

A Core Model for Choreographic Programming

Luís Cruz-Filipe and Fabrizio Montesi

University of Southern Denmark {lcf, fmontesi}@imada.sdu.dk

Abstract. Choreographic Programming is a paradigm for developing concurrent programs that are deadlock-free by construction, by 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.

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 [21]; (ii) combining different communication protocols in a single application is a major source of mistakes [20].


The paradigm of Choreographic Programming [22] 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” notation of security protocols. 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 [6, 25].

So far, work on choreographic programming focused on features of practical value – including web services [5], multiparty sessions [6, 8], modularity [24], and runtime adaptation [12]. 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.

However, 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 expressivity 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 [5, 6]. 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 on the right).



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. Thus, in a practical application of our work, CC would be the better candidate as frontend language for programmers, and MC could be used as an intermediate step in a compiler. A key technical advantage of this methodology is that it bypasses the need for the standard notion of merging [5], which is typically one of the most complicated steps in EPP. Our EPP for MC enjoys an elegant definition.

Structure of the paper. CC is defined in § 2. In § 3, we introduce Stateful Processes (SP), our target process model, and an EPP procedure from CC to SP. We show that CC and its set of choreography projections are Turing complete in § 4. In § 5, we show that all primitives of CC except for label selections are necessary to achieve Turing completeness; we then introduce MC (the fragment of CC without label selections) and prove both that it is Turing complete and that removing or weakening any of its primitives breaks this property. In § 6, we discuss the implications of our work for other choreography languages. Related work and discussion are given in § 7. Full definitions and proofs are in [9].

2 Core Choreographies

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

Syntax. The syntax of CC is as follows, where C ranges over choreographies.

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

We use two (infinite) disjoint sets of names: processes ($\mathbf{p}, \mathbf{q}, \dots$) and procedures (X, \dots). Processes run in parallel, and each process stores a value – a string of the form $\mathbf{s} \cdots \mathbf{s} \cdot \varepsilon$ – 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 – $\mathbf{p}.e \rightarrow \mathbf{q}$ – or a label selection – $\mathbf{p} \rightarrow \mathbf{q}[l]$. In $\mathbf{p}.e \rightarrow \mathbf{q}$, \mathbf{p} sends its local evaluation of expression e to \mathbf{q} , which stores the received value. Expressions are either the constant ε , the value of the sender (written as \mathbf{c}), or an application of the successor operator to \mathbf{c} . In $\mathbf{p} \rightarrow \mathbf{q}[l]$, \mathbf{p} communicates label l (either L or R) to \mathbf{q} . In a conditional $\text{if } \mathbf{p} \stackrel{\leftarrow}{=} \mathbf{q} \text{ then } C_1 \text{ else } C_2$, \mathbf{q} sends its value to \mathbf{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.

Semantics. 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 below and closed under structural precongurence \preceq .

$$\frac{v = e[\sigma(\mathbf{p})/\mathbf{c}]}{\mathbf{p}.e \rightarrow \mathbf{q}; C, \sigma \rightarrow C, \sigma[\mathbf{q} \mapsto v]} \text{ [C|Com]} \quad \frac{i = 1 \text{ if } \sigma(\mathbf{p}) = \sigma(\mathbf{q}), i = 2 \text{ o.w.}}{\text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } C_1 \text{ else } C_2, \sigma \rightarrow C_i, \sigma} \text{ [C|Cond]}$$

$$\frac{}{\mathbf{p} \rightarrow \mathbf{q}[l]; C, \sigma \rightarrow C, \sigma} \text{ [C|Sel]} \quad \frac{C_1, \sigma \rightarrow C'_1, \sigma'}{\text{def } X = C_2 \text{ in } C_1, \sigma \rightarrow \text{def } X = C_2 \text{ in } C'_1, \sigma'} \text{ [C|Ctx]}$$

These rules formalise the intuition presented earlier. In the premise of [C|Com], we write $e[\sigma(\mathbf{p})/\mathbf{c}]$ for the result of replacing \mathbf{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 . The key rule defining the structural precongurence is [C|Eta-Eta], allowing non-interfering actions to be executed in any order.

$$\text{[C|Eta-Eta]} \quad \text{if } \mathbf{pn}(\eta) \cap \mathbf{pn}(\eta') = \emptyset \text{ then } \eta; \eta' \equiv \eta'; \eta$$

Function $\mathbf{pn}(C)$ returns the set of all process names occurring in C , and $C \equiv C'$ stands for $C \preceq C'$ and $C' \preceq C$. The other rules for \preceq are standard, and support recursion unfolding and garbage collection of unused definitions.

