

Compositional Choreographies

Fabrizio Montesi¹ and Nobuko Yoshida²

¹ IT University of Copenhagen

² Imperial College London

Abstract. We propose a new programming model that supports a compositionality of choreographies. The key of our approach is the introduction of partial choreographies, which can mix global descriptions with communications among external peers. We prove that if two choreographies are composable, then the endpoints independently generated from each choreography are also composable, preserving their typability and deadlock-freedom. The usability of our framework is demonstrated by modelling an industrial use case implemented in a tool for Web Services, Jolie.

1 Introduction

Choreography-based programming is a powerful paradigm for designing communicating systems where the flow of communications is defined from a global viewpoint, instead of separately specifying the behaviour of each *endpoint* (peer). The local behaviour of the endpoints can then be automatically generated by means of *EndPoint Projection* (EPP). This paradigm has been used in standards [21,4] and language implementations [11,19,20,8]. Choreographies impact significantly the quality of software: they lower the chance for programming errors and ease their detection [16,6,7].

Previous works provide models for programming implementations of communicating systems with choreographies [6,7]. These models come with a type discipline for checking choreographies against protocol specifications given as session types [12], which are used to verify that the global behaviour of a choreography implements the expected communication flows. For example, a programmer may express a protocol using a *multiparty session type* [13] (or *global type*) such as the following one:

$$B \rightarrow C: \langle \text{string} \rangle; C \rightarrow B: \langle \text{int} \rangle; B \rightarrow T: \left\{ \begin{array}{l} \text{ok}: B \rightarrow T: \langle \text{string} \rangle; T \rightarrow B: \langle \text{date} \rangle, \\ \text{quit}: \text{end} \end{array} \right\}$$

Above, B, C and T are *roles* and abstractly represent endpoints in a system. In the protocol, a buyer B sends the name of a product to a catalogue C, which replies with the price for that product. Then, B notifies the transport role T of whether the price is accepted or not. In the first case (label *ok*), B sends also a delivery address to T and T replies with the expected delivery date. Otherwise (label *quit*), the protocol terminates immediately.

To the best of our knowledge, all previous choreography programming models (e.g., [7,6]) require the programmer to implement the behaviour of all roles in a protocol where it is used; e.g., it would not be possible to write the choreography of a system that uses the protocol above but gives the implementations only of roles C and T, to make

those reusable by other programs as software libraries through an API. This seriously hinders the applicability of choreographies in industrial settings, where the interoperability of different systems developed independently is the key. In particular, it is not currently possible to:

- use choreographies to develop software libraries that implement subsets of roles in protocols such that they can be reused from other systems;
- reuse an existing software library that implements subsets of roles in protocols from inside a choreography.

To tackle the issues above, we ask: *Can we design a choreography model in which the EPP of a choreography can be composed with other existing systems?* The main problem is that existing choreography models rely on the complete knowledge of the implementation details of all endpoints to ensure that the systems generated by EPP will behave correctly. This complete knowledge is not available when independently developed implementations of distributed protocols need to be composed. In order to answer our question, we build a model for developing *partial choreographies*. Partial choreographies implement the behaviour of subsets of the roles in the protocols they use. Endpoint implementations are then automatically generated from partial choreographies and composed with other systems, with the guarantee that their overall execution will follow the intended protocols and the behaviour of the originating choreographies.

Main Contributions. We provide the following contributions:

Compositional Choreographies. We introduce a new programming model for choreographies in which the implementation of some roles in protocols can be omitted (§ 3). These *partial choreographies* can then be composed with others through message passing. Our model allows to describe both choreographies with many participants or just a single endpoint. We provide a notion of EPP that produces correct endpoint code from a choreography, and we show that the EPP of a choreography preserves its compositional properties (§ 5). Our model introduces *shared channel mobility* to choreographies, which gains a dynamism when two protocols are composed.

Typing. We provide a type system for checking choreographies against protocol specifications given as multiparty session types [13]. The type system ensures that the composition of different programs implements the intended protocols correctly (§ 4), and that our EPP produces code that follows the behaviour of the originating choreographies. Our framework guarantees that the EPP is still typable (§ 5); therefore, the EPP is reusable as a “black box” composable with other systems and the result of the composition can be checked for errors by referring only to types.

Deadlock-freedom and Progress among Composed Choreographies. In the presence of partial choreographies, we prove that we can (i) capture the existing methodologies for deadlock-freedom in complete choreographies as in [6,7] and (ii) extend the notion of progress for incomplete systems investigated in [13] to choreographies (§ 5). Our results demonstrate for the first time that choreographies can be effectively used also as a tool for progress in a compositional setting, offering a new viewpoint for investigating progress and giving a fresh look to the results in [6,7].

Proofs, auxiliary definitions, and other resources are posted at [1], including an implementation of our use case (§ 2) with Jolie [15,18].

2 Motivations: A Use Case of Compositional Choreographies

We present motivations for this study by reporting a use case from our industry collaborators [14], and informally introducing our model. For clarity, we discuss only its most relevant parts. An extended version can be found at [1].

In our use case a buyer company needs to purchase a product from one of many available seller companies. The use case has two aspects that previous choreography models cannot handle: (i) the system of the buyer company is developed independently from those of the seller companies, and use the latter as software modules without revealing internal implementation details; (ii) depending on the desired product, the buyer company selects a suitable seller company at runtime. We address these issues with *partial choreographies*. A partial choreography implements a subset of the roles in a protocol, leaving the implementation of the other roles to an external system. External systems can be discovered at runtime. In our case, the buyer company will select a seller and then run the protocol from the introduction by implementing only the buyer role B, and rely on the external seller system to implement the other two roles C and T.

Buyer Choreography. We now define a choreography for the buyer company, C_B .

$$C_B = \begin{array}{l} 1. u[\bar{U}] \text{ starts } pd[PD] : a(k); u[\bar{U}].prod \rightarrow pd[PD].x : k; \\ 2. pd[PD] \text{ starts } r[R] : b(k'); pd[PD].x \rightarrow r[R].y : k'; r[R].find(y) \rightarrow pd[PD].z : k'; \\ 3. pd[B] \text{ req } C, T : z(k''); pd[B].x \rightarrow C : k''; C \rightarrow pd[B].price : k''; \\ 4. \text{ if } check(price)@pd \text{ then} \\ 5. \quad pd[B] \rightarrow T : k'' \oplus ok; pd[PD] \rightarrow u[\bar{U}] : k[del]; pd[PD] \rightarrow u[\bar{U}] : k\langle k''[B] \rangle; \\ 6. \quad u[B].addr \rightarrow T : k''; T \rightarrow u[B].ddate : k'' \\ 7. \text{ else} \\ 8. \quad pd[B] \rightarrow T : k'' \oplus quit; pd[PD] \rightarrow u[\bar{U}] : k[quit] \end{array}$$

Above, a purchase in the buyer company is initiated by a user process u . In Line 1, process u and the freshly created process pd (for purchasing department) start a session k by synchronising on shared channel a . Each process is annotated with the role it plays in the protocol that the session implements. Then, still in Line 1, u sends the product $prod$ the user wishes to buy to pd . In Line 2 pd starts a new session k' with a fresh process r (a service registry) through shared channel b . Then, pd forwards the product name to r , which replies with the shared channel of the seller to contact for the purchase.

