

Encoding Asynchrony in Choreographies

Luís Cruz-Filipe and Fabrizio Montesi*

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

ABSTRACT

Choreographies are widely used both for the specification and the programming of concurrent and distributed software architectures. Since many of such architectures use asynchronous communications, it is essential to understand how the behaviour described in a choreography can be correctly implemented in asynchronous settings. So far, this problem has been addressed by relying on additional technical machinery, such as ad-hoc syntactic terms, semantics, or equivalences. In this work, we show that such extensions are not needed for choreography languages that support primitives for process spawning and name mobility. Instead, we can just encode asynchronous communications in choreographies themselves, yielding a simpler approach.

CCS Concepts

•Theory of computation → Process calculi;

Keywords

Asynchrony; Choreography; Concurrency

1. INTRODUCTION

Programming concurrent and distributed systems is challenging, because it is difficult to program correctly the intended interactions among components executed concurrently (e.g., services). Empirical investigations of bugs in concurrent and distributed software [7, 8] reveal that most errors are due to: deadlocks; violations of atomicity intentions; or, violations of ordering intentions. The issue is particularly pressing in architectures where hundreds of components may interact via message passing, like microservices [5].

*Montesi was supported by CRC (Choreographies for Reliable and efficient Communication software), grant no. DFF-4005-00304 from the Danish Council for Independent Research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC'17, April 3-7, 2017, Marrakesh, Morocco

Copyright 2017 ACM 978-1-4503-4486-9/17/04...\$15.00

<http://dx.doi.org/xx.xxxx/xxxxxxx.xxxxxx>

To mitigate this problem, *choreographies* can be used as high-level formal specifications of the intended interactions among components [1, 2].

EXAMPLE 1. We use a choreography to define a scenario where a buyer, Alice (**a**), purchases a product from a seller (**s**) through her bank (**b**).

- | | |
|--------------------------------|--|
| 1. a.title → s ; | 4. if b $\stackrel{<}{=}$ a then |
| 2. s.price → a ; | 5. s.book → a ; |
| 3. s.price → b ; | 6. else 0 |

In Line 1, the term **a.title** → **s** denotes an interaction whereby **a** communicates the title of the book that Alice wishes to buy to **s**. The seller then sends the price of the book to both **a** and **b**. In Line 4, **a** sends the price she expects to pay to **b**, which confirms that it is the same amount requested by **s** (stored internally at **b**). If so, **s** sends the book to **a** (Line 5). Otherwise, the choreography terminates.

In addition to their clarity, choreographies enable new development methodologies. For example, in Choreographic Programming [9, 10], choreographies are compiled to compliant local implementations for the described components. In our example, the implementation inferred for Alice (**a**), would be: send the book title to **s**; receive the price from **s**; send the price to **b** for confirmation; await the success/failure notification from **b**; in case of success, receive the book from **s**.

In most software architectures, communications are asynchronous. Therefore, it is important to prove that the code generated by compiling a choreography implements it correctly under in such a setting. So far, such proofs have been developed by defining ad-hoc extensions to the syntax and semantics of the models used to represent choreographies or their compiled code [2, 6, 11].

In this paper, we show that choreography languages equipped with primitives for process spawning and name mobility are already powerful enough to capture asynchronous communications. The key idea is to use processes to represent messages in transit, allowing the sender to proceed immediately after having sent a message without having to synchronise with the receiver [12]. We present our result (sketch) as an endo-encoding in the new language of Dynamic Minimal Choreographies (DMC), an extension of the representative choreography calculus of Minimal Choreographies [4].

$$\begin{array}{c}
C ::= \mathbf{0} \mid \eta; C \mid \text{if } p \stackrel{\leftarrow}{=} q \text{ then } C_1 \text{ else } C_2 \mid \text{def } X(\bar{p}) = C_2 \text{ in } C_1 \mid X(\bar{p}) \qquad \eta ::= p.e \rightarrow q \mid p.r \rightarrow q \mid p \text{ starts } q \\
\hline
\frac{p \xrightarrow{G} q \quad e[\sigma(p)/*] \downarrow v}{G, p.e \rightarrow q; C, \sigma \rightarrow G, C, \sigma[q \mapsto v]} \text{ [C|Com]} \qquad \frac{i = 1 \text{ if } \sigma(p) = \sigma(q), \quad i = 2 \text{ otherwise}}{G, \text{if } p \stackrel{\leftarrow}{=} q \text{ then } C_1 \text{ else } C_2, \sigma \rightarrow G, C_i, \sigma} \text{ [C|Cond]} \\
\frac{}{G, p \text{ starts } q; C, \sigma \rightarrow G \cup \{p \leftrightarrow q\}, C, \sigma[q \mapsto \perp]} \text{ [C|Start]} \qquad \frac{p \xrightarrow{G} q \quad p \xrightarrow{G} r}{G, p.r \rightarrow q; C, \sigma \rightarrow G \cup \{q \rightarrow r\}, C, \sigma} \text{ [C|Intro]}
\end{array}$$

