Resolving Non-determinism in Choreographies ^{*}

Laura Bocchi¹, Hernán Melgratti², and Emilio Tuosto³

¹ Department of Computing, Imperial College London, UK

² Departamento de Computación, FCEyN, Universidad de Buenos Aires - Conicet, Argentina

³ Department of Computer Science, University of Leicester

Abstract. Resolving non-deterministic choices of choreographies is a crucial task. We introduce a novel notion of realisability for choreographies –called *whole-spectrum implementation*– that rules out deterministic implementations of roles that, no matter which context they are placed in, will never follow one of the branches of a non-deterministic choice. We show that, under some conditions, it is decidable whether an implementation is whole-spectrum. As a case study, we analyse the POP protocol under the lens of whole-spectrum implementation.

1 Introduction

The Context. A *choreography* describes the expected interactions of a system in terms of the message exchanged between its components (aka *roles*):

“Using the Web Services Choreography specification, a contract containing a global definition of the common ordering conditions and constraints under which messages are exchanged, is produced [...]. Each party can then use the global definition to build and test solutions that conform to it. The global specification is in turn realised by combination of the resulting local systems [...].”

The first part of the excerpt above taken from [15] envisages a choreography as a global contract regulating the exchange of messages; the last part identifies a distinctive element of choreographies: the global definition can be used to check the conformance of local components so to (correctly) realise the global contract. Choreographies allows for the combination of independently developed distributed components (e.g., services) while hiding implementation details. Moreover, the communication pattern specified in the choreography suffices to check each component.

For illustration, take a simple choreography, hereafter called *ATM*, involving the cash machine of a bank B and a customer C depicted as either of the following diagrams:

In the diagram on the left, the doubly stroked lines represent choices and the dashed lines connect interactions with the branches where they occur. On the right, *ATM* is expressed in terms of the conversation protocols of [12].

^{*} This work has been partially sponsored by Projects: EU 7FP Grant 295261 (MEALS), AN-PCyT BID-PICT-2008-00319, Ocean Observatories Initiative and EPSRC EP/K034413/1.

After successful authentication, B offers a deposit and an overdraft service to C. When opting for a deposit, C indicates the amount of money to be deposited. If C asks to overdraft then B can either grant or deny it; in the former case C will communicate the amount of money required.

On Realisations. A set of processes is a *realisation* of a choreography when the behaviour emerging from their concurrent execution matches the behaviour specified by the choreography. A choreography is *realisable* when it has a realisation.

A realisation of *ATM* can be given using two CCS-like processes [20] (augmented with internal $_ \oplus _$ and external $_ + _$ choice operators) for roles B and C:

$$\begin{aligned} T_B &= \text{login}.\overline{(\text{deposit.amount} + \text{overdraft}.\overline{(\text{ok.amount} \oplus \overline{\text{ko}})})} \\ T_C &= \overline{\text{login}}.\overline{(\text{deposit.amount} \oplus \text{overdraft}.\overline{(\text{ok.amount} + \text{ko})})} \end{aligned}$$

In words, T_B specifies that, after C logs in, B waits to interact either on deposit or on overdraft; in the latter case, B non-deterministically decides whether to grant or deny the overdraft; T_C is the dual of T_B . Note that *ATM* uses non-determinism to avoid specifying the criteria for B to grant or deny an overdraft. The use of non-determinism is also reflected in realisations, in fact T_B uses the internal choice operator $_ \oplus _$ to model the reaction when C requests an overdraft.

Choreographies can be interpreted either as *constraints* or as *obligations* of distributed interactions [19]. The former interpretation (aka *partial* [19] or weak [23]) admits a realisation if it exhibits a subset of the behaviour. For instance, take

$$T'_B = \text{login}.\overline{(\text{deposit.amount} + \text{overdraft}.\overline{\text{ko}})}$$

then T'_B and T_C form a partial realisation of *ATM* where requests of overdraft are consistently denied. On the contrary, when interpreting choreographies as obligations, a realisation is admissible if it is able to exhibit *all* interaction sequences (hence such realisations are also referred to as *complete* realisations [19]). For instance, T_B and T_C form a complete realisation of *ATM*.

The Problem. Choreographies typically yield non-deterministic specifications; here we explore the problem of resolving their non-determinism. In fact, despite being a valuable abstraction mechanism, non-determinism has to be implemented using deterministic constructs such as conditional branch statements.

Using again *ATM*, we illustrate that traditional notions of complete realisation are not fully satisfactory. The non-deterministic choice in T_B abstracts away from the actual conditions used in implementations to resolve the choice. This permits, e.g., different banks to adopt different policies depending, for instance, on the type of the clients' accounts. Consider the (deterministic) implementations B_1 and B_2 of T_B below (for brevity, each name refers to the interaction of *ATM* with the same initial):

$$\begin{aligned} B_i &::= l(c);(d();a(x);Q + o();P_i(c)) \quad \text{for } i = 1, 2 \quad (Q \text{ is immaterial}) \\ P_1(c) &::= \text{if } \text{check}(c) : \overline{\text{ok}.a(x)} \text{ else } \overline{\text{ko}} \quad \text{and} \quad P_2(c) ::= \overline{\text{ko}} \end{aligned}$$

The expression $\text{check}(c)$ in P_1 deterministically discriminates if the overdraft should be granted. Clearly both B_1 and B_2 can be used as implementations of T_B in *partial realisations* of the choreography.¹ (as e.g. in [11]).

¹ For instance, both B_1 and B_2 type-check against T_B considered as a session type due to the fact that subtyping for session types [13] is contra-variant with respect to internal choices (and covariant with respect to external choices).

Conversely, neither B_1 nor B_2 can be used in a *complete realisation*. This is straightforward for B_2 (unable to interact over ok after receiving an overdraft request), but not so evident for B_1 . Depending on the credentials c sent by the customer to login, $check(c)$ will evaluate either to `true` or to `false`. Therefore, B_2 will be unable to exhibit both branches. This will be the case for any possible deterministic implementation of *ATM*: only one branch will be matched. Consequently, there is not a complete, deterministic realisation for *ATM*.

We prefer B_1 to B_2 arguing that they are not equally appealing when interpreting choreographies as obligations. In fact, B_2 consistently precludes one of the alternatives while B_1 guarantees only one or the other alternative (provided that *check* is not the constant map) depending on the deterministic implementation of the role T_c .

Contributions and Synopsis. We introduce *whole-spectrum implementation* (WSI), a new interpretation of choreographies as interaction obligations. A WSI of a role R guarantees that, whenever the choreography allows R to make an internal choice, there is a context (i.e., an implementation of the remaining roles) for which (the implementation of) R chooses such alternative. We illustrate the use of WSI to analyse the POP2 protocol (i.e., choreography § 2.2, implementation § 3.1, and verification § 5.1).

We develop our results in a behavioural typing framework since types directly relate specifications to implementations, but our results can be established in different contexts (c.f. [3, Appendix F]). Our technical contributions are a formalisation of WSI and a sound type system that guarantees that typable processes form WSIs. For instance, our type system validates B_1 against T_B while it discards B_2 . Typing is decidable if so is the logic expressing internal conditions. We relate a denotational semantics of global types (featuring optional behaviours) to the operational semantics of local types (c.f. Thm. 3). Finally, the strong connection between local types and processes ensures that well-typed processes enjoy whole-spectrum implementability (c.f. Thm. 4).

2 Global and Local Types

Our types elaborate from [18] and use a more tractable form of iteration (discussed below). We fix a countably infinite set $\mathbb{C}$ of (*session channel*) *names* ranged over by $u, y, s, \dots$ and a countably infinite set $\mathbb{P}$ of (*participants*) *roles* ranged over by $p, q, r, \dots$ (with $\mathbb{C} \cap \mathbb{P} = \emptyset$). Basic data types, called *sorts*, (e.g., booleans `Bool`, integers `Int`, strings `Str`, record types, etc.) are assumed; $\mathbb{U}$ ranges over sorts.

Tuples are written in bold font and, abusing notation, we use them to represent their underlying set (e.g., if $\mathbf{y} = (y_1, y_2, y_3)$, we write $y_2 \in \mathbf{y}$ for $y_2 \in \{y_1, y_2, y_3\}$). Let $\#X$ denote the cardinality of a set X . Write $\{-/-\}$ for substitutions and in $\{\mathbf{y}/\mathbf{s}\}$ assume that $\mathbf{s}$ and $\mathbf{y}$ have the same length, that the components of $\mathbf{y}$ are pairwise disjoint, and that the i -th element of $\mathbf{y}$ is replaced by the i -th element of $\mathbf{s}$.

2.1 Types

A *global type term* (GTT, for short) G is derived by the following grammar:

$$G ::= p \rightarrow q : y \langle \mathbf{U} \rangle \mid G + G \mid G \mid G \mid G; G \mid G^{*f} \mid \text{end}$$

In words, a GTT can either be a single interaction, the non-deterministic ($_ + _$), parallel ($_ \mid _$), or sequential ($_ ; _$) composition of two GTTs, the iteration of a GTT ($_^*$), or the empty term. Hereafter, we tacitly assume $p \neq q$ in any interaction $p \rightarrow q : y \langle U \rangle$. As in [7], we adopt a form of iteration to statically check for WSI (see § 4); in G^{*f} , f injectively maps roles in G to pairs of channels and sorts; i.e., $f(p) = y \langle U \rangle$ is used to notify $p \in G$ when the iteration ends. We use $\text{cod}(f)$ to denote the set of channels appearing as first component in the image of f .

For a GTT G , $\text{ch}(G) \subseteq \mathbb{C}$ are the names, $\mathcal{P}(G)$ are the participants, and $\text{fst}(G)$ are the initially enabled input and output actions of each each participant in G ; e.g., in

$$G_f = p \rightarrow q : y \langle U \rangle; q \rightarrow s : z \langle U \rangle \quad (2.1)$$

$\text{ch}(G_f) = \{y, z\}$, $\mathcal{P}(G_f) = \{p, q, s\}$, and $\text{fst}(G_f) = \{(p, \bar{y}), (q, y), (s, z)\}$. Formal definitions of such maps are standard and relegated in [3, Appendix A].

A *global type* is defined by an equation $\mathcal{G}(\mathbf{y}) \triangleq G$ where $\mathbf{y} \subseteq \mathbb{C}$ are pairwise distinct names and $\text{ch}(G) \subseteq \mathbf{y}$. The syntax of global types explicitly mentions names as they are needed when typing processes to check if they form a WSI (c.f. § 5). We write $\mathcal{G}(\mathbf{y})$ when the defining equation of a global type is understood or its corresponding GTT is immaterial; we write $\mathcal{G}$ or G instead of $\mathcal{G}(\mathbf{y})$ when parameters are understood.

GTTs are taken up to *structural congruence*, defined as the smallest congruence $\equiv$ such that $_ ; _$, $_ \mid _$, and $_ + _$ form a monoid with identity end and $_ \mid _$ and $_ + _$ are commutative. Two global types $\mathcal{G}_1(\mathbf{y}_1) \triangleq G_1$ and $\mathcal{G}_2(\mathbf{y}_2) \triangleq G_2$ are structurally equivalent when $G_1 \equiv G_2\{\mathbf{y}_2/\mathbf{y}_1\}$, in which case we write $\mathcal{G}_1 \equiv \mathcal{G}_2$.