We refer to statements such as those in Lines 1-2 as complete, since they describe the behaviour of all participants, both sender and receiver(s). On the other hand, the continuation in Lines 3-8 is a partial choreography that relies on the selected external seller to implement the protocol shown in the introduction and perform the purchase.

The partial choreography in Lines 3-8 is depicted as a sequence chart in Fig. 1.a, where dashed lines indicate interactions with external participants. In Line 3 pd *requests* a synchronisation on the shared channel stored in its local variable z to create the new session k'' , declaring that it will play role B and that it expects the environment to implement roles C and T. Session k'' proceeds as specified by the protocol in the introduction. First, pd sends the product name stored in x through session k'' to the external process that is playing role C (the product catalogue executed by the seller company). Observe that here we do not specify the process name of the receiver, since that will be established by the external seller system. Then, pd waits to receive the price for the product from the external process playing role C in k'' . In Line 4, pd checks whether




Fig. 1. Sequence charts for buyer (a), seller (b), and their composition (c)

the price is acceptable; if so, in Line 5 pd tells the external process playing role T (the transport process executed by the seller company) and user u (which remains internal to the buyer choreography) to proceed with the purchase (labels ok and del respectively). Still in Line 5, pd *delegates* to u the continuation of session k'' in its place, as role B . In Line 6, the user sends her address to T and receives a delivery date. If the price is not acceptable, Line 7, then in Line 8 pd informs the others to quit the purchase attempt.

Seller Choreography and Composition. We define now a choreography for a seller that can be contacted by C_B . Let the find function in C_B return shared channel c for electronic products, and c' for other products; we refer to the choreographies of the respective seller companies as C_S and C'_S . Below, we define C_S (C'_S , omitted, is similar).

$$C_S = \begin{array}{l} 1. \text{acc } c[C], t[T] : c(k''); \quad B \rightarrow c[C].x_2 : k''; \quad c[C].price(x_2) \rightarrow B : k''; \\ 2. B \rightarrow t[T] : k'' \& \left\{ \begin{array}{l} ok : B \rightarrow t[T].daddr : k''; \quad t[T].time(daddr) \rightarrow B : k'' \\ quit : \mathbf{0} \end{array} \right. \end{array}$$

The choreography C_S , depicted as a sequence chart in Fig. 1.b, starts by *accepting* the creation of session k'' through shared channel c , offering to spawn two fresh processes c and t . Choreographies starting with an acceptance act as replicated, modelling typical always-available modules. The acceptance in Line 1 would synchronise with the request made by C_B in the case $z = c$. Afterwards, c expects to receive the product name from the process playing B in session k'' , and replies with the respective price. In Line 2, t (the process for the transport) waits for either label ok or $quit$. In the first case, t also waits for a delivery address and then sends back the expected time of arrival.

From the code of C_B and C_S and, graphically, from their respective sequence charts we can see that they are *compatible*: sending actions match receiving actions on the other side and vice versa. Our model can recognise this by using roles in protocols as interfaces between partial choreographies (§ 4). The code for buyer and seller companies can be composed in a network with the parallel operator $|$ as: $C = C_B | C_S | C'_S$. Parallel composition allows partial terms in different choreographies to communicate. In (§ 3, Semantics) we formalise a semantics for choreography composition. To give the intuition behind our semantics, let us consider the sequence charts in Fig. 1.a and Fig. 1.b; their composition will behave as the sequence chart in Fig. 1.c.

$C ::= \eta; C$	<i>(seq)</i>	$C_1 \mid C_2$	<i>(par)</i>
$\text{if } e @ \mathbf{p} \text{ then } C_1 \text{ else } C_2$	<i>(cond)</i>	$(\nu r) C$	<i>(res)</i>
$\text{rec } X(x @ \mathbf{p}, \tilde{k}, \tilde{\mathbf{p}}) = C_2 \text{ in } C_1$	<i>(rec)</i>	$X(x @ \mathbf{p}, \tilde{k}, \tilde{\mathbf{p}})$	<i>(call)</i>
$\mathbf{0}$	<i>(inact)</i>	$\mathbf{A} \rightarrow q : k \& \{l_i : C_i\}_{i \in I}$	<i>(branch)</i>
$\eta ::= p \text{ starts } \tilde{q} : a(k)$	<i>(start)</i>	$p.e \rightarrow q.x : k$	<i>(com)</i>
$p \rightarrow q : k[l]$	<i>(sel)</i>	$p \rightarrow q : k \langle k'[\mathbf{A}] \rangle$	<i>(del)</i>
$p \text{ req } \tilde{\mathbf{B}} : u(k)$	<i>(req)</i>	$\mathbf{acc} \tilde{q} : a(k)$	<i>(acc)</i>
$p.e \rightarrow \mathbf{B} : k$	<i>(com-s)</i>	$\mathbf{A} \rightarrow q.x : k$	<i>(com-r)</i>
$p \rightarrow \mathbf{B} : k \langle k'[\mathbf{C}] \rangle$	<i>(del-s)</i>	$\mathbf{A} \rightarrow q : k \langle k'[\mathbf{C}] \rangle$	<i>(del-r)</i>
$p \rightarrow \mathbf{B} : k \oplus l$	<i>(sel-s)</i>		
$p, q ::= \mathbf{p}[\mathbf{A}]$		$u ::= x \mid a$	

Fig. 2. Compositional Choreographies.

3 Compositional Choreographies

This section introduces our model for compositional choreographies, a calculus where complete and partial actions can be freely interleaved.

Syntax. Fig. 2 defines the syntax of our calculus. C is a choreography, η is a complete or partial action, p is a typed process identifier made by a process identifier \mathbf{p} and a role annotation \mathbf{A} , k is a session identifier, and a is a shared channel. A term $\eta; C$ denotes a choreography that may execute action η and then proceed as C . In the productions for η , terms *(start)*, *(com)*, *(sel)* and *(del)* are complete actions, whereas all the others are partial. In the productions for C , term *(branch)* is also partial.

Complete Actions. Term *(start)* initiates a session: process \mathbf{p} starts a new multiparty session through shared channel a and tags it with a fresh identifier k . \mathbf{p} is already running and dubbed *active process*, while \tilde{q} (which we assume nonempty) is a set of bound *service processes* that are freshly created. $\mathbf{A}, \tilde{\mathbf{B}}$ represent the respective roles played by the processes in session k . Term *(com)* denotes a communication where process \mathbf{p} sends, on session k , the evaluation of a first-order expression e to process q , which binds it to its *local variable* x . Expressions may be shared channel names, capturing shared channel mobility. In *(sel)*, \mathbf{p} communicates to q its selection of branch l . Term *(del)* models session mobility: process \mathbf{p} delegates to q through session k its role \mathbf{C} in session k' .

