

Certifying Choreography Compilation^{*}

Luís Cruz-Filipe^[0000–0002–7866–7484], Fabrizio Montesi^[0000–0003–4666–901X],
and Marco Peressotti^[0000–0002–0243–0480]

Department of Mathematics and Computer Science, University of Southern Denmark
Campusvej 55, 5230 Odense M, Denmark
{lcfilipe,fmontesi,peressotti}@imada.sdu.dk

Abstract. Choreographic programming is a paradigm for developing concurrent and distributed systems, where programs are *choreographies* that define, from a global viewpoint, the computations and interactions that communicating processes should enact. *Choreography compilation* translates choreographies into the local definitions of process behaviours, given as terms in a process calculus.

Proving choreography compilation correct is challenging and error-prone, because it requires relating languages in different paradigms (global interactions vs local actions) and dealing with a combinatorial explosion of proof cases. We present the first certified program for choreography compilation for a nontrivial choreographic language supporting recursion.

Keywords: Choreographic Programming · Formalisation · Compilation.

1 Introduction

Choreographic programming is an emerging programming paradigm where the desired communication behaviour of a system of communicating processes can be defined from a global viewpoint in programs known as *choreographies* [25]. Then, a provably-correct compiler can automatically generate executable code for each process, with the guarantee that executing these processes together implements the communications prescribed in the choreography [5,7]. The theory of such compilers is typically called EndPoint Projection (EPP).

Choreographies are inspired by the “Alice and Bob” notation for security protocols [27]. The key idea is to have a linguistic primitive for a communication from a participant to another: statement `Alice.e → Bob.x` reads “Alice evaluates expression `e` and sends the result to Bob, who stores it in variable `x`”. This syntax has two main advantages. First, the desired communications are syntactically manifest in a choreography, which makes choreographic programming suitable for making interaction protocols precise. Second, it disallows writing mismatched send and receive actions, so code generated from a choreography enjoys progress (the system never gets stuck) [7].

The potential of choreographic programming has motivated the study of choreographic languages and EPP definitions for different applications, including

^{*} Work partially supported by Villum Fonden, grant no. 29518.

self-adaptive systems [15], information flow [21], system integration [16], parallel algorithms [10], cyber-physical systems [17,23,24], and security protocols [17].

EPP involves three elements: the source choreography language, the target language, and the compiler. The interplay between these components, where a single instruction at the choreographic level might be implemented by multiple instructions in the target language, makes the theory of choreographic programming error-prone: for even simpler approaches, like abstract choreographies without computation, it has been recently discovered that a few key results published in peer-reviewed articles do not hold and their theories required adjustments [28], raising concerns about the soundness of these methods.

In this article, we present a certified program for EPP, which translates terms of a Turing complete choreographic language into terms of a distributed process calculus. Our main result is the formalisation of the hallmark result of choreographic programming, the “EPP Theorem”: an operational correspondence between choreographies and their endpoint projections. This is the first time that this result has been formalised in a theorem prover, increasing our confidence in the methodology of choreographies.

Structure and Related Work. Our formalisation is developed in Coq [4], and we assume some familiarity with it. We start from a previous formalisation [13] of a choreographic language (Core Choreographies [12]), which we recap in Sect. 2. This formalisation only deals with the choreographic language and its properties; in particular, the target calculus for EPP is not defined therein.

In Sect. 3, we define our target language: a distributed process calculus inspired by the informal presentation in [26]. The calculus has communication primitives that recall those commonly used for implementing choreography languages, e.g., as found in (multiparty) session types [20]. In Sect. 4, we define *merging* [5]: a partial operator that addresses the standard problem (for choreographies) of checking that each process implementation eventually agrees on the choice between alternative behaviours in protocols [3,5]. Building on merging, in Sect. 5, we define EPP. Then, in Sect. 6, we explore *pruning* [5]: a preorder induced by merging that plays a key role in the EPP Theorem, proved in Sect. 7.

Choreographies are used in industry for the specification and definition of web services and business processes [19,29]. These languages feature recursion or loops, which are not present in the only other formalisation work on choreographies that we are aware of [18]. We have validated the theory of EPP from [26], and made explicit several properties that are typically only implicitly assumed. Our results show that the ideas developed by researchers on choreographies, like merging, can be relied upon for languages of practical appeal.

2 Background

We use the choreographic language of [26], which is inspired by Core Choreographies (CC) [12]. So far, this is the only Turing complete choreographic language that has been formalised [13]. In this section, we recap this formalisation, which

our work builds upon. We refer the reader to [13] for a discussion of the design options both behind the choreographic language and the formalisation. Some of these are relevant for our development, and we explain them when needed.

CC is designed to model communication protocols involving several participants, called processes, each equipped with memory cells, identified by variables. Communications are of two kinds: *value communications*, where a process evaluates an expression using the values stored in its memory and sends the result to a (distinct) process; and *label selections*, where a process selects from different behaviours available at another process by means of an identifier (the label). Selections are used to communicate local choices made by a process to other processes. Recursive and infinite behaviour is achieved by defining procedures.

Syntax. The formalisation of CC is parametric on the types of processes (Pid , ranged over by p), variables (Var , ranged over by x), expressions (Expr , ranged over by e), values stored in memory (Value , ranged over by v), Boolean expressions (BExpr , ranged over by b) and procedure names (RecVar , ranged over by X). Equality on these types must be decidable. Labels, ranged over by l , are either **left** or **right**, which is common in choreographies and session types [6,8,30].

The syntax of choreographies is defined by the following BNF grammar.

```

eta ::= p.e -> q.x | p -> q[l]
C ::= eta; C | If p.b Then Ct Else Ce | Call X | RT_Call X ps C | End
    
```

An **eta** is a communication action, where $p.e \rightarrow q.x$ is a value communication¹ and $p \rightarrow q[l]$ is a label selection. **Eta** is the type of all communication actions.

Choreographies are ranged over by C . A choreography **eta; C**, can execute a communication **eta** and continue as C . A conditional **If p.b Then Ct Else Ce** evaluates the Boolean expression b in the memory of process p and continues as Ct or Ce , according to whether b evaluates to **true** or **false**. Choreography **Call X** is a call to a procedure X . Term **RT_Call X ps C** is a *runtime term*, discussed below. Term **End** is the terminated choreography. We write **Choreography** for the type of all choreographies.

This grammar is implemented as a Coq inductive type, e.g., **eta; C** stands for **Interaction eta C**, where **Interaction** : **Eta** \rightarrow **Choreography** \rightarrow **Choreography** is a constructor of type **Choreography**.

Example 1 (Distributed Authentication). The choreography **C1** below describes a multiparty authentication scenario where an identity provider **ip** authenticates a client **c**, to server **s** (we name subterms for later use).

```

C1 := c.credentials -> ip.x; If ip.(check x) Then C1t Else C1e
C1t := ip -> s[left]; ip -> c[left]; s.token -> c.t; End
C1e := ip -> s[right]; ip -> c[right]; End
    
```

C1 starts with c communicating its credentials, stored in the local variable **credentials**, to **ip**, which stores them in x . Then, **ip** checks if the credentials are

¹ For readability, we use notations closer to the usual mathematical ones than in formalisation, where they are slightly different due to Coq's restrictions on overloading.

valid or not by evaluating the local expression `check x`, and signals the result to `s` and `c` by selecting `left` when the credentials are valid (`C1t`) and `right` otherwise (`C1e`). In the first case, the server communicates a token (stored in its local variable `token`) to the client, otherwise the choreography ends.

Because the guard of the conditional is evaluated by `ip`, only this process knows which branch of the choreography to execute. This is an instance of the *knowledge of choice* problem. Label selections propagate this information to processes whose behaviour depends on this choice (in this case, both `s` and `c`). ◁

A *program* is a pair $(\text{Procedures } P, \text{Main } P) : \text{DefSet} * \text{Choreography}$. Elements of type `DefSet` (set of procedure definitions) map each `X` to a pair containing the processes used in `X` (`Vars P X`) and the choreography to be executed (`Procs P X`). These procedure definitions can then be called from each other and from `Main P` (using term `Call X`), allowing for the definition of recursive behaviour.

