Theorem 5 (Turing completeness of MC). *If $f : \mathbb{N}^n \rightarrow \mathbb{N}$ and $f \in \mathcal{R}$, then, for every k , $\{\!|f|\!\}_k^{\tilde{\mathbf{p}} \mapsto \mathbf{q}}$ implements f with input processes $\tilde{\mathbf{p}} = \mathbf{p}_1, \dots, \mathbf{p}_n$ and output process \mathbf{q} .*

Proof. The proof is by induction on the definition of the set of partial recursive functions. We use a stronger induction hypothesis – namely, that if $\sigma(\mathbf{p}_i) = \lceil x_i \rceil$ and $f(\tilde{x})$ is defined, then $\{\!|f|\!\}_k^{\tilde{\mathbf{p}} \mapsto \mathbf{q}}, \sigma \rightarrow^* \sigma'$ where $\sigma'(\mathbf{p}_i) = \lceil x_i \rceil$ and $\sigma'(\mathbf{q}) = \lceil f(\tilde{x}) \rceil$. The extra assumption that the input values are not changed during execution is essential for the inductive step. In the case where $f(\tilde{x})$ is not defined, we assume as before that $\{\!|f|\!\}_k^{\tilde{\mathbf{p}} \mapsto \mathbf{q}}, \sigma \not\rightarrow^* \mathbf{0}$.

1. For each base case, it is straightforward to compute the sequence of reductions from the rules and the definition of the corresponding choreography. We exemplify this with successor.

$$\{\!|S|\!\}_\ell^{\mathbf{p} \mapsto \mathbf{q}} : \mathbf{p} \cdot (\mathbf{s} \cdot \mathbf{c}) \rightarrow \mathbf{q}, \{ \mathbf{p} \mapsto \lceil x \rceil \} \rightarrow \mathbf{0}, \left\{ \begin{array}{l} \mathbf{p} \mapsto \lceil x \rceil \\ \mathbf{q} \mapsto \lceil x + 1 \rceil \end{array} \right\}$$

2. Let $h = C(f, g_1, \dots, g_k) : \mathbb{N}^n \rightarrow \mathbb{N}$. The result follows directly from the induction hypothesis and Lemma 1.

3. Let $h = R(f, g) : \mathbb{N}^{n+1} \rightarrow \mathbb{N}$. By induction hypothesis, choreographies $\{\{f\}\}_{\ell_f}^{p_1, \dots, p_n \mapsto q}$ and $\{\{g\}\}_{\ell_g}^{p_1, \dots, p_{n+2} \mapsto q}$ implement f and g , respectively, for all \tilde{p} , q , ℓ_f and ℓ_g . Again, assume first that $h(x_0, \tilde{x})$ is defined. Then:

$$\begin{aligned} \{\{h\}\}_{\ell}^{p_0, \tilde{p} \mapsto q} &: \text{def } T = (\dots) \text{ in } \{\{f\}\}_{\ell_f}^{\tilde{p} \mapsto q'} \circledast r_t. \varepsilon \rightarrow r_c; T, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil f(\tilde{x}) \rceil \end{array} \right\} \\ &\xrightarrow{IH} \text{def } T = (\dots) \text{ in } r_t. \varepsilon \rightarrow r_c; T, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil f(\tilde{x}) \rceil \end{array} \right\} \\ &\rightarrow \text{def } T = (\dots) \text{ in } T, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil h(0, \tilde{x}) \rceil \\ r_c \mapsto \lceil 0 \rceil \end{array} \right\} \end{aligned}$$

We now prove that

$$\text{def } T = (\dots) \text{ in } T, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil h(k, \tilde{x}) \rceil \\ r_c \mapsto \lceil k \rceil \end{array} \right\} \rightarrow^* \text{def } T = (\dots) \text{ in } T, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil h(k+1, \tilde{x}) \rceil \\ r_c \mapsto \lceil k+1 \rceil \end{array} \right\}$$

for all $k < x_0$. We only need to unfold T once, so we omit the $\text{def } T = (\dots) \text{ in}$ wrapper in the next reduction sequence.

Since $k < x_0$, the definition of T reduces to the else branch:

$$\begin{aligned} T &\rightarrow^* \{\{g\}\}_{\ell_g}^{r_c, q', \tilde{p} \mapsto r_t} \circledast r_t. c \rightarrow q'; r_c. c \rightarrow r_t; r_t. (s \cdot c) \rightarrow r_c; T, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil h(k, \tilde{x}) \rceil \\ r_c \mapsto \lceil k \rceil \end{array} \right\} \\ &\xrightarrow{IH} r_t. c \rightarrow q'; r_c. c \rightarrow r_t; r_t. (s \cdot c) \rightarrow r_c; T, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil h(k, \tilde{x}) \rceil \\ r_c \mapsto \lceil k \rceil \\ r_t \mapsto \lceil g(k, h(k, \tilde{x}), \tilde{x}) \rceil \end{array} \right\} \\ &\rightarrow^* T, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil h(k+1, \tilde{x}) \rceil \\ r_c \mapsto \lceil k+1 \rceil \\ r_t \mapsto \lceil k \rceil \end{array} \right\} \end{aligned}$$

which establishes the thesis, ignoring the value in r_t .

By induction on x_0 we obtain that

$$\begin{aligned} \{\{h\}\}_{\ell}^{p_0, \tilde{p} \mapsto q}, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \end{array} \right\} &\rightarrow^* \text{def } T = (\dots) \text{ in } T, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil h(x_0, \tilde{x}) \rceil \\ r_c \mapsto \lceil x_0 \rceil \end{array} \right\} \\ &\xrightarrow{(1)} \text{def } T = (\dots) \text{ in } q'. c \rightarrow q; \mathbf{0}, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil h(x_0, \tilde{x}) \rceil \\ r_c \mapsto \lceil x_0 \rceil \end{array} \right\} \\ &\rightarrow^* \text{def } T = (\dots) \text{ in } \mathbf{0}, \left\{ \begin{array}{l} p_i \mapsto \lceil x_i \rceil \\ q' \mapsto \lceil h(x_0, \tilde{x}) \rceil \\ r_c \mapsto \lceil x_0 \rceil \\ q \mapsto \lceil h(x_0, \tilde{x}) \rceil \end{array} \right\} \end{aligned}$$

and the last process is equivalent to $\mathbf{0}$. In (1) we used the fact that the contents of r_c and p_0 are both equal to $\lceil x_0 \rceil$.

If $h(x_0, \tilde{x})$ is not defined, there are two possible cases. If $f(\tilde{x})$ is not defined, then $\{\!\{f\}\!\}_{\ell_f}^{\tilde{p} \rightarrow \mathbf{q}'}$ diverges from any state where each \mathbf{p}_i contains $\lceil x_i \rceil$, whence so does $\{\!\{h\}\!\}_{\ell}^{p_0, \tilde{p} \rightarrow \mathbf{q}'}$ by Lemma 1 and rule C|Ctx. If $g(k, h(k, \tilde{x}), \tilde{x})$ is undefined for some $k < x_0$, then divergence is likewise obtained from the fact that $\{\!\{g\}\!\}_{\ell_g}^{r_c, \mathbf{q}', \tilde{p}}$ diverges from any state where r_c contains $\lceil k \rceil$, \mathbf{q}' contains $\lceil h(k, \tilde{x}) \rceil$, and \mathbf{p}_i contains $\lceil x_i \rceil$.

4. The case where $h = M(f) : \mathbb{N}^n \rightarrow \mathbb{N}$ is very similar, the auxiliary result now stating that

$$\text{def } T = (\dots) \text{ in } T, \left\{ \begin{array}{l} \mathbf{p}_i \mapsto \lceil x_i \rceil \\ r_c \mapsto \lceil k \rceil \end{array} \right\} \rightarrow^* \text{def } T = (\dots) \text{ in } T, \left\{ \begin{array}{l} \mathbf{p}_i \mapsto \lceil x_i \rceil \\ r_c \mapsto \lceil k + 1 \rceil \end{array} \right\}$$

as long as $f(\tilde{x}, k)$ is defined and different from 0.

The only new aspect is that non-termination may arise from the fact that $f(\tilde{x}, k)$ is defined and non-zero for every $k \in \mathbb{N}$, in which case we get an infinite reduction sequence

$$\begin{aligned} \{\!\{h\}\!\}_{\ell}^{\tilde{p} \rightarrow \mathbf{q}}, \left\{ \begin{array}{l} \mathbf{p}_i \mapsto \lceil x_i \rceil \\ r_c \mapsto \lceil 0 \rceil \end{array} \right\} &\rightarrow^* \text{def } T = (\dots) \text{ in } T, \left\{ \begin{array}{l} \mathbf{p}_i \mapsto \lceil x_i \rceil \\ r_c \mapsto \lceil 0 \rceil \end{array} \right\} \\ &\rightarrow^* \text{def } T = (\dots) \text{ in } T, \left\{ \begin{array}{l} \mathbf{p}_i \mapsto \lceil x_i \rceil \\ r_c \mapsto \lceil n \rceil \end{array} \right\} \\ &\rightarrow^* \dots \end{aligned}$$

□

Since MC is a fragment of CC, this result trivially implies Turing completeness of CC. Let $\text{SP}^{\text{CC}} = \{\llbracket C, \sigma \rrbracket \mid \llbracket C, \sigma \rrbracket \text{ is defined}\}$ be the set of the projections of all projectable choreographies in CC. By Theorem 4, all terms in SP^{CC} are deadlock-free. By Lemma 2, for every function f we can amend $\{\!\{f\}\!\}$ to an equivalent projectable choreography. Then SP^{CC} is Turing complete by Theorems 3 and 5.

Corollary 1 (Turing completeness of SP^{CC}). *Every partial recursive function is implementable in SP^{CC} .*

Proof. Let $f \in \mathcal{R}$. By Theorem 5, $C = \{\!\{f\}\!\}^{\tilde{p} \rightarrow \mathbf{q}}$ for any suitable \tilde{p} and \mathbf{q} implements f . By Lemma 2, $\text{Amend}(C)$ is projectable and operationally equivalent to C . Hence, by Theorem 3, $\llbracket \text{Amend}(C), \sigma \rrbracket$ is a term in SP that correctly implements f . □

We finish this section by showing how to optimize our encoding and obtain parallel process implementations of independent computations. If h is defined by composition from f and g_1, \dots, g_k , then in principle the computation of the g_i s could be completely parallelised. However, $\{\!\{\}\!\}$ does not fully achieve this, as $\{\!\{g_1\}\!\}, \dots, \{\!\{g_k\}\!\}$ share the processes containing the input. We define a modified variant $\{\!\{\}\!\}$ of $\{\!\{\}\!\}$ where, for $h = C(f, g_1, \dots, g_k)$, $\{\!\{h\}\!\}_{\ell}^{\tilde{p} \rightarrow \mathbf{q}}$ is

$$\left\{ \mathbf{p}_j \cdot \mathbf{c} \rightarrow \mathbf{p}_j^i \right\}_{1 \leq i \leq k, 1 \leq j \leq n} \circ \left\{ \{\!\{g_i\}\!\}_{\ell_i}^{\tilde{p}^i \rightarrow \mathbf{r}'_i} \right\}_{i=1}^k \circ \{\!\{f\}\!\}_{\ell_{k+1}}^{\mathbf{r}'_1, \dots, \mathbf{r}'_k \rightarrow \mathbf{q}}$$

with a suitably adapted label function ℓ . Now Theorem 2 applies, yielding:

Theorem 6. *Let $h = C(f, g_1, \dots, g_k)$. For all \tilde{p} and q , if $h(\tilde{x})$ is defined and σ is such that $\sigma(p_i) = \ulcorner x_i \urcorner$, then all the $\{\{g_i\}_{\ell_i}^{\tilde{p}^i \mapsto r'_i}\}$ run in parallel in $\{\{h\}^{\tilde{p} \mapsto q}\}$.*

Proof. By induction on k . For $k = 1$ there is nothing to prove. By definition of $\{\{\}, \{g_{k+1}\}_{\ell_{k+1}}^{\tilde{p}^{k+1} \mapsto r'_{k+1}}\}$ and $\{\{g_i\}_{\ell_i}^{\tilde{p}^i \mapsto r'_i}\}$ do not share any process names for $i \leq k$, hence Theorem 2 implies that $\{\{g_{k+1}\}_{\ell_{k+1}}^{\tilde{p}^{k+1} \mapsto r'_{k+1}}\}$ and $\left\{\{g_i\}_{\ell_i}^{\tilde{p}^i \mapsto r'_i}\right\}_{i=1}^k$ run in parallel in $\left\{\{g_i\}_{\ell_i}^{\tilde{p}^i \mapsto r'_i}\right\}_{i=1}^{k+1}$, and hence in $\{\{h\}^{\tilde{p} \mapsto q}\}$. \square

This parallelism is preserved by EPP, through Theorem 3.