Partial Actions. In term *(req)*, process \mathbf{p} is willing to start a new session k by synchronising through shared channel a with some other external processes. \mathbf{p} is willing to play role \mathbf{A} in the session and expects the other processes to play the other roles $\tilde{\mathbf{B}}$. *(req)* terms are supposed to synchronise with always-available *service processes*, modelled by term *(acc)*. In term *(acc)*, processes \tilde{q} are dynamically spawned whenever requested by a matching *(req)* term on the same shared channel a . Term *(com-s)* models the sending of a message from a process \mathbf{p} to an external process playing role \mathbf{B} in session k . Dually, in *(com-r)* process q receives a message intended for \mathbf{B} in session k from the external process playing role \mathbf{A} . *(del-s)* and *(del-r)* model, respectively, the sending and receiving of a delegation of role \mathbf{C} in session k' . *(sel-s)* models the sending of a selection of label

l . (*sel-s*) can synchronise with a (*branch*) term, which offers a choice on multiple labels. Once a label l_i is selected, (*branch*) proceeds by executing its continuation C_i .

Other Terms. In term (*cond*), process p evaluates condition e to choose the continuation C_1 or C_2 . Term (*res*) restricts the usage of a name r to a choreography C . r can be any name, i.e., a process identifier p , a session identifier k , or a shared channel a . Term (*par*) models the parallel composition of choreographies, allowing partial actions to interact through the network. The other terms are standard: terms (*rec*), (*call*) and (*inact*) model, respectively, a recursive procedure, a recursive call, and termination.

For clarity, we have annotated process identifiers with roles in all communications. Technically, this is necessary only for terms (*start*), (*req*) and (*acc*) since roles can be inferred from session identifiers in all other terms (cf. [7]).

Semantics. We give semantics to choreographies with a labelled transition system (Its), whose rules are defined in Fig. 3 and whose labels λ are defined as:

$$\lambda ::= \eta \mid A \rightarrow q : k\&l \mid \text{if}@p \mid (\nu r) \lambda$$

We distinguish between labels representing complete or partial actions with the respective sets CAct and PAct. CAct is the smallest set containing all η that are complete actions and the labels of the form $\text{if}@p$, closed under restrictions (νr) . PAct is the smallest set containing all η that are partial actions and the labels of the form $A \rightarrow q : k\&l$, similarly closed under restriction of names. We also use other auxiliary definitions. $\text{fc}(C)$ returns the set of all session/role pairs $k[A]$ such that k is free in C and there is a process performing an action as role A in session k in C . $\text{rc}(\lambda)$ is defined only for partial labels that are not (*req*) or (*acc*), and returns the session/role pair of the intended external sender or receiver of λ ; e.g., $\text{rc}(p.e \rightarrow B : k) = k[B]$. fn and bn denote the sets of free and bound names in a label or a term. $\text{snd}(\eta)$ returns the name of the sender process in η , and is undefined if η has no sender process (e.g., when η is a (*com-r*)). $\text{rcv}(\eta)$, instead, returns the session/role pair $k[A]$ where k is the session used in η and A is the role of the receiver (similarly for $\text{rcv}(\lambda)$). $\text{fc}(\lambda)$ is as $\text{fc}(C)$, but applied on labels.

We comment the rules. Rule $[\text{C}_{\text{ACT}}]$ handles actions that can be simply consumed. Rule $[\text{C}_{\text{START}}]$ starts a session with a global action, by restricting the names of the newly created session identifier k and processes \tilde{q} . Rule $[\text{C}_{\text{COM}}]$ handles the communication of a value by substituting, in the continuation C , the binding occurrence x under process identifier q with value v (evaluated from expression e). Similarly, rules $[\text{C}_{\text{COM-S}}]$ and $[\text{C}_{\text{COM-R}}]$ implement the respective partial sending and receiving actions of a communication. In rule $[\text{C}_{\text{BRANCH}}]$, process q receives a selection on a branching label and proceeds accordingly. Rules $[\text{C}_{\text{COND}}]$, $[\text{C}_{\text{RES}}]$, and $[\text{C}_{\text{CTX}}]$ are standard. Rule $[\text{C}_{\text{PAR}}]$ makes global actions observable and blocks partial actions if their counterpart is in the parallel branch C_2 . In rule $[\text{C}_{\text{EQ}}]$, the relation \mathcal{R} can either be the swapping relation \simeq_C , which swaps terms that describe the behaviour of different processes [7], or the structural congruence \equiv , which handles name restriction and recursion unfolding (see [1]).

Rule $[\text{C}_{\text{SYNC}}]$ is the main rule and enables two choreographies to perform compatible sending/receiving partial actions λ and λ' to interact and realise a global action, defined by $\lambda \circ \lambda'$. Function $\circ : \text{PAct} \times \text{PAct} \rightarrow \text{CAct}$ is formally defined by the rules below:

$$\begin{aligned} p[A] \rightarrow B : k\langle v \rangle \circ A \rightarrow q[B] : k\langle v \rangle &= p[A] \rightarrow q[B] : k\langle v \rangle \\ p[A] \rightarrow B : k\langle k'[C] \rangle \circ A \rightarrow q[B] : k\langle k'[C] \rangle &= p[A] \rightarrow q[B] : k\langle k'[C] \rangle \\ p[A] \rightarrow B : k \oplus l \circ A \rightarrow q[B] : k\&l &= p[A] \rightarrow q[B] : k[l] \end{aligned}$$