Executing procedure calls generates runtime terms. Choreography `Call X` can reduce by a process `p` entering `X`. It becomes `RT_Call X ps C`, where `C` is the definition of `X` and `ps` is the list of processes used in `X` (other than `p`). This term can then either reduce by another process entering the procedure (and being removed from `ps`), or by executing some action in `C` that does not involve processes in `ps` (and `C` is updated). When the last process in `ps` calls `X`, `RT_Call X ps C` reduces to `C`. The runtime term `RT_Call X ps C` is not meant to be used when writing a choreography, and in particular it should not occur in procedure definitions. A choreography that does not include any such term is called *initial*.

There are three kinds of restrictions when writing choreographies.

- (i) Intended use of choreographies. Interactions must have distinct processes (no self-communication), e.g., `p.e -> p.x` is disallowed.
- (ii) Intended use of runtime terms. All choreographies `Procs P X` must be initial. `Main P` may include subterms `RT_Call X ps C`, but `ps` must be nonempty and include only process names that occur in `Vars P X`.
- (iii) Design choices in the formalisation. Informally, `Vars X` contains the processes that are used in `Procs X`. This introduces constraints that are encapsulated in the definition of well-formed program.

The constraints in the last category are particularly relevant for the proof of the EPP Theorem, so we discuss them in that context in Sect. 7.

Example 2. Let `Defs:DefSet` map `FileTransfer` to the pair consisting of the process list `c :: s :: nil` and the following choreography.

```
s.(file, check) -> c.x;                (* send file and check data *)
If c.(crc(fst(x)) == snd(x))           (* cyclic redundancy check *)
  Then c -> s[left]; End                (* file received correctly, end *)
  Else c -> s[right]; Call FileTransfer (* errors detected, retry *)
```

`FileTransfer` describes a file transfer protocol between a server `s` and a client `c` using Cyclic Redundancy Checks (`crc`) to detect errors from a noisy channel. ◁

Semantics. The semantics of CC is defined as a labelled transition system using inductive types. It uses a *state*, which is a function mapping process variables to their values: `State:Pid → Var → Value`.

The semantics is structured in three layers. The first layer specifies transitions with the following relation, parameterised on a set of procedure definitions.

$$\begin{aligned} \text{CCC_To} : \text{DefSet} &\rightarrow \text{Choreography} \rightarrow \text{State} \\ &\rightarrow \text{RichLabel} \rightarrow \text{Choreography} \rightarrow \text{State} \rightarrow \text{Prop} \end{aligned}$$

Rules in this set include that $p.e \rightarrow q.x;C$ in state s transitions to C in state s' , where s' coincides with s except that $s'.q.x$ now stores the value obtained by evaluating e at p in state s .² `RichLabel` includes information about the executed term – the above communication is labelled `R_Com p v q x`. There are also rules for out-of-order execution: interactions involving distinct processes can be executed in any order, reflecting concurrency. For example, $p.e \rightarrow q.x; r \rightarrow s[\text{left}]$ can execute as a communication from p to q followed by a label selection from r to s , but also as the latter label selection followed by the former communication.

The second layer raises transitions to the level of programs, abstracting from unobservable details. It is defined by a single rule: if $\text{CCC_To Defs } C \text{ s t } C' \text{ s}'$, then $(\{| \text{Defs}; C |, s) \xrightarrow{[\text{forget } t]} (\{| \text{Defs}; C' |, s')$. Here, $\{| \text{Defs}; C |$ denotes the program built from `Defs` and C , and `forget` removes unobservable details from transition labels (e.g., `forget (R_Com p v q x)` is `L_Com p v q`). Labels for conditionals and procedure calls all simplify to `L_Tau p`, denoting an internal action at p . The third layer defines the reflexive and transitive closure of program transitions.

Important properties of CC formalised in [13] include deadlock-freedom by design (any program P such that `Main P` \neq `End` reduces), confluence (if P can execute two different sequences of actions, then the resulting programs can always reduce to a common program and state), and Turing completeness.

Example 3. Consider the program $\{|D; C1|\}$ where $C1$ is the choreography in Example 1 and $D:\text{DefSet}$ is arbitrary (there are no recursive calls in $C1$).

$$\{|D; C1|\}, \text{st1} \xrightarrow{[\text{L_Com } c \text{ ip } v1]} (\{|D; \text{If ip.}(\text{check } x) \text{ Then } C1t \text{ Else } C1e|\}, \text{st2})$$

where $v1$ is the evaluation of e at c in st1 and st2 updates the value of ip 's variable x accordingly. If `check x` is `true` at ip in st2 , then it continues as follows.

$$\begin{aligned} \{|D; \text{If ip.}(\text{check } x) \text{ Then } C1t \text{ Else } C1e|\}, \text{st2} &\xrightarrow{[\text{L_Tau ip}]} \{|D; C1t|\}, \text{st2} \\ &\xrightarrow{[\text{L_Sel ip } s \text{ left}; \text{L_Sel ip } c \text{ left}]}^* \{|D; s.\text{token} \rightarrow c.t; \text{End}\}, \text{st2} \\ &\xrightarrow{[\text{L_Com } s \text{ c } v2]} \{|D; \text{End}\}, \text{st3} \end{aligned}$$

where $v2$ is the evaluation of `token` at s in st2 and st3 updates st2 accordingly. Otherwise, it continues as follows.

$$\begin{aligned} \{|D; \text{If ip.}(\text{check } x) \text{ Then } C1t \text{ Else } C1e|\}, \text{st2} &\xrightarrow{[\text{L_Tau ip}]} \{|D; C1e|\}, \text{st2} \\ &\xrightarrow{[\text{L_Sel ip } s \text{ right}; \text{L_Sel ip } c \text{ right}]}^* \{|D; \text{End}\}, \text{st2} \end{aligned}$$

In compound transitions, the actions in the label are executed in order. ◁

Example 4. Let `Defs` as in Example 2 and C be the body of `FileTransfer`. Consider the program $\{| \text{Defs}; \text{Call FileTransfer} |\}$. The processes in the procedure `FileTransfer` can start a call in any order as exemplified by the transitions below.

² The semantics of CC only requires extensional equality of states, which is why s' is quantified over rather than directly defined from s .

$$\{\text{Defs}; \text{Call FileTransfer}\}, \text{st}) \xrightarrow{[\text{L_Tau } c]} (\{\text{Defs}; \text{RT_Call FileTransfer } s::\text{nil } C\}, \text{st}) \xrightarrow{[\text{L_Tau } s]} (\{\text{Defs}; C\}, \text{st})$$

$$\{\text{Defs}; \text{Call FileTransfer}\}, \text{st}) \xrightarrow{[\text{L_Tau } s]} (\{\text{Defs}; \text{RT_Call FileTransfer } c::\text{nil } C\}, \text{st}) \xrightarrow{[\text{L_Tau } c]} (\{\text{Defs}; C\}, \text{st})$$

The state st is immaterial. ◁

3 Stateful Processes

Our first contribution is formalising SP (Stateful Processes), the process calculus for implementing CC [26], which is inspired by the process calculi in [11,12].³

3.1 Syntax

SP is defined as a Coq functor with the same parameters as CC. Its syntax contains three ingredients: behaviours, defining the actions performed by individual processes; networks, consisting of several processes running in parallel; and programs, defining a set of procedure definitions that all processes can use.