5. Removing Selections

As we saw earlier, MC is the fragment of CC that does not contain selections, and MC choreographies can be amended into projectable CC choreographies. We now show that selections are not necessary to ensure projectability: we can encode those selections introduced by amendment using only conditionals and extra communications.

5.1. Selections as Value Communications

Selections are used pervasively for projectability in previous choreography languages, so the fact that they are technically unnecessary is both interesting and somewhat unexpected. This construction increases the size of the choreography exponentially, but both the number of processes and the size of the endpoint projections grow only by a linear factor.

We motivate our construction with an example.

Example 8. *Consider a choreography where p makes a choice depending on the value stored by q , and then r needs to be notified of the result (because, e.g., it is involved in further communications in one or both of the branches). As an example, we take C to be the choreography $\text{if } p \stackrel{\leq}{=} q \text{ then } p \rightarrow r[L]; C_1 \text{ else } p \rightarrow r[R]; C_2$, where r has different behaviours in C_1 and C_2 .*

In order to eliminate the label selections, r must be able to perform a conditional that is guaranteed to choose the same branch as taken by p . With this in mind, we introduce an auxiliary process p^ and add communications from p to p^* of ε (then branch) or $s \cdot c$ (else branch). Then r can recover this information by first setting its contents to ε and then comparing them with p^* ; this requires another auxiliary process r^* to store r 's value in the meantime. Furthermore, even though we know at a global level what the result of the comparison will be, the EPP (in particular, merging) demands that we consider both branches in both cases. We therefore rewrite C as follows.*

$$\begin{aligned} &\text{if } p \stackrel{\leq}{=} q \text{ then } \left(p.\varepsilon \rightarrow p^*; r.c \rightarrow r^*; r^*.\varepsilon \rightarrow r; \text{if } r \stackrel{\leq}{=} p^* \text{ then } (r^*.c \rightarrow r; C_1) \text{ else } (r^*.c \rightarrow r; C_2) \right) \\ &\quad \text{else } \left(p.s \cdot c \rightarrow p^*; r.c \rightarrow r^*; r^*.\varepsilon \rightarrow r; \text{if } r \stackrel{\leq}{=} p^* \text{ then } (r^*.c \rightarrow r; C_1) \text{ else } (r^*.c \rightarrow r; C_2) \right) \end{aligned}$$

$$\llbracket C \rrbracket^+ = \mathbf{p}.\varepsilon \rightarrow \mathbf{z}; \llbracket C \rrbracket$$

$$\begin{aligned} \llbracket \mathbf{0} \rrbracket &= \mathbf{0} & \llbracket \mathbf{p}.e \rightarrow \mathbf{q}; C \rrbracket &= \mathbf{p}.e \rightarrow \mathbf{q}; \llbracket C \rrbracket & \llbracket \text{def } X = C_2 \text{ in } C_1 \rrbracket &= \text{def } X = \llbracket C_2 \rrbracket \text{ in } \llbracket C_1 \rrbracket \\ \llbracket X \rrbracket &= X & \llbracket \text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } C_1 \text{ else } C_2 \rrbracket &= \text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } \llbracket C_1, C_2 \rrbracket_1 \text{ else } \llbracket C_1, C_2 \rrbracket_2 \end{aligned}$$

$$\begin{aligned} \llbracket C_1, C_2 \rrbracket &= \langle \llbracket C_1 \rrbracket, \llbracket C_2 \rrbracket \rangle \text{ if } C_1 \text{ and } C_2 \text{ do not begin with a selection} \\ \llbracket \mathbf{p} \rightarrow \mathbf{q}[\mathbf{L}]; C_1, \mathbf{p} \rightarrow \mathbf{q}[\mathbf{R}]; C_2 \rrbracket &= \\ &\left\langle \mathbf{q}.c \rightarrow \mathbf{q}^\bullet; \mathbf{p}.\varepsilon \rightarrow \mathbf{q}; \text{if } \mathbf{q} \stackrel{\leq}{=} \mathbf{z} \text{ then } \mathbf{q}^\bullet.c \rightarrow \mathbf{q}; \llbracket C_1, C_2 \rrbracket_1 \text{ else } \mathbf{q}^\bullet.c \rightarrow \mathbf{q}; \llbracket C_1, C_2 \rrbracket_2, \right. \\ &\quad \left. \mathbf{q}.c \rightarrow \mathbf{q}^\bullet; \mathbf{p}.sc \rightarrow \mathbf{q}; \text{if } \mathbf{q} \stackrel{\leq}{=} \mathbf{z} \text{ then } \mathbf{q}^\bullet.c \rightarrow \mathbf{q}; \llbracket C_1, C_2 \rrbracket_1 \text{ else } \mathbf{q}^\bullet.c \rightarrow \mathbf{q}; \llbracket C_1, C_2 \rrbracket_2 \right\rangle \end{aligned}$$

Figure 10: Elimination of selections from amended choreographies.

Observe that the behaviour of the processes not performing conditionals (\mathbf{p}^* and \mathbf{r}^*) is the same in all four branches, while \mathbf{p} and \mathbf{r} have two possible behaviours that are independent of each other's choices. This guarantees that merging will work for all projections. \square

Recall that the definition of amendment guarantees that selections only occur in branches of conditionals, and that they are always paired and in the same order. These properties are essential to our construction. The fragment of CC obtained by amending choreographies in MC can be inductively generated by

$$C ::= \mathbf{p}.e \rightarrow \mathbf{q}; C \mid \text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } S(\mathbf{p}, \tilde{r}, \mathbf{L}, C_1) \text{ else } S(\mathbf{p}, \tilde{r}, \mathbf{R}, C_2) \mid \text{def } X = C_2 \text{ in } C_1 \mid X \mid \mathbf{0}$$

where $S(\mathbf{p}, \tilde{r}, \ell, C)$ prepends selections of label ℓ from \mathbf{p} to all processes in the list \tilde{r} . Formally, S is defined as

$$S(\mathbf{p}, \emptyset, \ell, C) = C \quad S(\mathbf{p}, \mathbf{r} :: \tilde{r}, \ell, C) = \mathbf{p} \rightarrow \mathbf{r}[\ell]; S(\mathbf{p}, \tilde{r}, \ell, C)$$

Definition 10 (Selection elimination). *Let C be a choreography obtained by amending a choreography in MC. The encoding $\llbracket C \rrbracket^+$ of C in MC uses processes $\mathbf{p}, \mathbf{p}^\bullet$ for each $\mathbf{p} \in \text{pn}(C)$, plus a special process \mathbf{z} , and is defined in Figure 10.*

The auxiliary function $\llbracket C_1, C_2 \rrbracket$ is used to eliminate selections in conditionals, by simultaneously traversing both branches. As usual, we write $\llbracket C_1, C_2 \rrbracket_1$ to denote the first component of the resulting pair, and likewise for the second component.

This definition significantly exploits the structure of amended choreographies, where selections are always paired at the top of the two branches of conditionals. It follows from it that $|\text{pn}(\llbracket C \rrbracket^+)| = 2|\text{pn}(C)| + 1$ and that $|\llbracket C \rrbracket^+| \leq 2^{|C|}$. However, the EPP from MC to MP collapses all branches of conditionals, hence $\llbracket \llbracket C \rrbracket^+ \rrbracket_{\mathbf{q}^\bullet} \leq \llbracket \llbracket C \rrbracket^+ \rrbracket_{\mathbf{q}} \leq 3\llbracket C \rrbracket_{\mathbf{q}}$ for every $\mathbf{q} \in \text{pn}(C)$.

Theorem 7 (Selection elimination). *For every choreography $C \in MC$, $\llbracket (\text{Amend}(C)) \rrbracket$ is defined.*

For convenience, we split the proof of this result in several lemmas.

Lemma 3. *If $q \in \tilde{r}$, then $\llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_1 \llbracket q \rrbracket = \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_2 \llbracket q \rrbracket$.*

Proof. By induction on the length of \tilde{r} . If $\tilde{r} = \emptyset$, then the result is vacuously true. If \tilde{r} does not start with q , then the result follows trivially from the induction hypothesis. So consider the case where $\tilde{r} = q :: \tilde{r}'$. In this case, $\llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket$ unfolds to

$$\begin{aligned} & \langle q.c \rightarrow q^\bullet; p.\varepsilon \rightarrow q; \text{if } q \stackrel{\leq}{=} z \text{ then } q^\bullet.c \rightarrow q; \llbracket (S(p, \tilde{r}', L, C_1), S(p, \tilde{r}', R, C_2)) \rrbracket_1 \\ & \quad \text{else } q^\bullet.c \rightarrow q; \llbracket (S(p, \tilde{r}', L, C_1), S(p, \tilde{r}', R, C_2)) \rrbracket_2, \\ & q.c \rightarrow q^\bullet; p.sc \rightarrow q; \text{if } q \stackrel{\leq}{=} z \text{ then } q^\bullet.c \rightarrow q; \llbracket (S(p, \tilde{r}', L, C_1), S(p, \tilde{r}', R, C_2)) \rrbracket_1 \\ & \quad \text{else } q^\bullet.c \rightarrow q; \llbracket (S(p, \tilde{r}', L, C_1), S(p, \tilde{r}', R, C_2)) \rrbracket_2 \end{aligned}$$

and the endpoint projections of both choreographies for q become

$$\begin{aligned} & q^\bullet \langle c \rangle; p?; \text{if } c \stackrel{\leq}{=} z \text{ then } q^\bullet?; \llbracket (S(p, \tilde{r}', L, C_1), S(p, \tilde{r}', R, C_2)) \rrbracket_1 \llbracket q \rrbracket \\ & \quad \text{else } q^\bullet?; \llbracket (S(p, \tilde{r}', L, C_1), S(p, \tilde{r}', R, C_2)) \rrbracket_2 \llbracket q \rrbracket \end{aligned}$$

which are defined and identical. \square

Lemma 4. *If $\llbracket (C_1) \rrbracket \llbracket q \rrbracket = \llbracket (C_2) \rrbracket \llbracket q \rrbracket$ and $p \neq q \notin \tilde{r}$, then $\llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_1 \llbracket q \rrbracket = \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_2 \llbracket q \rrbracket$.*

Proof. By induction on the length of \tilde{r} . If $\tilde{r} = \emptyset$, then the result reduces to the hypothesis. Otherwise, $\llbracket (S(p, r :: \tilde{r}, L, C_1), S(p, r :: \tilde{r}, R, C_2)) \rrbracket$ unfolds to

$$\begin{aligned} & \langle r.c \rightarrow r^\bullet; p.\varepsilon \rightarrow r; \text{if } r \stackrel{\leq}{=} z \text{ then } r^\bullet.c \rightarrow r; \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_1 \\ & \quad \text{else } r^\bullet.c \rightarrow r; \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_2, \\ & r.c \rightarrow r^\bullet; p.sc \rightarrow r; \text{if } r \stackrel{\leq}{=} z \text{ then } r^\bullet.c \rightarrow r; \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_1 \\ & \quad \text{else } r^\bullet.c \rightarrow r; \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_2 \end{aligned}$$

and since $q \neq r$ there are three cases to consider.