We define the set of *ready participants* of G as follows.

$$\begin{aligned} \text{rdy}(p \rightarrow q : y \langle U \rangle) &= \{p\} & \text{rdy}(G + G') &= \text{rdy}(G \mid G') = \text{rdy}(G) \cup \text{rdy}(G') & \text{rdy}(\text{end}) &= \emptyset \\ \text{rdy}(G; G') &= \text{rdy}(G), \text{ if } \text{rdy}(G) \neq \emptyset & \text{rdy}(G; G') &= \text{rdy}(G'), \text{ if } \text{rdy}(G) = \emptyset & G^{*f} &= \text{rdy}(G) \end{aligned}$$

(note that for the GTT (2.1) $\text{rdy}(G_f) = \{p\}$). We extend $\mathcal{P}(_)$ and $\text{rdy}(_)$ to global types $\mathcal{G}(\mathbf{y}) \triangleq G$ by defining $\mathcal{P}(\mathcal{G}) = \mathcal{P}(G)$ and $\text{rdy}(\mathcal{G}) = \text{rdy}(G)$.

As customary in session types, we restrict the attention to *well-formed* global types in order to rule out specifications that cannot be implemented distributively. A global type is *well-formed* when it enjoys the following properties: *linearity*, *single threadness*, *single selector* [14], *knowledge of choice* [14, 7], and *single iteration controller*. All but the last condition are standard. The last condition is specific to our form of iteration; informally, it requires that in each iteration there is a unique participant that decides when to exit the loop (see [3, Appendix B] for its definition).

A *local type term* (LTT for short) T is derived by the following grammar:

$$T ::= \bigoplus_{i \in I} y_i ! U_i; T_i \mid \sum_{i \in I} y_i ? U_i; T_i \mid T_1; T_2 \mid T^* \mid \text{end}$$

An LTT is either an internal ($\bigoplus$) or external ($\sum$) guarded choice, the sequential composition of LTTs $_ ; _$, an iteration $_^*$, or the empty term end . The set $\text{ch}(T)$ of channels of T is standard (see [3, Appendix A]).

A *local type* is defined by an equation $\mathcal{T}(\mathbf{y}) \triangleq T$ where $\mathbf{y}$ are pairwise distinct names and $\text{ch}(T) \subseteq \mathbf{y}$. Hereafter, we write $\mathcal{T}(\mathbf{y})$ when the defining equation of a local type

is understood or its corresponding LTT is immaterial; we may write $\mathcal{T}$ or T instead of $\mathcal{T}(\mathbf{y})$ when parameters are understood. We overload $\equiv$ to denote the structural congruence over local types defined as the least congruence such that internal and external choice are associative, commutative and have end as identity, while $;$ is associative. In the following, we consider types up-to structural congruence.

The projection operation extracts the local types from a global type. For a well-formed GTT G and $\mathbf{r} \in \mathbb{P}$, $G \upharpoonright \mathbf{r}$ is the *projection* of G on $\mathbf{r}$ and it is defined homomorphically on $;$, $\oplus$, $\rightarrow$, $+$, $?$, and $!$, and as follows on the remaining constructs:

$$G \upharpoonright \mathbf{r} = \begin{cases} y!U & (\text{resp. } y?U) & \text{if } G = \mathbf{r} \rightarrow \mathbf{p} : y \langle U \rangle \quad (\text{resp. if } G = \mathbf{p} \rightarrow \mathbf{r} : y \langle U \rangle) \\ G_i \upharpoonright \mathbf{r} & & \text{if } G = G_1 \mid G_2 \text{ and } \mathbf{r} \notin \mathcal{P}(G_j) \text{ with } j \neq i \in \{1, 2\} \\ (G_1 \upharpoonright \mathbf{r})^*; b_1!U_1; \dots; b_n!U_n & & \text{if } G = G_1^{*f}, \text{cod}(f) = \{b_1 \langle U_1 \rangle, \dots, b_n \langle U_n \rangle\}, \text{ and } \mathbf{r} \in \text{rdy}(G_1) \\ (G_1 \upharpoonright \mathbf{r})^*; b?U & & \text{if } G = G_1^{*f}, f(\mathbf{r}) = b \langle U \rangle, \text{ and } \mathbf{r} \notin \text{rdy}(G_1) \\ \text{end} & & \text{if } G = \mathbf{p} \rightarrow \mathbf{q} : y \langle U \rangle \text{ and } \mathbf{r} \neq \mathbf{p}, \mathbf{q} \text{ or if } G = \text{end} \\ \text{end} & & \text{if } G = G_1^{*f} \text{ and } \mathbf{r} \notin \mathcal{P}(G_1) \text{ or } f(\mathbf{r}) \text{ is undefined} \end{cases}$$

Our projection is total on well-formed global types. All but the clauses for the projections of iteration in the definition of $;$ are straightforward (c.f. [14]). Each iteration has a unique participant $\mathbf{r} \in \text{rdy}(G_1)$ (by well-formedness) dictating when to stop the iteration, and a number of ‘passive’ participants. Projection sends messages from $\mathbf{r}$ to each passive participant to signal the termination of the iteration. The *projection* $G(\mathbf{y}) \upharpoonright \mathbf{r}$ of a global type $G(\mathbf{y}) \triangleq G$ with respect to $\mathbf{r}$ is a local type $\mathcal{T}(\mathbf{y}) \triangleq T$ where $T = G \upharpoonright \mathbf{r}$.

Example 1. Let $G = G_f^{*f}$, with G_f defined in (2.1), $f(\mathbf{q}) = b_1 \langle U_1 \rangle$ and $f(\mathbf{s}) = b_2 \langle U_2 \rangle$. Then, the projections of G are

$$G \upharpoonright \mathbf{p} = (y!U)^*; b_1!U_1; b_2!U_2 \quad G \upharpoonright \mathbf{q} = (y?U; z!U)^*; b_1?U_1 \quad G \upharpoonright \mathbf{s} = (z?U)^*; b_2?U_2$$

2.2 Running Example

We illustrate our approach on a real yet tractable protocol, the Post Office Protocol - Version 2 (POP2) [5] between a client and a mail server. We describe POP2 with the following choreography where $G_{\text{EXIT}} = S \rightarrow C : \text{BYE} \langle \rangle$:

$$\begin{aligned} G_{\text{POP}} &= C \rightarrow S : \text{QUIT} \langle \rangle; G_{\text{EXIT}} + C \rightarrow S : \text{HELO} \langle \text{Str} \rangle; G_{\text{MBOX}} \\ G_{\text{MBOX}} &= S \rightarrow C : \text{R} \langle \text{Int} \rangle; G_{\text{NMBR}} + S \rightarrow C : \text{E} \langle \rangle; G_{\text{EXIT}} \\ G_{\text{NMBR}} &= (C \rightarrow S : \text{FOLD} \langle \text{Str} \rangle; S \rightarrow C : \text{R} \langle \text{Int} \rangle) \\ &\quad + C \rightarrow S : \text{READ} \langle \text{Int} \rangle; S \rightarrow C : \text{R} \langle \text{Int} \rangle; G_{\text{SIZE}}^{*S \rightarrow \text{QUIT} \langle \rangle}; G_{\text{EXIT}} \\ G_{\text{SIZE}} &= (C \rightarrow S : \text{RETR} \langle \rangle; S \rightarrow C : \text{MSG} \langle \text{Data} \rangle; G_{\text{XFER}} \\ &\quad + C \rightarrow S : \text{READ} \langle \text{Int} \rangle; S \rightarrow C : \text{R} \langle \text{Int} \rangle)^{*S \rightarrow \text{FOLD} \langle \text{Str} \rangle}; S \rightarrow C : \text{R} \langle \text{Int} \rangle \\ G_{\text{XFER}} &= C \rightarrow S : \text{ACKS} \langle \rangle; S \rightarrow C : \text{R} \langle \text{Int} \rangle + C \rightarrow S : \text{ACKD} \langle \rangle; S \rightarrow C : \text{R} \langle \text{Int} \rangle \\ &\quad + C \rightarrow S : \text{NACK} \langle \rangle; S \rightarrow C : \text{R} \langle \text{Int} \rangle \end{aligned}$$

The protocol G_{POP} starts with C sending S either an empty message along channel QUIT to quit the session, or a string on channel HELO representing C ’s password. In the first case, the protocol ends as per G_{EXIT} while in the latter case the G_{MBOX} is executed.

In G_{MBOX} , the server S either sends the number of messages in the default mailbox or it signals an error and ends the session as per G_{EXIT} . In the former case, G_{NMBR} establishes that C repeatedly asks either (a) to enter a folder (sending the folder's name on FOLD) and then receiving back the number of messages in that folder, or (b) to request a message by sending its index along READ and then receiving back the length of the message. In case (a), the loop is immediately repeated after S 's reply, in case (b) the protocol continues as G_{SIZE} where another loop starts with C either (a) retrieving the message or (b) asking for another message (by interacting again on READ). For (a), C signals on RETR that it is ready to receive data that are sent by S on MSG (sort Data abstracts away the format of messages specified in [10]); after these interactions the choreography continues as G_{XFER} where the transmission is acknowledged by C with the interactions in G_{XFER} : ACKS keeps the message in the mailbox, ACKD deletes the message, NACK notifies that the message has not been received and must be kept in the mailbox; after any acknowledgement, S sends C the length of the next message. After some iterations in G_{SIZE} , C specifies a different folder and repeats G_{NMBR} .

The projection $T_S = G_{\text{POP}} \upharpoonright S$ of G_{POP} onto the server is below; $G_{\text{POP}} \upharpoonright C$ is dual.

$$\begin{aligned}
 T_S &= \text{QUIT?}; T_{\text{EXIT}} & T_{\text{NMBR}} &= (\text{FOLD?Str}; R!Int + \text{READ?Int}; R!Int; T_{\text{SIZE}})^*; \\
 &+ \text{HELO?Str}; T_{\text{MBOX}} & & \text{QUIT?}; T_{\text{EXIT}} \\
 T_{\text{MBOX}} &= R!Int; T_{\text{NMBR}} \oplus E!; T_{\text{EXIT}} & T_{\text{SIZE}} &= (\text{RETR?}; \text{MSG!Data}; T_{\text{XFER}} + \text{READ?Int}; R!Int)^*; \\
 T_{\text{EXIT}} &= \text{BYE!} & & \text{FOLD?Str}; R!Int \\
 & & T_{\text{XFER}} &= \text{ACKS?}; R!Int + \text{ACKD?}; R!Int + \text{NACK?}; R!Int
 \end{aligned}$$

The messages in T_S are as in G_{POP} and S iterates until either a signal on QUIT or on FOLD is sent by C .

In Ex. 2 we present, for illustrative purpose, a multiparty variant of G_{POP} where the authentication is outsourced.