Remark 1 (Label Selection). To the reader unfamiliar with choreographies, the role of selection – $\mathbf{p} \rightarrow \mathbf{q}[l]$ – may be unclear at this point. They are crucial in making choreographies projectable, as we anticipate with the choreography $\text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } (\mathbf{p}.c \rightarrow \mathbf{r}; \mathbf{0}) \text{ else } (\mathbf{r}.c \rightarrow \mathbf{p}; \mathbf{0})$. Here, \mathbf{p} checks whether its value is the same as that of \mathbf{q} . If so, \mathbf{p} communicates its value to \mathbf{r} ; otherwise, it is \mathbf{r} that communicates its value to \mathbf{p} . Recall that processes are assumed to run independently and share no data. Here, \mathbf{p} is the only process that knows which branch of the conditional should be executed. However, \mathbf{r} also needs to know this information, since it must behave differently. Intuitively, we need to propagate \mathbf{p} 's decision to \mathbf{r} , which is achieved with selections: $\text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } (\mathbf{p} \rightarrow \mathbf{r}[L]; \mathbf{p}.c \rightarrow \mathbf{r}; \mathbf{0}) \text{ else } (\mathbf{p} \rightarrow \mathbf{r}[R]; \mathbf{r}.c \rightarrow \mathbf{p}; \mathbf{0})$. Now, \mathbf{p} tells \mathbf{r} about its choice by sending a different label. This intuition will be formalised in our definition of EndPoint Projection in § 3. The first choreography we presented (without label selections) is not projectable, whereas the second one is.

Theorem 1. *If C is a choreography, then either $C \preceq \mathbf{0}$ (C has terminated) or, for all $\sigma, C, \sigma \rightarrow C', \sigma'$ for some C' and σ' (C can reduce).*

The semantics of CC suggests a natural definition of computation. We write \rightarrow^* for the transitive closure of \rightarrow and $C, \sigma \not\rightarrow^* \mathbf{0}$ for $C, \sigma \not\rightarrow^* \mathbf{0}, \sigma'$ for any σ' .

Definition 1. *A choreography C implements a function $f : \mathbb{N}^n \rightarrow \mathbb{N}$ with input processes $\mathbf{p}_1, \dots, \mathbf{p}_n$ and output process \mathbf{q} if, for all $x_1, \dots, x_n \in \mathbb{N}$ and for every state σ s.t. $\sigma(\mathbf{p}_i) = \lceil x_i \rceil$:*

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

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

Sequential composition and parallelism. The results in the remainder use choreographies with only one exit point (a single occurrence of $\mathbf{0}$). When C has a single exit point, we write $C \mathbin{\text{;}} C'$ for the choreography obtained by replacing $\mathbf{0}$ in C with C' . Then, $C \mathbin{\text{;}} C'$ behaves as a “sequential composition” of C and C' .

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

Structural precongruence gives $C \mathbin{\text{;}} C'$ fully parallel behaviour in some cases. Intuitively, C_1 and C_2 run in parallel in $C_1 \mathbin{\text{;}} C_2$ 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 2. *Let $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{q}}$ be disjoint. Then, $\tilde{\sigma}$ is an interleaving of $\widetilde{\sigma_1}$ and $\widetilde{\sigma_2}$ wrt $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{q}}$ 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 $\mathbf{p} \in \tilde{\mathbf{p}}$, and $\widetilde{\sigma'_2(\mathbf{q})} = \widetilde{\sigma_2(\mathbf{q})}$ for all $\mathbf{q} \in \tilde{\mathbf{q}}$;
- $\widetilde{\sigma(r)}$ is a constant sequence for all $r \notin \tilde{\mathbf{p}} \cup \tilde{\mathbf{q}}$.

Definition 3. *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 $C_1 \mathbin{\text{;}} C_2$ if: whenever $C_i \xrightarrow{\tilde{\sigma}_i}^* \mathbf{0}$, then $C_1 \mathbin{\text{;}} C_2 \xrightarrow{\tilde{\sigma}}^* \mathbf{0}$ 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. *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 $C_1 \mathbin{\text{;}} C_2$.*

Example 1. We present examples of choreographies in CC, writing them as macros (syntax shortcuts). We use the notation $\text{M}(params) \triangleq C$, where M is the name of the macro, *params* its parameters, and C its body.

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 L); otherwise, it asks them to do another step (selection of R). In each step,

the values of \mathbf{p} and r are incremented using \mathbf{t}_1 and \mathbf{t}_2 . The compositional usage of INC is allowed, as it has exactly one exit point.