- q is r^\bullet : then both endpoint projections become

$$r?; (r! \langle c \rangle); \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_1 \llbracket q^\bullet \rrbracket \sqcup (r! \langle c \rangle); \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_2 \llbracket q^\bullet \rrbracket$$

and by induction hypothesis the two processes being merged are identical, so the result is defined.

- q is z : then both endpoint projections become

$$r! \langle c \rangle; \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_1 \llbracket z \rrbracket \sqcup \llbracket (S(p, \tilde{r}, L, C_1), S(p, \tilde{r}, R, C_2)) \rrbracket_2 \llbracket z \rrbracket$$

and again by induction hypothesis the two processes being merged are identical, so the result is defined.

- q is another process: then both endpoint projections become simply

$$\llbracket (S(\mathbf{p}, \tilde{r}, L, C_1), S(\mathbf{p}, \tilde{r}, R, C_2)) \rrbracket_1 \sqcup \llbracket (S(\mathbf{p}, \tilde{r}, L, C_1), S(\mathbf{p}, \tilde{r}, R, C_2)) \rrbracket_2 \rrbracket_q$$

whence the induction hypothesis guarantees again that the two processes being merged are identical, so the result is defined. \square

Lemma 5. *For every choreography C in MC and every process r , $\llbracket (\text{Amend}(C)) \rrbracket_r$ is defined.*

Proof. By structural induction on $\text{Amend}(C)$. The only non-trivial case is that where $\text{Amend}(C)$ is $\text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } S(\mathbf{p}, \tilde{r}, L, C_1) \text{ else } S(\mathbf{p}, \tilde{r}, R, C_2)$, where we need to consider the possible cases for r . If $r = \mathbf{p}$, then the induction hypothesis establishes the thesis with induction over \tilde{r} . If $r \in \tilde{r}$, then Lemma 3 guarantees that both branches of the conditional will be equal, hence the endpoint projection is again defined. Finally, if $r \notin \tilde{r}$, then by definition of amendment $\llbracket \text{Amend}(C_1) \rrbracket_r = \llbracket \text{Amend}(C_2) \rrbracket_r$, whence Lemma 4 applies and establishes the thesis as in the previous case. \square

Proof (Theorem 7). Straightforward consequence of Lemma 5. \square

The operational semantics of C and $\llbracket \text{Amend}(C) \rrbracket$ are related by the following results, which are straightforward to prove by structural induction.

Lemma 6. *Choreographies C and $\llbracket \text{Amend}(C) \rrbracket^+$ are equivalent wrt $\text{pn}(C)$.*

Lemma 7. *If $C, \sigma \rightarrow C', \sigma'$ and σ^+ is such that $\sigma^+(\mathbf{p}) = \sigma(\mathbf{p})$ for $\mathbf{p} \in \text{pn}(C)$ and $\sigma^+(z) = \varepsilon$, then $\llbracket \text{Amend}(C) \rrbracket, \sigma^+ \rightarrow^* \llbracket C' \rrbracket, \sigma'^+$ for some σ'^+ similarly related to σ' . Furthermore, the latter reduction consists of only one step except for the case when the former uses rule $C|Cond$.*

Conversely, if $\llbracket \text{Amend}(C) \rrbracket, \sigma^+ \rightarrow C', \sigma'$, then $C, \sigma \rightarrow C'', \sigma''$ where $C', \sigma' \rightarrow^ \llbracket \text{Amend}(C'') \rrbracket, \sigma''^+$. Furthermore, the latter reduction is non-empty only in the case when the former uses rule $C|Cond$.*

Corollary 2. *With the notation of the previous lemma, if $C, \sigma \rightarrow^* C', \sigma'$, then $\llbracket \text{Amend}(C) \rrbracket^+, \sigma^+ \rightarrow^* \llbracket \text{Amend}(C') \rrbracket, \sigma'^+$.*

As a consequence, the set $\text{SP}^{\text{MC}} = \{\llbracket C, \sigma \rrbracket \mid \llbracket C, \sigma \rrbracket \text{ is defined}\}$ of projections of minimal choreographies is also Turing complete.

Corollary 3 (Turing completeness of SP^{MC}). *Every partial recursive function is implementable in SP^{MC} .*

5.2. Discussion

Figure 11 displays the connections among the different calculi that we studied in this work. We now discuss some consequences of these connections.

The results in this section show that label selection is not a necessary primitive in a choreography calculus, and thus we could take MC (rather than CC) as our core choreography language. Furthermore, the construction in § 4.2 shows

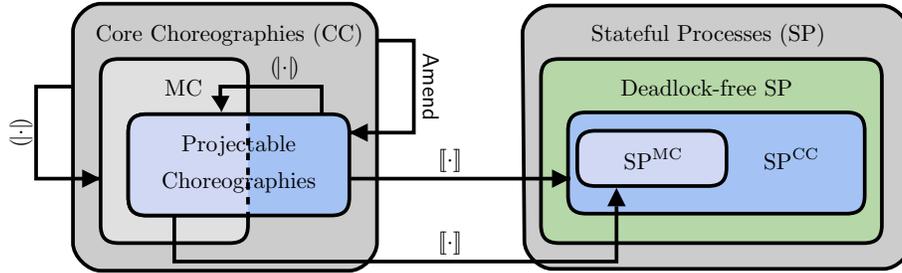


Figure 11: Summary of the different mappings between calculi.

that selections are not needed for implementing computable functions in CC; they are used only for obtaining projectable choreographies, via amendment.

There is however a strong argument for including label selections in a core choreography calculus. The advantages of eliminating selections are a simpler choreography language, a simpler definition of EPP (without merging), and a simpler process language (without selection and branching). The main drawback is that eliminating a selection needed for projectability makes the choreography exponentially larger and requires the addition of extra processes and communications; this significantly changes the structure of the choreography, potentially making it unreadable. Selections are also present in virtually all choreography models (Coppo et al., 2016; Carbone et al., 2012; Carbone and Montesi, 2013; Honda et al., 2008; Dalla Preda et al., 2015; Qiu et al., 2007), therefore we believe that a core model such as CC should have them (in addition to the drawback we mentioned).

The combination of our results on amendment and selection elimination suggests the viability of a particular implementation strategy for choreographic programming. Programmers could write choreographies without label selections, ignoring how information about control flow is propagated between processes. Then, our results could be used to translate these choreographies to process implementations in a simple language that does not include label communications (like SP^{MC}). This would simplify the target language, since it would not require primitives for selection and branching. The exponential growth of the intermediate choreography representation can be bypassed by using shared data structures for the syntax tree, since the generated choreographies contain a lot of duplicate terms.

However, such a methodology makes use of amendment, and relying exclusively on amendment removes an important ability provided in CC and all other standard choreography calculi: deciding at which point of execution selections should be performed. In more expressive languages than CC, processes can perform complex internal computations (Cruz-Filipe and Montesi, 2017c). For example, assume that p had to assign tasks to other two processes r and s based on a condition. In one case, r would run a slow task and s a fast one; otherwise, r would run a fast task and s a slow one. In this case, p should begin by sending a selection to the process with the slow task and then by sending it the

necessary data for its computation, before it sends the selection to the process with the fast task. (Note that it is only amendment that causes this issue, since our selection elimination procedure preserves the control flow specified by a choreography with selections.)

6. Minimality in Choreographies

We now discuss our choice of primitives for CC, showing that it is indeed a minimal core language for choreographic programming. We first show that if we remove or simplify any primitive from MC we are no longer able to compute all partial recursive functions using projectable choreographies. Since label selection can be encoded in MC, we also discuss why it should be included in a core language. Then we discuss the implications of our results for other choreography languages.

6.1. Minimality of MC

We proceed by analysing each primitive of MC. Recall that Turing completeness of MC is a pre-requisite for the Turing completeness of choreography projections. In most cases, simplifying MC yields a decidable termination problem (thus breaking Turing completeness). We start with the easiest terms.

Lemma 8. *Let C be a choreography with no exit points. Then C does not terminate.*

Proof. Straightforward by structural induction on C . □

Lemma 9. *Let C be a choreography with no communications. If C implements a function $f : \mathbb{N}^n \rightarrow \mathbb{N}$, then, for all inputs $\vec{x} \in \mathbb{N}^n$, either $f(\vec{x}) = x_i$ for some i or $f(\vec{x})$ is undefined.*

Proof. By the semantics of MC, only communication actions can change the state σ , hence structural induction on C shows that $C\sigma \not\rightarrow C'\sigma'$ with $\sigma' \neq \sigma$. The thesis is a consequence of the definition of function implementation. □

Observe further that the syntax of expressions is trivially minimal: ε (zero) is the only terminal, removing \mathbf{c} makes termination decidable (since values become statically defined), and likewise for \mathbf{s} (since no new values can be computed).

Lemma 10. *Let C be a choreography with no recursive definitions. Then C always terminates.*

Proof. Without recursive definitions, rule C|Unfold is never applicable, hence execution C always reduces the size of the choreography. □

Again we observe that recursive definitions are already severely restricted: MC supports only tail recursion and definitions are not parameterised.

Removing conditionals naturally also breaks Turing completeness.

Theorem 8. *Let C be a choreography with no conditionals. Then termination of C is decidable and independent of the initial state.*

Proof. The second part is straightforward, since rule C|Cond is the only rule whose conclusion depends on the state.

For the first part, we reduce termination to a decidable graph problem. Define $\mathcal{G}_C = \langle V, E \rangle$ to be the graph whose set of vertices V contains C and $\mathbf{0}$, and is closed under the following rules.

- if $\eta; C \in V$, then $C \in V$;
- if $\text{def } X = C_2 \text{ in } C_1 \in V$, then $C_1 \in V$;
- if $\text{def } X = C_2 \text{ in } \eta; C_1 \in V$, then $\text{def } X = C_2 \text{ in } C_1 \in V$;
- if $\text{def } X = C_2 \text{ in } \eta; X \in V$, then $\text{def } X = C_2 \text{ in } \eta; C_2 \in V$.

This set is finite: all rules add smaller choreographies to V , except the last one, which can only be applied once for each variable in C .

There is an edge between C_1 and C_2 iff $C_1, \sigma \rightarrow C_2, \sigma'$ for some σ, σ' without using rule C|Eta-Eta. This is decidable, as the possibility of a reduction does not depend on the state (as observed above). Also, if there is a reduction from C_1 , then there is always an edge from C_1 in the graph, as swapping communication actions cannot unblock execution.

Then C terminates iff there is a path from C to $\mathbf{0}$, which can be decided in finite time, as \mathcal{G}_C is finite. \square

More interestingly, limiting processes to evaluating only their own local values in conditions makes termination decidable. Intuitively, this is because a process can only hold a value at a time and thus no process can compare its current value to that of another process anymore.

Theorem 9. *If the conditional is replaced by $\text{if } p.c = v \text{ then } C_1 \text{ else } C_2$, where v is a value, and rule C|Cond by*

$$\frac{i = 1 \text{ if } \sigma(p) = v, \quad i = 2 \text{ otherwise}}{\text{if } p.c = v \text{ then } C_1 \text{ else } C_2, \sigma \rightarrow C_i, \sigma},$$

then termination is decidable.

Proof. We first show that termination is decidable for processes of the form $\text{def } X = C_2 \text{ in } X$ and comparison with 0. The proof is by induction on the number of recursive definitions in C_2 .

Consider first the case where C_2 has no recursive definitions, and let P be the set of all process names occurring in C_2 . We define an equivalence relation on states by

$$\sigma \equiv_P \sigma' \text{ iff } (\forall p \in P, \sigma(p) = \varepsilon \text{ iff } \sigma'(p) = \varepsilon).$$

The vertices of the graph are the $2^{|P|}$ equivalence classes of states wrt \equiv_P , plus \top . Note that \equiv_P is compatible with the transition relation excluding rule

C|Eta-Eta: for any choreography C using only process names in P , $\sigma_1 \equiv \sigma_2$ and $C, \sigma_i \rightarrow \sigma'_i$, then $\sigma'_1 \equiv \sigma'_2$.

