AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Self-adaptive multiparty sessions

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/154878> since

Published version:

DOI:10.1007/s11761-014-0171-9

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

UNIVERSITÀ DEGLI STUDI DI TORINO

This is an author version of the contribution published on:

Mario Coppo, Mariangiola Dezani, Betti Venneri
Self-adaptive multiparty sessions
SERVICE ORIENTED COMPUTING AND APPLICATIONS (2014)
DOI: 10.1007/s11761-014-0171-9

The definitive version is available at:

<http://link.springer.com/content/pdf/10.1007/s11761-014-0171-9>

Self-Adaptive Multiparty Sessions

Mario Coppo · Mariangiola Dezani-Ciancaglini · Betti Venneri

Abstract To model the notion of *self-adaptiveness* for *multiparty sessions*, we propose a formal framework, where participants can access and modify the *global state*, in such a way that the whole system can react promptly to unforeseen events by reconfiguring itself. The adaptation strategy is triggered by the overall communication choreography, represented by a global type. When the global type is dynamically updated, its projections define new *monitors*, which set-up novel communication protocols for the participants. The key result of this paper is that self-adaptations are performed in a type-safe way, while providing a high degree of flexibility. *Subject Reduction* and *Progress* properties are proven: any session executes all required communications in a type safe way and never gets stuck.

1 Introduction

The topic of self-adaptiveness emerged as a key research subject within various application domains, as a response to the growing complexity of software systems operating in many different scenarios and in highly dynamic environments. To manage this complexity at a reasonable cost, novel approaches are needed in which a system can promptly react to crucial changes by reconfiguring its behaviour autonomously and dynamically, in accordance with evolving policies and objectives.

As for a precise definition of *self-adaptivity*, this is still a debated question, due to the wide spectrum of the involved

features. In our opinion, a simple, but rather deep, characterisation is the one presented in [7]: *we define adaptation as the run-time modification of the control data ...and a component is self-adaptive if it is able to modify its own control data at run-time*. We follow [7] in claiming that we need to distinguish between standard data and control data: a change in the system behaviour is part of the application logic if it is based on standard data, it is an adaptation if it is based on control data.

This paper injects the above notion of *self-adaptivity* into the formal framework of *multiparty sessions* [28], where each participant can access and modify the *global state* representing those (control) data whose values are critical for planning the adaptation steps, in such a way that the whole system can react to changes in the global data by reconfiguring itself. A system comprises four active parties: *global types*, *monitors*, *processes*, and *adaptation functions*.

A global type represents the overall communication choreography [9]; its projections onto participants generate the monitors, which are essentially local types and set-up the communication protocols of the participants. The association of a monitor with a compliant process, dubbed *monitored process*, incarnates a participant where the process provides the implementation of the monitoring protocol. Notably, we exploit intersection types, union types and subtyping to make this compliance relation flexible. Processes are able to follow different incompatible computational paths. For instance, a process could contain both the code needed to buy a book and the one needed to arrange a friend meeting, the choice between the two being determined by the monitor controlling it.

The adaptation strategy is defined by global types and adaptation functions. The choreography decides *when* the adaptation takes place, since its monitors prescribe when some participants have to check global data, and then send a request of adaptation to the other participants together with

M. Coppo · M. Dezani-Ciancaglini
Dipartimento di Informatica, Università di Torino, corso Svizzera 185,
10149 Torino, Italy {coppo,dezani}@di.unito.it

B. Venneri
Dipartimento di Statistica, Informatica, Applicazioni, Università di
Firenze Viale Morgagni 65, 50134 Firenze, Italy betti.venneri@unifi.it

an adaptation function. The adaptation functions contain the dynamic evolution policy, since they prescribe *how* the system needs to reconfigure itself based on the changes of the critical data.

When an adaptation flag is thrown, new monitors are generated, according to a new choreography: indeed, the community involved in the session modifies both its set of participants and the internal communication patterns. Therefore, dynamic adaptations are essentially triggered by control data and monitors.

Most of the approaches to self-adaptive systems in the literature do not face the main challenge of including formal tools to ensure correctness of dynamic adaptations. Some approaches address this issue by providing verification techniques for testing properties of the performed adaptation (e.g., model checking in [25] and web services testbed in [36]). Differently, the focus of the present paper is on the formal properties of the proposed framework, which ensure that adaptations steps are performed in a correct way, being controlled by global types, monitors and process types. The key results are the proofs of Subject Reduction and Progress theorems: in any session, all outputs will eventually be consumed and all processes waiting for an input will eventually receive it.