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

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

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

Definition 3 and Theorem 2 straightforwardly generalise to an arbitrary number of processes. We provide an example of such parallel behaviour in Theorem 6.

3 Stateful Processes and EndPoint Projection

We present Stateful Processes (SP), our target process model, and show how to synthesise process implementations from choreographies in CC.

Syntax. The syntax of SP is reported below. 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.

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

Expressions and labels are as in CC. A send term $\mathbf{q}!(e); B$ 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 $\text{if } c \stackrel{\leftarrow}{=} \mathbf{q} \text{ then } B_1 \text{ else } B_2$, 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.

Semantics. The reduction rules for SP are mostly standard, from process calculi. The key difference from CC is that execution is now distributed over processes.

We report the key rules for synchronisation:

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

Rule [S|Com] follows the standard communication rule in process calculi. A process \mathfrak{p} executing a send action towards a process \mathfrak{q} can synchronise with a receive-from- \mathfrak{p} action at \mathfrak{q} ; in the reduct, \mathfrak{q} 's value is updated with the value sent by \mathfrak{p} , obtained by replacing the placeholder c in e with the value of \mathfrak{p} . Rule [S|Sel] is selection from session types [15], with the sender selecting one of the branches offered by the receiver. In rule [S|Cond], \mathfrak{p} (executing the conditional) acts as a receiver for the value sent by the process whose value it wants to read (\mathfrak{q}). All other rules are standard (see [9]), and use a structural precongruence that supports: recursion unfolding, garbage collection of terminated processes and unused definitions, and associativity and commutativity of parallel composition.

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

Definition 4. 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 $f(\vec{x})$ is defined, then $N(\vec{x}) \rightarrow^* \mathfrak{q} \triangleright_{\Gamma f(\vec{x})} \mathbf{0}$;
- if $f(\vec{x})$ is not defined, then $N(\vec{x}) \not\rightarrow^* \mathbf{0}$.

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

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

We first discuss the rules for projecting the behaviour of a single process \mathfrak{p} , a partial function $\llbracket C \rrbracket_{\mathfrak{p}}$ defined by the rules in Figure 1. 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 $\mathfrak{p}.e \rightarrow \mathfrak{q}$, we project a send action for the sender \mathfrak{p} , a receive action for the receiver \mathfrak{q} , or just the continuation otherwise. The rule for selection is similar. The rules for projecting recursive definitions and calls assume 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 [5]).

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. Merging is a homomorphic binary operator; for all terms but branchings it requires isomorphism, e.g.: $\mathfrak{q}!(e); B \sqcup \mathfrak{q}!(e); B' = \mathfrak{q}!(e); (B \sqcup B')$. Branching terms can have unmergeable continuations, as long as

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

Fig. 1. Minimal Choreographies, Behaviour Projection.

they are guarded by distinct labels. In this case, merge returns a larger branching including all options (merging branches with the same label):

$$\begin{aligned}
p\&\{l_i : B_i\}_{i \in J} \sqcup p\&\{l_i : B'_i\}_{i \in K} &= \\
p\&(\{l_i : (B_i \sqcup B'_i)\}_{i \in J \cap K} \cup \{l_i : B_i\}_{i \in J \setminus K} \cup \{l_i : B'_i\}_{i \in K \setminus J}) &
\end{aligned}$$

Merging explains the role of selections in CC, common in choreography models [2, 5, 6, 16, 12, 25]. Recall the choreographies from Remark 1. In the first one, 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). The second one is projectable, and the behaviour of r is $\llbracket C \rrbracket_r = p\&\{L : p?; \mathbf{0}, R : p!(c); \mathbf{0}\}$.

Definition 5. 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 : $\llbracket C, \sigma \rrbracket = \prod_{p \in \text{pn}(C)} p \triangleright_{\sigma(p)} \llbracket C \rrbracket_p$.

Since the σ s are 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*.

EPP guarantees the following operational correspondence.

Theorem 3. Let C be a projectable choreography. Then, for all σ :

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

Soundness: If $\llbracket C, \sigma \rrbracket \rightarrow N$, then $C, \sigma \rightarrow C', \sigma'$ for some σ' , with $\llbracket C', \sigma' \rrbracket \prec N$.

The *pruning relation* \prec [5, 6] deletes branches introduced by merging when no longer needed; $N \succ N'$ means $N' \prec N$. Pruning does not alter the behaviour of a network: eliminated branches are never selected, as shown in [5, 18, 12]. As a consequence of Theorems 1 and 3, choreography projections never deadlock.

Theorem 4. 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).