$$\begin{array}{l}
\llbracket^c\text{ACT}\rrbracket \quad \eta \notin \{(com), (com-s), (com-r), (start), (acc)\} \Rightarrow \eta; C \xrightarrow{\eta} C \\
\llbracket^c\text{START}\rrbracket \quad \eta = p \text{ starts } \widetilde{q[\mathbf{B}]} : a(k) \Rightarrow \eta; C \xrightarrow{\eta} (\nu k, \tilde{q}) C \\
\llbracket^c\text{COM}\rrbracket \quad \eta = p.e \rightarrow q[\mathbf{B}].x : k \Rightarrow \eta; C \xrightarrow{p \rightarrow q[\mathbf{B}]:k\langle v \rangle} C[v/x@q] \quad (e \downarrow v) \\
\llbracket^c\text{COM-S}\rrbracket \quad p.e \rightarrow \mathbf{B} : k; C \xrightarrow{p \rightarrow \mathbf{B}:k\langle v \rangle} C \quad (e \downarrow v) \\
\llbracket^c\text{COM-R}\rrbracket \quad \mathbf{A} \rightarrow q[\mathbf{B}].x : k; C \xrightarrow{\mathbf{A} \rightarrow q[\mathbf{B}]:k\langle v \rangle} C[v/x@q] \\
\llbracket^c\text{BRANCH}\rrbracket \quad \mathbf{A} \rightarrow q : k \& \{l_i : C_i\}_{i \in I} \xrightarrow{\mathbf{A} \rightarrow q:k \& l_j} C_j \quad (j \in I) \\
\llbracket^c\text{COND}\rrbracket \quad \text{if } e@p \text{ then } C_1 \text{ else } C_2 \xrightarrow{\text{if}@p} C_i \quad (i = 1 \text{ if } e \downarrow \text{true}, i = 2 \text{ otherwise}) \\
\llbracket^c\text{RES}\rrbracket \quad C \xrightarrow{\lambda} C' \Rightarrow (\nu r) C \xrightarrow{(\nu r)\lambda} (\nu r) C' \\
\llbracket^c\text{CTX}\rrbracket \quad C_1 \xrightarrow{\lambda} C'_1 \Rightarrow \text{rec } X(\widetilde{x@p}, \tilde{k}, \tilde{p}) = C_2 \text{ in } C_1 \xrightarrow{\lambda} \text{rec } X(\widetilde{x@p}, \tilde{k}, \tilde{p}) = C_2 \text{ in } C'_1 \\
\llbracket^c\text{PAR}\rrbracket \quad C_1 \xrightarrow{\lambda} C'_1 \Rightarrow C_1 \mid C_2 \xrightarrow{\lambda} C'_1 \mid C_2 \quad (\lambda \in \text{CACT} \vee \text{rc}(\lambda) \notin \text{fc}(C_2)) \\
\llbracket^c\text{EQ}\rrbracket \quad \mathcal{R} \in \{\equiv, \simeq_C\} \quad C_1 \mathcal{R} C'_1 \quad C'_1 \xrightarrow{\lambda} C'_2 \quad C'_2 \mathcal{R} C_2 \Rightarrow C_1 \xrightarrow{\lambda} C_2 \\
\llbracket^c\text{SYNC}\rrbracket \quad C_1 \xrightarrow{\lambda} C'_1 \quad C_2 \xrightarrow{\lambda'} C'_2 \Rightarrow C_1 \mid C_2 \xrightarrow{\lambda \circ \lambda'} C'_1 \mid C'_2 \\
\llbracket^c\text{P-START}\rrbracket \quad \left. \begin{array}{l} i \in [1, n] \quad \{\tilde{q}\} = \{\tilde{q}_1, \dots, \tilde{q}_n\} \\ \{\tilde{\mathbf{B}}\} = \{\tilde{\mathbf{B}}_1, \dots, \tilde{\mathbf{B}}_n\} \quad C'' = \prod_i C_i \\ C \xrightarrow{p \text{ req } \tilde{\mathbf{B}}:u(k)} C' \\ C_i = \text{acc } \widetilde{q[\mathbf{B}]_i} : a(k); C'_i \end{array} \right\} \Rightarrow C \mid C'' \xrightarrow{\lambda} (\nu k, \tilde{q}) (C' \mid \prod_i (C'_i)) \mid C'' \\ \quad (\lambda = p \text{ starts } \widetilde{q[\mathbf{B}]_1}, \dots, \widetilde{q[\mathbf{B}]_n} : a(k)) \\
\llbracket^c\text{ASYNC}\rrbracket \quad C \xrightarrow{\lambda} (\nu \tilde{r}) C' \Rightarrow \eta; C \xrightarrow{\lambda} (\nu \tilde{r}) \eta; C' \quad \left(\begin{array}{l} \text{snd}(\eta) \in \text{fn}(\lambda) \quad \tilde{r} = \text{bn}(\lambda) \\ \text{rcv}(\eta) \notin \text{fc}(\lambda) \quad \tilde{r} \notin \text{fn}(\eta) \\ \eta \notin \{(start), (acc)\} \end{array} \right)
\end{array}$$

Fig. 3. Semantics of Compositional Choreographies

Observe that if $\lambda \circ \lambda'$ is not defined (the actions are incompatible), then the rule cannot be applied. Similarly, $\llbracket^c\text{P-START}\rrbracket$ models a session start by synchronising a partial choreography that requests to start a session with other choreographies that can accept the request on the same shared channel. The choreographies accepting the request remain available afterwards, for reuse. Finally, rule $\llbracket^c\text{ASYNC}\rrbracket$ models asynchrony, allowing the sender process of an interaction η ($\text{snd}(\eta)$) to send a message and then proceed freely before the intended receiver actually receives it. In the rule, we require asynchrony to preserve the message ordering in a session wrt receivers with a causality check ($\text{rcv}(\eta) \notin \text{fc}(\lambda)$).

4 Typing Compositional Choreographies

We now present our typing discipline, which ensures that sessions in a choreography follow protocol specifications given as global types [13,3]. The key advances from

previous work [7] are: (i) introduction of the typing rules for partial choreographies and shared channel passing; and (ii) typing endpoints by local types, which offer transparent compositional properties for the behaviour of each process.

Global and Local Types from [13,3] are defined below:

$$\begin{aligned} G &::= A \rightarrow B : \langle U \rangle; G \mid A \rightarrow B : \{l_i : G_i\}_{i \in I} \mid \mu \mathbf{t}. G \mid \mathbf{t} \mid \text{end} \\ T &::= !A \langle U \rangle; T \mid ?A \langle U \rangle; T \mid \oplus A \{l_i : T_i\}_{i \in I} \mid \&A \{l_i : T_i\}_{i \in I} \mid \mu \mathbf{t}. T \mid \mathbf{t} \mid \text{end} \\ S &::= G \mid \text{int} \mid \text{bool} \cdots \quad U ::= S \mid T \end{aligned}$$

G is a global type. $A \rightarrow B : \langle U \rangle; G$ abstracts a communication from role A to role B with continuation G , where U is the type of the exchanged message. U can either be a sort type S (used for typing values or shared channels), or a local type T (used for typing session delegation). In $A \rightarrow B : \{l_i : G_i\}_{i \in I}$, role A selects one label l_i offered by role B and the global type proceeds as G_i . All other terms are standard.

T denotes a local type. $!A \langle U \rangle; T$ represents the sending of a message of type U to role A , with continuation T . Dually, $?A \langle U \rangle; T$ represents the receiving of a message of type U from role A . $\oplus A \{l_i : T_i\}_{i \in I}$ and $\&A \{l_i : T_i\}_{i \in I}$ abstract the selection and the offering of some branches. The other terms are standard.

To relate a global type to the behaviour of an endpoint, we project a global type G onto a local type that represents the behaviour of a single role. We write $\llbracket G \rrbracket_A$ to denote the projection of G onto the role A , which is defined following [10] (cf. [1]).

Type Checking. We now introduce our type checking discipline for checking choreographies against global types. We use two kinds of typing environments, the linear session typing environments Δ and the unrestricted service environments Γ :

$$\Delta ::= \Delta, k[A] : T \mid \emptyset \quad \Gamma ::= \Gamma, x @ \mathbf{p} : S \mid \Gamma, X : (\Gamma, \Delta) \mid \Gamma, \mathbf{p} : k[A] \mid \Gamma, a : G \langle A \mid \tilde{B} \mid \tilde{C} \rangle \mid \emptyset$$

Δ is standard [3], where $k[A] : T$ maps a local type T to a role A in a session k . In Γ , $x @ \mathbf{p} : S$ types variable x of process \mathbf{p} with type S . $X : (\Gamma, \Delta)$ types recursive procedure X . $\mathbf{p} : k[A]$ establishes that process \mathbf{p} owns role A in session k . $a : G \langle A \mid \tilde{B} \mid \tilde{C} \rangle$ types a shared channel a with global type G : A is the role of the active process that starts the session through a ; \tilde{B} are the roles of the service processes; \tilde{C} are the roles, in \tilde{B} , that a choreography implements for the shared channel a , enabling compositionality of services. Whenever we write $a : G \langle A \mid \tilde{B} \mid \tilde{C} \rangle$ in Γ , we assume that $\tilde{C} \subseteq \tilde{B}$, $A \notin \tilde{B}$, and that $A, \tilde{B} = \text{roles}(G)$. $\text{roles}(G)$ returns the set of roles in a global type G .