Example 2. A multiparty variant of POP2 is given by G'_{POP} below where S uses a third-party authentication service A :

$$\begin{aligned}
 G'_{\text{POP}} &= C \rightarrow S : \text{QUIT} \langle \rangle; G_{\text{EXIT}} + C \rightarrow S : \text{HELO} \langle \text{Str} \rangle; G'_{\text{MBOX}} \\
 G'_{\text{MBOX}} &= S \rightarrow A : \text{REQ} \langle \text{Str} \rangle; A \rightarrow S : \text{RES} \langle \text{Bool} \rangle; \\
 &S \rightarrow C : R \langle \text{Int} \rangle; G_{\text{NMBR}} + S \rightarrow C : E \langle \rangle; G_{\text{EXIT}}
 \end{aligned}$$

where, on RES , A sends the result of the authentication of C (G_{NMBR} and G_{EXIT} remain unchanged). The projection of G'_{POP} on S is

$$\begin{aligned}
 T'_S &= \text{QUIT?}; T_{\text{EXIT}} + \text{HELO?Str}; T_{\text{AUTH}} & \blacklozenge \\
 T_{\text{AUTH}} &= \text{REC!Str}; \text{RES?Bool}; T'_{\text{MBOX}} & T'_{\text{MBOX}} = R!Int; T_{\text{NMBR}} \oplus E!; T_{\text{EXIT}}
 \end{aligned}$$

2.3 Behaviour of Types

The semantics of local types is given in terms of *specifications*, namely pairs of partial functions Γ and Δ such that Γ maps session names to global types and names to sorts, and Δ maps tuples of session names to local types. We use $\Gamma \bullet \Delta$ to denote a specification and adopt the usual syntactic notations for environments:

$$\Gamma ::= \emptyset \mid \Gamma, u : \mathcal{G} \mid \Gamma, x : \mathcal{U} \qquad \Delta ::= \emptyset \mid \Delta, s : \mathcal{T}$$

$$\begin{array}{c}
\frac{\Gamma(u) \equiv \mathcal{G}(y)}{\Gamma \bullet \Delta \xrightarrow{\bar{u}^y} \Gamma \bullet \Delta, y : \mathcal{G}(y) \downarrow 0} [\text{TReq}] \\
\frac{v : U_j \quad s_j \in s \quad j \in I}{\Gamma \bullet \Delta, s : \bigoplus_{i \in I} s_i!U_i; \mathcal{T}_i \xrightarrow{\bar{s}_j^y} \Gamma \bullet \Delta, s : \mathcal{T}_j} [\text{TSend}] \\
\frac{\Gamma \bullet \Delta, s : \mathcal{T} \xrightarrow{\alpha} \Gamma \bullet \Delta, s : \mathcal{T}'}{\Gamma \bullet \Delta, s : \mathcal{T}; \mathcal{T}'' \xrightarrow{\alpha} \Gamma \bullet \Delta, s : \mathcal{T}'; \mathcal{T}''} [\text{TSeq}] \\
\Gamma \bullet \Delta, s : \mathcal{T}^* \xrightarrow{\tau} \Gamma \bullet \Delta, s : \text{end} [\text{TLoop1}]
\end{array}
\quad
\begin{array}{c}
\frac{\Gamma(u) \equiv \mathcal{G}(y)}{\Gamma \bullet \Delta \xrightarrow{u^y} \Gamma \bullet \Delta, y : \mathcal{G}(y) \downarrow i} [\text{TAcc}] \\
\frac{v : U_j \quad s_j \in s \quad j \in I}{\Gamma \bullet \Delta, s : \sum_{i \in I} s_i?U_i; \mathcal{T}_i \xrightarrow{s_j^y} \Gamma \bullet \Delta, s : \mathcal{T}_j} [\text{TRec}] \\
\frac{\Gamma \bullet \Delta_1 \xrightarrow{\tau} \Gamma \bullet \Delta'_1}{\Gamma \bullet \Delta_1, \Delta_2 \xrightarrow{\tau} \Gamma \bullet \Delta'_1, \Delta_2} [\text{TPar}] \\
\Gamma \bullet \Delta, s : \mathcal{T}^* \xrightarrow{\tau} \Gamma \bullet \Delta, s : \mathcal{T}; \mathcal{T}^* [\text{TLoop2}]
\end{array}$$

Fig. 1. Labelled transitions for specifications

as usual, when writing $\Delta, s : \mathcal{T}$, $s \notin \text{dom}(\Delta)$ is implicitly assumed (likewise for $\Gamma, _ : _$) and $\Delta_1, \Delta_2 \equiv \Delta_2, \Delta_1$.

The semantics of specifications is generated by the rules in Fig. 1 using the labels

$$\alpha ::= \bar{u}^y s \mid u_i s \mid \bar{s}v \mid sv \mid \tau \quad (2.3)$$

that respectively represent the request on u for the initialisation of a session among $n + 1$ roles, the acceptance of joining a session of u as the i -th role, the sending of a value on s , the reception of a value on s , and the silent step.

Intuitively, the rules of Fig. 1 specify how a single participant behaves in a session s and are instrumental for type checking processes. Rules [TReq] and [TAcc] allow a specification to initiate a new session by projecting (on 0 and i , resp.) the global type associated² to name u in Γ . By [TSend], if types are respected, a specification can send any value on one of the names in a branch of an internal choice. Dually, [TRec] accounts for the reception of a value. Note that values occur only on the label of the transitions and are not instantiated in the local types. Rule [TSeq] is trivial. Rule [TPar] allows part of a specification to make a transition. Finally, an iterative local type can either stop by rule [TLoop1] or arbitrarily repeat itself by rule [TLoop2].

3 Processes and Systems

As we will see (Def. 1 in § 4), global types are implemented by *systems*. Our systems exchange values specified by *expressions* having the following syntax:

$$e ::= x \mid v \mid e_1 \text{ op } e_2 \quad \ell ::= [e_1, \dots, e_n] \mid e_1..e_2$$

An expression e is either a variable, or a value, or else the composition of expressions (we assume that expressions are implicitly sorted and do not include names). Lists $[e_1, \dots, e_n]$ and numerical ranges $e_1..e_2$ are used for iteration; in the former case, all the items of a list have the same sort, in the latter case, both expressions are integers and

² The use of $\equiv$ in the premises caters for α -conversion of names y . Also, $\mathbb{P}$ is the set of natural numbers (0 is the initiator of sessions) and for readability, in examples we use names to denote participants.

the value of e_1 is smaller than or equal to the value of e_2 . The empty list is denoted as ε and the operations $\text{hd}(\ell)$ and $\text{tl}(\ell)$ respectively return the head and tail of ℓ (defined as usual). We write $\text{var}(e)$ and $\text{var}(\ell)$ for the set of variables of e and ℓ .

The syntax of processes and systems below relies on *queues* of basic values M and input-guarded non-deterministic sequential process N , respectively defined as

$$M ::= \mathbf{0} \mid \mathbf{v} \cdot M \quad N ::= \sum_{i \in I} y_i(x_i); P_i$$

where $i \neq j \in I \implies y_i \neq y_j$; we define $\mathbf{0} \triangleq \sum_{i \in \emptyset} y_i(x_i); P_i$.

The syntax of systems S and processes P is

$$\begin{aligned} P, Q ::= & u_i(\mathbf{y}).P \mid \bar{u}^n(\mathbf{y}).P \mid N \mid \bar{s}e \mid \text{if } e : P \text{ else } Q \\ & \mid P;P \mid \text{for } x \text{ in } \ell : P \mid \text{do } N \text{ until } b(x) \\ S ::= & P \mid (\text{vs})S \mid S \mid S \mid s : M \end{aligned}$$

All constructions but loops are straightforward. In $\text{for } x \text{ in } \ell : P$, the body P is executed for each element in ℓ , while $\text{do } N \text{ until } b(x)$ repeats N until a message on b is received. Intuitively, the former construct is executed by the (unique) role that decides when to exit the iteration while the latter construct is used by the “passive” roles in the loop (see § 2.1 and § 5). Given a process P , $\text{fv}(P)$ denotes the set of all variables appearing outside the scope of input prefixes in P . Also, we extend $\text{var}(_)$ to systems in the obvious way. In $(\text{vs})S$, names $\mathbf{s}$ are bound (the set $\text{fc}(S)$ of free session names of S is defined as expected); a system S is *closed* when $\text{fc}(S) = \emptyset$ and it is *initial* when S does not contain runtime constructs, namely new session $(\text{vs})S'$ and queues $s : M$. Formally, S is initial iff for each s and S' , if $S \equiv (\text{vs})S'$ then $\mathbf{s} \not\subseteq \text{fc}(S')$.

The structural congruence $\equiv$ is the least congruence over systems closed with respect to α -conversion, such that $_ \mid _$ and $_ + _$ are associative, commutative and have $\mathbf{0}$ as identity, $_;$ is associative and has $\mathbf{0}$ as identity, and the following axioms hold:

$$(\text{vs})\mathbf{0} \equiv \mathbf{0} \quad (\text{vs})(\text{vs}')S \equiv (\text{vs}')(\text{vs})S \quad (\text{vs})(S \mid S') \equiv S \mid (\text{vs})S', \text{ when } \mathbf{s} \not\subseteq \text{fc}(S)$$

The operational semantics of systems is in Fig. 2 where a store σ records the values assigned to variables, $e \downarrow \sigma$ is the evaluation of e (defined if $\text{var}(e) \subseteq \text{dom}(\sigma)$ and undefined otherwise), and $\sigma[x \mapsto \mathbf{v}]$ is the update of σ at x with $\mathbf{v}$. Labels are obtained by extending the grammar in (2.3) with the production $\alpha ::= e \vdash \alpha$ where e is a boolean expression used in conditional transitions $\langle S, \sigma \rangle \xrightarrow{e \vdash \alpha} \langle S', \sigma' \rangle$ representing the fact that $\langle S, \sigma \rangle$ has an α -transition to $\langle S', \sigma' \rangle$ provided that $e \downarrow \sigma$ actually holds. We may write α instead of $\text{true} \vdash \alpha$ and $e \wedge e' \vdash \alpha$ instead of $e \vdash (e' \vdash \alpha)$.

We comment on the rules in Fig. 2 where $\text{fc}(\alpha)$ is defined as $\text{fc}(\bar{u}^n \mathbf{s}) = \text{fc}(u_i \mathbf{s}) = \{u\}$, $\text{fc}(\bar{s} \mathbf{v}) = \text{fc}(s \mathbf{v}) = \{s\}$, and $\text{fc}(\tau) = \emptyset$. Rules [SReq] and [SAcc] are for requesting and accepting new sessions; in their continuations, newly created session names $\mathbf{s}$ replace $\mathbf{y}$. Rule [SRec] is for receiving messages in an early style approach (variables are assigned when firing input prefixes); note that the store is updated by recording that x is assigned $\mathbf{v}$. Rule [SSend] is for sending values. Rules [SThen] and [SElse] handle “if” statements as expected; their only peculiarity is that the guard is recorded on the label of the transition: this is instrumental for the correspondence between systems

$\frac{s \notin \text{fc}(P)}{\langle \bar{u}^n(\mathbf{y}).P, \sigma \rangle \xrightarrow{\bar{u}^n s} \langle P\{\mathbf{y}/s\}, \sigma \rangle} [\text{SReq}]$	$\frac{\ell \downarrow \sigma \neq \varepsilon \quad \langle P, \sigma[x \mapsto \text{hd}(\ell \downarrow \sigma)] \rangle \xrightarrow{e^+ \alpha} \langle P', \sigma' \rangle}{\langle \text{for } x \text{ in } \ell : P, \sigma \rangle \xrightarrow{e^+ \alpha} \langle P'; \text{for } x \text{ in } \text{tl}(\ell) : P, \sigma' \rangle} [\text{SFor}_2]$
$\frac{s \notin \text{fc}(P)}{\langle u_i(\mathbf{y}).P, \sigma \rangle \xrightarrow{u_i s} \langle P\{\mathbf{y}/s\}, \sigma \rangle} [\text{SAcc}]$	$\frac{e \downarrow \sigma = \text{true} \quad \langle P, \sigma \rangle \xrightarrow{e^+ \alpha} \langle P', \sigma' \rangle}{\langle \text{if } e : P \text{ else } Q, \sigma \rangle \xrightarrow{e^+ \alpha} \langle P', \sigma' \rangle} [\text{SThen}]$
$\langle s(x); P + N, \sigma \rangle \xrightarrow{sv} \langle P, \sigma[x \mapsto v] \rangle [\text{SRec}]$	$\frac{e \downarrow \sigma = \text{false} \quad \langle Q, \sigma \rangle \xrightarrow{e^+ \alpha} \langle Q', \sigma' \rangle}{\langle \text{if } e : P \text{ else } Q, \sigma \rangle \xrightarrow{e^+ \alpha} \langle Q', \sigma' \rangle} [\text{SElse}]$
$\frac{e \downarrow \sigma = v}{\langle \bar{s}e, \sigma \rangle \xrightarrow{\bar{s}v} \langle \mathbf{0}, \sigma \rangle} [\text{SSend}]$	$\langle \text{do } P \text{ until } b(x), \sigma \rangle \xrightarrow{bv} \langle \mathbf{0}, \sigma[x \mapsto v] \rangle [\text{SLoop}_1]$
$\frac{\langle P, \sigma \rangle \xrightarrow{e^+ \alpha} \langle P', \sigma' \rangle}{\langle P; Q, \sigma \rangle \xrightarrow{e^+ \alpha} \langle P'; Q, \sigma' \rangle} [\text{SSeq}]$	$\frac{\langle P, \sigma \rangle \xrightarrow{e^+ \alpha} \langle P', \sigma' \rangle \quad b \notin \text{fc}(\alpha)}{\langle \text{do } P \text{ until } b, \sigma \rangle \xrightarrow{e^+ \alpha} \langle P'; \text{do } P \text{ until } b, \sigma' \rangle} [\text{SLoop}_2]$
$\frac{\ell \downarrow \sigma = \varepsilon}{\langle \text{for } x \text{ in } \ell : P, \sigma \rangle \xrightarrow{\tau} \langle \mathbf{0}, \sigma \rangle} [\text{SFor}_1]$	$\frac{P \equiv P' \quad \langle P', \sigma \rangle \xrightarrow{e^+ \alpha} \langle Q', \sigma' \rangle \quad Q' \equiv Q}{\langle P, \sigma \rangle \xrightarrow{e^+ \alpha} \langle Q, \sigma' \rangle} [\text{SStruct}]$