Figure 1: Dynamic Minimal Choreographies, Syntax and Semantics.

2. LANGUAGE MODEL

We introduce Dynamic Minimal Choreographies (DMC), an extension of the calculus from [4].

The syntax of DMC is given in Figure 1 (top), where C ranges over choreographies. Processes (p, q, \dots) run in parallel, and each process stores a value in a local memory cell that can be read with the expression $*$. Term $\eta; C$ is an interaction between two processes, read “the system may execute η and proceed as C ”. In a value communication $p.e \rightarrow q$, p sends its local evaluation of expression e (whose syntax we leave undefined) to q , which stores the received value. In term $p \text{ starts } q$, process p starts a new process q , whose name is known only by p . Names can be communicated via term $p.r \rightarrow q$. In a conditional $\text{if } p \stackrel{\leftarrow}{=} q \text{ then } C_1 \text{ else } C_2$, q sends its value to p , which checks if the received value is equal to its own; the choreography proceeds as C_1 , if that is the case, or as C_2 , otherwise. In all these actions, the two interacting processes must be different. Definitions and invocations of (parametric) recursive procedures (X) are standard. The term $\mathbf{0}$ is the terminated choreography.

In the semantics of DMC, we use a graph of connections G [3], keeping track of which pairs of processes are allowed to communicate. This graph is directed, and an edge from p to q in G (written $p \xrightarrow{G} q$) means that p knows the name of q . In order for an actual message to flow between p and q , both processes need to know each other, which we write as $p \xleftrightarrow{G} q$. The semantics for DMC uses reductions of the form $G, C, \sigma \rightarrow G', C', \sigma'$, where G and G' are the connection graphs before and after executing C , respectively, and the total state function σ maps each process name to its value (values are denoted v, w, \dots). The complete rules are given in Figure 1 (bottom), closed under a structural precongruence that allows for unfolding of procedure calls, garbage collection, and swapping of independent actions (see [4]).

In the premise of [C|Com] , $e[\sigma(p)/*]$ denotes replacing $*$ with $\sigma(p)$ in e . In the reductum, $\sigma[q \mapsto v]$ denotes the updated state function σ where q now maps to v . In [C|Start] , the fresh process q is assigned a default value \perp . We write $G \cup G'$ for the graph obtained by merging G with G' .

The main limitation of DMC is that its semantics is synchronous. Indeed, in a real-world scenario implementation of Example 1, we would expect s to proceed immediately to sending its message in Line 3 after having sent the one in Line 2, without waiting for a to receive the latter. Capturing

this kind of asynchronous behaviour is the main objective of our development in the remainder of this paper.

3. ASYNCHRONY IN DMC

The calculus of Minimal Choreographies (MC) from [4] is the fragment of DMC that does not include process spawning and name mobility. In this fragment, we can omit procedure parameters by assuming that all procedures take all processes as arguments. In this section, we focus on MC and show that any MC choreography can be encoded in DMC in such a way that communication becomes asynchronous. More precisely, we provide a mapping $\{\!\cdot\!\} : \text{MC} \rightarrow \text{DMC}$ such that every communication action $p.e \rightarrow q \in C \in \text{MC}$ becomes split into a send/receive pair in $\{\!\{C}\!\} \in \text{DMC}$, with the properties that: p can continue executing without waiting for q to receive its message (and even send further messages to q); and messages from p to q are delivered in the same order as they were originally sent.