Behaviours. Behaviours are sequences of local actions – counterparts to the terms that can be written in CC – defined by the following BNF grammar.

$$B ::= \text{End} \mid p!e;B \mid p?x;B \mid p\oplus 1;B \mid p\&mB1//mBr \mid \text{If } b \text{ Then } Bt \text{ Else } Be \mid \text{Call } X$$

This grammar is again formalised as a Coq inductive type, called **Behaviour**. Terms **End**, **If b Then Bt Else Be** and **Call X** are as in CC. Value communications are split into $p!e; B$, which evaluates e , sends the result to p , and continues as B , and by $p?x; B$, which receives a value from p , stores it at x , and continues as B .⁴

Label selections are divided into $p\oplus 1; B$ – sending the label 1 to p and continuing as B – and $p \& mB1 // mBr$, where one of $mB1$ or mBr is chosen according to the label selected by p . Both $mB1$ and mBr have type `option Behaviour` (either `None` or `Some B`, where B is a behaviour): a process does not need to offer behaviours corresponding to all possible labels. Informally, branching terms are partial functions from labels to behaviours; we capitalise on the fact that there are only two labels to simplify the formalisation.

Many results about **Behaviours** are proved by structural induction, requiring inspection of subterms of branching terms. The induction principle automatically generated by Coq is not strong enough for this, and our formalisation includes a stronger result that we use in later proofs.

³ The choreography language CC is also inspired by the works [11,12], but formalising it in Coq benefitted substantially from adopting a labelled transition system semantics. This is discussed extensively in [13]. In this work, we made similar changes to the process calculus not only for similar reasons, but also to keep a close correspondence with the choreography language.

⁴ Processes communicate by name. In practice, names can be either process identifiers (cf. actors), network addresses, or session correlation data.

Networks. Networks have type $\text{Pid} \rightarrow \text{Behaviour}$. We define extensional equality of networks, $N == N'$, and show that it is an equivalence relation. To model practice, where networks are written as finite parallel compositions of behaviours, we introduce some constructions: $N|N'$ is the parallel composition of N and N' ; $p[B]$ is the network mapping p to B and all other processes to End ; and $N \sim p$ denotes network N where p is now mapped to End . The formalisation includes a number of lemmas about extensional equality, for example that updating the behaviours of two processes yields the same result independent of the order of the updates.

Programs. Finally, a **Program** is a pair $(\text{Procs } P, \text{Net } P) : \text{DefSetB} * \text{Network}$, where $\text{DefSetB} = \text{RecVar} \rightarrow \text{Behaviour}$ maps procedure names to Behaviours.

Programs, networks and behaviours should satisfy well-formedness properties similar to those of CC. However, these properties are automatically ensured when networks are automatically generated from choreographies, so we do not discuss them here. They are included in the formalisation.

Example 5. Consider the network $N = c[Bc] \mid s[Bs] \mid ip[Bip]$ composed by the behaviours below.

```
Bc := ip!credentials; ip & Some (s?t; End) // Some End
Bs := ip & Some (c!token; End) // Some End
```

```
Bip := c?x; Bip'
Bip' := If (check x) Then (s⊕left; c⊕left; End) Else (s⊕right; c⊕right; End)
```

This network implements the choreography in Example 1. ◁

3.2 Semantics

The semantics of SP is defined by a labelled transition system. Transitions for communications match dual actions in two processes, while conditionals and procedure calls simply run locally. We report the representative cases – the complete definition can be found in the source formalisation [14]. For readability, we first present them in the standard rule notation.

$$\frac{v := \text{eval_on_state } e \ s \ p \quad N \ p = q!e; B \quad N \ q = p?x; B' \quad N' == N \sim p \sim q \mid p[B] \mid q[B'] \quad s' == \text{update } s \ q \ x \ v}{(\{\text{Defs}; N\}, s) \xrightarrow{[L_Com \ p \ v \ q]} (\{\text{Defs}; N'\}, s')} \text{S_Com}$$

$$\frac{N \ p = q\oplus\text{left}; B \quad N \ q = p \ \& \ \text{Some } B1 \ // \ Br \quad N' == N \sim p \sim q \mid p[B] \mid q[B1] \quad s == s'}{(\{\text{Defs}; N\}, s) \xrightarrow{[L_Sel \ p \ q \ \text{left}]} (\{\text{Defs}; N'\}, s')} \text{S_LSel}$$

$$\frac{N \ p = \text{If } b \ \text{Then } B1 \ \text{Else } B2 \quad \text{beval_on_state } b \ s \ p = \text{true} \quad N' == N \sim p \mid p[B1] \quad s == s'}{(\{\text{Defs}; N\}, s) \xrightarrow{[L_Tau \ p]} (\{\text{Defs}; N'\}, s')} \text{S_Then}$$

$$\frac{N \text{ p} = \text{Call } X \quad N' == N \sim\sim \text{p} \mid \text{p}[\text{Defs } X] \quad \text{s} == \text{s}'}{(\{\text{Defs}; N\}, \text{s}) \xrightarrow{[\text{L_Tau } X \text{ p}]} (\{\text{Defs}; N'\}, \text{s}')} \text{S_Call}$$

Functions `eval_on_state` and `beval_on_state` evaluate a (Boolean) expression locally, given a state, while `update s q x v` updates state `s` by changing the value of `q`'s variable `x` to `v`. We write also `==` for extensional equality of states.

As for choreographies, the formalisation of these rules is done in two steps. The first step defines a transition relation parameterised on the set of procedure definitions, and includes richer transition labels (necessary for doing case analysis on transitions). This relation is defined as an inductive type, whose first defining clause is shown below.⁵

```

Inductive SP_To (Defs : DefSetB) :
  Network → State → RichLabel → Network → State → Prop :=
| S_Com N p e B q x B' N' s s' : let v := (eval_on_state e s p) in
  N p = (q!e; B) → N q = (p?x; B') → N' == N ~\~ p ~\~ q \mid p[B] \mid q[B'] →
  s' == (update s q x v) → SP_To Defs N s (R_Com p v q x) N' s'   (...)

```

This relation is then lifted to configurations just as for CC: if `SP_To Defs N s t N' s'`, then $(\{\text{Defs}; N\}, \text{s}) \xrightarrow{[\text{forget } t]} (\{\text{Defs}; N'\}, \text{s}')$. Closure under reflexivity and transitivity (with similar notation) is again defined as for choreographies.

Example 6. We illustrate the possible transitions of the network from Example 5. We abbreviate the behaviours of processes that do not change in a reduction to ... to make it more clear what parts of the network are changed. Furthermore, we omit trailing `Ends` in `Behaviours`.

The network starts by performing the transition

$$\begin{aligned} & \{\text{D}; \text{c}[\text{Bc}] \mid \text{s}[\text{Bs}] \mid \text{ip}[\text{Bip}]\}, \text{st1} \xrightarrow{[\text{L_Com } \text{c ip v1}]} \\ & \{\text{D}; \text{c}[\text{ip \& Some (s?t) // Some End}] \mid \text{s}[\dots] \mid \text{ip}[\text{Bip}']\}, \text{st2} \end{aligned}$$

where `v1` and `st2` are as in Example 3. If `eval_on_state (check x) st2 ip` is `true`, it continues as follows

$$\begin{aligned} & \{\text{D}; \text{c}[\text{ip \& Some (s?t) // Some End}] \mid \text{s}[\text{Bs}] \mid \text{ip}[\text{Bip}'], \text{st2}\} \\ & \xrightarrow{[\text{L_Tau ip}]} \{\text{D}; \text{c}[\dots] \mid \text{s}[\dots] \mid \text{ip}[\text{s}\oplus\text{left};\text{c}\oplus\text{left}], \text{st2}\} \\ & \xrightarrow{[\text{L_Sel ip s left}]} \{\text{D}; \text{c}[\dots] \mid \text{s}[\text{c!token}] \mid \text{ip}[\text{c}\oplus\text{left}], \text{st2}\} \\ & \xrightarrow{[\text{L_Sel ip c left}]} \{\text{D}; \text{c}[\text{s?t}] \mid \text{s}[\dots] \mid \text{ip}[\text{End}], \text{st2}\} \\ & \xrightarrow{[\text{L_Com s c v2}]} \{\text{D}; \text{c}[\text{End}] \mid \text{s}[\text{End}] \mid \text{ip}[\text{End}], \text{st3}\} \end{aligned}$$

where `v2` and `st3` are again as in Example 3. Otherwise, it continues as follows.