$\frac{s \notin \text{fc}(P_i) \quad Q_i = P_i\{\mathbf{y}_i/s\} \text{ for } i = 0, \dots, n}{\langle \bar{u}^n(\mathbf{y}_0).P_0 \mid u_1(\mathbf{y}_1).P_1 \mid \dots \mid u_n(\mathbf{y}_n).P_n, \sigma \rangle \xrightarrow{\tau} \langle (\text{vs})(Q_0 \mid \dots \mid Q_n \mid s : \mathbf{0}), \sigma \rangle} [\text{SInit}]$	$\frac{\langle P, \sigma \rangle \xrightarrow{e^+ \bar{s}v} \langle P', \sigma' \rangle}{\langle P \mid s : M, \sigma \rangle \xrightarrow{e^+ \tau} \langle P' \mid s : M \cdot v, \sigma' \rangle} [\text{SCom}_1]$
$\frac{\langle P, \sigma \rangle \xrightarrow{e^+ \bar{s}v} \langle P', \sigma' \rangle}{\langle P \mid s : M, \sigma \rangle \xrightarrow{e^+ \tau} \langle P' \mid s : M \cdot v, \sigma' \rangle} [\text{SCom}_1]$	$\frac{\langle P, \sigma \rangle \xrightarrow{e^+ sv} \langle P', \sigma' \rangle}{\langle P \mid s : v \cdot M, \sigma \rangle \xrightarrow{e^+ \tau} \langle P' \mid s : M, \sigma' \rangle} [\text{SCom}_2]$
$\frac{\langle S, \sigma \rangle \xrightarrow{e^+ \alpha} \langle S', \sigma' \rangle \quad s \notin \text{fc}(\alpha)}{\langle (\text{vs})S, \sigma \rangle \xrightarrow{e^+ \alpha} \langle (\text{vs})S', \sigma' \rangle} [\text{SNews}]$	$\frac{\langle S_1, \sigma \rangle \xrightarrow{e^+ \alpha} \langle S'_1, \sigma' \rangle \quad \text{var}(S_1) \cap \text{var}(S_2) = \emptyset}{\langle S_1 \mid S_2, \sigma \rangle \xrightarrow{e^+ \alpha} \langle S'_1 \mid S_2, \sigma' \rangle} [\text{SPar}]$

Fig. 2. Labelled transitions for processes (top) and systems (bottom)

and their types (c.f. § 6). Rules [SFor₁], [SFor₂], [SLoop₁], [SLoop₂] unfold the corresponding iterative program in an expected way. Except for session initialisation, the remaining rules are standard. Rule [SInit] allows n roles to synchronise with $\bar{u}^n(\mathbf{y}_0).P_0$; in the continuation of each role i , the bound names $\mathbf{y}_i$ are replaced with a tuple of freshly chosen session names for which the corresponding queues are created. Such queues are used to exchange values as prescribed by rules [SCom₁] and [SCom₂]. Rule [SNews] is standard and allows an action α to pass a restriction that does not involve the names of α . Rule [SInit] requires the synchronisation of all roles. Since processes are single-threaded, this is only possible when each process plays exactly one role in that session. Note that the semantics relies on a global store σ . However, the condition $\text{var}(S_1) \cap \text{var}(S_2) = \emptyset$ in rule [SPar] ensures that each program has its own local (logical) store (i.e., there is no confusion between local variables of different programs).

Note that, in a sequential composition $P; Q$, the store σ allows us to extend the scope of names bound in P by input prefixes to Q .

3.1 Running Examples

In Ex. 3 we give the implementation of T_S (i.e., participant S of G_{POP}) from § 2.2. To ease the presentation, we use the following auxiliary functions.

- $\text{auth} : \text{Str} \rightarrow \text{Bool}$ that is used for authenticating clients;
- $\text{fn} : \text{Str} \rightarrow \text{Int}$ that given a folder name returns the number of messages in that folder (we assume inbox to be the default folder);
- $\text{mn} : \text{Int} \rightarrow \text{Int}$ that given a message number returns its length (in bytes);
- $\text{data} : \text{void} \rightarrow \text{Data}$ that returns the current message;
- $\text{next} : \text{void} \rightarrow \text{Int}$ that returns the next message number;
- $\text{del} : \text{void} \rightarrow \text{Int}$ that returns the next message number and deletes the current message from the folder.

Let s_k denote the name in $\mathbf{s}$ corresponding to channel k in G_{POP} and likewise for G'_{POP} .

Example 3. The process P_{INIT} below implements POP2's server.

$$\begin{aligned}
 P_{\text{INIT}} &= u_{\mathbf{s}}(\mathbf{s}).P_{\mathbf{s}} & P_{\mathbf{s}} &= \text{squit}(); P_{\text{EXIT}} + s_{\text{HELO}}(x); P_{\text{MBOX}} & P_{\text{EXIT}} &= \overline{s_{\text{BYE}}} \\
 P_{\text{MBOX}} &= \text{if } \text{auth}(x) : \overline{s_{\text{Rfn}}}(\text{inbox}); P_{\text{NMBR}} \text{ else } \overline{s_{\text{E}}}; P_{\text{EXIT}} \\
 P_{\text{NMBR}} &= \text{do}(s_{\text{FOLD}}(x); \overline{s_{\text{Rfn}}}(x) + s_{\text{READ}}(x'); \overline{s_{\text{Rmn}}}(x'); P_{\text{SIZE}}) \text{ until } \text{squit}(); P_{\text{EXIT}} \\
 P_{\text{SIZE}} &= \text{do}(s_{\text{RETR}}(); \overline{s_{\text{MSG}}}\text{data}()); P_{\text{XFER}} + s_{\text{READ}}(x'); \overline{s_{\text{Rmn}}}(x')) \text{ until } s_{\text{FOLD}}(x); \overline{s_{\text{Rfn}}}(x) \\
 P_{\text{XFER}} &= s_{\text{ACKS}}(); \overline{s_{\text{Rmn}}}(\text{next}()) + s_{\text{ACKD}}(); \overline{s_{\text{Rmn}}}(\text{del}()) + s_{\text{NACK}}(); \overline{s_{\text{Rmn}}}(x')
 \end{aligned}$$

Firstly, P_{INIT} initiates a session of type G_{POP} as $\mathbf{S}$ then it behaves according to $T_{\mathbf{s}}$. The non-deterministic choice is resolved in the conditional statement of P_{MBOX} . ♦

Ex. 4 gives an implementation of the server $T'_{\mathbf{s}}$ of the multiparty variant of POP2.

Example 4. Let G'_{POP} be as in Ex. 2 and $P'_{\text{INIT}} = u_{\mathbf{s}}(\mathbf{s}).P'_{\mathbf{s}}$ where

$$\begin{aligned}
 P'_{\mathbf{s}} &= \text{squit}(); P_{\text{EXIT}} + s_{\text{HELO}}(x); P_{\text{AUTH}} \\
 P_{\text{AUTH}} &= \overline{s_{\text{REQ}}}(x); s_{\text{RES}}(y); P'_{\text{MBOX}} \\
 P'_{\text{MBOX}} &= \text{if } \text{auth}(x) \wedge y : \overline{s_{\text{Rfn}}}(\text{inbox}); P_{\text{NMBR}} \text{ else } \overline{s_{\text{E}}}; P_{\text{EXIT}}
 \end{aligned}$$

Here, P'_{INIT} resolves the non-deterministic choice in P'_{MBOX} by taking into account both the value returned by $\text{auth}(_)$ and the feedback of $\mathbf{A}$ stored in variable y . ♦

4 Whole-Spectrum Implementation

Definition 1 below introduces the notion of candidate implementation of a global type, that is a system consisting of one process for each role in the global type.

Definition 1 (Implementation). Given $\mathcal{G}(\mathbf{y}) \triangleq G$ s.t. $\mathcal{P}(\mathcal{G}) = \{p_1, \dots, p_n\}$ and a mapping $\mathbf{v}$ assigning a process to each $p \in \mathcal{P}(\mathcal{G})$, a $\mathbf{v}$ -implementation of $\mathcal{G}$ is a system $I_{\mathcal{G}}^{\mathbf{v}}$ such that either (i) $I_{\mathcal{G}}^{\mathbf{v}} \equiv \mathbf{v}(p_1) \mid \dots \mid \mathbf{v}(p_n)$ and $\mathbf{y} \cap \text{fc}(\mathbf{v}(p_1)) = \dots = \mathbf{y} \cap \text{fc}(\mathbf{v}(p_n)) = \emptyset$ or (ii) $I_{\mathcal{G}}^{\mathbf{v}} \equiv (\mathbf{v}\mathbf{y})(\mathbf{v}(p_1) \mid \dots \mid \mathbf{v}(p_n) \mid \mathbf{y} : \mathbf{M})$.

In case (i) the session that implements $\mathcal{G}$ is not initiated. For simplicity, we assume that roles do not use the channels defined by the global type before initiating the corresponding session (i.e., $\mathbf{y} \cap \text{fc}(\mathbf{v}(p_i)) = \emptyset$). This is not a limitation since channel names can always be renamed to avoid clashes. Case (ii) captures already initiated sessions; wlog, we assume that the system and the global type use the same session channels $\mathbf{y}$.