Typical scenarios that can benefit from our self-adaptation framework are those characterised by the following features:

- a community, established for a common task or mission, has many distributed entities which interact with each other according to a given operational plan,
- the complex dynamic environment can present crucial events which require the community to modify its plan dynamically,
- those critical events are observed in any separate component of the system: this component can be checked by the session participants, so that the whole system can react promptly by updating itself,
- the dynamic changes need to be rather flexible: in each new phase, other participants can be introduced or some of the old participants are no longer involved (temporarily or permanently),
- these dynamic changes need to be safe: interactions must proceed correctly to pursue the common task.

Example As an example of such a scenario, let us consider a company which has various productive units and sale organisations scattered around the world. Each factory has a number of machines and produces several products for nearby markets or for export. The state of the plants is checked periodically. Communications among factories and sellers exchange several data about products and prices, according to a given combination factory-seller for each product. The company chief supervises the whole organisation. In particular, she equipped the company with an adaptation policy, which gives potential alternative plans for moving produc-

tions and/or sales of a product to different entities. All the interactions among these participants run under the control of the monitors that are originated from a global type. Finally, a global state contains crucial data, for instance the performance of machineries, plants and sale organisations. Unforeseen circumstances, such as the catastrophic event of a fire incapacitating an entire plant, can require the company organisation to update itself: new production and sale plans have to be adopted to maintain uninterrupted supply to customers.

We simplify the above scenario in the case of a Company which has two factories, **iF** (Italian factory) and **aF** (American factory), and two sellers, **iS** (Italian seller) and **aS** (American seller). In order to give a preliminary intuition of our system, we use a simplified and incomplete syntax (w.r.t. the formal presentation of next section). In Section 5 we will enrich and formalise this example.

To show how self-adaptation works, we consider the case when a fire incapacitates a factory. The global state contains either **OK** or **KO** for each of the two plants. When both plants are **OK** the interaction takes place according to the following global type:

$$G_1 = \begin{cases} \text{iS} \rightarrow \text{iF} : (\text{String}, \text{Int}). \\ \text{aS} \rightarrow \text{aF} : (\text{String}, \text{Int}). \\ \text{Ada} \rightarrow \{\text{iS}, \text{iF}, \text{aS}, \text{aF}\} : \text{check} \end{cases}$$

Each seller requires to the corresponding factory a certain amount (of type Int) of an item (of type String), then the chief Ada sends a checking flag to all, as an alert for a possible adaptation. When the Italian factory is **OK**, while the American factory is **KO**, the global type is:

$$G_2 = \begin{cases} \text{Ada} \rightarrow \text{Bob} : \text{String}. \\ \text{iS} \rightarrow \text{iF} : (\text{String}, \text{Int}). \\ \text{aS} \rightarrow \text{iF} : (\text{String}, \text{Int}). \\ \text{Ada} \rightarrow \{\text{iS}, \text{iF}, \text{aS}, \text{Bob}\} : \text{check} \end{cases}$$

where Ada sends to Bob a contract (type String) for rebuilding the plant and both sellers send their requests to the Italian factory. In the symmetric case, the global type G_3 prescribes that both sellers send their requests to the American factory. Finally, when both factories are **KO**, Ada just closes down the business by sending the label *bye* to both sellers:

$$G_4 = \text{Ada} \rightarrow \{\text{iS}, \text{aS}\} : \text{bye.end}$$

The processes in Table 1 are implementations of the monitors generated by projection from all the above global types. For instance, the monitor of **aS** from G_1 and G_3 is

$$\text{aF}!(\text{String}, \text{Int}).\text{Ada}?check,$$

where ! represents output and ? represents input. The monitor of the American seller from G_2 is similar:

$$\text{iF}!(\text{String}, \text{Int}).\text{Ada}?check.$$

$Seller = \mu X.y!(item, amount).y?check.X + y?bye$
 $Factory = \mu X.y?(x, w).if \dots then y?check.X + y?(x, w).y?check.X$
 $\quad \quad \quad \text{else write } \textcolor{red}{KO}.(y?check + y?(x, w).y?check)$
 $Ada = \mu X.y!check(F).X + y!contract.y!check(F).X + y!bye$
 $Bob = \mu X.y?(z).if \dots then write \textcolor{green}{OK}.y?check else y?check.X$

Table 1 Processes for the Company example

Its monitor from G_4 is simpler: $Ada?bye.end$. The process code for the seller has only two alternative behaviours, since processes do not mention senders and receivers. The seller can send on channel y item and amount, receive the *check* and then restart. Otherwise, he can receive *bye* and stop.