We can write $\Gamma, \mathbf{p} : k[A]$ only if \mathbf{p} is not associated to any other role in session k in Γ (a process may only play one role per session). A process \mathbf{p} may however appear more than once in a same Γ , allowing processes to run multiple sessions. As usual, we require all other kinds of occurrences in environments to have disjoint identifiers.

A typing judgement $\Gamma \vdash C \triangleright \Delta$ establishes that a choreography C is well-typed. Intuitively, C is well-typed if shared channels are used according to Γ and sessions are used according to Δ . Δ gives the session types of the free sessions in C . Following the design idea that services should always be available, shared by other models [6,7], we assume that all (*acc*) terms in a choreography are not guarded by other actions. A selection of the rules defining our typing judgement is reported in Fig. 4.

We comment the typing rules. Rule $\lceil^T \text{START} \rceil$ types a (*start*); $a : G \langle A \mid \tilde{B} \mid \tilde{B} \rangle$ checks that the choreography should implement all roles in protocol G ; processes \tilde{q} are checked to

$$\begin{array}{c}
\text{[}^T\text{]}_{\text{START}} \frac{\Gamma, a : G\langle A|\tilde{B}|\tilde{B}\rangle, \Gamma' \vdash C \triangleright \Delta, \Delta' \quad r[c] \in p[A], \tilde{q}[\tilde{B}] \Leftrightarrow (r : k[c] \in \Gamma' \wedge k[c] : \llbracket G \rrbracket_c \in \Delta') \quad \tilde{q} \notin \Gamma}{\Gamma, a : G\langle A|\tilde{B}|\tilde{B}\rangle \vdash p[A] \text{ starts } \tilde{q}[\tilde{B}] : a(k); C \triangleright \Delta} \\
\text{[}^T\text{]}_{\text{SEL}} \frac{j \in I \quad \Gamma \vdash p : k[A], q : k[B] \quad \Gamma \vdash C \triangleright \Delta, k[A] : T_j, k[B] : T'_j}{\Gamma \vdash p[A] \rightarrow q[B] : k[l_j]; C \triangleright \Delta, k[A] : \oplus B\{l_i : T_i\}_{i \in I}, k[B] : \& A\{l_i : T'_i\}_{i \in I}} \\
\text{[}^T\text{]}_{\text{REQ}} \frac{\Gamma \vdash x @ p : G\langle A|\tilde{B}|\emptyset\rangle \quad \Gamma, p : k[A] \vdash C \triangleright \Delta, k[A] : \llbracket G \rrbracket_A}{\Gamma \vdash p[A] \text{ req } \tilde{B} : x(k); C \triangleright \Delta} \quad \text{[}^T\text{]}_{\text{PAR}} \frac{\Gamma, \Gamma_i \vdash C_i \triangleright \Delta_i}{\Gamma, \Gamma_1 \circ \Gamma_2 \vdash C_1 | C_2 \triangleright \Delta_1, \Delta_2} \\
\text{[}^T\text{]}_{\text{ACC}} \frac{\Gamma, a : G\langle D|\tilde{B}|\emptyset\rangle, \Gamma' \vdash C \triangleright \Delta, \Delta' \quad r[c] \in \tilde{q}[\tilde{A}] \Leftrightarrow (r : k[c] \in \Gamma' \wedge k[c] : \llbracket G \rrbracket_c \in \Delta') \quad \tilde{q} \notin \Gamma}{\Gamma, a : G\langle D|\tilde{B}|\tilde{A}\rangle \vdash \text{acc } \tilde{q}[\tilde{A}] : a(k); C \triangleright \Delta} \\
\text{[}^T\text{]}_{\text{COM-S}} \frac{\Gamma \vdash e @ p : S \quad \Gamma \vdash p : k[A] \quad \Gamma \vdash C \triangleright \Delta, k[A] : T \quad q : k[B] \notin \Gamma}{\Gamma \vdash p[A].e \rightarrow B : k; C \triangleright \Delta, k[A] : !B(S); T} \quad \text{[}^T\text{]}_{\text{ZERO}} \frac{\text{cosha}(\Gamma) \quad \Delta \text{ end only}}{\Gamma \vdash \mathbf{0} \triangleright \Delta} \\
\text{[}^T\text{]}_{\text{BRANCH}} \frac{i \in I \quad \Gamma \vdash C_i \triangleright \Delta, k[A] : T_i \quad I \subseteq J \quad p : k[A] \notin \Gamma}{\Gamma \vdash A \rightarrow q[B] : k \& \{l_i : C_i\}_{i \in I} \triangleright \Delta, k[B] : \& B\{l_j : T_j\}_{j \in J}}
\end{array}$$

Fig. 4. Typing Rules for Compositional Choreographies (selection)

be fresh ($\tilde{q} \notin \Gamma$); the continuation C is checked by updating Γ' and Δ' respectively with the process ownerships for their roles in k and the local types for their behaviour in k . $\text{[}^T\text{]}_{\text{SEL}}$ deals with selection, checking that the selected label l_j is specified in the local types. In rule $\text{[}^T\text{]}_{\text{REQ}}$, we check that the choreography requesting the services is not responsible for implementing them, to avoid deadlocks due to the lack of services in parallel required by rule $\text{[}^C\text{]}_{\text{P-START}}$, and that the requesting process behaves as expected by its role in the protocol. Conversely, $\text{[}^T\text{]}_{\text{ACC}}$ types an (*acc*) term by ensuring that all the roles for which the choreography is responsible are implemented (the other checks are similar to $\text{[}^T\text{]}_{\text{START}}$). This *distribution* of the responsibilities for implementing the different roles in a protocol is handled by rule $\text{[}^T\text{]}_{\text{PAR}}$, using the role distribution function $\Gamma_1 \circ \Gamma_2$. Formally, $\Gamma_1 \circ \Gamma_2$ is defined as the union of Γ_1 and Γ_2 except for the typing of shared channels with the same name, which are merged with the following rule:

$$a : G\langle A|\tilde{B}|\tilde{C}\rangle = a : G\langle A|\tilde{B}|\tilde{D}\rangle \circ a : G\langle A|\tilde{B}|\tilde{E}\rangle \quad (\tilde{C} = \tilde{D} \uplus \tilde{E})$$

In rule $\text{[}^T\text{]}_{\text{ZERO}}$ we check that all responsibilities have been implemented and that the sessions in Δ have been executed. Specifically, predicate $\text{cosha}(\Gamma)$ checks that for every $a : G\langle A|\tilde{B}|\tilde{C}\rangle$ in Γ either (i) $\tilde{C} = \tilde{B}$, meaning that a was used only internally with (*start*) terms; or (ii) $\tilde{C} = \emptyset$, meaning that a is used compositionally in collaboration with other choreographies and all roles that the current choreography is responsible for (\tilde{C}) have been implemented correctly with (*acc*) terms. Rules $\text{[}^T\text{]}_{\text{COM-S}}$ and $\text{[}^T\text{]}_{\text{BRANCH}}$ type respectively a sending action and a branching. They are very similar to their complete versions since local types allow us to look at the behaviour of processes independently. They also check that the counterpart for the partial action is not in the continuation, by ensuring that there is not process q such that q plays the other role for session k in Γ , which could obviously lead to a deadlock because process p would not have another process to communicate with in parallel as required by rule $\text{[}^C\text{]}_{\text{SYNC}}$.

