

Merging Multiparty Protocols in Multiparty Choreographies

Marco Carbone*

IT University of Copenhagen
Copenhagen, Denmark
carbonem@itu.dk

Fabrizio Montesi

IT University of Copenhagen
Copenhagen, Denmark
fmontesi@itu.dk

Choreography-based programming is a powerful paradigm for defining communication-based systems from a global viewpoint. A choreography can be checked against multiparty protocol specifications, given as behavioural types, that may be instantiated indefinitely at runtime. Each protocol instance is started with a synchronisation among the involved peers.

We analyse a simple transformation from a choreography with a possibly unbounded number of protocol instantiations to a choreography instantiating a single protocol, which is the merge of the original ones. This gives an effective methodology for obtaining new protocols by composing existing ones. Moreover, by removing all synchronisations required for starting protocol instances, our transformation reduces the number of communications and resources needed to execute a choreography.

1 Introduction

Communication-based programming is a widespread design paradigm where the entities of a system communicate exclusively by means of message passing. Communication-based systems are employed in many areas, from multi-core programming [4] to service-oriented and cloud computing [14, 5, 1, 10]. In such systems, entities engage in communication flows where their input and output operations must follow a specific order. The structure of each flow is defined by a protocol.

Choreography-based programming is an emerging methodology for defining communication-based systems in terms of *global descriptions*. A global description gives a global view of how messages are exchanged during execution, in contrast with the methodologies where the code for each entity is defined separately. Global descriptions have been studied as models [8, 7, 13, 11], as standards [18, 2], and as language implementations [12, 16, 15].

In [9] we propose a language where both the abstract and the concrete descriptions of a system are given in global terms. Programmers can use *choreographies* for defining the concrete behaviour of a system and then check them against *protocols*, given as *global types* [11]. The language allows for instantiating different protocols multiple times, checking that each instantiation respects the corresponding protocol type. In the sequel, we introduce a small example for explaining the basic mechanisms of a choreography and how it integrates with a protocol specification.

Our example implements two protocols, called G_a and G_b . In G_a , a user U sends a message to a client application C together with some authentication credentials (encoded as a string). This is expressed by the global type:

$$G_a = U \rightarrow C : \mathbf{string}$$

where U and C are called the *roles* of the protocol. The above behavioural type simply expresses that for executing protocol G_a , whoever plays role U needs to send a message of type **string** to a party playing

*Research supported by the Danish Agency for Science, Technology and Innovation.

role C. Protocol G_b is a little bit more complex:

$$G_b = C \rightarrow F : \mathbf{string}; \quad F \rightarrow C : \{ok : C \rightarrow F : \mathbf{file}, \quad quit : \mathbf{end}\}$$

Here, the roles interacting are a client C and a file server F. In G_b , the client sends a string to the file server F (some authentication credentials) which can reply with either label *ok* or label *quit*. In the first case, the client will send a file to be stored to the file server, whereas in the other case the protocol will just terminate.

Although the global types above give a good abstraction of the protocols that a programmer wishes to use, they give no information on how they can be combined in an implementation. A possible implementation can be given with the choreography C defined below:

$$\begin{aligned}
C = & \quad 1. \quad \text{rec } X \text{ in} \quad c[C], u[U] \mathbf{start } a(k); \\
& \quad 2. \quad u[U].\text{password}() \rightarrow c[C].\text{pwd} : k; \\
& \quad 3. \quad c[C], f[F] \mathbf{start } b(k'); \\
& \quad 4. \quad c[C].\text{pwd} \rightarrow f[F].y : k'; \\
& \quad 5. \quad \text{if } \text{check}(y)@f \text{ then } f[F] \rightarrow c[C] : k'[ok]; \\
& \quad 6. \quad \quad \quad c[C].\text{file} \rightarrow f[F].z : k' \\
& \quad 7. \quad \quad \quad \text{else } f[F] \rightarrow c[C] : k'[quit]; X
\end{aligned} \tag{1}$$

We briefly describe the choreography above. In Line 1, the operation $c[C], u[U] \mathbf{start } a(k)$ starts protocol G_a where c and u denote two executing threads willing to implement protocol G_a playing roles C and U respectively. The two symbols a and k denote the name of the protocol G_a and a session identifier k (which functions as a binder). In Line 2, we can see a communication over session k where the user u , playing role U, sends the return value of some internal function `password()` to c . In Line 3, thread c and another thread f start protocol G_b , similarly to the start of protocol G_a . Line 4 contains a new interaction where c forwards the password pwd to f . In Line 5 we have an if-then-else implementing the abstract branching given in the protocol description of G_b . Note that in the then-branch f communicates the choice of *ok* to c , whereas it uses label *quit* in the else-branch.