$$\begin{aligned} & \{\text{D}; \text{c}[\text{ip \& Some (s?t) // Some End}] \mid \text{s}[\text{Bs}] \mid \text{ip}[\text{Bip}'], \text{st2}\} \\ & \xrightarrow{[\text{L_Tau ip}]} \{\text{D}; \text{c}[\dots] \mid \text{s}[\dots] \mid \text{ip}[\text{s}\oplus\text{right};\text{c}\oplus\text{right}], \text{st2}\} \\ & \xrightarrow{[\text{L_Sel ip s right}]} \{\text{D}; \text{c}[\dots] \mid \text{s}[\text{End}] \mid \text{ip}[\text{c}\oplus\text{right}], \text{st2}\} \\ & \xrightarrow{[\text{L_Sel ip c right}]} \{\text{D}; \text{c}[\text{End}] \mid \text{s}[\text{End}] \mid \text{ip}[\text{End}], \text{st2}\} \end{aligned}$$

⁵ Coq's type inference mechanism allows us to omit types of most universally quantified variables and parameters. We abuse this possibility to lighten the presentation.

The labels in these reductions are exactly as in Example 3. ◁

Transitions are compatible with network equality and state equivalence.

Lemma SPP_To_eq : $\forall P \ s1 \ t1 \ P' \ s2 \ s1' \ s2',$
 $s1 == s1' \rightarrow s2 == s2' \rightarrow (P, s1) \xrightarrow{[t1]} (P', s2) \rightarrow (P, s1') \xrightarrow{[t1]} (P', s2').$

Lemma SPP_To_Network_eq : $\forall P1 \ P1' \ P2 \ s \ s' \ t1, (\text{Net } P1 == \text{Net } P1') \rightarrow$
 $(\text{Procs } P1 = \text{Procs } P1') \rightarrow (P1, s) \xrightarrow{[t1]} (P2, s') \rightarrow (P1', s) \xrightarrow{[t1]} (P2, s').$

These results are instrumental in some of the later proofs, and are proven for the three levels of reductions. The following, related, result is also important.

Lemma SPP_To_Defs_stable : $\forall \text{Defs}' \ N \ N' \ t1 \ s \ s',$
 $\{|\text{Defs } N, s|\} \xrightarrow{[t1]} \{|\text{Defs}' \ N', s'|\} \rightarrow \text{Defs} = \text{Defs}'.$

We also prove that transitions are completely determined by the label.

Lemma SPP_To_deterministic : $\forall P \ s \ t1 \ P' \ s' \ P'' \ s'',$
 $(P, s) \xrightarrow{[t1]} (P', s') \rightarrow (P, s) \xrightarrow{[t1]} (P'', s'') \rightarrow$
 $(\text{Net } P' == \text{Net } P'') \wedge \text{Procs } P' = \text{Procs } P'' \wedge (s' == s'').$

Finally, we show that the semantics of SP is confluent. Although this is not required for our main theorem, it is a nice result that confirms our expectations.

The formalisation of SP consists of 37 definitions, 80 lemmas, and 2000 lines.

4 Merging

Intuitively, process implementations are generated from choreographies recursively, by projecting each action in the choreography to the corresponding process action – for example, a value communication $p.e \rightarrow q.x$ should be projected as a send action $q!e$ at p , as a receive action $p?x$ at q , and not projected to any other processes. However, this causes a problem with conditionals. Projecting a choreography **If** $p.b$ **Then** Ct **Else** Ce for any process other than p , say q , requires combining the projections obtained for Ct and Ce , such that q can “react” to whichever choice p will make. This combination is called *merging* [5].

Merge is typically defined as a partial function mapping pairs of behaviours to behaviours that returns a behaviour combining all possible executions of the two input behaviours (if possible). For SP, two behaviours can be merged if they are structurally similar, with the possible exception of branching terms: we can merge branchings that offer options on distinct labels. For example, merging $p \& (\text{Some } B) // \text{None}$ with $p \& \text{None} // (\text{Some } B')$ yields $p \& (\text{Some } B) // (\text{Some } B')$, allowing **If** $p??b$ **Then** $(p \rightarrow q[\text{left}]; q.e \rightarrow p; \text{End})$ **Else** $(p \rightarrow q[\text{right}]; \text{End})$ to be projected for q as $p \& \text{Some } (p!e; \text{End}) // \text{Some } \text{End}$.

Since functions in Coq are total, formalising merging poses a problem. Furthermore, assigning type $\text{Behaviour} \rightarrow \text{Behaviour} \rightarrow \text{option } \text{Behaviour}$ to merging causes ambiguity, since branching terms have subterms of type $\text{option } \text{Behaviour}$.⁶

⁶ Essentially because it is not possible to distinguish if the behaviour assigned to a label is **None** because it was not defined, or because a recursive call to merge failed.

Instead, we define a type `XBehaviour` of extended behaviours, defined as `Behaviour` with an extra constructor `XUndefined : XBehaviour`. Thus, an `XBehaviour` is a `Behaviour` that may contain `XUndefined` subterms.

The connection between `Behaviour` and `XBehaviour` is established by means of two functions: `inject : Behaviour → XBehaviour`, which isomorphically injects each constructor of `Behaviour` into the corresponding one in `XBehaviour`, and `collapse : XBehaviour → Behaviour` that maps all `XBehaviours` with `XUndefined` as a subterm to `XUndefined`. The most relevant properties of these functions are:

Lemma `inject_elim` : $\forall B, \exists B', \text{inject } B = B' \wedge B' \neq \text{XUndefined}$.

Lemma `collapse_inject` : $\forall B, \text{collapse } (\text{inject } B) = \text{inject } B$.

Lemma `collapse_char''` : $\forall B, \text{collapse } B = \text{XUndefined} \rightarrow \forall B', B \neq \text{inject } B'$.

Lemma `collapse_∃` : $\forall B, \text{collapse } B \neq \text{XUndefined} \rightarrow \exists B', B = \text{inject } B'$.

Using this type, we first define `XMerge` on `XBehaviours` as below, where we report only representative cases. (`Pid_dec` and `Expr_dec` are lemmas stating decidability of equality on `Pid` and `Expr`, allowing us to do case analysis.)

```

Fixpoint Xmerge (B1 B2:XBehaviour) : XBehaviour := match B1, B2 with
| XEnd, XEnd ⇒ XEnd
| XSend p e B, XSend p' e' B' ⇒ if Pid_dec p p' && Expr_dec e e'
  then match Xmerge B B' with XUndefined ⇒ XUndefined
    | _ ⇒ XSend p e (Xmerge B B') end
  else XUndefined
| XBranching p B1 Br, XBranching p' B1' Br' ⇒ if Pid_dec p p'
  then let BL := match B1 with None ⇒ B1' | Some B ⇒
    match B1' with None ⇒ B1 | Some B' ⇒ Some (Xmerge B B') end end
    in let BR := match Br with None ⇒ Br' | Some B ⇒
    match Br' with None ⇒ Br | Some B' ⇒ Some (Xmerge B B') end end
    in match BL, BR with Some XUndefined, _ ⇒ XUndefined
      | _, Some XUndefined ⇒ XUndefined | _, _ ⇒ XBranching p BL BR end
  else XUndefined

```

Using `XMerge` we can straightforwardly define merging.

Definition `merge B1 B2` := `Xmerge (inject B1) (inject B2)`.