The edges in the graph are defined as follows. There is an edge from $[\sigma]$ to $[\sigma']$ if $C_2, \sigma \rightarrow X, \sigma'$, and there is an edge from $[\sigma]$ to \top if $C_2, \sigma \rightarrow \mathbf{0}, \sigma'$ or $C_2, \sigma \rightarrow Y, \sigma'$ for some $Y \neq X$. This is constructible, as reductions in C_2 are always finite, and well-defined, as alternative reduction paths always end in the same state.

Since reductions are deterministic and \equiv_P is compatible with reduction, every node has exactly one edge leaving from it, except for \top . Therefore, we can decide if $\text{def } X = C_2 \text{ in } X$ terminates from an initial state σ by simply following the path starting at σ and returning Yes if we reach \top and No if we pass some node twice. This procedure terminates, as the graph is finite.

For the inductive step, proceed as above but add an extra node to the graph, labeled \perp . When constructing the edges in the graph, if C_2 reduces to a variable Y different than X , we split into two cases. If Y is not bound in C_2 , we proceed as in the previous case. If Y is bound, then we apply the induction hypothesis to the choreography $\text{def } Y = C_Y \text{ in } Y$ (where $Y = C_Y$ is the same as in C_2) to decide whether the reduction from Y will terminate; if this is not the case, we add an edge to \perp , otherwise we proceed with the simulation. At the end, we return No in the case that the path followed leads to \perp .

The general case follows, as C has the same behaviour as $\text{def } X = C \text{ in } X$ for some X not occurring in C .

If we allow comparisons with other values, the strategy is the same, but the relation \equiv_P has to be made finer. The key observation is that only a finite number of values can be used in comparisons, so we can identify states if they only differ on processes whose contents are larger than all values used in conditionals. \square

Summarising, simplifying MC in any of the ways described above makes it no longer a representative model of choreographic programming.

6.2. CC as a Core Language

In the Appendix, we formally present an embedding of CC into the choreography model from (Carbone and Montesi, 2013), which we refer to in this work as Channel Choreographies (ChC). ChC is a very rich choreography language designed to be projected to a variant of the session-typed π -calculus (Coppo et al., 2016), which we refer to as Channel Processes (ChP). Communications in ChP are based on channels, instead of process names as in SP. This layer of indirection means that a process performing an I/O action does not know [which other process it is going to communicate with](#), and that there can be race conditions on the usage of channels. ChC comes with a typing discipline for checking that the usage of channels specified in a choreography does not cause errors in the process code generated by EPP. In particular, we show that our embedding always yields well-typed and projectable ChC choreographies.

Channel Choreographies. Our formal translation from CC to ChC shows that many primitives of ChC are not needed to achieve Turing completeness, including: asynchronous communications, creation of sessions and processes, channel mobility, parameterised recursive definitions, arbitrary local computation, unbounded memory cells at processes, and multiparty sessions. While useful in practice, these primitives come at the cost of making the formal treatment of ChC technically involved. In particular, ChC (as well as its implementation Chor) requires a sophisticated type system, linearity analysis, and definition of EPP to ensure correctness of projected processes. These features are not needed in CC.

Other Choreography Languages. The language WS-CDL from W3C (W3C WS-CDL Working Group, 2004) and the formal models inspired by it – e.g., (Carbone et al., 2012) – are very similar to ChC, and a similar translation from CC could be formally developed, with similar implications as above. The same applies to the choreography language developed in (Dalla Preda et al., 2015), which adds higher-order features to choreographies to achieve runtime adaptation. Finally, the language of compositional choreographies presented in (Montesi and Yoshida, 2013) is an extension of ChC, and therefore our translation applies directly. This implies that adding modularity to choreographies does not add any computational power, as expected.

Process Languages. Our embedding of CC in ChC identifies a fragment of ChP, via EPP, that is also Turing complete. This fragment is isomorphic to value-passing CCS (Milner, 1980): since we only have one channel, we can interpret the constructs $k[A]!B\langle e \rangle$ and $k[B]?A(x)$ as sending and receiving over a channel with name kAB . We thus obtain a deadlock-free and Turing complete fragment of value-passing CCS. [Deciding whether a given CCS process lies in this fragment is undecidable by Rice’s Theorem, but it is possible to define a procedure that establishes deadlock-freedom for a large class of such processes. Both results are thoroughly discussed in \(Cruz-Filipe et al., 2017\), where we also present a procedure to extract a choreography that represents a particular network in SP.](#)

Since ChC has also been translated to the Jolie programming language (Gabrielli et al., 2015; Montesi et al., 2014), our reasoning also applies to the latter and, in general, to service-oriented languages based on message correlation. Namely, our results identify a deadlock-free and Turing complete fragment of Jolie.

7. Related Work and Discussion

Choreographies. [The origins of choreographic programming \(Montesi, 2013\) stem from the efforts of using choreographic descriptions for the specification of interactions among web services. In particular, the Web Services Choreography Description Language by the W3C, WS-CDL for short, is a choreography language for describing the observable interaction behaviour of web services from](#)

a global viewpoint (W3C WS-CDL Working Group, 2004). The usefulness of formal methods was recognised early on in the development of WS-CDL, and spawned a successful line of research based on choreographies. This led to the introduction of the notion of EndPoint Projection (EPP), which has been formalised in process calculi (Carbone et al., 2007; Qiu et al., 2007), and also adopted in implementations (Red Hat, 2017).

Already in the early works on choreographies, it was evident that choices at the choreography level (like our conditionals) play an important role for EPP: their wrong programming may lead to projections that behave incorrectly, as we discussed in § 2.3 and § 3.3. The solution of checking whether the other processes could distinguish which branch they should execute based on the labels that they receive – or, equivalently, operation or method names, as in service-oriented or object-oriented programming respectively – was given together with some early proposals of EPP (Carbone et al., 2007; Lanese et al., 2008). The idea of merging was introduced in (Carbone et al., 2007), for a different calculus than CC.

In the first choreography languages (W3C WS-CDL Working Group, 2004; Carbone et al., 2007), the construct for performing a communication includes both a carried value and a label. This corresponds to having both value communication and label selections in a single construct. It is a choice motivated by practice, since that is how invocations work in service-oriented computing (where labels are the operations offered by a service) and object-oriented programming (where labels are the names of the methods offered by an object).

Very soon afterwards, however, a series of choreography models with separate constructs for value communications and label selections (as we have in CC) started emerging. This transition was influenced by the interaction with the research line on session types for process calculi, which have the same distinction (Honda et al., 1998). There are two main reasons for having this distinction: it makes the model more foundational, since each construct is more primitive; and it allows for studying the two primitives separately, which is useful since labels are statically defined whereas communicated values are computed at runtime (for example, the fact that the two labels are distinct can be statically computed is what allows merging to be defined). The interaction between session types and choreographies spawned a prolific research area where choreographies are used as types for protocols in process calculi and concurrent languages (Hüttel et al., 2016; Ancona et al., 2016). The seminal work in this direction is the theory of Multiparty Session Types (Honda et al., 2008, 2016). The theory of session types have been recently found to be in a propositions-as-types correspondence with linear logic (Caires and Pfenning, 2010; Wadler, 2012), where the action of making a choice (corresponding to a label selection) is again distinct. This correspondence extends naturally to choreographies used as types (Carbone et al., 2017).

Thus, selections are important for practical reasons – for example, they model the selection of an operation, or method, from the interface offered by a service, or an object – and for theoretical reasons – they are primitives identified in foundational theories like session types and linear logic. Indeed, all current

implementations of choreography languages have them (Yoshida et al., 2013; Chor, 2017; Dalla Preda et al., 2017).

Our encoding from CC to MC and Corollary 3 show, for the first time, that as far as the computational expressivity of choreography projections is concerned we need only consider value communications. Technically, the communication of a choice can be simulated by communicating appropriate values and then using local conditionals to understand which choice was transmitted by the sender. This is possible because in choreographic programming we have complete control of the local computations performed by each process, differently from works where choreographies are used as types (which abstract from computation).

Which model between CC and MC should we then use when studying choreographies?

If the aim is to determine whether a given choreography language is Turing complete, then MC is the obvious choice. For example, if a language has (an equivalent version of) MC as a fragment, then it is obviously Turing complete. Also, the simple EPP for MC gives a procedure for constructing a process implementation for all possible computations.

Another common aim is to observe the communications enacted by a choreography, either statically (e.g., for verification) or at runtime (e.g., monitoring). Analysing the flow of choice communications is a key element of all behavioural type systems for choreographies to date (Carbone et al., 2012; Carbone and Montesi, 2013; Montesi and Yoshida, 2013; Carbone et al., 2014; Giallorenzo, 2016; Chor, 2017; Red Hat, 2017), which aim at checking that a choreography correctly implements some protocol specifications. For these, CC is a representative model that can be used as foundations to start from. For other kinds of analyses, like information flow in choreographies (Luch Lafuente et al., 2015), selections are not strictly necessary and MC may be the better choice. Similar considerations apply to monitoring, since typically that involves checking whether the actions performed by a choreographed system respect some specifications: if the specifications contain information about explicit choice communications, then selections are necessary (CC), otherwise MC may offer a simpler base model. In general, CC is preferable for all frameworks that have mechanisms where operations (or methods) can be selected out of the interface offered by some component (like a service or an object).

Choreographic Programming and Applications of our Development. Essentially, choreographic programming applies the ideas of choreographies and EPP to synthesise correct-by-construction implementations of concurrent processes. Languages for choreographic programming are typically more complicated than choreography languages for specifications. For example, we have choreographic programming languages for: service-oriented computing (Carbone and Montesi, 2013), including notions such as dynamic networks; adaptable computing (Dalla Preda et al., 2017), including runtime code updates; and cyber-physical systems (López et al., 2016; López and Heussen, 2017), including broadcasts and failures.

Thus, so far, most expressivity results have been investigated for choreographic specifications. For example, we know of a strong characterisation result for multiparty session types: a variant of multiparty session types corresponds to communicating finite state machines (Brand and Zafiropulo, 1983) that respect the property of multiparty compatibility (Deniélou and Yoshida, 2013). By contrast, for choreographies used as concrete implementations (our interest here), this question has barely been scratched before this work: session-typed choreographies with finite traces correspond to proofs in multiplicative-additive linear logic (Carbone et al., 2014). The language in (Carbone et al., 2014) does not include any constructs for programming repetitive behaviour. To the best of our knowledge, MC is the first choreography language to be identified as minimally Turing complete.

Just as it happened for process calculi, the field of choreographic programming is evolving into a workshop of different languages developed for different purposes (Montesi, 2015). In this context, it makes sense to develop new notions by following a minimalistic approach, such that they remain as easy to adopt in other choreography models as possible. The results that we have presented here in extended form have already been useful in this sense. We briefly report on some developments.

In (Cruz-Filipe and Montesi, 2017d), we showed that CC is unable to encode asynchronous communication, and then discussed which primitives to add in order to be able to mimic this behaviour in a synchronous semantics. Thanks to CC, in addition to the technical convenience of working in a setting that is as simple as possible, our study again identifies a minimal set of primitives that must be present in a choreography language to build such an encoding, and the encoding itself can be structurally extended to most choreography languages.

In (Cruz-Filipe and Montesi, 2017b), we formally defined what a “good” asynchronous semantics for choreographies is, and show that it is possible to equip MC with such a semantics, obtaining a version of MC with an asynchronous reduction semantics. Again, our construction is simple yet general and is easy to adopt in more sophisticated choreography languages.