The choreography in (1) interleaves the two protocols G_a and G_b in a particular order decided by the programmer. The local behaviour of each thread (implementation) can be then automatically generated by means of *EndPoint Projection* [9, 8]. We observe that (1), since it executes two protocols, has two operations for protocol initiation (Lines 1 and 3) in the body of a recursion. At the endpoint level starts are implemented through synchronisations between peers [6, 11], which may be computationally expensive in a distributed system. Now, we ask:

Can we remove the synchronisation points introduced by start operations in a choreography? And, what are the consequences?

In this paper, we analyse a straightforward transformation on choreographies that cancels out start operations. E.g., the choreography C in (1) could be transformed into:

$$\begin{aligned}
& \quad 1. \quad c[C], u[U], f[F] \mathbf{start } c(k); \\
& \quad 2. \quad \text{rec } X \text{ in} \quad u[U].\text{password}() \rightarrow c[C].\text{pwd} : k; \\
& \quad 4. \quad c[C].\text{pwd} \rightarrow f[F].y : k; \\
& \quad 5. \quad \text{if } \text{check}(y)@f \text{ then } f[F] \rightarrow c[C] : k[ok]; \\
& \quad 6. \quad \quad \quad c[C].\text{file} \rightarrow f[F].z : k \\
& \quad 7. \quad \quad \quad \text{else } f[F] \rightarrow c[C] : k[quit]; X
\end{aligned} \tag{2}$$

Although (2) has a single start operation, we observe that it is semantically related to (1), since all data communications performed in (1) are also performed in (2) and viceversa. Moreover, since the single

synchronisation point in (2) is no longer under recursion, we conjecture that this aspect may greatly improve the execution of choreographies in asynchronous settings.

We also observe that (2) does no longer implement the two binary protocols G_a and G_b , but it subsumes a new three-party protocol G_c obtained by composing the former two:

$$G_c = \text{rec } \mathbf{t} \text{ in } \mathbf{U} \rightarrow \mathbf{C} : \mathbf{string}; \mathbf{C} \rightarrow \mathbf{F} : \mathbf{string}; \mathbf{F} \rightarrow \mathbf{C} : \{\text{ok} : \mathbf{C} \rightarrow \mathbf{F} : \mathbf{file}, \text{quit} : \mathbf{t}\}$$

Note that because of the recursive behaviour appearing in C (but not in G_a and G_b), we need to include some recursive behaviour in the new type (hence the recursion $\text{rec } \mathbf{t}$). Observe also that G_c is a *multiparty* protocol, i.e. it considers more than two participants. Because of this, we believe that such a transformation could be used for creating new protocols. In fact, it may happen that such a choreography represents a pattern that the programmer may want to reuse in other programs. Unfortunately, there is no way to reuse such a pattern in a safe way other than copying and editing the code. By using the transformation hinted above, we could abstract the behaviour of a choreography and make it reusable in other programs.

In the remainder of the paper, we try to lay the foundations of this idea by giving a formalisation of the concept into a simplified version of the global calculus with multiparty protocols [9].