We show `merge` to be idempotent, commutative and associative. This last proof illustrates the major challenge of this stage of the formalisation: it requires a triple induction with 512 cases, of which 84 cannot be solved automatically. These had to be divided in further subcases of different levels of complexity. The final case, when all behaviours are branching terms, requires six (!) nested inductions to generate the 64 possible combinations of defined/undefined branches, which had to be done by hand. The total proof is over 500 lines.

The largest set of lemmas about `merge` deals with inversion results, such as:

Lemma `merge_inv_Send` : $\forall B B' p e X, \text{merge } B B' = \text{XSend } p e X \rightarrow$
 $\exists B1 B1', B = p ! e; B1 \wedge B' = p ! e; B1' \wedge \text{merge } B1 B1' = X$.

The similar result for branching terms is much more complex, since there are several cases for each branch (if it is `None`, then both `B` and `B'` must have `None` in the corresponding branch, otherwise it may be from `B`, from `B'`, or a merge of

both), and its proof suffers from the same problems as the proof of associativity (thankfully, not to such a dramatic level). Automation works better here, and the effect of the large number of subcases is mostly felt in the time required by the `auto` tactic. Still, the formalisation of merging consists of 6 definitions, 43 lemmas, and 2550 lines – giving an average proof length of over 50 lines.

5 EndPoint Projection

The next step is defining EndPoint Projection (EPP): a partial function that maps programs in CC to programs in SP. The target instance of SP has the same parameters as CC, except that the set of procedure names is `RecVar * Pid` – each procedure is implemented from the point of view of each process in it.

Partiality of EPP stems from the problem of choreography realisability [3]. A choreography such as `If p.b Then (q.e -> r.x) Else End` cannot be implemented without additional communications between `p`, `q` and `r`, since the latter processes need to know the result of the conditional to decide whether to communicate (see also Example 1). We say that this choreography is not *projectable* [12].

We define EPP in several layers. First, we define a function `bproj : DefSet → Choreography → Pid → XBehaviour` projecting the behaviour of a single process. Intuitively, `bproj Defs C p` attempts to construct `p`'s behaviour as specified by `C`; the parameter `Defs` is used for procedure calls, whose projections depend on whether `p` participates in the procedure. Returning an `XBehaviour` instead of an `option Behaviour` gives information about where exactly merging fails (the location of `XUndefined` subterms), which can be used for debugging, providing information to programmers, or for automatic repair of choreographies [2,12].

We show some illustrative cases of the definition of `bproj`.

```

Fixpoint bproj (Defs:DefSet) (C:Choreography) (r:Pid) : XBehaviour := match C with
| p.e -> q.x; C' =>
    if Pid_dec p r then XSend q e (bproj Defs C' r)
    else if Pid_dec q r then XRecv p x (bproj Defs C' r)
    else bproj Defs C' r
| p -> q[left]; C' =>
    if Pid_dec p r then XSel q left (bproj Defs C' r)
    else if Pid_dec q r then XBranching p (Some (bproj Defs C' r)) None
    else bproj Defs C' r
| If p.b Then C1 Else C2 =>
    if Pid_dec p r then XCond b (bproj Defs C1 r) (bproj Defs C2 r)
    else Xmerge (bproj Defs C1 r) (bproj Defs C2 r)
| CCBase.Call X => if In_dec P.eq_dec r (fst (Defs X)) then XCall (X,r) else XEnd
(...)
    
```

The next step is generating projections for all relevant processes. We take the set of processes as a parameter, and collapse all individual projections.

```

Definition epp_list (Defs:DefSet) (C:Choreography) (ps:list Pid)
: list (Pid * XBehaviour) := map (fun p => (p, collapse (bproj Defs C p))) ps.
    
```

A choreography C is projectable wrt Defs and ps if $\text{epp_list Defs } C \text{ ps}$ does not contain $X\text{Undefined}$, and Defs:DefSet is projectable wrt a set of procedure names Xs if $\text{snd (Defs } X)$ is projectable wrt Defs and $\text{fst (Defs } X)$ for each X in Xs . Projectability of programs is a bit more involved, and we present its Coq formalisation before discussing it.

Definition $\text{projectable } Xs \text{ ps } P :=$
 $\text{projectable_C (Procedures } P) \text{ ps (Main } P) \wedge \text{projectable_D } Xs \text{ (Procedures } P) \wedge$
 $(\forall p, \text{In } p \text{ (CCC_pn (Main } P) (\text{fun } _ \Rightarrow \text{nil}))} \rightarrow \text{In } p \text{ ps}) \wedge$
 $(\forall p X, \text{In } X \text{ } Xs \rightarrow \text{In } p \text{ (fst (Procedures } P \text{ } X))} \rightarrow \text{In } p \text{ ps}) \wedge$
 $(\forall p X, \text{In } X \text{ } Xs \rightarrow \text{In } p \text{ (CCC_pn (snd (Procedures } P \text{ } X)) (\text{fun } _ \Rightarrow \text{nil}))} \rightarrow \text{In } p \text{ ps}).$

The first two conditions simply state that $\text{Main } P$ and $\text{Procedures } P$ are projectable wrt the appropriate parameters. The remaining conditions state that the sets ps and Xs include all processes used in P and all procedures needed to execute P . (These sets are not necessarily computable, since Xs is not required to be finite. However, in practice, these parameters are known – so it is easier to include them in the definition.) Function CCC_pn returns the set of processes used in a choreography, given the sets of processes each procedure call is supposed to use.

We now define compilation of projectable choreographies (epp_C), projectable sets of procedure definitions (epp_D), and projectable programs (epp). These definitions depend on proof terms whose structure needs to be explored, and are done interactively; afterwards, we show them to be independent of the proof terms, and that they work as expected. We give a few examples.

Lemma $\text{epp_C_wd} : \forall \text{Defs } C \text{ ps } H \text{ H}', (\text{epp_C Defs ps } C \text{ H}) == (\text{epp_C Defs ps } C \text{ H}')$.

Lemma $\text{epp_C_Com_p} : \forall \text{Defs ps } C \text{ p } e \text{ q } x \text{ HC HC}', \text{In } p \text{ ps} \rightarrow$
 $\text{epp_C Defs ps (p.e} \rightarrow \text{q.x; } C) \text{ HC } p = \text{q!e; epp_C Defs ps } C \text{ HC}' \text{ p.}$

Lemma $\text{epp_C_Cond_r} : \forall \text{Defs ps } p \text{ b } C1 \text{ } C2 \text{ HC HC1 HC2 } r, p \neq r \rightarrow$
 $\text{inject (epp_C Defs ps (If } p.\text{b Then } C1 \text{ Else } C2) \text{ HC } r)$
 $= \text{merge (epp_C Defs ps } C1 \text{ HC1 } r) (\text{epp_C Defs ps } C2 \text{ HC2 } r).$

Projectability of C does not imply projectability of choreographies that C can transition to. This is due to the way runtime terms are projected: $\text{RT_Call } X \text{ ps } C'$ is projected as a call to (X, p) if p is in ps , and as the projection of C' otherwise. Our definition of projectability allows in principle for C to be unprojectable for a process in ps , which would make it unprojectable after transition. That this situation does not arise is a consequence of the intended usage of runtime terms: initially C' is obtained from the definition of a procedure, and ps is the set of processes used in this procedure. Afterwards ps only shrinks, while C' may change due to execution of actions outside ps . We capture these conditions in the notion of strong projectability, whose representative case is:

Fixpoint $\text{str_projectable Defs (C:Choreography) (r:Pid) : Prop} :=$
 $\text{match } C \text{ with } | \text{RT_Call } X \text{ ps } C \Rightarrow \text{str_projectable Defs } C \text{ } r \wedge$
 $(\forall p, \text{In } p \text{ ps} \rightarrow \text{In } p \text{ (fst (Defs } X))}$
 $\wedge X\text{more_branches (bproj Defs (snd (Defs } X)) \text{ } p) (\text{bproj Defs } C \text{ } p)) \quad (\dots)$

The relation $X\text{more_branches}$, explained in the next section, is a semantic characterisation of how the projection of $\text{bproj Defs snd (Defs } X) \text{ } p$ may change due to execution of actions not involving p in $\text{snd (Defs } X)$.