Typing Expressiveness. Our typing system exploits the global information given by complete terms and seamlessly falls back to typical session typing when dealing with partial actions. In particular, $[\cdot]^T_{\text{SEL}}$ judges that a choice in a protocol is implemented correctly even if only one of the branches is actually followed. This is sound because we are typing a complete term, and therefore we know that the other branches are not used. This expressiveness is typical of choreography-based models [6,7]. However, such a global knowledge is not available in a partial choreography. For example, in rule $[\cdot]^T_{\text{BRANCH}}$ we cannot know which branch will be selected by the sender and we must therefore require that the receiver process supports at least all the branches specified by the corresponding local type, as in standard session typing for endpoints [12,13].

Properties. We conclude this section by presenting the expected main properties of our type system. Below, to state session fidelity, we use the transition of local types $\Delta \xrightarrow{\alpha} \Delta'$ (defined as [13] and fully given in [1]), where α types a partial or complete action. $\alpha \vdash \lambda$ judges that the label λ is for the same session as α and respects its roles and carried type. We also extend our typing judgement with the extra environment Σ , for handling session ownerships with asynchronous delegations at runtime (see [1]).

Theorem 1 (Typing Soundness). *Let $\Gamma; \Sigma \vdash C \triangleright \Delta$. Then,*

- (Subject Swap) $C \simeq_C C'$ implies $\Gamma; \Sigma \vdash C' \triangleright \Delta$.
- $C \xrightarrow{\lambda} C'$ implies that there exists Δ' such that
 - (Subject Reduction) $\Gamma'; \Sigma' \vdash C' \triangleright \Delta'$ for some Γ', Σ' ;
 - (Session Fidelity) if λ is a communication on session k , then $\Delta \xrightarrow{\alpha} \Delta'$ with $\alpha \vdash \lambda$; else, $\Delta = \Delta'$.

5 Properties of Compositional Choreographies

This section states the main properties of our framework wrt the execution of actual systems composed by endpoints.

Endpoint Projection (EPP) generates correct endpoint code from a choreography. Formally, by endpoint code we refer to choreographies that do not contain complete actions. To define the complete EPP, we first define how the behaviour of a single process in a choreography can be projected. We denote this *process projection* of a process p in a choreography C with $\llbracket C \rrbracket_p$. Selected rules of process projection are given below:

$$\begin{array}{l}
 \llbracket p[A] \text{ starts } \widetilde{q[B]} : a(k); C \rrbracket_r \\
 = \begin{cases} p[A] \text{ req } \widetilde{B} : a(k); \llbracket C \rrbracket_r & \text{if } r = p \\ \text{acc } r[C] : a(k); \llbracket C \rrbracket_r & \text{if } r[C] \in \widetilde{q[B]} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
 \llbracket p[A].e \rightarrow B : k; C \rrbracket_r \\
 = \begin{cases} p[A].e \rightarrow B : k; \llbracket C \rrbracket_r & \text{if } r = p \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
 \llbracket \text{if } e @ p \text{ then } C_1 \text{ else } C_2 \rrbracket_r \\
 = \begin{cases} \text{if } e @ p \text{ then } \llbracket C_1 \rrbracket_r \text{ else } \llbracket C_2 \rrbracket_r & \text{if } r = p \\ \llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r & \text{otherwise} \end{cases} \\
 \llbracket p[A].e \rightarrow q[B].x : k; C \rrbracket_r \\
 = \begin{cases} p[A].e \rightarrow B : k; \llbracket C \rrbracket_r & \text{if } r = p \\ A \rightarrow q[B].x : k; \llbracket C \rrbracket_r & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
 \llbracket A \rightarrow q[B].x : k; C \rrbracket_r \\
 = \begin{cases} A \rightarrow q[B].x : k; \llbracket C \rrbracket_r & \text{if } r = p \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
 \llbracket A \rightarrow q[B] : k \& \{l_i : C_i\}_{i \in I} \rrbracket_r \\
 = \begin{cases} A \rightarrow q[B] : k \& \{l_i : \llbracket C_i \rrbracket_r\}_{i \in I} & \text{if } r = q \\ \bigsqcup_{i \in I} \llbracket C_i \rrbracket_r & \text{otherwise} \end{cases}
 \end{array}$$

Process projection follows the structure of the originating choreography. In a *(start)*, we project the active process p to a request and the service processes \tilde{q} to (always-available) accepts. In a *(com)*, the sender is projected to a partial sending action and the receiver to a partial receiving action. The projections of *(sel)* and *(del)*, omitted, follow the same principle. Above we also report the rule for projecting *(com-s)* and *(com-r)* to exemplify how we treat partial choreographies: these are simply projected as they are for their respective process, following the structure of the choreography. The projections of conditionals and partial branchings are the only special cases. In a conditional, we project it as it is for the process evaluating the condition, but for all other we merge their behaviours with the *merging* partial operator \sqcup [6]. $C \sqcup C'$ is defined only for partial choreographies that define the behaviour of a single process and returns a choreography isomorphic to C and C' up to branching, where all branches with distinct labels are also included. We use \sqcup also in the projection of *(branch)* terms, where we require the behaviour of all processes not receiving the selection to be merged. As an example, the process projection for process u in the choreography C_B from our example in § 2 is:

$$\llbracket C_B \rrbracket_u = \begin{array}{l} u[\mathbb{U}] \text{ req PD} : a(k); u[\mathbb{U}].\text{prod} \rightarrow \text{PD} : k; \\ \text{PD} \rightarrow u[\mathbb{U}] : k \& \left\{ \begin{array}{l} \text{del} : \text{PD} \rightarrow u[\mathbb{U}] : k\langle k''[\mathbb{B}] \rangle; u[\mathbb{B}].\text{addr} \rightarrow \text{T} : k''; \\ \text{T} \rightarrow u[\mathbb{U}].\text{ddate} : k'', \\ \text{quit} : \mathbf{0} \end{array} \right. \end{array}$$

Using process projection, we can now define the EPP of a whole system. Since different service processes may be started through *(start)* terms on the same shared channel and play the same role, we use \sqcup for merging their behaviours into a single service. We identify these processes with the service grouping operator $\llbracket C \rrbracket_A^a$, which computes the set of all service process names in a start or a request in C on shared channel a playing role A . Formally, EPP is the endofunction $\llbracket C \rrbracket$ defined in the following.