2 Formalisation and Results

Calculus, Semantics and Types. We formalise our choreographies with a simplification of the Global Calculus (GC) [9]. Fig. 1 reports the syntax of GC. In the Figure, τ is a thread (running process); p, q, \dots

$$\begin{array}{ll}
C ::= & \eta; C \quad (seq) \\
& | \text{if } e @ \tau \text{ then } C_1 \text{ else } C_2 \quad (cond) \\
& | \text{rec } X \text{ in } C \quad (rec) \\
& | X \quad (call) \\
& | (\nu k) C \quad (res) \\
& | 0 \quad (inact) \\
\eta ::= & \tau_1[p_1], \dots, \tau_n[p_n] \mathbf{start} a(k) \quad (start) \\
& | \tau_1[p].e \rightarrow \tau_2[q].x : k \quad (com) \\
& | \tau_1[p] \rightarrow \tau_2[q] : k[l] \quad (sel)
\end{array}$$

Figure 1: Global Calculus, syntax.

are roles; a is a public channel; k is a session channel; x is a placeholder for values; and l is a *label* for branching. e denotes a first-order expression on values (integers, strings, \dots), whose syntax we leave unspecified. We read (seq) as: do η and then proceed as C . η represents an interaction between some threads. Term $(start)$ denotes the initiation of a multiparty session (protocol): threads τ_i wish to start the multiparty session a and tag it with a fresh session channel (identifier) k , which is bound in the choreography continuation. The threads τ_i are ordinary threads running in parallel. The p_i 's denote the roles played by the threads in the session. In-session communication is denoted by the term (com) where thread τ_1 sends the evaluation of expression e to thread τ_2 which binds it to variable x in the choreography continuation. In term (sel) , τ_1 communicates to τ_2 her wish to select branch l . In term $(cond)$, thread τ makes an internal choice between branches C_1 and C_2 by evaluating e . (rec) and $(call)$ model standard recursion. (res) , used only at runtime, allows to bind session channel k in C . We use $(\nu k_1, \dots, k_n)$ as an abbreviation for $(\nu k_1) \dots (\nu k_n)$. 0 is the empty choreography.

$$\begin{array}{l}
\llbracket^C\rrbracket_{\text{START}} \quad \tau_1[p_1], \dots, \tau_n[p_n] \mathbf{start} \ a(k); C \rightarrow (vk) C \\
\llbracket^C\rrbracket_{\text{COM}} \quad \tau_1[p].e \rightarrow \tau_2[q].x : k; C \rightarrow C[v/x] \quad (e \downarrow v) \\
\llbracket^C\rrbracket_{\text{SEL}} \quad \tau_1[p] \rightarrow \tau_2[q] : k[l]; C \rightarrow C \\
\llbracket^C\rrbracket_{\text{IF}} \quad \text{if } e @ \tau \text{ then } C_1 \text{ else } C_2 \rightarrow C_i \quad (i = 1 \text{ if } e \downarrow \text{true}, i = 2 \text{ otherwise}) \\
\llbracket^C\rrbracket_{\text{CTX}} \quad C \rightarrow C' \Rightarrow \text{rec } X \text{ in } C \rightarrow \text{rec } X \text{ in } C' \\
\llbracket^C\rrbracket_{\text{RES}} \quad C \rightarrow C' \Rightarrow (vk) C \rightarrow (vk) C' \\
\llbracket^C\rrbracket_{\text{STRUCT}} \quad C_1 \equiv C'_1 \quad C'_1 \rightarrow C'_2 \quad C'_2 \equiv C_2 \Rightarrow C_1 \rightarrow C_2
\end{array}$$

Figure 2: Global Calculus, semantics.

The semantics of the global calculus is a reduction relation \rightarrow which just reduces the size of a choreography wrt prefixing. Formally, \rightarrow is the smallest relation satisfying the rules reported in Fig. 2. In $\llbracket^C\rrbracket_{\text{STRUCT}}$, structural congruence \equiv is standard: it handles alpha-renaming and expansion of recursive calls. Given that there are some fresh names created, the semantics may introduce restriction operators (Rule $\llbracket^C\rrbracket_{\text{START}}$). For instance, the term

$$C = c[C], u[U] \mathbf{start} \ a(k); u[U].\text{password}() \rightarrow c[C].\text{pwd} : k; C'$$

would have the following reduction chain:

$$C \rightarrow (vk) u[U].\text{password}() \rightarrow c[C].\text{pwd} : k; C' \rightarrow (vk) C'[\text{password}()/\text{pwd}]$$

In [9], we develop a type theory for our choreography model exploiting global types for representing protocols (as we did in the Introduction). A type system checks that the protocol instances in a choreography follow the given global types. As an example, we can see that protocols G_a and G_b are correctly used by the choreography C given in the Introduction. Fig. 3 reports the syntax for global types.