Projectability and strong projectability coincide for initial choreographies. Furthermore, we state and prove lemmas that show that strong projectability of C imply strong projectability of any choreography that C can transition to.

6 Pruning

The key ingredient for our correspondence result is a relation on behaviours usually called *pruning* [5,7]. Pruning relates two behaviours that differ only in that one offers more options in branching terms than the other; we formalise this relation with the more suggestive name of `more_branches` (in line with [26]), and we include some illustrative cases of its definition.

```

Inductive more_branches : Behaviour → Behaviour → Prop :=
| MB_Send p e B B' : more_branches B B' → more_branches (p ! e; B) (p ! e; B')
| MB_Branching_None_None p mB1 mBr :
  more_branches (p & mB1 // mBr) (p & None // None)
| MB_Branching_Some_Some p B1 B1' Br Br' :
  more_branches B1 B1' → more_branches Br Br' →
  more_branches (p & Some B1 // Some Br) (p & Some B1' // Some Br')      (...)
    
```

The need for pruning arises naturally when one considers what happens when a choreography C executes a conditional at a process p . In the continuation, only one of the branches is kept. However, no process other than p knows that this action has been executed; therefore, if the projection of C executes the corresponding action, both behaviours are still available for all processes other than p . Given how merging is defined, this means that these processes' behaviours may contain more branches than those of the projection of the choreography after reduction.

Pruning is naturally related to merging, as stated in the following lemmas.

Lemma `more_branches_char` : $\forall B B', \text{more_branches } B B' \leftrightarrow \text{merge } B B' = \text{inject } B$.

Lemma `merge_more_branches` :

$\forall B1 B2 B, \text{merge } B1 B2 = \text{inject } B \rightarrow \text{more_branches } B B1$.

Pruning is also reflexive and transitive. Further, if two behaviours have an upper bound according to pruning, then their merge is defined and is their lub.

Lemma `more_branches_merge` : $\forall B B1 B2, \text{more_branches } B B1 \rightarrow$

$\text{more_branches } B B2 \rightarrow \exists B', \text{merge } B1 B2 = \text{inject } B' \wedge \text{more_branches } B B'$.

Finally, two behaviours with fewer branches than two mergeable behaviours are themselves mergeable.

Lemma `more_branches_merge_extend` : $\forall B1 B2 B1' B2' B,$

$\text{more_branches } B1 B1' \rightarrow \text{more_branches } B2 B2' \rightarrow \text{merge } B1 B2 = \text{inject } B \rightarrow$

$\exists B', \text{merge } B1' B2' = \text{inject } B' \wedge \text{more_branches } B B'$.

These two results are key ingredients to the cases dealing with conditionals in the proof of the EPP Theorem. They require extensive case analysis (512 cases for the last lemma, of which 81 are not automatically solved by Coq's `inversion`

tactic even though most are contradictory). Analogous versions of some of these lemmas also need to be extended to $\mathbf{XBehaviours}$, which is straightforward.

Pruning extends pointwise to networks, which we denote as $\mathbf{N} \gg \mathbf{N}'$. The key result is that, due to how the semantics of SP is defined, pruning a network cannot add new transitions.

Lemma $\text{SP_To_more_branches_N} : \forall \text{Defs } \mathbf{N1} \ \mathbf{s} \ \mathbf{N2} \ \mathbf{s}' \ \text{Defs}' \ \mathbf{N1}' \ \mathbf{t1},$
 $\text{SP_To } \text{Defs } \mathbf{N1} \ \mathbf{s} \ \mathbf{t1} \ \mathbf{N2} \ \mathbf{s}' \rightarrow \mathbf{N1}' \gg \mathbf{N1} \rightarrow (\forall \ \mathbf{X}, \text{Defs } \mathbf{X} = \text{Defs}' \ \mathbf{X}) \rightarrow$
 $\exists \ \mathbf{N2}', \text{SP_To } \text{Defs}' \ \mathbf{N1}' \ \mathbf{s} \ \mathbf{t1} \ \mathbf{N2}' \ \mathbf{s}' \wedge \mathbf{N2}' \gg \mathbf{N2}.$

The reciprocal of this result only holds for choreography projections, and is proven after the definition of EPP.

The formalisation of pruning includes 3 definitions, 25 lemmas, and 950 lines. Some of these results are used for defining of EPP (previous section), but we delayed their presentation as its motivation is clearer after seeing those definitions.

Example 7. The choreography $\mathbf{C1}$ in Example 1 is projectable, yielding the network in Example 5: $\text{bproj } \text{Defs } \mathbf{C1} \ \mathbf{c} = \text{inject } \mathbf{Bc}$, $\text{bproj } \text{Defs } \mathbf{C1} \ \mathbf{s} = \text{inject } \mathbf{Bs}$, and $\text{bproj } \text{Defs } \mathbf{C1} \ \mathbf{ip} = \text{inject } \mathbf{Bip}$.

If we remove selections from $\mathbf{C1}$, the resulting choreography is not projectable on process \mathbf{c} : projecting the conditional requires merging the projections at \mathbf{c} of the two branches (now simply $\mathbf{s.token} \rightarrow \mathbf{c.t}; \text{End}$ and End), which fails since $\text{bproj } \text{Defs } \mathbf{C5} \ \mathbf{c} = \text{inject } (\mathbf{s}?t; \text{End})$ and $\text{bproj } \text{Defs } \text{End} \ \mathbf{c} = \text{inject } \text{End}$, but $\text{merge } (\mathbf{s}?t; \text{End}) \ \text{End} = \mathbf{XUndefined}$. Likewise, merging would fail for \mathbf{s} . \triangleleft

7 EPP Theorem

We now prove the operational correspondence between choreographies and their projections, in two directions: if a choreography can make a transition, then its projection can make the same transition; and if the projection of a choreography can make a transition, then so can the choreography. The results of the transitions are not directly related by projection, since choreography transitions may eliminate some branches in the projection; thus, establishing the correspondence for multi-step transitions requires some additional lemmas on pruning.

Preliminaries. Both directions of the correspondence depend on a number of results relating choreography transitions and their projections. These results follow a pattern: the results for communications state a precise correspondence; the ones for conditionals include pruning in their conclusions; and the ones for procedure calls require additional hypotheses on the set of procedure definitions.

Lemma $\text{CCC_To_bproj_Sel_p} : \forall \text{Defs } \mathbf{C} \ \mathbf{s} \ \mathbf{C}' \ \mathbf{s}' \ \mathbf{p} \ \mathbf{q} \ \mathbf{l}, \text{str_projectable } \text{Defs } \mathbf{C} \ \mathbf{p} \rightarrow$
 $\text{CCC_To } \text{Defs } \mathbf{C} \ \mathbf{s} \ (\text{CCBase.TL.R_Sel } \mathbf{p} \ \mathbf{q} \ \mathbf{l}) \ \mathbf{C}' \ \mathbf{s}' \rightarrow$
 $\exists \ \mathbf{Bp}, \text{bproj } \text{Defs } \mathbf{C} \ \mathbf{p} = \mathbf{XSel} \ \mathbf{q} \ \mathbf{l} \ \mathbf{Bp} \wedge \text{bproj } \text{Defs } \mathbf{C}' \ \mathbf{p} = \mathbf{Bp}.$