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 6. 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 $\text{if } p \stackrel{\leftarrow}{=} q \text{ then } C_1 \text{ else } C_2$ 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 if $p \stackrel{\leftarrow}{=} q$ then C_1 else C_2 in C is replaced with:

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

From the definitions of Amend, EPP and the semantics of CC, we get:

Lemma 2. 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 σ , $C, \sigma \rightarrow^* C', \sigma'$ iff $\text{Amend}(C), \sigma \rightarrow^* \text{Amend}(C'), \sigma'$.

Example 2. Applying Amend to the first choreography in Remark 1 yields the second choreography in the same remark. 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 $\text{Amend}(C) = C$.

4 Turing completeness of CC and SP

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 [17], 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 CC is Turing complete regardless of projectability, since any unprojectable choreography can be amended to one 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).

Our proof is in line with other traditional proofs of computational completeness [11, 17, 27], where data and programs are distinct. This differs from other proofs of similar results for, e.g., π -calculus [26] and λ -calculus [1], 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).

Partial Recursive Functions. Our definition of the class of partial recursive functions \mathcal{R} is slightly simplified, but equivalent to, that in [17], where it is shown to be the class of computable functions. \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 $z < y$.

Encoding Partial Recursive Functions in CC. All functions in \mathcal{R} can be implemented in CC, in the sense of Definition 1. Given $f : \mathbb{N}^n \rightarrow \mathbb{N}$, we denote its implementation by $\llbracket f \rrbracket_{\tilde{p} \rightarrow \mathfrak{q}}$, 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 ($\llbracket f \rrbracket_{\ell}^{\tilde{p} \rightarrow \mathfrak{q}}$). 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 $A_1 \mathfrak{;} \dots \mathfrak{;} A_n$. We omit the selections needed for projectability, as they can be inferred by amendment; we will discuss this aspect formally later.

The encoding of the base cases is straightforward.

$$\llbracket Z \rrbracket_{\ell}^{\tilde{p} \rightarrow \mathfrak{q}} = p.\varepsilon \rightarrow \mathfrak{q} \quad \llbracket S \rrbracket_{\ell}^{\tilde{p} \rightarrow \mathfrak{q}} = p.(s \cdot c) \rightarrow \mathfrak{q} \quad \llbracket P_m^n \rrbracket_{\ell}^{\tilde{p} \rightarrow \mathfrak{q}} = p_m.c \rightarrow \mathfrak{q}$$

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

$$\llbracket h \rrbracket_{\ell}^{\tilde{p} \rightarrow \mathfrak{q}} = \left\{ \llbracket g_i \rrbracket_{\ell_i}^{\tilde{p} \rightarrow r'_i} \right\}_{i=1}^k \mathfrak{;} \llbracket f \rrbracket_{\ell_{k+1}}^{r'_1, \dots, r'_k \rightarrow \mathfrak{q}}$$

where $r'_i = r_{\ell+i-1}$, $\ell_1 = \ell + k$ and $\ell_{i+1} = \ell_i + \pi(g_i)$. Each auxiliary process 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 $\mathfrak{;} \llbracket g_{i+1} \rrbracket$ is substituted for the (unique) exit point of $\llbracket g_i \rrbracket$, and $\llbracket f \rrbracket$ is substituted for the exit point of $\llbracket g_k \rrbracket$. The resulting choreography also has only one exit point (that of $\llbracket f \rrbracket$). 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 1 for brevity:

$$\begin{aligned} \llbracket h \rrbracket_{\ell}^{p_0, \dots, p_n \mapsto q} = & \text{def } T = \text{if } (r_c \stackrel{\leftarrow}{=} p_0) \text{ then } (q'.c \rightarrow q; \mathbf{0}) \\ & \text{else } \llbracket g \rrbracket_{\ell_g}^{r_c, q', p_1, \dots, p_n \mapsto r_t} \ ; \ r_t.c \rightarrow q'; \text{INC}(r_c, r_t) \ ; \ T \\ & \text{in } \llbracket f \rrbracket_{\ell_f}^{p_1, \dots, p_n \mapsto q'} \ ; \ r_t.\varepsilon \rightarrow r_c; \ T \end{aligned}$$

where $q' = r_{\ell}$, $r_c = r_{\ell+1}$, $r_t = r_{\ell+2}$, $\ell_f = \ell + 3$ and $\ell_g = \ell_f + \pi(f)$. Process r_c is a counter, q' stores intermediate results, and r_t is temporary storage; T checks the value of r_c and either outputs the result or recurs. Note that $\llbracket h \rrbracket$ has only one exit point (after the communication from r to q), as the exit points of $\llbracket f \rrbracket$ and $\llbracket g \rrbracket$ 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 r_c and compute successive values of f , stored in 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.