In (Cruz-Filipe et al., 2017), we considered the problems of (i) deciding whether a process implementation can be described by a choreography, and (ii) synthesising such a choreography in the affirmative case; we showed that, for CC and SP, both problems are solvable in exponential time. Since all languages for choreographic programming introduced so far include CC, the technical challenges identified for these problems are present in all such languages, and the identified exponential complexity is the current lower bound for all of them.

We have also used CC as a basis to obtain expressive choreography languages for more practical purposes. In (Cruz-Filipe and Montesi, 2017c), we studied how to extend CC with general sequential composition – a feature that is unavailable in most choreography calculi equipped with recursion – together with other commonly occurring primitives that are useful in practice (process spawning and name passing). The resulting calculus is simple, yet expressive enough to capture different parallel computing algorithms, like parallel versions of gaussian elimination and fast fourier transform (Cruz-Filipe and Montesi,

2016). These two works illustrate how concepts defined for CC (for example, the notion of “running in parallel”, Definition 4) are naturally applicable in more complex settings.

Our results support the claim that there are substantial advantages to gain from first studying choreographic programming in itself, abstracting from features that are specific to other models (like channels in process calculi), and then applying the obtained insights to particular scenarios. We believe that CC is a useful step in this direction, and that it will serve as a stepping stone for the future developments of the paradigm of choreographic programming in general.

Full β -reduction and Nondeterminism. Execution in CC is nondeterministic due to the swapping of communications allowed by the structural precongruence \preceq . This recalls full β -reduction for λ -calculus, where sub-terms can be evaluated whenever possible. However, the two mechanisms are actually different. Consider the choreography $C \triangleq p.c \rightarrow q; q.\varepsilon \rightarrow r; \mathbf{0}$. If CC supported full β -reduction, we should be able to reduce the second communication before the first one, since there is no data dependency between the two. Formally, for some $\sigma: C, \sigma \rightarrow p.c \rightarrow q; \mathbf{0}, \sigma[r \mapsto \varepsilon]$. However, this reduction is disallowed by our semantics: rule C|Eta-Eta cannot be applied because q is present in both communications. This difference is a key feature of choreographies, stemming from their practical origins: controlling sequentiality by establishing causalities using process identifiers is important for the implementation of business processes (W3C WS-CDL Working Group, 2004). For example, imagine that the choreography C models a payment transaction and that the message from q to r is a confirmation that p has sent its credit card information to q ; then, it is a natural requirement that the second communication happens only after the first. Note that we would reach the same conclusions even if we adopted an asynchronous messaging semantics for SP, since the first action by q is a blocking input.

While execution in CC can be nondeterministic, computation results are deterministic as in many other choreography languages (Carbone and Montesi, 2013; Carbone et al., 2014; Montesi and Yoshida, 2013): if a choreography terminates, the result will always be the same regardless of how its execution is scheduled, recalling the Church–Rosser Theorem for the λ -calculus (Church and Rosser, 1936). Nondeterministic computation is not necessary for our results. Nevertheless, it can be easily added to CC. Specifically, we could augment CC with the syntax primitive $C_1 \oplus^p C_2$ and the reduction rule $C_1 \oplus^p C_2 \rightarrow C_i$ for $i = 1, 2$. Extending SP with an internal choice $B_1 \oplus B_2$ and our definition of EPP is straightforward: in SP, we would also allow $B_1 \oplus B_2 \rightarrow B_i$ for $i = 1, 2$, and define $\llbracket C_1 \oplus^p C_2 \rrbracket_r$ to be $\llbracket C_1 \rrbracket_r \oplus \llbracket C_2 \rrbracket_r$ if $r = p$ and $\llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r$ otherwise.

Merging and Amendment. Amendment was first studied by Lanese et al. (2013) for a simple language with finite traces (thus not Turing complete). Our definition is different, since it uses merging for the first time.

We could define our amendment procedure in different ways, e.g., by propagating selections from a process to another as a chain, rather than from one process to all the others. This would not influence our results.

Actors and Asynchrony. Processes in SP communicate by using direct references to each other, recalling actor systems. However, there are notable differences: communications are synchronous and inputs specify the intended sender. The first difference comes from minimality: asynchrony would add possible behaviours to CC, which are unnecessary to establish Turing completeness. We leave an investigation of asynchrony in CC to future work. The second difference arises because CC is a choreography calculus, and communication primitives in choreographies typically express both sender and receiver.

Register Machines. The computational primitives in CC recall those of the Unlimited Register Machine (URM) (Cutland, 1980), but CC and URM differ in two main aspects. First, URM programs contain go-to statements, while CC supports only tail recursion. Second, in the URM there is a single sequential program manipulating the cells, whereas in CC computation is distributed among the various cells (the processes), which operate concurrently.

Simulating the URM is an alternative way to prove Turing completeness of CC. However, our proof using partial recursive functions is more direct and gives an algorithm to implement any function in CC, given its proof of membership in \mathcal{R} . It also yields the natural interpretation of parallelisation stated in Theorem 6. Similarly, we could establish Turing completeness of CC using only a bounded number of processes. However, such constructions encode data using Gödel numbers, which is not in the spirit of our declarative notion of function implementation. They also restrict concurrency, breaking Theorem 6.

Acknowledgements

We thank Hugo Torres Vieira and Gianluigi Zavattaro for their useful comments. Montesi was supported by CRC (Choreographies for Reliable and efficient Communication software), grant no. DFF-4005-00304 from the Danish Council for Independent Research.

References

- Ancona, D., Bono, V., Bravetti, M., Campos, J., Castagna, G., Deniélou, P., Gay, S. J., Gesbert, N., Giachino, E., Hu, R., Johnsen, E. B., Martins, F., Mascardi, V., Montesi, F., Neykova, R., Ng, N., Padovani, L., Vasconcelos, V. T., Yoshida, N., 2016. Behavioral types in programming languages. *Foundations and Trends in Programming Languages* 3 (2-3), 95–230.
- Barendregt, H., 1984. *The Lambda Calculus: Its Syntax and Semantics*, 2nd Edition. North Holland.
- Brand, D., Zafiropulo, P., Apr. 1983. On communicating finite-state machines. *J. ACM* 30 (2), 323–342.
URL <http://doi.acm.org/10.1145/322374.322380>

- Caires, L., Pfenning, F., 2010. Session types as intuitionistic linear propositions. In: CONCUR. Vol. 6269 of LNCS. Springer, pp. 222–236.
- Carbone, M., Honda, K., Yoshida, N., 2007. Structured communication-centred programming for web services. In: Proc. of ESOP. Vol. 4421 of LNCS. Springer-Verlag, pp. 2–17.
- Carbone, M., Honda, K., Yoshida, N., 2012. Structured communication-centered programming for web services. ACM Trans. Program. Lang. Syst. 34 (2), 8.
- Carbone, M., Montesi, F., 2013. Deadlock-freedom-by-design: multiparty asynchronous global programming. In: Giacobazzi, R., Cousot, R. (Eds.), POPL. ACM, pp. 263–274.
- Carbone, M., Montesi, F., Schürmann, C., 2014. Choreographies, logically. In: Baldan, P., Gorla, D. (Eds.), CONCUR. Vol. 8704 of LNCS. Springer, pp. 47–62.
- Carbone, M., Montesi, F., Schürmann, C., Yoshida, N., 2017. Multiparty session types as coherence proofs. Acta Inf. 54 (3), 243–269.
- Chor, 2017. Programming Language. <http://www.chor-lang.org/>.
- Church, A., Rosser, J., 1936. Some properties of conversion. Transactions of the American Mathematical Society 39 (3), 472–482.
URL <http://www.jstor.org/stable/1989762>
- Coppo, M., Dezani-Ciancaglini, M., Yoshida, N., Padovani, L., 2016. Global progress for dynamically interleaved multiparty sessions. Mathematical Structures in Computer Science 26 (2), 238–302.
URL <http://dx.doi.org/10.1017/S0960129514000188>
- Cruz-Filipe, L., Larsen, K. S., Montesi, F., 2017. The paths to choreography extraction. In: Esparza, J., Murawski, A. S. (Eds.), FoSSaCS. Vol. 10203 of LNCS. Springer, pp. 424–440.
- Cruz-Filipe, L., Montesi, F., 2016. Choreographies in practice. In: Albert, E., Lanese, I. (Eds.), FORTE 2016. Vol. 9688 of LNCS. Springer, pp. 114–123.
- Cruz-Filipe, L., Montesi, F., 2017a. A core model for choreographic programming. In: Kouchnarenko, O., Khosravi, R. (Eds.), FACS. Vol. 10231 of LNCS. Springer, pp. 17–35.
- Cruz-Filipe, L., Montesi, F., 2017b. On asynchrony and choreographies. In: Proceedings of ICE. To appear.
- Cruz-Filipe, L., Montesi, F., 2017c. Procedural choreographic programming. In: Bouajjani, A., Silva, A. (Eds.), FORTE 2017. Vol. 10321 of LNCS. Springer, pp. 92–107.

- Cruz-Filipe, L., Montesi, F., 2017d. That’s enough: Asynchrony with standard choreography primitives. In: Caleiro, C., Dionísio, F., Gouveia, P., Mateus, P., Rasga, J. (Eds.), *Logic and Computation: Essays in Honour of Amílcar Sernadas*. Vol. 33 of *Tributes*. College Publications, pp. 125–142.
- Cutland, N., 1980. *Computability: an Introduction to Recursive Function Theory*. Cambridge University Press.
- Dalla Preda, M., Gabbrielli, M., Giallorenzo, S., Lanese, I., Mauro, J., 2015. Dynamic choreographies – safe runtime updates of distributed applications. In: Holvoet, T., Viroli, M. (Eds.), *COORDINATION*. Vol. 9037 of *LNCS*. Springer, pp. 67–82.
URL http://dx.doi.org/10.1007/978-3-319-19282-6_5
- Dalla Preda, M., Gabbrielli, M., Giallorenzo, S., Lanese, I., Mauro, J., 2017. Dynamic choreographies: Theory and implementation. *Logical Methods in Computer Science* 13 (2).
- Deniélou, P.-M., Yoshida, N., 2013. Multiparty compatibility in communicating automata: Characterisation and synthesis of global session types. In: Fomin, F., Freivalds, R., Kwiatkowska, M., Peleg, D. (Eds.), *ICALP (II)*. Vol. 7966 of *LNCS*. Springer, pp. 174–186.
- Gabbrielli, M., Giallorenzo, S., Montesi, F., 2015. Applied choreographies. *CoRR* abs/1510.03637.
URL <http://arxiv.org/abs/1510.03637>
- Giallorenzo, S., 2016. Real-world choreographies. Ph.D. thesis, University of Bologna, Italy.
- Honda, K., Vasconcelos, V., Kubo, M., 1998. Language primitives and type disciplines for structured communication-based programming. In: Hankin, C. (Ed.), *ESOP*. Vol. 1381 of *LNCS*. Springer, pp. 122–138.
- Honda, K., Yoshida, N., Carbone, M., 2008. Multiparty asynchronous session types. In: Necula, G., Wadler, P. (Eds.), *POPL*. ACM, pp. 273–284.
- Honda, K., Yoshida, N., Carbone, M., 2016. Multiparty asynchronous session types. *J. ACM* 63 (1), 9.
URL <http://doi.acm.org/10.1145/2827695>
- Hüttel, H., Lanese, I., Vasconcelos, V. T., Caires, L., Carbone, M., Deniélou, P., Mostrous, D., Padovani, L., Ravara, A., Tuosto, E., Vieira, H. T., Zavattaro, G., 2016. Foundations of session types and behavioural contracts. *ACM Comput. Surv.* 49 (1), 3:1–3:36.
- Kleene, S., 1952. *Introduction to Metamathematics*. North-Holland Publishing Co.