Lemma $\text{CCC_To_bproj_Call_p} : \forall \text{Defs } \mathbf{C} \ \mathbf{s} \ \mathbf{C}' \ \mathbf{s}' \ \mathbf{p} \ \mathbf{X} \ \mathbf{Xs}, \text{str_projectable } \text{Defs } \mathbf{C} \ \mathbf{p} \rightarrow$
 $(\forall \ \mathbf{Y}, \text{In } \mathbf{Y} \ \mathbf{Xs} \rightarrow \text{str_projectable } \text{Defs } (\text{snd } (\text{Defs } \mathbf{Y})) \ \mathbf{p}) \rightarrow$
 $(\forall \ \mathbf{Y}, \text{set_incl_pid } (\text{CCC_pn } (\text{snd } (\text{Defs } \mathbf{Y}))) \ (\text{fun } \mathbf{X} \Rightarrow \text{fst } (\text{Defs } \mathbf{X})))$

$$\begin{aligned}
 & (\text{fst } (\text{Defs } Y)) \rightarrow \\
 & \text{In } X \text{ Xs} \rightarrow \text{CCC_To Defs } C \text{ s } (\text{CCBase.TL.R_Call } X \text{ p}) \text{ C' s'} \rightarrow \\
 & \text{bproj Defs } C \text{ p} = \text{XCall } (X, \text{p}) \\
 & \wedge \text{Xmore_branches } (\text{bproj Defs } (\text{snd } (\text{Defs } X)) \text{ p}) (\text{bproj Defs } C' \text{ p}).
 \end{aligned}$$

These lemmas are simple to prove by induction on C . The tricky part is getting the hypotheses strong enough that the thesis holds, and weak enough that any well-formed program will satisfy them throughout its entire execution.

From these results, it follows that projectability is preserved by choreography reductions. This property is needed even to state the EPP Theorem, since we can only compute projections of projectable programs.

Lemma CCC_To_projectable : $\forall P \text{ Xs ps}$,
 $\text{Program_WF } X \text{ s } P \rightarrow \text{well_ann } P \rightarrow \text{projectable } X \text{ s } P \rightarrow$
 $(\forall \text{ p, In p ps} \rightarrow \text{str_projectable } (\text{Procedures } P) (\text{Main } P) \text{ p}) \rightarrow$
 $(\forall \text{ p, In p } (\text{CCC_pn } (\text{Main } P) (\text{Vars } P)) \rightarrow \text{In p ps}) \rightarrow$
 $(\forall \text{ p } X, \text{In } X \text{ Xs} \rightarrow \text{In p } (\text{Vars } P \text{ X}) \rightarrow \text{In p ps}) \rightarrow$
 $\forall \text{ s t1 } P' \text{ s'}, (P, \text{s}) \xrightarrow{[\text{t1}]} (P', \text{s}') \rightarrow \text{projectable } X \text{ s } P'.$

Some of the hypotheses from the previous lemmas are encapsulated in the first two conditions: well-formedness of P , which ensures that any runtime term $\text{RT_Call } X \text{ qs } C$ in $\text{Main } P$ only includes processes in qs that are declared to be used by X (this trivially holds if $\text{Main } P$ is initial, and is preserved throughout execution); and well-annotation of P , i.e., the processes used in any procedure are a subset of those it declares.⁷ The remaining hypotheses state, as before, that X and ps include all processes and procedures relevant for executing $\text{Main } P$.

Similarly, we prove that strong projectability is preserved by transitions.

Completeness. In the literature, completeness of EPP is proven by induction on the derivation of the transition performed by the choreography C . For each case, we look at how a transition for C can be derived, and show that the projection of C can make the same transition to a network with more branches than the projection of C' . The proof is lengthy, but poses no surprises.

Lemma EPP_Complete : $\forall P \text{ Xs ps, Program_WF } X \text{ s } P \rightarrow \text{well_ann } P \rightarrow \forall \text{ HP}$,
 $(\forall \text{ p, In p ps} \rightarrow \text{str_projectable } (\text{Procedures } P) (\text{Main } P) \text{ p}) \rightarrow$
 $(\forall \text{ p, In p } (\text{CCC_pn } (\text{Main } P) (\text{Vars } P)) \rightarrow \text{In p ps}) \rightarrow$
 $(\forall \text{ p } X, \text{In } X \text{ Xs} \rightarrow \text{In p } (\text{Vars } P \text{ X}) \rightarrow \text{In p ps}) \rightarrow$
 $\forall \text{ s t1 } P' \text{ s'}, (P, \text{s}) \xrightarrow{[\text{t1}]} (P', \text{s}') \rightarrow$
 $\exists \text{ N t1'}, (\text{epp } X \text{ s } P \text{ HP, s}) \xrightarrow{[\text{t1}']} (\text{N}, \text{s}')$
 $\wedge \text{Procs } \text{N} = \text{Procs } (\text{epp } X \text{ s } P \text{ HP}) \wedge \forall \text{ H, Net } \text{N} \gg \text{Net } (\text{epp } X \text{ s } P' \text{ H}).$

By combining with the earlier results on pruning, we immediately obtain the generalisation for multi-step transitions.

Soundness. Soundness is proven by case analysis on the transition made by the network, and then by induction on the choreography inside each case. For convenience, we split this proof in separate proofs, one for each transition. We

⁷ Equality is not necessary, and it would make this property harder to prove.

omit the statements of these lemmas, since they include a number of technical hypotheses (similar to those in e.g. `SP_To_bproj_Com`, which is used in the proof of the case of communication, but more complex). By contrast with completeness, all these lemmas are complex to prove: each case requires around 300 lines of Coq code. The proofs have similar structure, but are still different enough that adapting them can not be done mechanically.

The last ingredient is a lemma of practical interest on procedure names: each process only uses “its” copy of the original procedure names. This lemma is not only crucial in the proof of the next theorem, but also interesting in itself: it shows that the set of procedure definitions can be fully distributed among the processes with no duplications.

Lemma `SP_To_bproj_Call_name` : $\forall \text{Defs Defs}' \text{ps C HC s N}' \text{s}' \text{p X}$,
 $\text{SP_To Defs}' (\text{epp_C Defs ps C HC}) \text{s} (\text{R_Call X p}) \text{N}' \text{s}' \rightarrow$
 $\exists Y, X = (Y, \text{p}) \wedge X_Free Y \text{C}$.

All these ingredients are combined in the proof of soundness of EPP.

Lemma `EPP_Sound` : $\forall P Xs \text{ps}$, `Program_WF Xs P` \rightarrow `well_ann P` \rightarrow $\forall \text{HP}$,
 $(\forall \text{p}, \text{In p ps} \rightarrow \text{str_projectable (Procedures P) (Main P) p}) \rightarrow$
 $(\forall \text{p}, \text{In p (CCC_pn (Main P) (Vars P))} \rightarrow \text{In p ps}) \rightarrow$
 $(\forall \text{p X}, \text{In X Xs} \rightarrow \text{In p (Vars P X)} \rightarrow \text{In p ps}) \rightarrow$
 $\forall \text{s t1 N}' \text{s}'$, $(\text{epp Xs ps P HP, s}) \text{---[t1]---} (\text{N}', \text{s}') \rightarrow$
 $\exists P' \text{t1}'$, $(\text{P}, \text{s}) \text{---[t1]---} (\text{P}', \text{s}') \wedge \forall \text{H}, \text{Net N}' \gg \text{Net (epp Xs ps P' H)}$.

Generalising this result to multi-step transitions requires showing that pruning does not eliminate possible transitions of a network. This is in general not true, but it holds when the pruned network is the projection of a choreography.

Lemma `SP_To_more_branches_N_epp` : $\forall \text{Defs N1 s N2 s}' \text{t1 Defs}' \text{ps C HC}$,
 $\text{N1} \gg \text{epp_C Defs}' \text{ps C HC} \rightarrow \text{SP_To Defs N1 s t1 N2 s}' \rightarrow$
 $\exists \text{N2}'$, $\text{SP_To Defs (epp_C Defs}' \text{ps C HC}) \text{s t1 N2}' \wedge \text{N2} \gg \text{N2}'$.

The formalisation of EPP and the proof of the EPP theorem consists of 13 definitions, 110 lemmas, and 4960 lines of Coq code. The proof of the EPP Theorem and related lemmas make up for around 75% of this size.