$$\begin{aligned} \llbracket h \rrbracket_{\ell}^{p_1, \dots, p_{n+1} \mapsto q} = & \text{def } T = \llbracket f \rrbracket_{\ell_f}^{p_1, \dots, p_n, r_c \mapsto q'} \ ; \ r_c.\varepsilon \rightarrow r_z; \\ & \text{if } (r_z \stackrel{\leftarrow}{=} q') \text{ then } (r_c.c \rightarrow q; \mathbf{0}) \text{ else } (\text{INC}(r_c, r_z) \ ; \ T) \\ & \text{in } r_z.\varepsilon \rightarrow r_c; \ T \end{aligned}$$

where $q' = r_{\ell}$, $r_c = r_{\ell+1}$, $r_z = r_{\ell+2}$, $\ell_f = \ell + 3$ and $\ell_g = \ell_f + \pi(f)$. In this case, the whole if-then-else is inserted at the exit point of $\llbracket f \rrbracket$; the only exit point of this choreography is again after communicating the result to q .

Definition 7. Let $f \in \mathcal{R}$. The encoding of f in CC is $\llbracket f \rrbracket^{\tilde{p} \mapsto q} = \llbracket f \rrbracket_0^{\tilde{p} \mapsto q}$.

Main Results. We prove that our construction is sound by induction.

Theorem 5. If $f : \mathbb{N}^n \rightarrow \mathbb{N}$ and $f \in \mathcal{R}$, then, for every k , $\llbracket f \rrbracket_k^{\tilde{p} \mapsto q}$ implements f with input processes $\tilde{p} = p_1, \dots, p_n$ and output process q .

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 $\llbracket f \rrbracket$ to an equivalent projectable choreography. Then SP^{CC} is Turing complete by Theorems 3 and 5.

Corollary 1. Every partial recursive function is implementable in SP^{CC} .

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, $\llbracket \cdot \rrbracket$ does not fully achieve this, as $\llbracket g_1 \rrbracket, \dots, \llbracket g_k \rrbracket$ share the processes containing the input. We define a modified variant $\{\!\!\{ \cdot \}\!\!\}$ of $\llbracket \cdot \rrbracket$ where, for $h = C(f, g_1, \dots, g_k)$, $\{\!\!\{ h \}\!\!\}_{\ell}^{\tilde{p} \mapsto q}$ is

$$\{p_j.c \rightarrow p_j^i\}_{1 \leq i \leq k, 1 \leq j \leq n} \ ; \ \left\{ \{\!\!\{ g_i \}\!\!\}_{\ell_i}^{\tilde{p}^i \mapsto r_i'} \right\}_{i=1}^k \ ; \ \{\!\!\{ f \}\!\!\}_{\ell_{k+1}}^{r_1', \dots, r_k' \mapsto q}$$

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 $\{g_i\}_{\ell_i}^{\tilde{p}^i \rightarrow r'_i}$ run in parallel in $\{h\}^{\tilde{p} \rightarrow q}$.*

This parallelism is preserved by EPP, through Theorem 3.

5 Minimality in Choreography Languages

We now discuss our choice of primitives for CC, showing it to be a good candidate core language for choreographic programming. We analyse each primitive of CC. Recall that Turing completeness of CC is a pre-requisite for the Turing completeness of choreography projections. In many cases, simplifying CC yields a decidable termination problem (thus breaking Turing completeness). We discuss these cases first, and then proceed to a discussion on label selection.

Minimality in CC. Removing or simplifying the following primitives makes termination decidable.

- Exit point – $\mathbf{0}$: without it, no choreography terminates.
- Value communication – $p.e \rightarrow q$: without it, values of processes cannot be changed, and termination becomes decidable. The syntax of expressions is also minimal: ε (zero) is the only terminal; without c values become statically defined, while without s no new values can be computed; in either case, termination is decidable.
- Recursion – $\text{def } X = C_2 \text{ in } C_1$ and X : without it, all choreographies trivially terminate. The terms are minimal: they support only tail recursion and definitions are not parameterised.

Theorem 7. *Let C be a choreography with no conditionals. Then, termination of C is decidable and independent of the initial state.*

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

Label selection. The argument for including label selections in CC is of a different nature. As the construction in § 4 shows, selections are not needed for implementing computable functions in CC; they are used only for obtaining projectable choreographies, via amendment. We now show that we can encode selections introduced by amendment using the other primitives of CC, thereby eliminating the need for them from a purely computational point of view.