Let C be a choreography in MC. In order to encode C in DMC, we use a function $M : \mathcal{P}^2 \rightarrow \mathbb{N}$, where $\mathcal{P} = \text{pn}(C)$ is the set of process names in C . Intuitively, $\{\!\{C}\!\}$ use a countable set of auxiliary processes $\{\text{pq}^i \mid p, q \in \mathcal{P}, i \in \mathbb{N}\}$, where pq^i holds the i th message from p to q .

First, we setup initial channels for communications between all processes occurring in C .

$$\{\!\{C}\!\} = \{p \text{ start } \text{pq}^0; p : q \leftarrow \text{pq}^0\}_{p, q \in \mathcal{P}, p \neq q}; \{\!\{C}\!\}_{M_0}$$

Here, $M_0(p, q) = 0$ for all p and q . For simplicity, we write pq^M for $\text{pq}^{M(p, q)}$ and pq^{M+} for $\text{pq}^{M(p, q)+1}$. The definition of $\{\!\{C}\!\}_M$ is given in Figure 2.

We write \bar{M} for $\{\text{pq}^M \mid p, q \in \mathcal{P}, p \neq q\}$, where we assume that the order of the values of M is fixed. In recursive definitions, we reset M to M_0 ; note that the parameter declarations act as binders, so these process names are still fresh.

In order to encode $p.e \rightarrow q$, p uses the auxiliary process pq^M to store the value it wants to send to q . Then, p creates a fresh process (to use in the next communication) and sends its name to pq^M . Afterwards, p is free to proceed with execution. In turn, pq^M communicates q 's name to the new process, which now is ready to receive the next message from p . Finally, pq^M waits for q to be ready to receive both the value being communicated and the name of the process that will store the next value.

$$\begin{aligned}
\{\{p.e \rightarrow q; C\}\}_M &= p.e \rightarrow pq^M; p \text{ start } pq^{M+}; p : pq^M \leftrightarrow pq^{M+}; \\
&\quad pq^M.q \rightarrow pq^{M+}; pq^M.pq^{M+} \rightarrow q; pq^M.* \rightarrow q; \{\{C\}\}_{M[(p,q) \rightarrow M(p,q)+1]} \\
\{\{ \text{if } p \stackrel{\leq}{=} q \text{ then } C_1 \text{ else } C_2 \}\}_M &= q.* \rightarrow qp^M; q \text{ start } qp^{M+}; q : qp^M \leftrightarrow qp^{M+}; qp^M.p \rightarrow qp^{M+}; qp^M.qp^{M+} \rightarrow p; \\
&\quad \text{if } p \stackrel{\leq}{=} qp^M \text{ then } \{\{C_1\}\}_{M[(p,q) \rightarrow M(p,q)+1]} \text{ else } \{\{C_2\}\}_{M[(p,q) \rightarrow M(p,q)+1]} \\
\{\{0\}\}_M &= 0 \quad \{\{X\}\}_M = X(\overline{M}) \quad \{\{\text{def } X = C_2 \text{ in } C_1\}\}_M = \text{def } X(\overline{M_0}) = \{\{C_2\}\}_{M_0} \text{ in } \{\{C_1\}\}_M
\end{aligned}$$

Figure 2: Encoding MC in DMC.

The behaviours of the choreographies C and $\{\{C\}\}$ are closely related, as formalised in the following theorems.

THEOREM 1. *Let $p \in \text{pn}(C)$ and $pq \in \text{pn}(\{\{C\}\}) \setminus \text{pn}(C)$. If $G, \{\{C\}\}, \sigma \rightarrow^* G', C_1, \sigma_1 \rightarrow G', C_2, \sigma_2$ where in the last transition a value v is sent from p to pq , then there exist G'', C_3, σ_3, C_4 and σ_4 such that $G', C_2, \sigma_2 \rightarrow^* G'', C_3, \sigma_3 \rightarrow G'', C_4, \sigma_4$ and in the last transition the same value v is sent from pq to some process $q \in \text{pn}(C)$.*

THEOREM 2. *If $G, \{\{C\}\}_M, \sigma \rightarrow^* G_1, C_1, \sigma_1$, then there exist C', σ' and σ'' such that $G, C, \sigma \rightarrow^* G, C', \sigma'$, and $G_1, C_1, \sigma_1 \rightarrow^* G', \{\{C'\}\}_M, \sigma''$, and σ' and σ'' coincide on the values stored at $\text{pn}(C)$.*