The control data can be modified by the factory, writing $\textcolor{red}{KO}$ when it is incapacitated, and by Bob, writing $\textcolor{green}{OK}$ when he accomplished the rebuilding task. The adaptation function F in Ada's process gives the new global type when applied to the pair (*state* $\textcolor{blue}{iF}$, *state* $\textcolor{red}{aF}$), i.e.

$$\begin{array}{ll}
 F(\textcolor{green}{OK}, \textcolor{green}{OK}) = G_1 & F(\textcolor{green}{OK}, \textcolor{red}{KO}) = G_2 \\
 F(\textcolor{red}{KO}, \textcolor{green}{OK}) = G_3 & F(\textcolor{red}{KO}, \textcolor{red}{KO}) = G_4
 \end{array}$$

A process can implement several different monitors also thanks to the external choice constructor. For instance, the process *Seller* can fill all the monitors that are generated by projecting the above global types onto the participants $\textcolor{blue}{iS}$ and $\textcolor{red}{aS}$.

Let us consider the system choreographed by G_1 with the global data $(\textcolor{green}{OK}, \textcolor{green}{OK})$. The American factory changes its state to $\textcolor{red}{KO}$ and then, when the chief checks the global data, the function F generates the adaptation step which produces the global type G_2 . After this adaptation Bob is a new participant, while the American factory is out. Then the American seller, as prescribed by his monitor, sends his requests to the Italian factory. When process Bob writes $\textcolor{green}{OK}$ for the American factory and the Italian factory is still $\textcolor{green}{OK}$, the global type produced by the adaptation step is again G_1 . Then the American factory comes back into the scene.

The present paper is a revised and extended version of [12]. Key additions with respect to [12] are a full formalisation of the safety properties and more examples which illustrate characterising features of our framework.

Structure of the paper Sections 2 and 3 present the syntax of our calculus and its type system, respectively. The formal semantics is given in Section 4. Examples in Section 5 enlighten key technical points of our approach. Formal properties are the content of Section 6. Related works are discussed in Section 7. Section 8 concludes.

2 Syntax

Global types Following a widely common approach, the set-up of protocols starts from global types. Global types establish overall communication schemes. In our setting they

also control the reconfiguration phase, in which a system adapts itself to new environmental conditions.

Let L be a set of *labels*, ranged over by ℓ , which mark the exchanged values as in [19] and Λ be a set of *flags*, ranged over by λ , which transmit the adaptation information. We assume to have some basic *sorts*, ranged over by S , i.e.

$$S ::= \text{Bool} \mid \text{Int} \mid \dots$$

Definition 1 *Global types* are defined by:

$$G ::= p \rightarrow \Pi : \{\ell_i(S_i).G_i\}_{i \in I} \mid p \rightarrow \Pi : \{\lambda_i\}_{i \in I} \mid \text{end}$$

In writing $\{\ell_i(S_i).G_i\}_{i \in I}$ and $\{\lambda_i\}_{i \in I}$ we implicitly assume that $\ell_i \neq \ell_j$ and $\lambda_i \neq \lambda_j$ for all $i \neq j$. There are only two kinds of communications: value exchange and adaptation flag exchange. Each value exchange is characterised by a label which allows to represent choices. The sender is p , while Π is the set of the receivers, which does not contain p and cannot be the empty set. The participants of a global type G are all the senders and the receivers in G , ranged over by $p, q, \dots$. We denote by $\text{pa}(G)$ the set of all participants in G .

Global types can terminate in two ways: either with the usual end or with the exchange of adaptation flags. In the latter case the adaptation flags are sent by a participant to all the other ones. Adaptation flags can be seen as synchronisation points, interleaved in a conversation, at which different interaction paths can be taken. In the global types syntax there is no recursion operator, but recursive protocols can be obtained by reconfiguring the system with the same global type. Recursion can then be considered as a particular case of reconfiguration. A recursion operator is included instead in the syntax of processes. For instance, all processes in the company example of the Introduction are recursive and they implement monitors which are projections of non-recursive global types.

Notably, we do not allow parallel composition of global types, which is quite common in the literature [28, 9, 3, 10]. As a matter of fact many papers [28, 9, 3] require that two global types can be put in parallel only if their sets of participants are disjoint, so parallel composition can be expressed by interleaving. Without this condition parallel composition of global types requires some care [10]. This issue is orthogonal to the present framework, where each participant, in all reconfiguration steps, follows one global type only (see Tables 8 and 9).