We denote by Minimal Choreographies (MC) the fragment of CC that does not contain label selections. We can therefore view amendment as a function from

$$\begin{aligned}
\langle \mathbf{0} \rangle &= \mathbf{0} & \langle \mathbf{p}.e \rightarrow \mathbf{q}; C \rangle &= \mathbf{p}.e \rightarrow \mathbf{q}; \langle C \rangle & \langle \text{def } X = C_2 \text{ in } C_1 \rangle &= \text{def } X = \langle C_2 \rangle \text{ in } \langle C_1 \rangle \\
\langle X \rangle &= X & \langle \text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } C_1 \text{ else } C_2 \rangle &= \text{if } \mathbf{p} \stackrel{\leq}{=} \mathbf{q} \text{ then } \langle C_1, C_2 \rangle_1 \text{ else } \langle C_1, C_2 \rangle_2
\end{aligned}$$

$$\begin{aligned}
\langle C_1, C_2 \rangle &= \langle \langle C_1 \rangle, \langle C_2 \rangle \rangle \text{ if } C_1 \text{ and } C_2 \text{ do not begin with a selection} \\
\langle \mathbf{p} \rightarrow \mathbf{q}[L]; C_1, \mathbf{p} \rightarrow \mathbf{q}[R]; C_2 \rangle &= \\
&\left\langle \mathbf{q}.c \rightarrow \mathbf{q}^\bullet; \mathbf{p}.\varepsilon \rightarrow \mathbf{q}; \text{if } \mathbf{q} \stackrel{\leq}{=} \mathbf{z} \text{ then } \mathbf{q}^\bullet.c \rightarrow \mathbf{q}; \langle C_1, C_2 \rangle_1 \text{ else } \mathbf{q}^\bullet.c \rightarrow \mathbf{q}; \langle C_1, C_2 \rangle_2, \right. \\
&\quad \left. \mathbf{q}.c \rightarrow \mathbf{q}^\bullet; \mathbf{p}.sc \rightarrow \mathbf{q}; \text{if } \mathbf{q} \stackrel{\leq}{=} \mathbf{z} \text{ then } \mathbf{q}^\bullet.c \rightarrow \mathbf{q}; \langle C_1, C_2 \rangle_1 \text{ else } \mathbf{q}^\bullet.c \rightarrow \mathbf{q}; \langle C_1, C_2 \rangle_2 \right\rangle
\end{aligned}$$

Fig. 2. Elimination of selections from amended choreographies.

MC into the subset of projectable CC choreographies. 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 (see Definition 6). The fragment of CC obtained by amending choreographies in MC is thus:

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

Term $S(\mathbf{p}, \tilde{r}, l, C)$ denotes a series of selections of label l from \mathbf{p} to all processes in the list \tilde{r} . Formally, $S(\mathbf{p}, \emptyset, l, C) = C$ and $S(\mathbf{p}, \mathbf{r} :: \tilde{r}, l, C) = \mathbf{p} \rightarrow \mathbf{r}[l]; S(\mathbf{p}, \tilde{r}, l, C)$.

Definition 8. Let C be obtained by amending a choreography in MC. The encoding $\langle C \rangle^+$ of C in MC uses processes $\mathbf{p}, \mathbf{p}^\bullet$ for each $\mathbf{p} \in \text{pn}(C)$ and a special process \mathbf{z} , and is defined as $\langle C \rangle^+ = \mathbf{p}.\varepsilon \rightarrow \mathbf{z}; \langle C \rangle$, with $\langle C \rangle$ defined in Figure 2.

The definition of $\langle C \rangle$ exploits the structure of amended choreographies, where selections are always paired at the top of the two branches of conditionals. It is immediate that $|\text{pn}(\langle C \rangle^+)| = 2|\text{pn}(C)| + 1$ (the extra process is \mathbf{z}). Let $|C|$ be the size of the syntax tree of C . Then, $|\langle C \rangle^+| \leq 2^{|C|}$, and in the worst case we get exponential growth. However, EPP collapses all branches of conditionals, hence projections do not grow exponentially: $|\llbracket \langle C \rangle^+ \rrbracket_{\mathbf{q}^\bullet}| \leq |\llbracket \langle C \rangle^+ \rrbracket_{\mathbf{q}}| \leq 3|\llbracket C \rrbracket_{\mathbf{q}}|$ for every $\mathbf{q} \in \text{pn}(C)$.

Theorem 9. For every choreography C in MC, $\llbracket \langle \text{Amend}(C) \rangle \rrbracket$ is defined.