Definition 1 (Endpoint Projection). Let $C \equiv (\nu \tilde{a}, \tilde{k}, \tilde{p}) C_f$, where C_f does not contain (res) terms. Then, the EPP of C is:

$$\llbracket C \rrbracket = (\nu \tilde{a}) \left((\nu \tilde{k}, \tilde{p}) \left(\prod_{p \in \text{fn}(C_f)} \llbracket C_f \rrbracket_p \right) \mid \prod_{a, A} \left(\bigsqcup_{p \in \llbracket C_f \rrbracket_A^a} \llbracket C_f \rrbracket_p \right) \right)$$

The EPP of a choreography C is the parallel composition of (i) the projections of all active processes and (ii) the merged projections of all service processes started under same shared channel and role. EPP respects the following Lemma, which shows that our model can adequately capture not only typical complete choreographies, but also scale down to describing the behaviour of a single endpoint.

Lemma 1 (Endpoint Choreographies). Let C be restriction-free, contain only partial terms, and be well-typed. If one of the following two conditions apply, then $C = \llbracket C \rrbracket$.

1. $C = \text{acc } q[\mathbb{B}] : a(k); C'$ and q is the only free process name in C' ;
2. otherwise, C has only one free process name.

We refer to choreographies that respect one of the two conditions above as *endpoint choreographies*. They implement either the behaviour of a single always-available service process (1), or that of a single free process (2). The EPP for these choreographies is the identity since they already model the behaviour of only one endpoint.

The projection of services may lead to undesirable behaviour if service roles for shared channels are not distributed correctly. For example, if we put the choreography C_B from § 2 in parallel with a choreography with a conflicting service on shared channel b for role R (which is internally implemented in C_B) we obtain a race condition, *even if protocols are correctly implemented*. Consider the following choreography:

$$C_R = \mathbf{acc} \ h[R] : b(k'); \text{PD} \rightarrow h[R].x : k'; \ h[R].c \rightarrow \text{PD} : k'$$

If we put the projection of C_B in parallel with that of C_R , we get a race condition between the service processes r and h for role R on shared channel b . Hence, the projection of process pd may synchronise with the service offered by C_R for creating session k' , instead of that by the projection of service process r in C_B . Consequently, C_B may not follow its intended behaviour. The distribution of service roles performed by our type system avoids this kind of situations. Observe that normal session typing cannot help us in detecting these problems, because the service process h correctly implements the same communication behaviour for session k' as service process r .

Main Theorems. We can now present our main theorems. We build our results on the foundation that the EPP of a choreography is still typable. As in previous work [16,6,7], we need to consider that in the projection of complete choreographies, due to merging, some projected processes may still offer branches that the original complete choreography has discarded with a conditional. Therefore, we state our type preservation result below under the *minimal typing* of choreographies \vdash_{\min} , in which the branches in rules $[\text{T}_{\text{SEL}}]$ and $[\text{T}_{\text{BRANCH}}]$ are typed using the respective minimal branch types.

Theorem 2 (EPP Type Preservation). *Let $\Gamma \vdash_{\min} C \triangleright \Delta$. Then, $\Gamma \vdash_{\min} \llbracket C \rrbracket \triangleright \Delta$.*

By Theorem 2, it follows that Theorem 1 applies also to the EPP of a choreography. We use this result to prove that EPP correctly implements the behaviour of the originating choreography, by establishing a formal relation between their respective semantics.

Theorem 3 (EPP Theorem). *Let $C \equiv (\nu \tilde{a}, \tilde{k}, \tilde{p}) C_f$, where C_f is restriction-free, be well-typed. Then,*

1. (Completeness) $C \xrightarrow{\lambda} C'$ implies $\llbracket C \rrbracket \xrightarrow{\lambda} \succ \llbracket C' \rrbracket$.
2. (Soundness) $\llbracket C \rrbracket \xrightarrow{\lambda} C'$ implies $C \xrightarrow{\lambda} C''$ and $\llbracket C'' \rrbracket \prec C'$.

Above, the *pruning relation* $C \prec C'$ is a strong typed bisimilarity [6] such that C has some unused branches and always-available accepts. \succ is a shortcut for \prec interpreted in the opposite direction.

Deadlock-freedom and Progress. We introduce our results on deadlock-freedom and progress mentioned in the Introduction. First, we define deadlock-freedom:

Definition 2 (Deadlock-freedom). *We say that choreography C is deadlock-free if either (i) $C \equiv \mathbf{0}$ or (ii) there exist C' and λ such that $C \xrightarrow{\lambda} C'$ and C' is deadlock-free.*

In our semantics (Fig. 3) complete terms can always be executed; therefore, choreographies that do not contain partial terms, or *complete choreographies*, are deadlock-free:

Theorem 4 (Deadlock-freedom for Complete Choreographies). *Let C be a complete choreography and contain no free variable names. Then, C is deadlock-free.*

By Theorems 3 and 4 we can obtain, as a corollary, that the EPP of well-typed complete choreographies never deadlock.

Corollary 1 (Deadlock-freedom for EPP). *Let C be a complete choreography, contain no free variable names, and be well-typed. Then, $\llbracket C \rrbracket$ is deadlock-free.*

Our model can also be used to talk of deadlock-freedom *compositionally*. In a compositional setting, a choreography may get stuck because of partial actions that need to be executed in parallel composition with other choreographies. We say that a choreography can *progress* if it can be composed with another choreography such that (i) all free names can be restricted and the resulting system is still well-typed, ensuring that protocols are implemented correctly; and (ii) the composition is deadlock-free. Differently from deadlock-freedom for complete choreographies, progress for partial choreographies does not follow directly from the semantics. For example, the following choreography does not have the progress property:

$$A \rightarrow q[B] : k; p[A].e \rightarrow B : k$$

Above, q is waiting for a message on session k from A , but that role is implemented by process p in the continuation. Thus, the two partial actions will never synchronise. As shown in § 4, our type system takes care of checking that roles in sessions or services are distributed correctly, avoiding cases such as this one and ensuring progress. In general, if a well-typed choreographies does not contain inner (*par*) terms we know that it can progress, since role distribution ensures that there exists a compatible environment.

Theorem 5 (Progress for Partial Choreographies). *Let C be a choreography, be well-typed, and contain no (*par*) terms. Then, there exists C' such that $(\nu \tilde{r})(C \mid C')$ with $\tilde{r} = \text{fn}(C \mid C')$, is well-typed and deadlock-free.*

By Theorems 2 and 5, it follows as a corollary that also the EPP of a well-typed choreography can progress:

Corollary 2 (Progress for EPP). *Let C contain no free variable names, be well-typed, and contain no (*par*) terms. Then, there exists C' such that $(\nu \tilde{r})(\llbracket C \rrbracket \mid C')$ with $\tilde{r} = \text{fn}(C \mid C')$, is well-typed and deadlock-free.*

Correctness of Choreography Composition. We end this section by presenting results that allow to reason about the composition of choreographies.

Lemma 2 (Compositional EPP). *Let $C = C_1 \mid C_2$ be well-typed. Then, $\llbracket C \rrbracket \equiv \llbracket C_1 \rrbracket \mid \llbracket C_2 \rrbracket$.*

By combining Lemma 2 with the Theorems shown so far, we get the following corollary, which summarises the properties for well-typed compositions of choreographies.