- Kouzapas, D., Yoshida, N., 2013. Globally governed session semantics. In: CONCUR. pp. 395–409.
- Lanese, I., Guidi, C., Montesi, F., Zavattaro, G., 2008. Bridging the gap between interaction- and process-oriented choreographies. In: Cerone, A., Gruner, S. (Eds.), SEFM. IEEE, pp. 323–332.
- Lanese, I., Montesi, F., Zavattaro, G., 2013. Amending choreographies. In: Ravara, A., Silva, J. (Eds.), WWV 2013. Vol. 123 of EPTCS. pp. 34–48.
- Leesatapornwongsa, T., Lukman, J. F., Lu, S., Gunawi, H. S., 2016. TaxDC: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In: ASPLOS. ACM, pp. 517–530.
- Lluch Lafuente, A., Nielson, F., Nielson, H. R., 2015. Discretionary information flow control for interaction-oriented specifications. In: Martí-Oliet, N., Ölveczky, P. C., Talcott, C. (Eds.), Logic, Rewriting, and Concurrency: Essays Dedicated to José Meseguer on the Occasion of His 65th Birthday. Springer International Publishing, Cham, pp. 427–450.
URL https://doi.org/10.1007/978-3-319-23165-5_20
- López, H. A., Heussen, K., 2017. Choreographing cyber-physical distributed control systems for the energy sector. In: Seffah, A., Penzenstadler, B., Alves, C., Peng, X. (Eds.), Proceedings of the Symposium on Applied Computing, SAC 2017, Marrakech, Morocco, April 3-7, 2017. ACM, pp. 437–443.
URL <http://doi.acm.org/10.1145/3019612.3019656>
- López, H. A., Nielson, F., Nielson, H. R., 2016. Enforcing availability in failure-aware communicating systems. In: Albert, E., Lanese, I. (Eds.), Formal Techniques for Distributed Objects, Components, and Systems - 36th IFIP WG 6.1 International Conference, FORTE 2016, Held as Part of the 11th International Federated Conference on Distributed Computing Techniques, DisCoTec 2016, Heraklion, Crete, Greece, June 6-9, 2016, Proceedings. Vol. 9688 of Lecture Notes in Computer Science. Springer, pp. 195–211.
URL https://doi.org/10.1007/978-3-319-39570-8_13
- Lu, S., Park, S., Seo, E., Zhou, Y., 2008. Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. In: ASPLOS. ACM, pp. 329–339.
- Milner, R., 1980. A Calculus of Communicating Systems. Vol. 92 of LNCS. Springer, Berlin.
- Montesi, F., 2013. Choreographic programming. Ph.D. thesis, IT University of Copenhagen, http://fabriziomontesi.com/files/choreographic_programming.pdf.
- Montesi, F., 2015. Kickstarting choreographic programming. In: WS-FM. Vol. 9421 of Lecture Notes in Computer Science. Springer, pp. 3–10.

- Montesi, F., Guidi, C., Zavattaro, G., 2014. Service-oriented programming with Jolie. In: Bouguettaya, A., Sheng, Q., Daniel, F. (Eds.), *Web Services Foundations*. Springer, pp. 81–107.
- Montesi, F., Yoshida, N., 2013. Compositional choreographies. In: D’Argenio, P., Melgratti, H. (Eds.), *CONCUR*. Vol. 8052 of LNCS. Springer, pp. 425–439.
- Needham, R. M., Schroeder, M. D., Dec. 1978. Using encryption for authentication in large networks of computers. *Commun. ACM* 21 (12), 993–999.
URL <http://doi.acm.org/10.1145/359657.359659>
- Qiu, Z., Zhao, X., Cai, C., Yang, H., 2007. Towards the theoretical foundation of choreography. In: Williamson, C., Zurko, M., Patel-Schneider, P., Shenoy, P. (Eds.), *WWW*. ACM, pp. 973–982.
- Red Hat, 2017. SAVARA Testable Architecture. <http://savara.jboss.org/>.
- Sangiorgi, D., Walker, D., 2001. *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press.
- Turing, A., 1937. Computability and λ -definability. *J. Symb. Log.* 2 (4), 153–163.
- W3C WS-CDL Working Group, 2004. Web services choreography description language version 1.0. <http://www.w3.org/TR/2004/WD-ws-cdl-10-20040427/>.
- Wadler, P., 2012. Propositions as sessions. In: Thiemann, P., Findler, R. B. (Eds.), *ICFP*. ACM, pp. 273–286.
- Yoshida, N., Hu, R., Neykova, R., Ng, N., 2013. The scribble protocol language. In: *Trustworthy Global Computing - 8th International Symposium, TGC 2013, Buenos Aires, Argentina, August 30-31, 2013, Revised Selected Papers*. pp. 22–41.

Appendix A. CC as a Core Language: Channel Choreographies

CC is representative of the body of previous work on choreographic programming, where choreographies are used for implementations, for example (Carbone et al., 2012; Carbone and Montesi, 2013; Chor, 2017; Montesi and Yoshida, 2013; Dalla Preda et al., 2015; W3C WS-CDL Working Group, 2004). All the primitives of CC (and therefore of MC) can be encoded in such languages. Thus, we obtain a notion of function implementation for these languages, induced by that for CC, for which they are Turing complete.

In this section we make this claim precise for the model in (Carbone and Montesi, 2013), which we refer to in this work as Channel Choreographies (ChC). ChC is designed to be projected to a variant of the session-typed π -calculus (Coppo et al., 2016), which we refer to as Channel Processes (ChP). Communications in ChP are based on channels, instead of process names as in

$$\begin{aligned}
C ::= & \eta; C \quad | \quad \text{if } p.(e = e') \text{ then } C_1 \text{ else } C_2 \quad | \quad \mathbf{0} \\
& | \quad \text{def } X(\tilde{D}) = C_2 \text{ in } C_1 \quad | \quad X\langle \tilde{E} \rangle \quad | \quad \boxed{\nu r} C \\
\eta ::= & \boxed{\widetilde{p[A]} \text{ start } \widetilde{q[B]} : a(k)} \quad | \quad p[A].e \rightarrow q[B].x : k \\
& | \quad p[A] \rightarrow q[B] : k[l] \quad | \quad \boxed{p[A] \rightarrow q[B] : k\langle k'[C] \rangle} \\
D ::= & p(\tilde{x}, \tilde{k}) \quad \quad \quad E ::= p(\tilde{e}, \tilde{k})
\end{aligned}$$

Figure A.12: Channel Choreographies, Syntax.

SP. This layer of indirection means that a process performing an I/O action does not know [which other process it is going to communicate with](#), and that there can be race conditions on the usage of channels. ChC comes with a typing discipline for checking that the usage of channels specified in a choreography does not cause errors in the process code generated by EPP.

Appendix A.1. Channel Choreographies

Syntax. We report the full syntax of ChC in Figure A.12. Several terms are unnecessary for our translation; we $\boxed{\text{box}}$ such terms in our presentation of the syntax. In the original presentation of ChC, expressions e may contain any basic values (integers, strings, etc.) or computable functions, making the language trivially Turing complete. Also, labels l range over an infinite set. Here, for our development, we need only to consider expressions of the form ε or $s \cdot x$, and labels L and R (as in CC). The major difference between CC and ChC is the usage of *public channels* a and *session channels* k . Public channels are used to create new processes and channels at runtime, whereas session channels are used for point-to-point communications between processes. We only need a single session channel in our development.

An interaction η in ChC can be either a start, a value communication, a selection, or a delegation. In a start term $\widetilde{p[A]} \text{ start } \widetilde{q[B]} : a(k)$, the processes \tilde{p} on the left synchronise at the public channel a in order to create a new private session k and spawn some new processes \tilde{q} (k and \tilde{q} are bound to the continuation). Each process is annotated with the role it plays in the created session. Roles are ranged over by A, B, C, \dots . They are used in the typing discipline of ChC to check whether sessions are used according to protocol specifications, given as multiparty session types (Honda et al., 2016).

In a value communication $p[A].e \rightarrow q[B].x : k$, process p sends its evaluation of expression e over session k to process q , which stores the result in its local variable x ; the name x appearing under q is bound to the continuation. Differently from CC, where each process has only one memory cell accessed through the placeholder c , in ChC each process has an unbounded number of cells (variables). Selections in ChC, of the form $p[A] \rightarrow q[B] : k[l]$, are very

$$\begin{array}{c}
\frac{}{\overline{\mathfrak{p}[\mathbf{A}].v \rightarrow \mathfrak{q}[\mathbf{B}].x : k; C \rightarrow C[v/x@q]}} \text{Ch|Com} \quad \frac{}{\overline{\mathfrak{p}[\mathbf{A}] \rightarrow \mathfrak{q}[\mathbf{B}] : k[l]; C \rightarrow C}} \text{Ch|Sel} \\
\frac{}{\overline{\mathfrak{p}[\mathbf{A}] \rightarrow \mathfrak{q}[\mathbf{B}] : k\langle k'[\mathbf{C}] \rangle; C \rightarrow C}} \text{Ch|Del} \quad \frac{}{\overline{\mathfrak{p}[\mathbf{A}] \text{ start } \mathfrak{q}[\mathbf{B}] : a(k); C \rightarrow (\nu \tilde{\mathfrak{q}}, k) C}} \text{Ch|Start} \\
\frac{i = 1 \text{ if } v = w, \ i = 2 \text{ otherwise}}{\text{if } \mathfrak{p}.(v = w) \text{ then } C_1 \text{ else } C_2 \rightarrow C_i} \text{Ch|Cond} \\
\frac{C_1 \preceq C_2 \quad C_2 \rightarrow C'_2 \quad C'_2 \preceq C'_1}{C_1 \rightarrow C'_1} \text{Ch|Struct} \\
\frac{C_1 \rightarrow C'_1}{\text{def } X(\tilde{D}) = C_2 \text{ in } C_1 \rightarrow \text{def } X(\tilde{D}) = C_2 \text{ in } C'_1} \text{Ch|Ctx}
\end{array}$$

Figure A.13: Channel Choreographies, Semantics.

similar to those in CC: the only difference is that we also have to write which role each process plays and the session used for communicating. In a delegation term $\mathfrak{p}[\mathbf{A}] \rightarrow \mathfrak{q}[\mathbf{B}] : k\langle k'[\mathbf{C}] \rangle$, process \mathfrak{p} delegates its role \mathbf{C} in session k' to process \mathfrak{q} ; delegation in ChC is a typed form of channel mobility, inspired by the π -calculus.

In a conditional, process \mathfrak{p} chooses a continuation based on whether the expressions e and e' evaluate to the same value according to its own local state. The restriction term $(\nu r) C$ is standard and binds the scope of r (which can be either a process name \mathfrak{p} or a session channel name k) to C . Finally, in the definition of a recursive procedure, the parameters \tilde{D} indicate which processes are used in the body of the procedure and which variables and sessions are used by each process. In the invocation of a procedure $X\langle \tilde{E} \rangle$, each process can pass generic expressions as parameters to itself.

Semantics. ChC was originally presented with an asynchronous semantics (Carbone and Montesi, 2013). We first present our results using only the (simpler) synchronous variant of the semantics of ChC, and defer the discussion of the general asynchronous case to the end of this section. This semantics is given in terms of a reduction relation, presented in Figure A.13. Rule Ch|Com is the key rule, where the value sent from a process \mathfrak{p} is received by a process \mathfrak{q} . Technically, this is modelled by replacing variable x with v in the continuation C , but only when it appears under the process name \mathfrak{q} (the *smart substitution* $C[v/x@q]$). Rule Ch|Cond models an internal choice: \mathfrak{p} chooses a continuation depending on whether the two values v and w are the same. Rules Ch|Del and Ch|Start implement the informal semantics of delegation and start described earlier; we do not use them in our development. The other rules are similar to those of CC. The structural precongruence \preceq is defined as expected, following the same intuition as that for CC. In particular, it supports swapping two terms

$$\begin{aligned}
P, Q ::= & \quad k[\mathbf{A}]!_{\mathbf{B}}\langle e \rangle; P \mid k[\mathbf{B}]?_{\mathbf{A}}(x); P \mid k[\mathbf{A}]!_{\mathbf{B}} \oplus l; P \mid k[\mathbf{B}]?_{\mathbf{A}} \& \{l_i : P_i\}_{i \in I} \\
& \mid P \mid Q \quad \mid \text{if } e = e' \text{ then } P \text{ else } Q \quad \mid \text{def } X(\tilde{x}, \tilde{k}) = Q \text{ in } P \quad \mid X(\tilde{e}, \tilde{k}) \quad \mid \mathbf{0}
\end{aligned}$$

Figure A.14: Channel Processes, Syntax (selection).

whenever they involve disjoint process names.