It is straightforward to prove that C and $\langle \text{Amend}(C) \rangle$ behave exactly in the same way when we only observe communications between the original processes – except that label selections are replaced by regular messages.

Lemma 3. 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 $\sigma^+(\mathbf{z}) = \varepsilon$, then $\llbracket \text{Amend}(C) \rrbracket, \sigma^+ \rightarrow^* \llbracket C' \rrbracket, \sigma'^+$ for some σ'^+ similarly related to σ' . Conversely, if $\llbracket \text{Amend}(C) \rrbracket, \sigma^+ \rightarrow C', \sigma'$, then $C, \sigma \rightarrow C'', \sigma''$ where $C', \sigma' \rightarrow^* \llbracket \text{Amend}(C'') \rrbracket, \sigma''^+$.

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 $\text{SP}^{\text{MC}} = \{\llbracket C, \sigma \rrbracket \mid C \text{ in MC and } \llbracket C, \sigma \rrbracket \text{ is defined}\}$ of projections of minimal choreographies is also Turing complete.

Corollary 3. *Every partial recursive function is implementable in SP^{MC} .*

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 a simpler one:

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

The advantages of eliminating selections are thus 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 [2, 5, 6, 16, 12, 25], therefore we believe that a core model such as CC should have them (in addition to the drawback we mentioned).

Our results suggest the viability of a particular implementation strategy for choreographic programming. Programmers could write choreographies without label selections, and then our results could be used to translate these choreographies to process implementations in a simple language that does not include label communications, thus simplifying the target language. 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, this implementation 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 [10]. 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.

6 CC and other languages

CC is representative of the body of previous work on choreographic programming, where choreographies are used for implementations, for example [5, 6, 8, 24, 12, 28]. All the primitives of CC 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. For the model in [6], we formalise this result in [9]. Below, we discuss the significance of our results for the cited languages.

Differently from CC, other choreography languages typically use channel-based communications (as in the π -calculus [26]). Communications via process references as in CC can be encoded by assigning a dedicated channel to each pair of processes [9]. For example, the calculus in [6], which we refer to as Channel Choreographies (ChC), features an EPP targeting the session-based π -calculus [2]. ChC is a fully-fledged calculus aimed at real-world application that has been implemented as a choreographic programming framework (the Chor language [8]). Our formal translation from CC to ChC (given in [9]) 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, multiparty sessions. While useful in practice, these primitives come at the cost of making the formal treatment of ChC very technically involved. In particular, ChC (and 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. Using our encoding from CC to ChC, we can repeat the argument in § 4 to characterise a fragment of the session-based π -calculus from [2] that contains only deadlock-free terms and is Turing complete. ChC has also been translated to the Jolie programming language [14, 23], whence our reasoning also applies to the latter and, in general, to service-oriented languages based on message correlation.

The language WS-CDL from W3C [28] and the formal models inspired by it (e.g., [5]) 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 [12], which adds higher-order features to choreographies in terms of runtime adaptation. Finally, the language of compositional choreographies presented in [24] 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.

7 Related Work and Discussion

Register Machines. The computational primitives in CC recall those of the Unlimited Register Machine (URM) [11], 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.

Multiparty Sessions and Types. The communication primitives in CC recall those of protocols for multiparty sessions, e.g., in Multiparty Session Types (MPST) [16] and conversation types [4]. These protocols are not meant for computation, as in choreographic programming (and CC); rather, they are types used to verify that sessions (e.g., π -calculus channels) are used accordingly to their respective protocol specifications. For such formalisms, we know of a strong characterisation result: a variant of MPST corresponds to communicating finite state machines [3] that respect the property of multiparty compatibility [13]. To the best of our knowledge, this is the first work studying the expressivity of choreographic programming (choreographies for implementations).

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 [28]. 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 [6, 7, 24]: 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). 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 in [19] for a simple language with finite traces (thus not Turing complete). Our definition is different, since it uses merging for the first time.

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.