Corollary 3 (Compositional Choreographies). *Let $C \mid C'$ be well-typed. Then,*

1. (*EPP Type Preservation*) $\llbracket C \rrbracket \mid \llbracket C' \rrbracket$ is well-typed.
2. (*Completeness*) $C \mid C' \xrightarrow{\lambda} C''$ implies $\llbracket C \rrbracket \mid \llbracket C' \rrbracket \xrightarrow{\lambda} \succ \llbracket C'' \rrbracket$.
3. (*Soundness*) $\llbracket C \rrbracket \mid \llbracket C' \rrbracket \xrightarrow{\lambda} C''$ implies $C \xrightarrow{\lambda} C'''$ and $\llbracket C''' \rrbracket \prec C''$.

Our corollary above formally addresses the issues mentioned in the Introduction. Choreographies (C and C' in the corollary) can be developed independently and then their respective projections can be composed.

6 Related Work

Previous works have tackled the problem of defining a formal model for choreographies and giving a correct EPP [7,6]. The main difference wrt our work is compositionality: previous models can only capture closed systems, and do not treat a methodology for composing choreographies. A major difficulty wrt composition given by the approach in [7] is that the EPP of a choreography could be untypable with known type systems for session types. Typability of EPP is important to achieve composition, since a programmer may need to reuse a choreography *after* it has been projected. [7] is the only previous work providing an asynchronous semantics for multiparty sessions in choreographies; however, asynchrony is modelled in two different ways in the choreography model and the endpoint model, raising complexity. As a consequence, the EPP Theorem in [7] has a more complex formulation with weak transitions and confluence, whereas ours can be formulated in a stronger form where EPP mimics its original choreography step by step. [6] preserves typability of projections but does not handle neither asynchrony nor multiparty sessions; instead, they type choreographies with binary sessions. We have shown that choreographies can be made compositional by introducing partial terms to perform message passing with the environment, and that it is possible to ensure typability of EPP in a multiparty and asynchronous setting. This is the first work introducing a compositional multiparty session typing for choreographies, exploiting the projection of global types onto local types. Finally, neither of [7,6] handles shared channel passing, and does not treat how to handle delegation in a compositional setting, where sessions may be delegated to external or internal processes.

Multiparty session types have been previously used for typing endpoint programs [13,3,9]. In our setting, endpoint programs can be captured as special cases of partial choreographies. Our global types are taken from [3]. Differently from our framework, these works capture asynchronous communications with dedicated processes that model order-preserving message queues. An approach more similar to ours can be found in the notion of *delayed input* presented in [17]. [3] defines a type system for progress by building additional restrictions on top of standard multiparty session typing; our model yields a simpler analysis, since we can rely on the fact that complete terms in a choreography do not get stuck. Nevertheless, [3] can capture sessions started by more than one active thread. We leave an extension of our model in this direction as future work.

In [2] the authors use a concept similar to our partial choreographies for protocol specifications, to allow a single process to implement more than one role in a protocol. Differently from our approach, these are not fully-fledged system implementations but abstract behavioural types, which are then used to type check endpoint code. In our setting, the techniques in [2] can be seen as a more flexible way of handling the projection from global types to local types. An extension of our type system to allow for a process to play more than one role in a session as in [2,9] is an interesting future work.

The relationship between choreographies and endpoints has been explored in, among others, [5,16,13,6,7]. Our work distinguishes itself by adopting the same calculus for describing choreographies and endpoints, simplifying the technical development.

Acknowledgements. Yoshida has been partially supported by the Ocean Observatories Initiative and EPSRC EP/K011715/1, EP/K034413/1 and EP/G015635/1.

References

1. Additional Resources, <http://www.itu.dk/people/fabr/papers/compchor/>
2. Baltazar, P., Caires, L., Vasconcelos, V.T., Vieira, H.T.: A Type System for Flexible Role Assignment in Multiparty Communicating Systems. In: Proc. of TGC (2012)
3. Bettini, L., Coppo, M., D'Antoni, L., De Luca, M., Dezani-Ciancaglini, M., Yoshida, N.: Global progress in dynamically interleaved multiparty sessions. In: van Breugel, F., Chechik, M. (eds.) CONCUR 2008. LNCS, vol. 5201, pp. 418–433. Springer, Heidelberg (2008), Long version at <http://www.di.unito.it/~dezani/papers/cdy12.pdf>
4. Business Process Model and Notation, <http://www.omg.org/spec/BPMN/2.0/>
5. Busi, N., Gorrieri, R., Guidi, C., Lucchi, R., Zavattaro, G.: Choreography and orchestration conformance for system design. In: Ciancarini, P., Wiklicky, H. (eds.) COORDINATION 2006. LNCS, vol. 4038, pp. 63–81. Springer, Heidelberg (2006)
6. Carbone, M., Honda, K., Yoshida, N.: Structured communication-centered programming for web services. ACM Trans. Program. Lang. Syst. 34(2), 8 (2012)
7. Carbone, M., Montesi, F.: Deadlock-freedom-by-design: multiparty asynchronous global programming. In: POPL, pp. 263–274 (2013)
8. Chor. Programming Language, <http://www.chor-lang.org/>
9. Deniérou, P.-M., Yoshida, N.: Dynamic multirole session types. In: Proc. of POPL, pp. 435–446. ACM (2011)
10. Deniérou, P.-M., Yoshida, N., Bejleri, A., Hu, R.: Parameterised multiparty session types. LMCS 8(4) (2012)
11. Honda, K., Mukhamedov, A., Brown, G., Chen, T.-C., Yoshida, N.: Scribbling interactions with a formal foundation. In: Natarajan, R., Ojo, A. (eds.) ICDCIT 2011. LNCS, vol. 6536, pp. 55–75. Springer, Heidelberg (2011)
12. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type disciplines for structured communication-based programming. In: Hankin, C. (ed.) ESOP 1998. LNCS, vol. 1381, pp. 122–138. Springer, Heidelberg (1998)
13. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: Proc. of POPL, vol. 43(1), pp. 273–284. ACM (2008)
14. italianaSoftware. <http://www.italianasoftware.com/>.
15. Jolie. Java Orchestration Language Interpreter Engine, <http://www.jolie-lang.org/>
16. Lanese, I., Guidi, C., Montesi, F., Zavattaro, G.: Bridging the gap between interaction- and process-oriented choreographies. In: Proc. of SEFM, pp. 323–332. IEEE (2008)
17. Merro, M., Sangiorgi, D.: On asynchrony in name-passing calculi. Mathematical Structures in Computer Science 14(5), 715–767 (2004)
18. Montesi, F., Guidi, C., Zavattaro, G.: Composing Services with JOLIE. In: Proc. of ECOWS, pp. 13–22 (2007)
19. PI4SOA (2008), <http://www.pi4soa.org>
20. Savara. JBoss Community, <http://www.jboss.org/savara/>
21. W3C WS-CDL Working Group. Web services choreography description language version 1.0 (2004), <http://www.w3.org/TR/ws-cdl-10/>