As expected, ChC offers a deadlock-freedom-by-design property in the style of Theorem 1 (Carbone and Montesi, 2013).

Appendix A.2. Channel Processes (ChP)

We now present Channel Processes (ChP), the target language that choreographies in ChC can be projected to. We discuss only the terms used in our work (see (Carbone and Montesi, 2013) for a complete presentation).

Syntax. The relevant part of the syntax of processes (P, Q) is reported in Figure A.14. Binding occurrences are denoted by the usage of round parentheses. Terms $\bar{a}[\tilde{\mathbf{A}}](k)$, $a[\mathbf{A}](k)$ and $!a[\mathbf{A}](k)$ are used to start a new session k by synchronising on the public channel a , and model respectively: the process requesting the creation of the session (responsible for playing the first role in $\tilde{\mathbf{A}}$); a process accepting to play role \mathbf{A} in the session; and, finally, a replicated process that will spawn a fresh process for playing role \mathbf{A} . In the first line we have the terms for in-session communications. In term $k[\mathbf{A}]!_{\mathbf{B}}\langle e \rangle; P$, as role \mathbf{A} on session k , we send the value of expression e to \mathbf{B} on the same session; then, we proceed as P . Dually, term $k[\mathbf{B}]?_{\mathbf{A}}(x); P$ receives a message for role \mathbf{B} from role \mathbf{A} on session k and stores it in variable x . Terms $k[\mathbf{A}]!_{\mathbf{B}} \oplus l$ and $k[\mathbf{B}]?_{\mathbf{A}} \& \{l_i : P_i\}_{i \in I}$ model, respectively, branch selection and offering. Finally, terms $k[\mathbf{A}]!_{\mathbf{B}}\langle k'[\mathbf{C}] \rangle$ and $k[\mathbf{B}]?_{\mathbf{A}}\langle k'[\mathbf{C}] \rangle$ capture channel mobility. The other terms are the standard parallel composition, procedure definition, procedure call, conditional (restricted to checking for equality), and terminated process.

Semantics. As before, we discuss only the synchronous semantics of ChP. We discuss only the communication rules, shown in Figure A.15, as all the other rules are standard – see (Kouzapas and Yoshida, 2013). As in typical calculi for multiparty sessions equipped with roles, each role in a session is a distinct communication endpoint. Therefore, a send action on a session k from a role \mathbf{A} towards a role \mathbf{B} synchronises with a receive action on the same session k by the target role \mathbf{B} wishing to receive from the sender role \mathbf{A} .

Appendix A.3. Endpoint Projection and Typing

As for CC, the Endpoint Projection from ChC to ChP is defined by first defining how to project the behaviour of a single process. The projection of a process \mathbf{p} from a choreography C , written $\llbracket C \rrbracket_{\mathbf{p}}$, is inductively defined on the structure of C in a similar way as the behaviour projection given in § 3.3.

$$\begin{array}{c}
\frac{}{k[\mathbf{A}]!\mathbf{B}\langle v \rangle; P \mid k[\mathbf{B}]?\mathbf{A}(x); Q \rightarrow P \mid Q[v/x]} \text{CP|Com} \\
\frac{(j \in I)}{k[\mathbf{A}]!\mathbf{B} \oplus l_j; P \mid k[\mathbf{B}]?\mathbf{A}\&\{l_i : Q_i\}_{i \in I} \rightarrow P \mid Q_j} \text{CP|Sel}
\end{array}$$

Figure A.15: Channel Processes, Semantics (selection).

$$\begin{array}{c}
\llbracket \mathbf{p}[\mathbf{A}].e \rightarrow \mathbf{q}[\mathbf{B}].x : k; C \rrbracket_r = \begin{cases} k[\mathbf{A}]!\mathbf{B}\langle e \rangle; \llbracket C \rrbracket_r & \text{if } r = \mathbf{p} \\ k[\mathbf{B}]?\mathbf{A}(x); \llbracket C \rrbracket_r & \text{if } r = \mathbf{q} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \text{if } \mathbf{p}.(e = e') \\ \text{then } C_1 \text{ else } C_2 \rrbracket_r = \begin{cases} \text{if } e = e' \\ \text{then } \llbracket C_1 \rrbracket_r \text{ else } \llbracket C_2 \rrbracket_r & \text{if } r = \mathbf{p} \\ \llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r & \text{otherwise} \end{cases}
\end{array}$$

Figure A.16: Channel Choreographies, EndPoint Projection (relevant cases).

The complete EPP procedure from ChC to ChP is technically involved, because the start term $\mathbf{p}[\mathbf{A}] \text{ start } \mathbf{q}[\mathbf{B}] : a(k)$ found in ChC enables the reuse of the same services exposed at a public channel a for spawning processes with potentially different behaviour. However, since start terms and restriction of names are unnecessary for our development, we can use a much simpler definition – see (Carbone and Montesi, 2013) for the general case. We report the rules for projecting value communications and conditionals in Figure A.16. The merging operator $P \sqcup Q$ works as in CC: it is isomorphic to P and Q aside from input branches with distinct labels, which are instead included in a larger input branching.

Appendix A.4. Typing ChC

Differently from CC, the EPP of a choreography in ChC does not always yield correct results. Consider the following choreography:

$$C = \mathbf{p}[\mathbf{A}] \rightarrow \mathbf{q}[\mathbf{B}] : k[\mathbf{L}]; \mathbf{q}[\mathbf{B}].\varepsilon \rightarrow \mathbf{p}[\mathbf{A}].x : k; \mathbf{r}[\mathbf{A}] \rightarrow \mathbf{q}[\mathbf{B}] : k[\mathbf{L}]$$

The choreography C above always terminates by reaching $\mathbf{0}$ (by using rules Ch|Sel, then Ch|Com, and then Ch|Sel again). However, its EPP (albeit defined)

$$\begin{aligned}
G &::= \mathbf{A} \rightarrow \mathbf{B} : \langle \mathbf{nat} \rangle; G \mid \mathbf{A} \rightarrow \mathbf{B} : \{l_i : G_i\}_{i \in I} \mid \mu \mathbf{t}; G \mid \mathbf{t} \mid \mathbf{end} \\
S &::= \mathbf{nat} \mid \mathbf{string} \mid \dots \qquad l ::= L \mid R
\end{aligned}$$

Figure A.17: Global Types, Syntax.

may get stuck:

$$\begin{aligned}
\llbracket C \rrbracket &= \overbrace{k[\mathbf{A}]! \mathbf{B} \oplus L; k[\mathbf{A}]? \mathbf{B}(x)}^{\llbracket C \rrbracket_{\mathbf{p}}} \mid \overbrace{k[\mathbf{A}]! \mathbf{B} \oplus L}^{\llbracket C \rrbracket_{\mathbf{r}}} \\
&\quad \mid \underbrace{k[\mathbf{B}]? \mathbf{A} \& \{L : k[\mathbf{B}]! \mathbf{A}(\varepsilon); k[\mathbf{B}]? \mathbf{A} \& \{L : \mathbf{0}\}\}}_{\llbracket C \rrbracket_{\mathbf{q}}}
\end{aligned}$$

Above, we have a race between the projections of process \mathbf{p} and process \mathbf{r} for the selection of label L offered by process \mathbf{q} . This is because both \mathbf{p} and \mathbf{r} play the same role \mathbf{A} in session k and therefore the receiver (the projection of process \mathbf{q}) cannot distinguish them. In the case where the race is won by the projection of process \mathbf{r} , not only do we obtain a reduction not defined by the originating choreography, but we even get into a deadlocked situation:

$$\llbracket C \rrbracket \rightarrow k[\mathbf{A}]! \mathbf{B} \oplus L; k[\mathbf{A}]? \mathbf{B}(x) \mid k[\mathbf{B}]! \mathbf{A}(\varepsilon); k[\mathbf{B}]? \mathbf{A} \& \{L : \mathbf{0}\}$$

To avoid such situations, ChC comes with a typing discipline based on multi-party session types that guarantees the absence of races.

A typing judgement for CC has the form $\Gamma; \Theta \vdash C \triangleright \Delta$, where Δ types the usage of sessions, Θ the ownership of roles by processes, and Γ variables and public channels.

Formally, the typing environment Γ contains variable typings of the form $x @ \mathbf{p} : S$, typing variable x at \mathbf{p} with data type S (which can only be \mathbf{nat} in our case). An environment Θ contains ownership typings of the form $\mathbf{p} : k[\mathbf{A}]$, read “process \mathbf{p} owns role \mathbf{A} in k ” (when writing $\Theta, \mathbf{p} : k[\mathbf{A}]$, it is assumed that no other process owns the same role for the same session in Θ). The environment Δ contains session typings of the form $k : G$, where G is a global type (Honda et al., 2016). The syntax of global types is given in Figure A.17. A global type G abstracts a communication between two roles in a session. A value communication is abstracted by $\mathbf{A} \rightarrow \mathbf{B} : \langle \mathbf{nat} \rangle$ (we restrict values to be natural numbers). A global type $\mathbf{A} \rightarrow \mathbf{B} : \{l_i : G_i\}_{i \in I}$ allows any selection from \mathbf{A} to \mathbf{B} of one of the labels l_i , provided that then the session proceeds as specified by the corresponding continuation G_i . The other terms are for recursion ($\mu \mathbf{t}$ and \mathbf{t}) and termination (\mathbf{end}).

We discuss the most relevant typing rules for ChC, given in Figure A.18. Rule T|Com checks that, in a value communication on session k , the sender and receiver processes own their respective roles in session k ($\Theta \vdash \mathbf{p} : k[\mathbf{A}], \mathbf{q} : k[\mathbf{B}]$), that the protocol for session k expects a communication for their respective roles

$$\begin{array}{c}
\frac{\Gamma \vdash e @ \mathbf{p} : S \quad \Theta \vdash \mathbf{p} : k[\mathbf{A}], \mathbf{q} : k[\mathbf{B}] \quad \Gamma, x @ \mathbf{q} : S; \Theta \vdash C \triangleright \Delta, k : G}{\Gamma; \Theta \vdash \mathbf{p}[\mathbf{A}].e \rightarrow \mathbf{q}[\mathbf{B}].x : k; C \triangleright \Delta, k : \mathbf{A} \rightarrow \mathbf{B} : \langle S \rangle; G} \text{T|Com} \\
\\
\frac{\Theta \vdash \mathbf{p} : k[\mathbf{A}], \mathbf{q} : k[\mathbf{B}] \quad j \in I \quad \Gamma; \Theta \vdash C \triangleright \Delta, k : G_j}{\Gamma; \Theta \vdash \mathbf{p}[\mathbf{A}] \rightarrow \mathbf{q}[\mathbf{B}] : k[l_j]; C \triangleright \Delta, k : \mathbf{A} \rightarrow \mathbf{B} : \{l_i : G_i\}_{i \in I}} \text{T|Sel} \\
\\
\frac{\Gamma; \Theta \vdash C_1 \triangleright \Delta \quad \Gamma; \Theta \vdash C_2 \triangleright \Delta}{\Gamma; \Theta \vdash \text{if } \mathbf{p}.(e = e') \text{ then } C_1 \text{ else } C_2 \triangleright \Delta} \text{T|Cond} \\
\\
\frac{\Delta \text{ end only} \quad \Gamma' \subseteq \Gamma}{\Gamma, X(\tilde{D}) : (\Gamma'; \Theta; \Delta'); \Theta \vdash X\langle \tilde{D} \rangle \triangleright \Delta, \Delta'} \text{T|Call} \\
\\
\frac{\Gamma, X(\tilde{D}) : (\Gamma'; \Theta'; \Delta'); \Theta \vdash C_1 \triangleright \Delta \quad \Gamma' \subseteq \Gamma \quad \Gamma', X(\tilde{D}) : (\Gamma'; \Theta'; \Delta'); \Theta' \vdash C_2 \triangleright \Delta' \quad \Theta' \subseteq \Theta}{\Gamma; \Theta \vdash \text{def } X(\tilde{D}) = C_2 \text{ in } C_1 \triangleright \Delta} \text{T|Def}
\end{array}$$

Figure A.18: Channel Choreographies, Typing Rules (selection).