We characterise WSI as a relation between the execution traces of a global type $\mathcal{G}$ and its implementations $I_{\mathcal{G}}^{\mathbf{v}}$. An execution trace of $I_{\mathcal{G}}^{\mathbf{v}}$ is a sequence of input and output actions decorated with the role that performs them (in symbols $\langle p, s!U \rangle$ and $\langle p, s?U \rangle$).

$$\begin{array}{c}
\text{Let } \langle \mathfrak{t}(\mathbf{p}), \sigma \rangle \xrightarrow{e' \vdash \alpha} \langle \mathfrak{t}'(\mathbf{p}), \sigma' \rangle \text{ stand for } \langle \mathfrak{t}(\mathbf{p}), \sigma \rangle \xrightarrow{e' \vdash \alpha} \langle P, \sigma' \rangle \text{ and } \mathfrak{t}' = \mathfrak{t}[\mathbf{p} \mapsto P] \\
\\
\frac{\langle I_G^1, \sigma \rangle \xrightarrow{e' \vdash \tau} \langle I_G^1, \sigma' \rangle \quad \langle \mathfrak{t}(\mathbf{p}), \sigma \rangle \xrightarrow{e' \vdash \alpha} \langle \mathfrak{t}'(\mathbf{p}), \sigma' \rangle}{\text{fc}(\alpha) \cap \mathbf{y} \neq \emptyset \quad r \in \mathcal{R}_u(\langle I_G^1, \sigma' \rangle) \quad \text{obj}(\alpha) : \mathbf{U}} [\text{RRInt}] \quad \frac{\langle I_G^1, \sigma \rangle \not\rightarrow}{\varepsilon \in \mathcal{R}_u(\langle I_G^1, \sigma \rangle)} [\text{RREnd}] \\
\\
\frac{\langle I_G^1, \sigma \rangle \xrightarrow{e' \vdash \alpha} \langle I_G^1, \sigma' \rangle \quad u \notin \text{fc}(\alpha) \quad \langle \mathfrak{t}(\mathbf{p}), \sigma \rangle \xrightarrow{e' \vdash \beta} \langle \mathfrak{t}'(\mathbf{p}), \sigma' \rangle}{\text{fc}(\beta) \cap \mathbf{y} = \emptyset \quad r \in \mathcal{R}_u(\langle I_G^1, \sigma' \rangle)} [\text{RRExt}] \\
\\
r \in \mathcal{R}_u(\langle I_G^1, \sigma \rangle)
\end{array}$$

Fig. 3. Runs of implementations

Definition 2 (Runs of implementations). Let I_G^1 be an implementation of $\mathcal{G}(\mathbf{y}) \triangleq G$. The set $\mathcal{R}_u(\langle I_G^1, \sigma \rangle)$ of runs of I_G^1 initiated on u with store σ is the least set closed with respect to the rules in Fig. 3. We write $\mathcal{R}_u(I_G^1)$ for $\mathcal{R}_u(\langle I_G^1, \emptyset \rangle)$. The runs of a set of implementations $\mathbb{I}$ is $\mathcal{R}_u(\mathbb{I}) = \bigcup_{I \in \mathbb{I}} \mathcal{R}_u(I)$.

Rules in Fig. 3 rely on the semantics of Fig. 2. In rule [RRInt] (where $\text{obj}(\alpha) = \mathbf{v}$ for $\alpha = \bar{s}\mathbf{v}$ or $\alpha = s\mathbf{v}$), a system reduces when some process $\mathfrak{t}(\mathbf{p})$ in the implementation interacts over a session channel (i.e., α is either $\bar{y}\mathbf{v}$ or $y\mathbf{v}$ with $y \in \mathbf{y}$). Since the action α performed by $\mathfrak{t}(\mathbf{p})$ involves a session channel of the global type, an event α associated to the role $\mathbf{p}$ is added to the trace. Note that the actual value of the message α is substituted by its type, i.e., $\alpha\{\text{obj}(\alpha)/\mathbf{U}\}$ in place of α . Rule [RREnd] is straightforward. Rule [RRExt] accounts for a computation step that does not involve session channels, i.e., an internal transition τ in a role, a communication over a channel not in $\mathbf{y}$, or a session initiation. This rule allows a process to freely initiate sessions over channels different from u (i.e., sessions that do not corresponds to the global type G). On the contrary, when a role attempts to initiate a session over u , rule [RRExt] requires all roles in the implementation to initiate the session (this behaviour is imposed by the premise $u \notin \text{fc}(\alpha)$). We assume that any role in the implementation will execute exactly one action over the channel u which also matches the role assigned by $\mathfrak{t}$. Nested sessions are handled by assuming that all sessions are created over different channels that have the same type. This is just a technical simplification analogous to the possibility of having annotations to indicate the particular instance of the session under analysis.

For global types, we deviate from standard definition of traces [9,7] and use, for technical convenience, *annotated traces* that distinguish mandatory from optional actions. We write $[r]$ to denote the optional sequence r . Moreover, we consider an asynchronous communication model as in [16] and a trace implicitly denotes the equivalence class of all traces obtained by permuting causally independent actions.

Definition 3 (Runs of a global type). Given a global type term G , the set $\mathcal{R}(G)$ denotes the runs allowed by G and is defined as the least set closed under the rules in Fig. 4.

The first four rules are straightforward. Rule [RGPar] considers just the sequential composition of the traces corresponding to the two parallel branches (recall that a trace denotes an equivalence class of executions). The traces of an iterative type G^{*f} are given

$$\begin{array}{c}
\frac{}{\varepsilon \in \mathcal{R}(\text{end})} [\text{RGEnd}] \qquad \frac{}{\langle p, s!U \rangle \langle q, s?U \rangle \in \mathcal{R}(p \rightarrow q : s(U))} [\text{RGComm}] \\
\frac{r \in \mathcal{R}(G_1) \cup \mathcal{R}(G_1)}{r \in \mathcal{R}(G_1 + G_2)} [\text{RGCh}] \qquad \frac{r_1 \in \mathcal{R}(G_1) \quad r_2 \in \mathcal{R}(G_2)}{r_1 r_2 \in \mathcal{R}(G_1 ; G_2)} [\text{RGSeq}] \\
\frac{r_1 \in \mathcal{R}(G_1) \quad r_2 \in \mathcal{R}(G_2)}{r_1 r_2 \in \mathcal{R}(G_1 \mid G_2)} [\text{RGPAr}] \qquad \frac{r_1 \in \mathcal{R}(G)}{r_1 \in \tilde{\mathcal{R}}(G^{*f})} [\text{RG*1}] \quad \frac{r_1 \in \tilde{\mathcal{R}}(G^{*f}) \quad r_2 \in \mathcal{R}(G)}{[r_1]r_2 \in \tilde{\mathcal{R}}(G^{*f})} [\text{RG*2}] \\
\frac{r \in \tilde{\mathcal{R}}(G^{*f}) \quad \text{rdy}(G) = \{p\} \quad \mathcal{P}(G) = \{p, p_1, \dots, p_n\} \quad \forall 1 \leq i \leq n : f(p_i) = s_i(U_i)}{r \langle p, s_1!U_1 \rangle \dots \langle p, s_n!U_n \rangle \langle p_1, s_1?U_1 \rangle \dots \langle p_n, s_n?U_n \rangle \in \mathcal{R}(G^{*f})} [\text{RGIter}]
\end{array}$$

Fig. 4. Runs of a global type

by the rule [RGIter]; the set $\tilde{\mathcal{R}}(G^{*f})$ in the premise contains the traces of the unfolding of G^{*f} defined by the rules [RG*1] and [RG*2]. Optional events are introduced when unfolding an iterative type (rule [RG*2]). The main motivation is that an iterative type G^{*f} denotes an unbounded number of repetitions G (i.e., an infinite number of traces). Note that $\tilde{\mathcal{R}}(G^{*f}) = \{r_1, [r_1]r_2, [[r_1]r_2]r_3, \dots\}$ with $r_i \in \mathcal{R}(G)$. When implementing an iterative type, we will allow the implementation to perform just a finite number of iterations (but we require at least once iteration). Annotation of optional events are instrumental to the comparison of traces associated with iterative types (which is defined below). Rule [RGIter] adds the events associated to the termination of an iteration: (i) the ready role p sends the termination signal to any other role by using the dedicated channels specified by f (i.e., $\langle p, s_1!U_1 \rangle \dots \langle p, s_n!U_n \rangle$), and (ii) the waiting roles receive the termination message (i.e., $\langle p_1, s_1?U_1 \rangle \dots \langle p_n, s_n?U_n \rangle$). As for parallel composition, we just consider one of the possible interleavings for the receive events (that can actually happen in any order).

We use the operator $\prec$ to compare annotated traces, which is defined as the least preorder satisfying the following rules

$$[r] \prec \varepsilon \quad \varepsilon \prec r \quad r \prec r' \implies [r] \prec [r'] \quad r_1 \prec r'_1 \wedge r_2 \prec r'_2 \implies r_1 r_2 \prec r'_1 r'_2$$

Basically, $r \prec r'$ means that r' matches all mandatory actions of r and all optional actions in r' are also optional in r . Let R_1 and R_2 be two sets of annotated traces, we write $R_1 \subseteq R_2$ if $r \in R_1$ implies $\exists r' \in R_2$ such that $r \prec r'$.

Definition 4 (Whole-spectrum implementation). A set $\mathbb{I}$ of implementations covers a global type G with respect to u iff $\mathcal{R}(G) \subseteq \mathcal{R}_u(\mathbb{I})$. A process P is a whole-spectrum implementation of $p_i \in \mathcal{P}(G) = \{p_0, \dots, p_n\}$ when there exists a set $\mathbb{I}$ of implementations that covers G with respect to u s.t. $I_G^\perp \in \mathbb{I}$ implies $\mathfrak{t}(p_i) = P$.

A whole-spectrum implementation (WSI) of a role p_i is a process P such that any expected behaviour of the global type can be obtained by putting P into a proper context. For iteration types, the comparison of annotated traces implies that the implementation has to be able to perform the iteration body at least once but possibly many times.

Remark 1. A set of implementations covering a global type $\mathcal{G}$ can exhibit more behaviour than the runs of $\mathcal{G}$. Nonetheless, we use WSI with the usual soundness requirement (given in § 6) to characterise valid implementations.

5 Typing Rules

We now give a typing system to guarantee that well-typed systems are a WSI of their global type. Systems are typed by judgements of the form $C \sqsubseteq \Gamma \vdash S \triangleright \Delta * \Gamma'$ stipulating that, under condition C and environment Γ , system S is typed as Δ and yields Γ' (where environments Γ , Γ' and Δ are as in § 2.3). Condition C is called *context assumption*; it is a logical formula derivable by the grammar

$$C ::= e \mid \neg C \mid C \wedge C \quad \text{where } e \text{ is of type bool}$$

that identifies the assumptions on variables taken by processes in S . The map Γ' extends Γ with the sorts for the names bound in S . This is needed to correctly type $P; Q$ where in fact a free name of Q could be bound in P .

Due to space limits, Fig. 5 gives only the typing rules to validate processes (the rules for systems are adapted from [14] and detailed in [3, Appendix C]). Condition $C \not\vdash \perp$ is implicitly assumed among the hypothesis of each rule of Fig. 5. Rule [VReq] types session requests of the form $\bar{u}^n(\mathbf{y}).P$; its premise checks that P can be typed by extending Δ with the mapping from session names $\mathbf{y}$ to the projection of the global type $\Gamma(u)$ on the 0-th role. Dually, rule [VAcc] types the acceptance of a session request as i -th role. Rule [VRec] types an external choice $P = \sum_{i \in I} y_i(x_i); P_i$ checking that each branch P_i can be typed against the respective continuation of the type, $\Delta, \mathbf{y} : \mathcal{T}_i$ (once Γ is updated with the type assignment on the bound name x_i); rule [VRec] cannot be applied (making the validation fail) when the names in $\text{fv}(P) \cup \text{fc}(P)$ are not mapped to the same sorts in all environments Γ_i . Rule [VSend] is trivial.