$$\begin{array}{l}
G ::= \quad p \rightarrow q : S; G \quad (com) \\
\quad | \quad p \rightarrow q : \{l_i : G_i\}_{i \in I} \quad (choice) \\
\quad | \quad \text{end} \quad (inact) \\
\quad | \quad \text{rec } \mathbf{t} \text{ in } G \quad (rec) \\
\quad | \quad \mathbf{t} \quad (call) \\
S ::= \quad \mathbf{bool} \mid \mathbf{int} \mid \mathbf{string} \mid \mathbf{file} \mid \dots \quad (sort)
\end{array}$$

Figure 3: Global Types, syntax.

Choreography Transformation. We can now present our transformation for choreographies. Formally, we define a function $\{\{C\}\}^k$ that transforms a choreography C into another choreography which implements the same behaviour of C using a single session k . $\{\{C\}\}^k$ is inductively defined by the following rules. Below we assume, without any loss of generality, that all session channels started in C are different, i.e. there are no two subterms of the form $\tau_1[p_1], \dots, \tau_n[p_n] \mathbf{start} \ a(k)$ in C with the same k .

$$\begin{array}{l}
\{\{\tau_1[p_1], \dots, \tau_n[p_n] \mathbf{start} \ a(k'); C'\}\}^k = \{\{C'\}\}^k \\
\{\{\tau_1[p].e \rightarrow \tau_2[q].x : k'; C\}\}^k = \tau_1[\tau_1].e \rightarrow \tau_2[\tau_2].x : k; \{\{C\}\}^k \\
\{\{\tau_1[p] \rightarrow \tau_2[q] : k'[l]; C\}\}^k = \tau_1[\tau_1] \rightarrow \tau_2[\tau_2] : k[l]; \{\{C\}\}^k \\
\{\{\text{if } e @ \tau \text{ then } C_1 \text{ else } C_2\}\}^k = \text{if } e @ \tau \text{ then } \{\{C_1\}\}^k \text{ else } \{\{C_2\}\}^k \quad \{\{\text{rec } X \text{ in } C\}\}^k = \text{rec } X \text{ in } \{\{C\}\}^k
\end{array}$$

We briefly comment the rules above. (*start*) terms are simply removed. In interactions the role of each thread is annotated with the thread name, in order to maintain the distinction between roles with the same name played by different threads. All other terms are preserved.

We can give an example of the transformation by applying it to the choreography C using protocols G_a and G_b in the Introduction. We obtain:

$$\{\{C\}\}^k = \begin{array}{l} 1. \text{ rec } X \text{ in } u[u].\text{password}() \rightarrow c[c].\text{pwd} : k; \quad c[c].\text{pwd} \rightarrow f[f].y : k; \\ 2. \quad \quad \quad \text{if } \text{check}(y)@f \text{ then } f[f] \rightarrow c[c] : k[\text{ok}]; \quad c[c].\text{file} \rightarrow f[f].z : k \\ 3. \quad \quad \quad \text{else } f[f] \rightarrow c[c] : k[\text{quit}]; X \end{array}$$

Observe that the result is the same to the transformation example we have shown in the Introduction, up to renaming of roles and the first start operation for starting k . We omit how to automatically generate the latter, since it can be done through a very simple traversal of the structure of $\{\{C\}\}^k$, tracking the roles of each thread in the interactions.

Results. Hereby, we present some of the properties enjoyed by our simple transformation.

Our first result is about the correctness of the behaviour of the transformation result. Specifically, the transformation does not introduce any additional behaviour (soundness) and it preserves the original behaviour up to removal of start terms and renaming of roles (completeness).

Theorem 2.1 (Correctness). *Let C be a choreography and k a session channel name. Then,*

- (Soundness) $\{\{C\}\}^k \rightarrow C'$ for some C' implies that there exists C'' such that $C \rightarrow C''$ and $C' = \{\{C''\}\}^k$
- (Completeness) $C \rightarrow C'$ for some C' implies that there exists C'' such that $C' \rightarrow^* C''$ and $\{\{C\}\}^k \rightarrow \{\{C''\}\}^k$

The result above can also be stated in a stronger form in terms of bisimilarity, using the labelled semantics reported in [9]. We chose this form for the sake of brevity.