8 Discussion and Conclusion

We have successfully formalised a translation from a Turing-complete choreographic language into a process calculus and proven its correctness in terms of an operational correspondence. This formalisation showed that the proof techniques used in the literature are correct, and identified only missing minor assumptions about runtime terms that trivially hold when these are used as intended.

To the best of our knowledge, this is the first time such a correspondence has been formalised for a full-fledged (Turing-complete) choreographic language. Comparable work includes a preliminary presentation on a certified compiler from choreographies to CakeML [18], which however deals only with finite behaviours [22]. In the related realm of multiparty session types (where choreogra-

phies do not include computation), a similar correspondence result has also been developed independently [9].

The complexity of the formalisation, combined with the similarities between several of the proofs, means that future extensions would benefit from exploiting semi-automatic generation of proof scripts.

Combining these results with those from [13] would yield a proof that SP is also Turing complete. Unfortunately, the choreographies used in the proof of Turing completeness in [13] are not projectable, but they can be made so automatically, by means of an amendment (repair) procedure [12]. In future work, we plan to formalise amendment in order to obtain this result.

References

1. 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, Lecture Notes in Computer Science, vol. 9688. Springer (2016)
2. Basu, S., Bultan, T.: Automated choreography repair. In: Stevens, P., Wasowski, A. (eds.) Procs. FASE. Lecture Notes in Computer Science, vol. 9633, pp. 13–30. Springer (2016). https://doi.org/10.1007/978-3-662-49665-7_2
3. Basu, S., Bultan, T., Ouederni, M.: Deciding choreography realizability. In: Field, J., Hicks, M. (eds.) Procs. POPL. pp. 191–202. ACM (2012). <https://doi.org/10.1145/2103656.2103680>
4. Bertot, Y., Castéran, P.: Interactive Theorem Proving and Program Development. Texts in Theoretical Computer Science, Springer (2004)
5. Carbone, M., Honda, K., Yoshida, N.: Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.* **34**(2), 8:1–8:78 (2012). <https://doi.org/10.1145/2220365.2220367>
6. Carbone, M., Lindley, S., Montesi, F., Schürmann, C., Wadler, P.: Coherence generalises duality: A logical explanation of multiparty session types. In: Desharnais, J., Jagadeesan, R. (eds.) Procs. CONCUR. LIPIcs, vol. 59, pp. 33:1–33:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016)
7. Carbone, M., Montesi, F.: Deadlock-freedom-by-design: multiparty asynchronous global programming. In: Giacobazzi, R., Cousot, R. (eds.) Procs. POPL. pp. 263–274. ACM (2013). <https://doi.org/10.1145/2429069.2429101>
8. Carbone, M., Montesi, F., Schürmann, C.: Choreographies, logically. *Distributed Comput.* **31**(1), 51–67 (2018). <https://doi.org/10.1007/s00446-017-0295-1>
9. Castro-Perez, D., Ferreira, F., Gheri, L., Yoshida, N.: Zooid: a DSL for certified multiparty computation: from mechanised metatheory to certified multiparty processes. In: Freund, S.N., Yahav, E. (eds.) Procs. PLDI. pp. 237–251. ACM (2021). <https://doi.org/10.1145/3453483.3454041>
10. Cruz-Filipe, L., Montesi, F.: Choreographies in practice. In: Albert and Lanese [1], pp. 114–123. https://doi.org/10.1007/978-3-319-39570-8_8
11. Cruz-Filipe, L., Montesi, F.: Procedural choreographic programming. In: Bouajjani, A., Silva, A. (eds.) Procs. FORTE. Lecture Notes in Computer Science, vol. 10321, pp. 92–107. Springer (2017). https://doi.org/10.1007/978-3-319-60225-7_7
12. Cruz-Filipe, L., Montesi, F.: A core model for choreographic programming. *Theor. Comput. Sci.* **802**, 38–66 (2020). <https://doi.org/10.1016/j.tcs.2019.07.005>

13. Cruz-Filipe, L., Montesi, F., Peressotti, M.: Formalising a Turing-complete choreographic language in Coq. In: Cohen, L., Kaliszzyk, C. (eds.) *Procs. ITP. LIPIcs*, vol. 193, pp. 15:1–15:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2021). <https://doi.org/10.4230/LIPIcs.ITP.2021.15>
14. Cruz-Filipe, L., Montesi, F., Peressotti, M.: A formalisation of a Turing-complete choreographic language in Coq. <https://doi.org/10.5281/zenodo.4548709>
15. Dalla Preda, M., Gabbrielli, M., Giallorenzo, S., Lanese, I., Mauro, J.: Dynamic choreographies: Theory and implementation. *Log. Methods Comput. Sci.* **13**(2) (2017). [https://doi.org/10.23638/LMCS-13\(2:1\)2017](https://doi.org/10.23638/LMCS-13(2:1)2017)
16. Giallorenzo, S., Lanese, I., Russo, D.: Chip: A choreographic integration process. In: Panetto, H., Debruyne, C., Proper, H.A., Ardagna, C.A., Roman, D., Meersman, R. (eds.) *Procs. OTM, part II. Lecture Notes in Computer Science*, vol. 11230, pp. 22–40. Springer (2018). https://doi.org/10.1007/978-3-030-02671-4_2
17. Giallorenzo, S., Montesi, F., Peressotti, M.: Choreographies as objects. *CoRR abs/2005.09520* (2020)
18. Gomez-Londono, A., Aman Pohjola, J.: Connecting choreography languages with verified stacks. In: *Procs. of the Nordic Workshop on Programming Theory*. pp. 31–33 (2018)
19. Object Management Group: Business Process Model and Notation. <http://www.omg.org/spec/BPMN/2.0/> (2011)
20. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. *J. ACM* **63**(1), 9 (2016). <https://doi.org/10.1145/2827695>, also: *POPL*, pages 273–284, 2008
21. Lluch-Lafuente, A., Nielson, F., Nielson, H.R.: Discretionary information flow control for interaction-oriented specifications. In: Martí-Oliet, N., Ölveczky, P.C., Talcott, C.L. (eds.) *Logic, Rewriting, and Concurrency. Lecture Notes in Computer Science*, vol. 9200, pp. 427–450. Springer (2015). https://doi.org/10.1007/978-3-319-23165-5_20
22. Londoño, A.G.: *Choreographies and cost semantics for reliable communicating systems* (2020)
23. López, H.A., Heussen, K.: Choreographing cyber-physical distributed control systems for the energy sector. In: Seffah, A., Penzenstadler, B., Alves, C., Peng, X. (eds.) *Procs. SAC*. pp. 437–443. ACM (2017). <https://doi.org/10.1145/3019612.3019656>
24. López, H.A., Nielson, F., Nielson, H.R.: Enforcing availability in failure-aware communicating systems. In: Albert and Lanese [1], pp. 195–211. https://doi.org/10.1007/978-3-319-39570-8_13
25. Montesi, F.: *Choreographic Programming*. Ph.D. Thesis, IT University of Copenhagen (2013)
26. Montesi, F.: *Introduction to choreographies* (2021), accepted for publication by Cambridge University Press
27. Needham, R.M., Schroeder, M.D.: Using encryption for authentication in large networks of computers. *Commun. ACM* **21**(12), 993–999 (1978). <https://doi.org/10.1145/359657.359659>
28. Scalas, A., Yoshida, N.: Less is more: multiparty session types revisited. *Proc. ACM Program. Lang.* **3**(POPL), 30:1–30:29 (2019). <https://doi.org/10.1145/3290343>
29. W3C: *WS Choreography Description Language*. <http://www.w3.org/TR/ws-cdl-10/> (2004)
30. Wadler, P.: Propositions as sessions. *J. Funct. Program.* **24**(2–3), 384–418 (2014). <https://doi.org/10.1017/S095679681400001X>