Rules [VThen] and [VElse] handle the cases in which the guard of the conditional statement is either a tautology or a contradiction. Rule [VCond] ensures that both branches can be selected by fixing a proper assumption (i.e., both $C \wedge e$ and $C \wedge \neg e$ are consistent). Note that C is augmented with the condition e (resp. $\neg e$) for typing the ‘then’-branch (resp. ‘else’-branch). The resulting type is $\Delta_1 \bowtie \Delta_2$ defined in Fig. 6. The merge $\Delta_1 \bowtie \Delta_2$ is defined only when Δ_1 and Δ_2 are *compatible*, namely iff

$$\forall \mathbf{s}_1 \in \text{dom}(\Delta_1), \mathbf{s}_2 \in \text{dom}(\Delta_2) : \mathbf{s}_1 \cap \mathbf{s}_2 \neq \emptyset \implies \mathbf{s}_1 = \mathbf{s}_2$$

For $\mathbf{s} \notin \text{dom}(\Delta_1) \cap \text{dom}(\Delta_2)$, the merging behaves as the union of environments Δ_1 and Δ_2 , otherwise it returns the merging of the local types $\mathcal{T}_1 = \Delta_1(\mathbf{s})$ and $\mathcal{T}_2 = \Delta_2(\mathbf{s})$; in turn, $\mathcal{T}_1 \bowtie \mathcal{T}_2$ yields an internal choice of $\mathcal{T}_1$ and $\mathcal{T}_2$, but for a common sequence of outputs. Rule [VFor1] assigns the type $\mathcal{T}^*$ to a for loop when its body P has type $\mathcal{T}$ under C extended with $x \in \ell$, and the environment Γ extended with $x : \mathcal{U}$. Rule [VFor2] is for empty lists. By rule [VLoop], the type of a loop is $\mathcal{T}^*; b? \mathcal{U}$ when its body P has type $\mathcal{T}$ and b is the channel used to receive the termination signal. Notice that the environments of the rules [VFor1] and [VLoop] include only one session (respectively $\mathbf{y} : \mathcal{T}^*$ and $\mathbf{y} : \mathcal{T}^*; b? \mathcal{U}$), hence the body can only perform actions within a single session.

$$\begin{array}{c}
\frac{\Gamma(u) \equiv \mathcal{G}(\mathbf{y}) \quad C \sqsubseteq \Gamma \vdash P \triangleright \Delta, \mathbf{y} : \mathcal{G}(\mathbf{y}) \upharpoonright 0 * \Gamma'}{C \sqsubseteq \Gamma \vdash \bar{u}^n(\mathbf{y}).P \triangleright \Delta * \Gamma'} [\text{VReq}] \\
\frac{\Gamma(u) \equiv \mathcal{G}(\mathbf{y}) \quad C \sqsubseteq \Gamma \vdash P \triangleright \Delta, \mathbf{y} : \mathcal{G}(\mathbf{y}) \upharpoonright i * \Gamma'}{C \sqsubseteq \Gamma \vdash u_i(\mathbf{y}).P \triangleright \Delta * \Gamma'} [\text{VAcc}] \\
\frac{P = \sum_{i \in I} y_i(x_i); P_i \quad \forall i : y_i \in \mathbf{y} \text{ and } C \sqsubseteq \Gamma, x_i : U_i \vdash P_i \triangleright \Delta, \mathbf{y} : \mathcal{T}_i * \Gamma_i \quad \Gamma' = \bigcap_{i \in I} \Gamma_i \quad \text{fv}(P) \cup \text{fc}(P) \subseteq \text{dom}(\Gamma')}{C \sqsubseteq \Gamma \vdash P \triangleright \Delta, \mathbf{y} : \sum_{i \in I} y_i ? U_i; \mathcal{T}_i * \Gamma'} [\text{VRec}] \\
\frac{\Gamma(e) = \text{U} \quad y \in \mathbf{y}}{C \sqsubseteq \Gamma \vdash \bar{y}e \triangleright \mathbf{y} : y!U * \Gamma} [\text{VSend}] \quad \frac{\Delta(s) = \text{end} \quad \forall s \in \text{dom}(\Delta)}{C \sqsubseteq \Gamma \vdash \mathbf{0} \triangleright \Delta * \Gamma} [\text{VEnd}] \\
\frac{C \sqsubseteq \Gamma \vdash P_1 \triangleright \Delta_1 * \Gamma_1 \quad C \sqsubseteq \Gamma_1 \vdash P_2 \triangleright \Delta_2 * \Gamma_2}{C \sqsubseteq \Gamma \vdash P_1; P_2 \triangleright \Delta_1; \Delta_2 * \Gamma_2} [\text{VSeq}] \\
\frac{\Gamma(e) = \text{bool} \quad C \wedge e \not\vdash \perp \quad C \wedge \neg e \vdash \perp \quad C \wedge e \sqsubseteq \Gamma \vdash P \triangleright \Delta * \Gamma'}{C \sqsubseteq \Gamma \vdash \text{if } e : P \text{ else } Q \triangleright \Delta * \Gamma'} [\text{VThen}] \\
\frac{\Gamma(e) = \text{bool} \quad C \wedge e \vdash \perp \quad C \wedge \neg e \not\vdash \perp \quad C \wedge \neg e \sqsubseteq \Gamma \vdash Q \triangleright \Delta * \Gamma'}{C \sqsubseteq \Gamma \vdash \text{if } e : P \text{ else } Q \triangleright \Delta * \Gamma'} [\text{VElse}] \\
\frac{\Gamma(e) = \text{bool} \quad C \wedge e \not\vdash \perp \quad C \wedge \neg e \not\vdash \perp \quad C \wedge e \sqsubseteq \Gamma \vdash P \triangleright \Delta_1 * \Gamma_1 \quad C \wedge \neg e \sqsubseteq \Gamma \vdash Q \triangleright \Delta_2 * \Gamma_2}{C \sqsubseteq \Gamma \vdash \text{if } e : P \text{ else } Q \triangleright \Delta_1 \bowtie \Delta_2 * \Gamma_1 \cap \Gamma_2} [\text{VCond}] \\
\frac{\Gamma(\ell) = [\text{U}] \quad C \vdash \ell \neq \varepsilon \quad C \wedge x \in \ell \sqsubseteq \Gamma, x : \text{U} \vdash P \triangleright \mathbf{y} : \mathcal{T} * \Gamma'}{C \sqsubseteq \Gamma \vdash \text{for } x \text{ in } \ell : P \triangleright \mathbf{y} : \mathcal{T}^* * \Gamma'} [\text{VFor1}] \\
\frac{C \vdash \ell = \varepsilon}{C \sqsubseteq \Gamma \vdash \text{for } x \text{ in } \ell : P \triangleright \mathbf{y} : \text{end} * \Gamma'} [\text{VFor2}] \\
\frac{C \sqsubseteq \Gamma \vdash N \triangleright \mathbf{y} : \mathcal{T} * \Gamma'}{C \sqsubseteq \Gamma \vdash \text{do } N \text{ until } b(x) \triangleright \mathbf{y} : \mathcal{T}^*; b?U * \Gamma', x : \text{U}} [\text{VLoop}]
\end{array}$$

Fig. 5. Typing rules for processes

Iterations involving messages over multiple sessions could not be checked compositionally since the conformance of a process to a local type would not be sufficient to ensure the correct coordination of a ‘for’-iteration with the corresponding ‘loop’-iterations. Rule [VEnd] types idle processes with a Δ that maps each session s to the end type. Rule [VSeq] checks sequential composition. Here $\Delta_1; \Delta_2$ is the pointwise sequential composition of Δ_1 and Δ_2 , i.e., $(\Delta_1; \Delta_2)(s) = \mathcal{T}_1; \mathcal{T}_2$ where $\mathcal{T}_i = \Delta_i(s)$ if $s \in \text{dom}(\Delta_i)$ and $\mathcal{T}_i = \text{end}$ otherwise, for $i = 1, 2$. Note that P_2 is typed under the environment Γ_1 , which contains the names bound by the input prefixes of P_1 .

The following result ensures that type checking is decidable (it follows from the obvious recursive algorithm and decidability of the underlying logic).

$$\begin{aligned}
(\Delta_1 \bowtie \Delta_2)(s) &= \begin{cases} \Delta_1(s) & \text{if } s \in \text{dom}(\Delta_1) \setminus \text{dom}(\Delta_2) \\ \Delta_2(s) & \text{if } s \in \text{dom}(\Delta_2) \setminus \text{dom}(\Delta_1) \\ \Delta_1(s) \bowtie \Delta_2(s) & \text{if } s \in \text{dom}(\Delta_1) \cap \text{dom}(\Delta_2) \end{cases} \\
\mathcal{T}_1 \bowtie \mathcal{T}_2 &= \begin{cases} \mathcal{T}_1 \oplus \mathcal{T}_2 & \text{if } \mathcal{T}_1 = y_1!U_1; \mathcal{T}'_1, \mathcal{T}_2 = y_2!U_2; \mathcal{T}'_2, y_1 \neq y_2 \\ y!U; (\mathcal{T}'_1 \bowtie \mathcal{T}'_2) & \text{if } \mathcal{T}_1 = y!U; \mathcal{T}'_1, \mathcal{T}_2 = y!U; \mathcal{T}'_2 \\ \perp & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 6. Composition of types

Theorem 1. *Given C, Γ, Γ', S and Δ , then the provability of $C \sqsubseteq \Gamma \vdash S \triangleright \Delta * \Gamma'$ is decidable.*

Our proof system discerns between B_1 and B_2 in the introduction (i.e., only B_1 is validated) due to the rules for conditional statements and to the lack of a rule for type refinement. In fact, after a few verification steps on B_1 (resp. B_2) we would reach the following scenario: $P_1(c) = \text{if } c : \overline{s_{OK}}; s_{AMOUNT}(x) \text{ else } \overline{s_{KO}}$ (resp. $P_2(c) = \overline{s_{KO}}$) and $\Delta = s : OK!; AMOUNT? \oplus KO!$. The verification of $P_1(c)$ terminates successfully after an application of [VCond]. In the case of $P_2(c)$ the only rule for a sending process, [VSend], cannot be applied against a type with a choice.

5.1 Running Examples

We now apply our typing to different implementations of POP2.

Example 5. The first few verification steps of P_{INIT} from Ex. 3 are shown below. By rule [VAcc], the newly created session is added to the session environment, then the verification of the external choice is split by [VRec] into the verification of each branch. As we omit the whole derivation, just assume that P_{INIT} yields $\Gamma' = \Gamma, x : \text{Str}$.

$$\frac{\frac{\text{true} \sqsubseteq \Gamma \vdash P_{EXIT} \triangleright s : T_{EXIT} * \Gamma' \quad \text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash P_{MBOX} \triangleright s : T_{MBOX} * \Gamma'}{\Gamma(u) \equiv G_{POP}(s) \quad \text{true} \sqsubseteq \Gamma \vdash s_{QUIT}(); P_{EXIT} + s_{HELO}(x); P_{MBOX} \triangleright s : T_S * \Gamma'} [\text{VRec}]}{\text{true} \sqsubseteq \Gamma \vdash u_S(s).P_S \triangleright \emptyset * \Gamma'} [\text{VAcc}]$$

Consider the second branch P_{MBOX} ; assuming that $\text{true} \wedge \text{auth}(x)$ is neither a tautology nor a falsum, we apply [VCond] (if it was, the verification would terminate unsuccessfully as the only possible branch would not validate against the choice in T_{MBOX}).

$$\frac{\frac{\text{auth}(x) \sqsubseteq \Gamma, x : \text{Str} \vdash \overline{s_{RMN}}(\text{inbox}); P_{NMBR} \triangleright s : R!Int; T_{NMBR} * \Gamma' \quad \neg \text{auth}(x) \sqsubseteq \Gamma, x : \text{Str} \vdash \overline{s_E}; P_{EXIT} \triangleright s : E!; T_{EXIT} * \Gamma'}{\text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash P_{MBOX} \triangleright s : T_{MBOX} * \Gamma'} [\text{VCond}]$$

The rest is trivial observing $s : T_{MBOX} = s : R!Int; T_{NMBR} \bowtie s : E!; T_{EXIT}$. ◆

Ex. 6 types the multiparty variant given in Ex. 4.