Our second result is about typing: there is a relationship between the typing of a choreography and its transformation. Intuitively, $\{\{C\}\}^k$ can be typed using a composition of the types of C . The following definition formalises this composition. We remind the reader that a global type can always be regarded to as a standard regular tree representation [17]. In the sequel, the function $\text{paths}(G)$ denotes the set of paths in the regular tree representation of a global type G . Moreover, the function interleave applied to a set of paths returns the set of all their possible interleaves. $*$ is the standard Kleene star, denoting closure of paths under repetition.

Definition 2.1 (Mesh Global Types). *Given a set of global types $\{G_1, \dots, G_n\}$, we define $\text{mesh}(\{G_1, \dots, G_n\})$, called the mesh of G_1, \dots, G_n , as the closure under α -renaming of the set*

$$\{ G \mid p \in \text{paths}(G) \text{ only if } p \in \text{interleave}(p_1^*, \dots, p_m^*) \text{ for some } p_j \in \bigcup_{1 \leq i \leq n} \text{paths}(G_i) \}$$

The mesh of a set of global types is the set of all the global types whose paths are the interleaving of some repetitions of the paths of the original types. We can now state our second main result: the transformation of a well-typed choreography is still well-typed and its type is in the mesh of the original types. Below, $a_1 : G_1, \dots, a_n : G_n \vdash C \triangleright k_{n+1} : G_{n+1}, \dots, k : G_m$ refers to the type system found in [9]. Intuitively, C is well-typed if it follows the protocols described by G_i in each session to be started through a_i and each running session k_i .

Theorem 2.2 (Transformation Typing). *Let C be a choreography such that*

$$a_1 : G_1, \dots, a_n : G_n \quad \vdash \quad C \quad \triangleright \quad k_{n+1} : G_{n+1}, \dots, k : G_m$$

Then, for every session channel name k there exists $G \in \text{mesh}(\{G_1, \dots, G_m\})$ such that

$$\emptyset \quad \vdash \quad \{\{C\}\}^k \quad \triangleright \quad k : G$$

Considering again our example, we can type its transformation $\{\{C\}\}^k$ with the following global type G .

$$G = \text{rec } \mathbf{t} \text{ in } \quad \mathbf{u} \rightarrow \mathbf{c} : \mathbf{string}; \mathbf{c} \rightarrow \mathbf{f} : \mathbf{string}; \mathbf{f} \rightarrow \mathbf{c} : \{\mathbf{ok} : \mathbf{c} \rightarrow \mathbf{f} : \mathbf{file}, \quad \mathbf{quit} : \mathbf{t}\}$$

Observe that type G is a nontrivial composition of the types G_a and G_b that we have shown in our introduction. Indeed we can observe that \mathbf{c} is a single role even though thread \mathbf{c} plays two protocols in the original choreography.

3 Conclusions and Further Developments

We have shown how a choreography implementing different sessions of different types can be transformed into a choreography with a single session implementing a single global type. Furthermore, the type of the latter is a composition of the original types.

Our transformation is useful for eliciting the abstract behaviour of a system that implements many protocols (given as types). A programmer may design a choreography and then check if the global abstract behaviour of its implementation is the expected one. Even more importantly, a software architect could exploit our transformation in order to design new standard protocols by extracting them from a choreography. Interestingly, our transformation could also be used for extracting global types out of binary session types once a global implementation is given [8].

Another potential benefit of our transformation lies in resource control. Our transformation removes protocol starts but preserves behaviour. This has two implications. First, all the synchronisations required for starting a protocol instance at the endpoint level are no longer required. This may help in improving the performance of a system. Second, in practice it is usually the case that threads (or processes) can be dynamically spawned at runtime whenever a new session is created. We believe that our transformation can be extended to transform a choreography with an unbounded number of threads and sessions (due to recursion) to a choreography with a finite number of these resources. This would help in managing the resource consumption of complex distributed systems, leading to applications, e.g., in the fields of embedded systems and optimisation. For example, one could design an ad-hoc system optimised for a choreography with a predetermined number of threads.