EXAMPLE 2. *We show the result of applying this transformation to Lines 1–3 of Example 1. The numbers refer to the line numbers in the original example.*

$$\begin{aligned}
& a \text{ start } as^0; a : as^0 \leftrightarrow s; \\
& s \text{ start } sa^0; s : sa^0 \leftrightarrow a; \\
& s \text{ start } sb^0; s : sb^0 \leftrightarrow b; \\
1. & a.\text{title} \rightarrow as^0; a \text{ start } as^1; a : as^1 \leftrightarrow as^0; \\
& as^0.as^1 \rightarrow s; as^0.s \rightarrow as^1; as^0.* \rightarrow s; \\
2. & s.\text{price} \rightarrow sa^0; s \text{ start } sa^1; s : sa^1 \leftrightarrow sa^0; \\
& sa^0.sa^1 \rightarrow a; sa^0.a \rightarrow sa^1; sa^0.* \rightarrow a; \\
3. & s.\text{price} \rightarrow sb^0; s \text{ start } sb^1; s : sb^1 \leftrightarrow sb^0; \\
& sb^0.sb^1 \rightarrow b; sb^0.b \rightarrow sb^1; sb^0.* \rightarrow b; \\
& \dots
\end{aligned}$$

The first three lines initialize three channels: from a to s ; from s to a ; and from s to b . Then one message is passed in each of these channels, as dictated by the encoding. All communications are asynchronous in the sense explained above, as in each case the main sender process sends its message to an intermediary (as^0 , sa^0 or sb^0 , respectively), who eventually delivers it to the recipient. Moreover, causal dependencies are kept: in Step 2, s can only send its message to sa^0 after receiving the message sent by a in Step 1. However, in Step 3 s can send its message to sb^0 without waiting for a to receive the previous message, as the action $s.\text{price} \rightarrow sa^0$ can swap with the three actions immediately preceding it. We briefly illustrate Theorems 1 and 2 in this example. Theorem 1 guarantees that the action $a.\text{title} \rightarrow as^0$ is eventually followed by a communication of title from as^0 to some other process in the original choreography (in this case, s). Theorem 2 implies that if $a.\text{title} \rightarrow as^0$ is executed, then it must be

“part” of an action in the original choreography (in this case, $a.\text{title} \rightarrow s$), and furthermore it is possible to find an execution path that will execute the remaining actions generated from that one (the remaining five actions in Step 1).

Our construction can be extended to the whole language of DMC, but we omit this for space constraints.

4. REFERENCES

- [1] M. Carbone, K. Honda, and N. Yoshida. Structured communication-centered programming for web services. *ACM Trans. Prog. Lang. Syst.*, 34(2):8, 2012.
- [2] 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.
- [3] L. Cruz-Filipe and F. Montesi. A language for the declarative composition of concurrent protocols. Submitted for publication.
- [4] L. Cruz-Filipe and F. Montesi. A core model for choreographic programming. Accepted for publication at FACS’16. <http://arxiv.org/abs/1510.03271>.
- [5] N. Dragoni, S. Giallorenzo, A. Lluch-Lafuente, M. Mazzara, F. Montesi, R. Mustafin, and L. Safina. Microservices: yesterday, today, and tomorrow. *CoRR*, abs/1606.04036, 2016.
- [6] K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. *J. ACM*, 63(1):9, 2016.
- [7] T. Leesatapornwongsa, J. Lukman, S. Lu, and H. Gunawi. TaxDC: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In *ASPLoS*, pages 517–530. ACM, 2016.
- [8] S. Lu, S. Park, E. Seo, and Y. Zhou. Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. *ACM SIGARCH Computer Architecture News*, 36(1):329–339, 2008.
- [9] F. Montesi. *Choreographic Programming*. Ph.D. thesis, IT University of Copenhagen, 2013. http://fabriziomontesi.com/files/choreographic_programming.pdf.
- [10] F. Montesi. Kickstarting choreographic programming. In *WS-FM*, volume 9421 of *LNCS*, pages 3–10. Springer, 2016.
- [11] F. Montesi and N. Yoshida. Compositional choreographies. In *CONCUR*, volume 8052 of *LNCS*, pages 425–439. Springer, 2013.
- [12] D. Sangiorgi and D. Walker. *The π -calculus: a Theory of Mobile Processes*. Cambridge Univ. Press, 2001.