Example 6. Assume $\Gamma(u) \equiv G'_{POP}$. The first steps are as in Ex. 5 by rules [VAcc] and [VRec]. We focus on the second branch that in this case is P_{AUTH} and apply [VSeq].

$$\frac{\text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash \overline{s_{REQ}}x \triangleright s : REQ!Str * \Gamma' \quad \text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash s_{RES}(y); P_{MBOX} \triangleright s : RES?Bool; T_{AUTH} * \Gamma'}{\text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash \overline{s_{REQ}}x; s_{RES}(y); P'_{MBOX} \triangleright s : T_{AUTH} * \Gamma'} [\text{VSeq}]$$

We show the verification of the first branch

$$\frac{\frac{}{\text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash \mathbf{0} \triangleright s : \text{end} * \Gamma'} [\text{VEnd}]}{\text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash \overline{s_{\text{REQ}}} x \triangleright s : \text{REQ!Str} * \Gamma'} [\text{VSend}]$$

and the successive steps for the second branch:

$$\frac{\text{true} \sqsubseteq \Gamma, x : \text{Str}, y : \text{Bool} \vdash P'_{\text{MBOX}} \triangleright s : T'_{\text{MBOX}} * \Gamma'}{\text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash s_{\text{RES}}(y); P'_{\text{MBOX}} \triangleright s : \text{RES?Bool}; T'_{\text{MBOX}} * \Gamma'} [\text{VRec}]$$

The verification of P'_{MBOX} proceeds with an application of [VCond], [VIf] or [VElse] depending on $\text{auth}() \wedge y$. If $\text{auth}()$ is not a contradiction then [VCond] can be applied as the condition depends from the context (that is the administrator). This leads to a successful validation. In this case, unlike in Ex. 5, the implementation is whole-spectrum even if $\text{auth}()$ is a tautology. If $\text{auth}()$ is a falsum then [VElse] is applied and the process will not validate against the type which has a choice. ♦

Ex. 7 deals with a process implementing two interleaved sessions. Ex. 7 shows that the verification scales to more complex processes that compose different protocols.

Example 7. We give a process that, upon request, engages as a server in a session G_{POP} (§ 2.2), and as a client in a session G_{ADMIN} to outsource the authentication. Instead of embedding in the same session this extra interactions with the administrator, as we did in Ex. 2, we represent the multiparty interaction as two interleaved sessions.

$$G_{\text{ADMIN}} = C \rightarrow A : \text{REQ} \langle \text{Str} \rangle; A \rightarrow C : \text{RES} \langle \text{Bool} \rangle \quad T_C = \text{REQ!Str}; \text{RES?Bool}$$

In G_{ADMIN} , the client C sends the administrator A a password and A replies along RES . T_C is the projection on G_{ADMIN} on C . We assume $\Gamma(u) \equiv G_{\text{POP}}$ and $\Gamma(v) \equiv G_{\text{ADMIN}}$. Process P_{INIT} starts, upon request, a session of type G_{POP} and then requests to start a session of type G_{ADMIN} . We omit the definition of processes P_{NMBR} and P_{EXIT} which are as in Ex. 3.

$$\begin{array}{ll} P''_{\text{INIT}} = u_S(s). \overline{v}^C(t). P''_S & P''_S = s_{\text{QUIT}}(); P_{\text{EXIT}} + s_{\text{HELO}}(x); P_{\text{AUTH}} \\ P_{\text{AUTH}} = \overline{t_{\text{REQ}}} x; t_{\text{RES}}(y); P''_{\text{MBOX}} & P''_{\text{MBOX}} = \text{if } y : \overline{s_{\text{RfN}}}(\text{inbox}); P_{\text{NMBR}} \text{ else } \overline{s_E}; P_{\text{EXIT}} \end{array}$$

The authentication is delegated to the administrator in session t via the message along t_{REQ} . Session s continues using the information in y , which stores the last message received in session t . The first verification steps are by rules [VAcc], [VReq] and [VRec].

$$\frac{\frac{\text{true} \sqsubseteq \Gamma \vdash P_{\text{EXIT}} \triangleright s : T_{\text{EXIT}}, t : T_S * \Gamma' \quad \text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash P_{\text{AUTH}} \triangleright s : T_{\text{MBOX}}, t : T_S * \Gamma'}{\Gamma(u) \equiv G_{\text{POP}}(s) \quad \Gamma(v) \equiv G_{\text{POP}}(t) \quad \text{true} \sqsubseteq \Gamma \vdash P''_S \triangleright s : T_S, t : T_S * \Gamma'} [\text{VRec}]}{\text{true} \sqsubseteq \Gamma \vdash u_S(s). \overline{v}^S(t). P''_S \triangleright \mathbf{0} * \Gamma'} [\text{VAcc}, [\text{VReq}]]$$

The verification of the second branch P_{AUTH} proceeds with one application of [VSeq] where $(s : T_{\text{MBOX}}, t : T_S) \equiv (s : \text{end}, t : \text{REQ!Str}); (s : T_{\text{MBOX}}, t : \text{RES?Bool})$.

Let $\text{push}([_], s!v) = s!v[_]$ and $\text{push}(s_1!v_1; \dots; s_n!v_n[_], s!v) = s_1!v_1; \dots; s_n!v_n; s!v[_]$

$$\begin{array}{c}
 \Gamma \bullet \Delta, s : s!v; \mathbb{M} @ p \xrightarrow{\bar{s}v} \Gamma \bullet \Delta, s : \mathbb{M} @ p \text{ [TQueue]} \\
 \frac{\Gamma \bullet s : \mathcal{T} \xrightarrow{\bar{s}v} \Gamma \bullet s : \mathcal{T}' \quad \mathbb{M}'[_] = \text{push}(\mathbb{M}[_], s!v)}{\Gamma \bullet \Delta, s : \mathbb{M}[\mathcal{T}] @ p \xrightarrow{\tau} \Gamma \bullet \Delta, s : \mathbb{M}'[\mathcal{T}'] @ p} \text{ [TCom1]} \\
 \frac{\Gamma \bullet s : \mathcal{T} \xrightarrow{\bar{s}v} \Gamma \bullet s : \mathcal{T}'}{\Gamma \bullet \Delta, s : \mathbb{M}[\mathcal{T}] @ p \xrightarrow{\tau} \Gamma \bullet \Delta, s : \mathbb{M}'[\mathcal{T}'] @ p} \text{ [TCom2]} \\
 \frac{u \in \text{dom}(\Gamma) \quad \Gamma(u) \equiv \mathcal{G}(s) \triangleq G \quad \mathcal{P}(G) = \{p_1, \dots, p_n\}}{\Gamma \bullet \Delta \xrightarrow{\tau} \Gamma \bullet \Delta, s : (G|p_1) @ p_1, \dots, s : (G|p_n) @ p_n} \text{ [TInit]}
 \end{array}$$

Fig. 7. Additional labelled transitions (to those of Fig. 1) for runtime specifications

$$\frac{\begin{array}{c} \text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash \overline{t_{\text{REQ}}x} \triangleright \Delta * s : \text{end}, t : \text{REQ!Str} \\ \text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash t_{\text{RES}}(y); P''_{\text{MBOX}} \triangleright s : T_{\text{MBOX}}, t : \text{RES?Bool} * \Gamma' \end{array}}{\text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash \overline{t_{\text{REQ}}x}; t_{\text{RES}}(y); P''_{\text{MBOX}} \triangleright s : T_{\text{MBOX}}, t : T_S * \Gamma'} \text{ [VSeq]}$$

Focusing on the second branch we apply [VRec]

$$\frac{\text{true} \sqsubseteq \Gamma, x : \text{Str}, y : \text{Bool} \vdash P''_{\text{MBOX}} \triangleright s : T_{\text{MBOX}}, t : \text{end} * \Gamma'}{\text{true} \sqsubseteq \Gamma, x : \text{Str} \vdash t_{\text{RES}}(y); P''_{\text{MBOX}} \triangleright s : T_{\text{MBOX}}, t : \text{RES?Bool} * \Gamma'} \text{ [VRec]}$$

Rule [VCond] can be applied since condition y of the conditional statement in P''_{MBOX} is neither a tautology nor a contradiction. The rest is as in Ex. 5. $\blacklozenge$

6 Properties of the Type System

Runtime Types. The properties of our type system are stated in terms of the behaviour of local types. As in [14], *runtime types* extend local types with *message contexts* $\mathbb{M}$ of the form $s_1!v_1; \dots; s_n!v_n[_]$ with $n \geq 0$, namely $\mathbb{M}$ is a sequence of outputs followed by a hole $[_]$. To model asynchrony, we stipulate the equality

$$s_1!v_1; s_2!v_2; \mathbb{M} \approx s_2!v_2; s_1!v_1; \mathbb{M} \quad \text{if } s_1 \neq s_2$$

A *runtime type* is either a message context $\mathbb{M}$ or a “type in context”, that is a term $\mathbb{M}[\mathcal{T}]$. We extend environments so to map session names s to runtime types of roles in s ; we write $\Delta, s : \mathbb{M}[\mathcal{T}] @ p$ to specify that (1) the runtime type of $p \in \mathbb{P}$ in s is $\mathbb{M}[\mathcal{T}]$ and (2) that for any $s : \mathbb{M}'[\mathcal{T}'] @ q$ in Δ we have $q \neq p$.

The semantics of runtime types is obtained by adding the rules in Fig. 7 to those in Fig. 1. Rule [TQueue] removes a message from of a queue. Rules [TCom1] and [TCom2] establish how runtime specifications send and receives messages (the transition in their premises are derived from the rules in Fig. 1). Rule [TInit] initiates a new session by mapping the new session s to the projections of the global type assigned by Γ .

Soundness. The typing rules in § 5 ensure the semantic conformance of processes with the behaviour prescribed by their types. Here, we define conformance in terms of *conditional simulation* that relates states and specifications. Our definition is standard, except for input actions, for which specifications have to simulate only inputs of messages with the expected type (i.e., systems are not responsible when receiving ill-typed messages).

Define $\xRightarrow{\alpha} = \xRightarrow{\tau}^* \xrightarrow{\alpha}$. Let $\Gamma \bullet \Delta \xRightarrow{s}$ shorten $\exists \Delta' \exists v : \Gamma \bullet \Delta \xRightarrow{sv} \Gamma \bullet \Delta'$.

Definition 5 (Conditional simulation). A relation $\mathbb{R}$ between states and specifications is a conditional simulation iff for any $(\langle S, \sigma \rangle, \Gamma \bullet \Delta) \in \mathbb{R}$, if $\langle S, \sigma \rangle \xrightarrow{el-\alpha} \langle S', \sigma' \rangle$ then

1. if $\alpha = sv$ then $\Gamma \bullet \Delta \xrightarrow{s} \Gamma \bullet \Delta'$ and if $\Gamma \bullet \Delta \xrightarrow{sv} \Gamma \bullet \Delta'$ then there is $\Gamma \bullet \Delta'$ such that $\Gamma \bullet \Delta \xrightarrow{sv} \Gamma \bullet \Delta'$ and $(\langle S', \sigma' \rangle, \Gamma \bullet \Delta') \in \mathbb{R}$
2. otherwise, $\Gamma \bullet \Delta \xrightarrow{\alpha} \Gamma \bullet \Delta'$ and $(\langle S', \sigma' \rangle, \Gamma \bullet \Delta') \in \mathbb{R}$.

We write $\langle S, \sigma \rangle \preceq \Gamma \bullet \Delta$ if there is a conditional simulation $\mathbb{R}$ s.t. $(\langle S, \sigma \rangle, \Gamma \bullet \Delta) \in \mathbb{R}$.