($k : \mathbf{A} \rightarrow \mathbf{B} : \langle S \rangle; G$), and that the expression sent by the sender has the expected type S . Rule T|Sel checks that a selection uses one of the labels expected by the protocol for the session ($j \in I$). Rule T|Cond is standard, requiring both branches to have the same typing; observe that different communication behaviour in the two branches may still occur, because of rule T|Sel. Rules T|Call and T|Def type, respectively, recursive calls and recursive procedures. These rules are simplified compared to the presentation in (Carbone and Montesi, 2013), taking into account that our encoding always calls procedures with exactly the same arguments (processes and variables) as they are declared.

Well-typedness is preserved by reductions. Furthermore, using this type system we get an operational correspondence result for EPP from ChC to ChP.

Theorem 10 (Operational Correspondence (ChC \leftrightarrow ChP) (Carbone and Montesi, 2013)). *Let C be a well-typed channel choreography without start subterms (terms of the form $\widetilde{\mathbf{p}[\mathbf{A}]} \text{ start } \widetilde{\mathbf{q}[\mathbf{B}]}$: $a(k)$) and such that its endpoint projection $\llbracket C \rrbracket$ is defined. Then:*

- (Completeness) $C \rightarrow C'$ implies $\llbracket C \rrbracket \rightarrow \succ \llbracket C' \rrbracket$;
- (Soundness) $\llbracket C \rrbracket \rightarrow P$ implies $C \rightarrow C'$ and $\llbracket C' \rrbracket \prec P$.

where \prec is the pruning relation defined in (Carbone and Montesi, 2013).

As for CC, the EPP of a well-typed channel choreography never deadlocks.

Appendix A.5. Embedding CC into ChC

Defining an embedding from CC to ChC is nontrivial, as the communication primitives of CC and ChC are different. In CC, messages are passed directly between processes: each process knows whom it is sending to or receiving from in each communication step; in ChC, communication is between roles in a session channel. To translate core choreographies into channel choreographies, we therefore assign to each process a role syntactically identical to its name, and perform all communication over a fixed channel k .

Conditional terms are also not directly translatable, as ChC evaluates guards in a single process. For this reason, each translated process uses two variables: x , storing its internal value, and y , used exclusively for temporary storage of a value required for a test.

For recursion, we recall that $\text{pn}(C)$ returns the set of process names in C .

Definition 11 (Embedding of CC in ChC). *The embedding of a core choreography C in ChC is $\{C\}$, inductively defined as follows.*

$$\begin{aligned} \{p.e \rightarrow q; C\} &= p[p].e[x/c] \rightarrow q[q].x : k; \{C\} \\ \{p \rightarrow q[l]; C\} &= p[p] \rightarrow q[q] : k[l]; \{C\} \\ \left\{ \text{if } p \stackrel{<}{=} q \text{ then } C_1 \text{ else } C_2 \right\} &= q[q].x \rightarrow p[p].y : k; \\ &\quad \text{if } p.(x = y) \text{ then } \{C_1\} \text{ else } \{C_2\} \\ \{\text{def } X = C_2 \text{ in } C_1\} &= (\text{def } X(*) = \{C_2\}) [|C_2|/*] \text{ in } \{C_1\} \\ \{X\} = X\langle * \rangle &\quad \{\mathbf{0}\} = \mathbf{0} \end{aligned}$$

where $|A| = \{p(\{x, y\}, k) \mid p \in \text{pn}(A)\}$.

Lemma 11. *Let C and C' be core choreographies. Then $C \preceq C'$ if and only if $\{C\} \preceq \{C'\}$.*

Proof. For the direct implication, observe that all structural precongruence rules in CC become valid instances of precongruence in ChC when mapped by $\{\cdot\}$. Conversely, given a structural precongruence rule in ChC, if its arguments are in the image of $\{\cdot\}$, then the rule can be pulled back to a valid precongruence in CC. \square

In order to compare the semantics of core and channel choreographies, we need to take the state into account. This is done by viewing each state as a substitution, replacing all free occurrences of x with the actual content of the process it belongs to.

Definition 12 (Substitution induced by state). *Let C be a core choreography and σ be a state. The substitution σ_C is defined as $\sigma_C = [\sigma(p)/x@p \mid p \in \text{pn}(C)]$, and the embedding of C in ChC via σ is the channel choreography $\{C\}_\sigma = \sigma_C(\{C\})$.*

Below, \rightarrow^+ denotes a chain of one or more applications of \rightarrow , and $\rightarrow^?$ denotes identity or one application of \rightarrow .

Theorem 11 (Operational Correspondence (CC \leftrightarrow ChC)). *Let C be a choreography in CC. Then, for all σ :*

- (Completeness) $C, \sigma \rightarrow C', \sigma'$ implies $\{C\}_\sigma \rightarrow^+ \{C'\}_{\sigma'}$;
- (Soundness) $\{C\}_\sigma \rightarrow C'$ implies $C, \sigma \rightarrow C^*, \sigma^*$ and $C' \rightarrow^? \{C^*\}_{\sigma^*}$.

Proof. • (Completeness) We analyze the possible cases for the rule justifying $C \xrightarrow{\eta} C'$. The only non-trivial case is rule C|Com. Suppose C is $\mathbf{p}.e \rightarrow \mathbf{q}; C'$. If $C, \sigma \rightarrow C', \sigma'$, then $\{C\}_\sigma = \mathbf{p}.\sigma_C(e[x/c]) \rightarrow \mathbf{q}.x : k; \{C'\}_{\sigma'}$ where $\sigma'' = \sigma_C \setminus \{\sigma(\mathbf{q})/x@q\}$ can make a transition to $\{C'\}_{\sigma''} [e[\sigma_C(\mathbf{p})/c]/x@q]$, which coincides with $\{C'\}_{\sigma'}$.

- (Soundness) The proof is again by case analysis on the transition from $\{C\}_\sigma$ to C' , noting that this cannot involve delegation or start actions. Most cases are straightforward, except when the transition is a communication obtained from translating a conditional. In this case, C' must execute the full conditional action, and $\{C^*\}_{\sigma^*}$ reduces to C' by application of rule Ch|Cond. □

We now define a notion of function implementation in ChC. Since the semantics of ChC does not have state, this definition is slightly different than that for CC. The embedding of CC is then a Turing complete fragment of ChC, which we show to be projectable.

Definition 13 (Implementation in ChC). *A channel choreography C implements a function $f : \mathbb{N}^n \rightarrow \mathbb{N}$ with input variables $\mathbf{p}_1.z_1, \dots, \mathbf{p}_n.z_n$ and output variable $\mathbf{q}.z$ if, for all $x_1, \dots, x_n \in \mathbb{N}$:*

- if $f(\tilde{x})$ is defined, then $C[\ulcorner x_i \urcorner / z_i @ \mathbf{p}_i] \rightarrow^* \mathbf{0}$, and \mathbf{q} receives exactly one message with $\ulcorner f(\tilde{x}) \urcorner$ as the value transmitted;
- if $f(\tilde{x})$ is not defined, then $C[\ulcorner x_i \urcorner / z_i @ \mathbf{p}_i] \not\rightarrow^* \mathbf{0}$, and \mathbf{q} never receives any messages.

Theorem 12 (Soundness). *If $f : \mathbb{N}^n \rightarrow \mathbb{N}$ is a partial recursive function, then $\{\{\{f\}\}^{\tilde{\mathbf{p}} \rightarrow \mathbf{q}}\}$ implements f with input variables $\tilde{\mathbf{p}}.x$ and output variable $\mathbf{q}.x$.*

Proof. Consequence of the proof of Theorem 5 and of Theorem 11: as the only free variables in $\{\{\{f\}\}^{\tilde{\mathbf{p}} \rightarrow \mathbf{q}}\}$ are $\mathbf{p}_i.x$ for $1 \leq i \leq n$, $\{\{\{f\}\}^{\tilde{\mathbf{p}} \rightarrow \mathbf{q}}\} [\ulcorner x_i \urcorner / z_i @ \mathbf{p}_i]$ coincides with $\{\{\{f\}\}^{\tilde{\mathbf{p}} \rightarrow \mathbf{q}}\}_\sigma$ whenever σ contains $\ulcorner x_i \urcorner$ at each process \mathbf{p}_i . □

We conclude by combining our results to characterise a Turing-complete and deadlock-free fragment of ChP.

Let ChP^{ChC} be the smallest fragment of ChP containing the projections of all typable and projectable choreographies in ChC, formally: $\text{ChP}^{\text{ChC}} = \{\llbracket C \rrbracket \mid \llbracket C \rrbracket \text{ is defined}\}$. By Theorem 10, all terms in ChP^{ChC} are deadlock-free.

We now show that ChP^{ChC} is also Turing powerful. The development is similar to that for SP^{CC} (§ 4.3), but we need two additional steps. First, the operational correspondence theorem for the EPP of ChC (Theorem 10) needs the projected channel choreography to be well-typed. Fortunately, this is always the case for the channel choreographies obtained by embedding amended CC terms.

Lemma 12. *Let C be a core choreography and σ a state. Then $C' = \{\text{Amend}(C)\}_\sigma$ implies $\Gamma; \Theta \vdash C' \triangleright \Delta$ for some Γ , Θ and Δ .*

Proof. Choosing Θ is trivial, as each process has its own role. For Γ , we assign type **nat** to all variables. Finally, $\Delta = k : G$, where G is inferred by abstracting the communications in C . The inductive construction of the latter is always possible since we applied **Amend**, so we can type each conditional with either the same global type or a branching global type with two labels. \square

Second, we need to know that the embedding of a projectable core choreography is also projectable in ChC.

Lemma 13. *If C is projectable, then $\{C\}_\sigma$ is projectable for any σ .*

Proof. The thesis follows from the fact that: if the projections of two choreographies are mergeable in CC, then the projections of their embeddings into ChC are mergeable. This is proven by structural induction. \square

Using these results, the proof of Corollary 1 can be adapted to yield the following property.

Corollary 4 (Turing completeness of ChP^{ChC}). *Every partial recursive function is implementable in ChP^{ChC} .*

We thus characterise a fragment of the session-based π -calculus from (Coppo et al., 2016) that contains only deadlock-free terms and is Turing complete.

Appendix A.6. Asynchronous semantics

The original semantics of ChC is made asynchronous by means of an additional rule that allows some transitions protected by a prefix to be executed, yielding more possible reduction sequences. In this scenario, the operational correspondence between CC and ChC is no longer as strong as stated in Theorem 11; in particular, the result for soundness now reads

$$\text{(Soundness)} \quad \{C\}_\sigma \rightarrow^+ C' \text{ implies } C, \sigma \rightarrow^+ C^*, \sigma^* \text{ and } C' \rightarrow^* \{C^*\}_{\sigma^*}$$

However, the reduction relation in the asynchronous setting, restricted to the language fragment we consider (embeddings of core choreographies, which in particular are typable) is still confluent, and it includes all synchronous executions. Therefore, Theorem 12 remains valid in this general case.