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

1. H. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. North Holland, 2nd edition, 1984.
2. L. Bettini, M. Coppo, L. D’Antoni, M. De Luca, M. Dezani-Ciancaglini, and N. Yoshida. Global progress in dynamically interleaved multiparty sessions. In F. van Breugel and M. Chechik, editors, *CONCUR*, volume 5201 of *LNCS*, pages 418–433. Springer, 2008.
3. D. Brand and P. Zafropulo. On communicating finite-state machines. *J. ACM*, 30(2):323–342, April 1983.
4. L. Caires and H. Torres Vieira. Conversation types. *Theor. Comput. Sci.*, 411(51-52):4399–4440, 2010.
5. M. Carbone, K. Honda, and N. Yoshida. Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.*, 34(2):8, 2012.
6. M. Carbone and F. Montesi. Deadlock-freedom-by-design: multiparty asynchronous global programming. In R. Giacobazzi and R. Cousot, editors, *POPL*, pages 263–274. ACM, 2013.
7. M. Carbone, F. Montesi, and C. Schürmann. Choreographies, logically. In P. Baldan and D. Gorla, editors, *CONCUR*, volume 8704 of *LNCS*, pages 47–62. Springer, 2014.
8. Chor. Programming Language. <http://www.chor-lang.org/>.
9. L. Cruz-Filipe and F. Montesi. Choreographies, computationally. *CoRR*, abs/1510.03271, 2015.

10. Luís Cruz-Filipe and Fabrizio Montesi. Choreographies, divided and conquered. *CoRR*, abs/1602.03729, 2016. Submitted for publication.
11. N.J. Cutland. *Computability: an Introduction to Recursive Function Theory*. Cambridge University Press, 1980.
12. M. Dalla Preda, M. Gabbrielli, S. Giallorenzo, I. Lanese, and J. Mauro. Dynamic choreographies – safe runtime updates of distributed applications. In T. Holvoet and M. Viroli, editors, *COORDINATION*, volume 9037 of *LNCS*, pages 67–82. Springer, 2015.
13. P.-M. Deniélou and N. Yoshida. Multiparty compatibility in communicating automata: Characterisation and synthesis of global session types. In F.V. Fomin, R. Freivalds, M.Z. Kwiatkowska, and D. Peleg, editors, *ICALP (II)*, volume 7966 of *LNCS*, pages 174–186. Springer, 2013.
14. M. Gabbrielli, S. Giallorenzo, and F. Montesi. Applied choreographies. *CoRR*, abs/1510.03637, 2015.
15. K. Honda, V. Vasconcelos, and M. Kubo. Language primitives and type disciplines for structured communication-based programming. In C. Hankin, editor, *ESOP*, volume 1381 of *LNCS*, pages 122–138. Springer, 1998.
16. K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. In G.C. Necula and P. Wadler, editors, *POPL*, pages 273–284. ACM, 2008.
17. S.C. Kleene. *Introduction to Metamathematics*. North-Holland Publishing Co., 1952.
18. I. Lanese, C. Guidi, F. Montesi, and G. Zavattaro. Bridging the gap between interaction- and process-oriented choreographies. In A. Cerone and S. Gruner, editors, *SEFM*, pages 323–332. IEEE, 2008.
19. I. Lanese, F. Montesi, and G. Zavattaro. Amending choreographies. In A. Ravara and J. Silva, editors, *WWV 2013*, volume 123 of *EPTCS*, pages 34–48, 2013.
20. Tanakorn Leesatapornwongsa, Jeffrey F. Lukman, Shan Lu, and Haryadi S. Gunawi. TaxDC: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In *ASPLOS*, pages 517–530. ACM, 2016.
21. Shan Lu, Soyeon Park, Eunsoo Seo, and Yuanyuan Zhou. Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. In *ASPLOS*, pages 329–339. ACM, 2008.
22. F. Montesi. *Choreographic Programming*. Ph.D. thesis, IT University of Copenhagen, 2013. http://fabriziomontesi.com/files/choreographic_programming.pdf.
23. F. Montesi, C. Guidi, and G. Zavattaro. Service-oriented programming with Jolie. In A. Bouguettaya, Q.Z. Sheng, and F. Daniel, editors, *Web Services Foundations*, pages 81–107. Springer, 2014.
24. F. Montesi and N. Yoshida. Compositional choreographies. In P.R. D’Argenio and H.C. Melgratti, editors, *CONCUR*, volume 8052 of *LNCS*, pages 425–439. Springer, 2013.
25. Z. Qiu, X. Zhao, C. Cai, and H. Yang. Towards the theoretical foundation of choreography. In C.L. Williamson, M.E. Zurko, P.F. Patel-Schneider, and P.J. Shenoy, editors, *WWW*, pages 973–982. ACM, 2007.
26. D. Sangiorgi and D. Walker. *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
27. A.M. Turing. Computability and λ -definability. *J. Symb. Log.*, 2(4):153–163, 1937.
28. W3C WS-CDL Working Group. Web services choreography description language version 1.0. <http://www.w3.org/TR/2004/TR/2004/WD-ws-cdl-10-20040427/>, 2004.