By (1), only inputs of S with the expected type have to be matched by $\Gamma \bullet \Delta$ (recall rule [TRec] in Fig. 1), while it is no longer expected to conform to the specification after an ill-typed input (i.e., not allowed by $\Gamma \bullet \Delta$).

Def. 6 establishes consistency for stores in terms of preservation of variables' sorts.

Definition 6 (Consistent store). Given an environment Γ , a context assumption C , and a state $\langle S, \sigma \rangle$ with $\text{var}(S) \subseteq \text{dom}(\sigma)$, store σ is consistent for S with respect to Γ and C iff $\forall x \in \text{dom}(\sigma)$, $\sigma(x) : \Gamma(x)$, and $C \downarrow \sigma = \text{true}$.

Theorem 2 (Subject reduction). Assume that

$$C \sqsubseteq \Gamma \vdash S \triangleright \Delta * \Gamma' \quad \text{and} \quad \langle S, \sigma \rangle \xrightarrow{el-\alpha} \langle S', \sigma' \rangle$$

with σ consistent for S with respect to Γ and C . Then

1. if $\alpha = sv$ then $\Gamma \bullet \Delta \xrightarrow{s} \Gamma \bullet \Delta'$ and if $\Gamma \bullet \Delta \xrightarrow{sv} \Gamma \bullet \Delta'$ then there is $\Gamma \bullet \Delta'$ such that $\Gamma \bullet \Delta \xrightarrow{sv} \Gamma \bullet \Delta'$ with $v : \mathbb{U}$ and $C \wedge e \sqsubseteq \Gamma, x : \mathbb{U} \vdash S' \triangleright \Delta' * \Gamma''$ for some x and some $\Gamma'' \supseteq \Gamma'$
2. otherwise $\Gamma \bullet \Delta \xrightarrow{\alpha} \Gamma \bullet \Delta'$ and $C \wedge e \sqsubseteq \Gamma \vdash S' \triangleright \Delta' * \Gamma''$ for some $\Gamma'' \supseteq \Gamma'$.

Corollary 1 (Soundness). If $C \sqsubseteq \Gamma \vdash S \triangleright \Delta * \Gamma'$ then $\langle S, \sigma \rangle \preceq \Gamma \bullet \Delta$ for all σ consistent store for S with respect to Γ and C .

WSI by typing

We show that well-typed processes are WSIs (Def. 4). First, we relate the runs of a global type with those of its corresponding runtime types. Then, we state the correspondence between the runs of runtime types and well-typed implementations.

Definition 7 (Runs of runtime types). The set $\mathcal{R}_s(\Delta)$ denotes the runs of events over the channels in s generated by Δ , and is inductively defined by the rules in Fig. 8.

Rule [RTCom] builds the runs for two communicating types. Rules [RTIt1] and [RTIt2] unfold the runs of an iterative type. Note that the mandatory actions of runs associated to recursive types are those requiring at least one execution of the iteration body, while additional executions are optional. The remaining rules are self-explanatory. (The correspondence between the operational and denotational semantics is in [3, Appendix E].)

Thm. 3 ensures that well-formed global types are covered by their projections, while Thm. 4 states that the set of well-typed implementation covers its specification.

Theorem 3. For any global type $\mathcal{G}(s)$, $\mathcal{R}(\mathcal{G}(s)) \subseteq \mathcal{R}_s(\{s : (\mathcal{G}(s) \upharpoonright p) @ p\}_{p \in \mathcal{P}(\mathcal{G}(s))})$.

Theorem 4. Let $\mathcal{G}(s) \triangleq G$ be a global type. Fix $p \in \mathcal{P}(G)$ and P a well-typed implementation of p . Define

$$\mathbb{I}_{p,P} = \{I_G^1 \upharpoonright \upharpoonright(p) = P, \forall q \in \mathcal{P}(G) : \text{true} \sqsubseteq \Gamma, u : \mathcal{G}(s) \vdash \upharpoonright(q) \triangleright \Delta, s : \mathcal{G}(s) \upharpoonright q * \Gamma'\}$$

then, $\mathcal{R}_s(\{s : (\mathcal{G}(s) \upharpoonright p) @ p\}_{p \in \mathcal{P}(G)}) \subseteq \mathcal{R}_u(\mathbb{I}_{p,P})$.

$$\begin{array}{c}
\frac{r \in \mathcal{R}_S(\Delta, s : \mathbb{M}[T_k] @ p, \mathbb{M}'[T'_k] @ q) \quad k \in J}{\langle p, s_k ! U_k \rangle \langle q, s_k ? U_k \rangle r \in \mathcal{R}_S(\Delta, s : \mathbb{M}[s_k ! U_k; T_k] @ p, \mathbb{M}'[\sum_{j \in J} s_j ? U_j; T'_j] @ q)} [\text{RTCom}] \\
\\
\frac{r \in \mathcal{R}_S(\Delta, s : T_i; T_j @ p)}{r \in \mathcal{R}_S(\Delta, s : T_i^*; T_j @ p)} [\text{RTIt1}] \quad \frac{rr' \in \mathcal{R}_S(\Delta, s : T_i; T_i^*; T_j @ p) \quad r' \in \mathcal{R}_S(\Delta, s : T_i; T_j @ p)}{[r]r' \in \mathcal{R}_S(\Delta, s : T_i^*; T_j @ p)} [\text{RTIt2}] \\
\\
\frac{r \in \mathcal{R}_S(\Delta) \quad s \neq \mathbf{r}}{r \in \mathcal{R}_S(\Delta, \mathbf{r} : T)} [\text{RTPar}] \quad \frac{r \in \mathcal{R}_S(\Delta)}{r \in \mathcal{R}_S(\Delta, s : \text{end} @ p)} [\text{RTEnd1}] \quad \varepsilon \in \mathcal{R}_S(\emptyset) [\text{RTEnd2}] \\
\\
\frac{r \in \mathcal{R}_S(\Delta, s : s_j ! U_j; T_j @ p) \quad j \in I}{r \in \mathcal{R}_S(\Delta, s : \bigoplus_{i \in I} s_i ! U_i; T_i @ p)} [\text{RTCh}]
\end{array}$$

Fig. 8. Runs of runtime local types

7 Conclusion and Related Work

WSI forbids implementations of a role that persistently avoid the execution of some alternative branches in a choreography. Although WSI is defined as a relation between the traces of a global type and those of its candidate implementations, it can be checked by using multiparty session types. Technically, we show that (i) the sets of the projections of a global type $\mathcal{G}$ preserves all the traces in $\mathcal{G}$ (Thm. 3); and (ii) any trace of a local type can be mimicked by a well-typed implementation, if interacting in a proper context (Thm. 4). The soundness of our type system (Corollary 1) ensures that well-typed implementations behave as prescribed by the choreography.

We are currently working on the extension of WSI to other models of choreography as e.g. those based on automata [12], which poses the classical question about the decidability of the notion of realisability (see [1]). To the best of our knowledge, the only proposal dealing with complete (i.e., exhaustive) realisations in a behavioural context is [7] but this approach focuses on non-deterministic implementation languages. Our type system is more restrictive than [14,2,4,8,6]. We do not consider subtyping because the liberal elimination of internal choices prevents WSI. The investigation of suitable forms of subtyping for WSIs is scope for future work.

WSI coincides with projection realisability [17,22,7] when implementation languages feature non-deterministic internal choices. On the contrary, WSI provides a finer criterion to distinguish deterministic implementations, as illustrated by the motivating example in the introduction. To some extent our proposal is related to the fair subtyping approach in [21], where refinement is studied under the fairness assumption: Fair subtyping differs from usual subtyping when considering infinite computations but WSI differs from partial implementation also when considering finite computations.

The static verification of WSI requires a form of recursion more restrictive than the one in [14,2], where the number of iterations is limited. This restriction is on the lines of [7] that also considers finite traces. The extension of our theory with a more general form of iteration is scope for future work.

Acknowledgements. We thank the reviewers for their insightful and helpful comments.

References

1. Basu, S., Bultan, T., Ouederni, M.: Deciding choreography realizability. In: POPL (2012)
2. 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)
3. Bocchi, L., Melgratti, H., Tuosto, E.: Extended version of this paper (2014)
<http://publicaciones.dc.uba.ar/Publications/2014/BMT14c/>
4. Bravetti, M., Zavattaro, G.: A theory of contracts for strong service compliance. MSCS 19(3) (2009)
5. Butler, M., Postel, J., Chase, D., Goldberger, J., Reynolds, J.: Post office protocol - version 2. RFC 918 (February 1985), <http://tools.ietf.org/html/rfc937>
6. Caires, L., Vieira, H.T.: Conversation types. In: Castagna, G. (ed.) ESOP 2009. LNCS, vol. 5502, pp. 285–300. Springer, Heidelberg (2009)
7. Castagna, G., Dezani-Ciancaglini, M., Padovani, L.: On global types and multi-party session. LMCS 8(1) (2012)
8. Castagna, G., Padovani, L.: Contracts for mobile processes. In: Bravetti, M., Zavattaro, G. (eds.) CONCUR 2009. LNCS, vol. 5710, pp. 11–228. Springer, Heidelberg (2009)
9. Chen, T.-C., Honda, K.: Specifying stateful asynchronous properties for distributed programs. In: Koutny, M., Ulidowski, I. (eds.) CONCUR 2012. LNCS, vol. 7454, pp. 209–224. Springer, Heidelberg (2012)
10. Crocker, D.: Standard for the format of arpa internet text messages. RFC 822 (February 1982), www.ietf.org/rfc/rfc0822.txt
11. Dezani-Ciancaglini, M., de'Liguoro, U.: Sessions and session types: An overview. In: Laneve, C., Su, J. (eds.) WS-FM 2009. LNCS, vol. 6194, pp. 1–28. Springer, Heidelberg (2010)
12. Fu, X., Bultan, T., Su, J.: Realizability of conversation protocols with message contents. Int. J. Web Service Res. 2(4), 68–93 (2005)
13. Gay, S., Hole, M.: Subtyping for Session Types in the Pi-Calculus. Acta Inf. 42(2/3), 191–225 (2005)
14. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: POPL (2008)
15. Kavantzaz, N., Burdett, D., Ritzinger, G., Fletcher, T., Lafon, Y. (2004),
<http://www.w3.org/TR/2004/WD-ws-cdl-10-20041217>
16. Lamport, L.: Time, clocks, and the ordering of events in a distributed system. CACM 21(7), 558–564 (1978)
17. Lanese, I., Guidi, C., Montesi, F., Zavattaro, G.: Bridging the gap between interaction-and process-oriented choreographies. In: SEFM (2008)
18. Lange, J., Tuosto, E.: Synthesising choreographies from local session types. In: Koutny, M., Ulidowski, I. (eds.) CONCUR 2012. LNCS, vol. 7454, pp. 225–239. Springer, Heidelberg (2012)
19. Lohmann, N., Wolf, K.: Decidability results for choreography realization. In: Kappel, G., Maamar, Z., Motahari-Nezhad, H.R. (eds.) Service Oriented Computing. LNCS, vol. 7084, pp. 92–107. Springer, Heidelberg (2011)
20. Milner, R.: Communication and Concurrency. Prentice Hall (1989)
21. Padovani, L.: Fair subtyping for multi-party session types. In: De Meuter, W., Roman, G.-C. (eds.) COORDINATION 2011. LNCS, vol. 6721, pp. 127–141. Springer, Heidelberg (2011)
22. Salaün, G., Bultan, T.: Realizability of choreographies using process algebra encodings. In: Integrated Formal Methods (2009)
23. Su, J., Bultan, T., Fu, X., Zhao, X.: Towards a theory of web service choreographies. In: Damas, M., Heckel, R. (eds.) WS-FM 2007. LNCS, vol. 4937, pp. 1–16. Springer, Heidelberg (2008)