The formalisation presented in this paper is only an initial step towards a more complete theory. Specifically, we did not deal with some useful features such as channel passing and thread spawning. We plan to investigate these features in future work. We also plan to give a concrete implementation of our transformation for the Chor language [9, 3], and use it to benchmark our theory through practical scenarios.

References

- [1] *Advanced Message Queueing Protocol (AMQP)*. <http://www.amqp.org/>.

- [2] *Business Process Model and Notation*. <http://www.omg.org/spec/BPMN/2.0/>.
- [3] *Chor. Choreography programming language*. <http://www.chor-lang.org/>.
- [4] *Message Passing Interface (MPI)*. <http://www.mpi-forum.org/>.
- [5] *WS-BPEL OASIS Web Services Business Process Execution Language*. <http://docs.oasis-open.org/wsbpel/2.0/wsbpel-specification-draft.html>.
- [6] Lorenzo Bettini, Mario Coppo, Loris D'Antoni, Marco De Luca, Mariangiola Dezani-Ciancaglini & Nobuko Yoshida (2008): *Global Progress in Dynamically Interleaved Multiparty Sessions*. In: *CONCUR, LNCS 5201*, Springer, pp. 418–433, doi:10.1007/978-3-540-85361-9_33.
- [7] N. Busi, R. Gorrieri, C. Guidi, R. Lucchi & G. Zavattaro (2006): *Choreography and orchestration conformance for system design*. In: *Proc. of Coordination'06, LNCS 4038*, Springer-Verlag, pp. 63–81, doi:10.1007/11767954_5.
- [8] Marco Carbone, Kohei Honda & Nobuko Yoshida (2012): *Structured Communication-Centered Programming for Web Services*. *ACM Trans. Program. Lang. Syst.* 34(2), p. 8, doi:10.1145/2220365.2220367.
- [9] Marco Carbone & Fabrizio Montesi (2011): *Typed Multiparty Global Programming*. Technical Report TR-2011-149, IT University of Copenhagen. Available at <http://www.itu.dk/people/fabr/multichor>.
- [10] Claudio Guidi, Paolo Anedda & Tullio Vardanega (2012): *Towards a New Paas Architecture Generation*. In: *CLOSER*, pp. 279–282.
- [11] K. Honda, N. Yoshida & M. Carbone (2008): *Multiparty Asynchronous Session Types*. In: *POPL'08*, 43(1), ACM Press, New York, NY, USA, pp. 273–284, doi:10.1145/1328438.1328472.
- [12] Kohei Honda, Aybek Mukhamedov, Gary Brown, Tzu-Chun Chen & Nobuko Yoshida (2011): *Scribbling Interactions with a Formal Foundation*. In: *ICDCIT, LNCS 6536*, Springer, pp. 55–75, doi:10.1007/978-3-642-19056-8_4.
- [13] I. Lanese, C. Guidi, F. Montesi & G. Zavattaro (2008): *Bridging the Gap between Interaction- and Process-Oriented Choreographies*. In Antonio Cerone & Stefan Gruner, editors: *Proc. of SEFM'08, IEEE*, pp. 323–332, doi:10.1109/SEFM.2008.11.
- [14] F. Montesi, C. Guidi & G. Zavattaro (2007): *Composing Services with JOLIE*. In: *Proc. of ECOWS 2007*, pp. 13–22, doi:10.1109/ECOWS.2007.19.
- [15] Nicholas Ng, Nobuko Yoshida & Kohei Honda (2012): *Multiparty Session C: Safe Parallel Programming with Message Optimisation*. In: *TOOLS (50)*, pp. 202–218, doi:10.1007/978-3-642-30561-0_15.
- [16] PI4SOA (2008): <http://www.pi4soa.org>.
- [17] Benjamin C. Pierce (2002): *Types and Programming Languages*. MIT Press, MA, USA.
- [18] W3C WS-CDL Working Group (2004): *Web Services Choreography Description Language Version 1.0*. <http://www.w3.org/TR/2004/WD-ws-cdl-10-20040427/>.