Monitors Monitors can be viewed as local types that are obtained as projections of global types onto individual participants, as in the standard approach of [28] and [2]. The only syntactic differences are the presence of the adaptation flags and the absence of recursion and delegation. In our calculus, however, monitors are more than types: they have an active role in system dynamics, since they guide communications and adaptations.

$$\begin{aligned}
& (p \rightarrow \Pi : \{\ell_i(S_i).G_i\}_{i \in I}) \upharpoonright q = \\
& \begin{cases} p? \{\ell_i(S_i).G_i\}_{i \in I} & \text{if } q \in \Pi \\ \Pi! \{\ell_i(S_i).G_i\}_{i \in I} & \text{if } q = p \\ G_{i_0} \upharpoonright q & \text{where } i_0 \in I \text{ if } q \neq p \text{ and } q \notin \Pi \\ & \text{and } G_i \upharpoonright q = G_j \upharpoonright q \text{ for all } i, j \in I \end{cases} \\
& (p \rightarrow \Pi : \{\lambda_i\}_{i \in I}) \upharpoonright q = \begin{cases} p? \{\lambda_i\}_{i \in I} & \text{if } q \in \Pi \\ \Pi! \{\lambda_i\}_{i \in I} & \text{if } q = p \\ \text{end} & \text{if } q \neq p \text{ and } q \notin \Pi \end{cases} \\
& \text{end} \upharpoonright p = \text{end}
\end{aligned}$$

Table 2 Projection of a global type onto a participant

Definition 2 The set of *monitors* is defined by:

$$\begin{aligned}
\mathcal{M} ::= & \quad p? \{\ell_i(S_i)..\mathcal{M}_i\}_{i \in I} \quad | \quad \Pi! \{\ell_i(S_i)..\mathcal{M}_i\}_{i \in I} \quad | \\
& \quad p? \{\lambda_i\}_{i \in I} \quad | \quad \Pi! \{\lambda_i\}_{i \in I} \quad | \\
& \quad \text{end}
\end{aligned}$$

The constructs in the first line correspond to input and output actions, respectively. An input monitor $p? \{\ell_i(S_i)..\mathcal{M}_i\}_{i \in I}$ fits with a process that can receive, for each $i \in I$, a value of sort S_i , labeled by ℓ_i , having as continuation a process which agrees with $\mathcal{M}_i$. This corresponds to an external choice. Dually an output monitor $\Pi! \{\ell_i(S_i)..\mathcal{M}_i\}_{i \in I}$ fits with a process which can send (by an internal choice) a value of sort S_i , distinguished by the label ℓ_i for each $i \in I$, and then continues as prescribed by $\mathcal{M}_i$.

The projection of global types onto participants is given in Table 2. A projection is undefined when two participants not involved, as sender or receiver, in a choice have different projections in different branchings (condition $G_i \upharpoonright q = G_j \upharpoonright q$ for all $i, j \in I$). Monitors are the results of such projections.

A global type G is *well formed* if its projections are defined for all participants and all occurrences of

$$p \rightarrow \Pi : \{\lambda_i\}_{i \in I}$$

are such that $\Pi \cup \{p\} = \text{pa}(G)$: i.e. all participants are involved in each flag exchange. In the following we assume that all global types are well formed.

Processes Processes represent code that is associated to monitors in order to implement participants.

Differently from session calculi [27, 23, 24, 28, 2, 33, 19, 11, 34, 3], processes do not specify the participants involved in sending and receiving actions. The associated monitors determine senders and receivers. Processes represent flexible code that can be associated to different monitors to incarnate different participants. Besides communicating, processes can access the global state to read or change it.

The communication actions of processes are performed through *channels*. Each process owns a unique channel. We use y to denote this channel in the user code. As usual, the

user channel y will be replaced at run time by a session channel $s[p]$ (where s is the session name and p is the current participant). Let c denote a user channel or a session channel. We could avoid to write explicitly the channel in the user syntax, but not in the run-time syntax. We have chosen to write all channels to simplify the definition of processes.

Definition 3 *Processes* are defined by:

$$\begin{aligned}
P ::= & \quad \mathbf{0} \quad | \quad \text{op}.P \quad | \quad X \quad | \quad \mu X.P \quad | \\
& \quad c? \ell(x).P \quad | \quad c! \ell(e).P \quad | \\
& \quad c?(\lambda, T).P \quad | \quad c!(\lambda(F), T).P \quad | \\
& \quad \text{if } e \text{ then } P \text{ else } P \quad | \quad P+P
\end{aligned}$$

The syntax of processes is rather standard, in particular the operator $+$ represents external choice. Notice that internal and external choices with many branches can easily be encoded in our calculus, which gains in simplicity. In writing processes we assume the precedence: prefix, external choice, recursion. Note that in the sending and receiving actions the involved participants are missing. For instance, $c! \ell(e).P$ denotes a process which sends via the channel c the label ℓ and the value of the expression e and then has P as continuation. Notably, a system has a global state (see Definition 5) and the op operator represents an action on this global state, for instance a “read” or “write” operation. We leave unspecified the kind of actions since we are only interested in the dynamic changes of this state, which plays the role of the control data for the self-reconfiguration of the whole system.

Types, which are statically assigned to processes, will be formally introduced in Section 3 (Definition 7). Types are mainly aimed at checking the matching between processes and monitors. It is convenient to include a type annotation in the syntax of the adaptation flag. The input flag $c?(\lambda, T).P$ represents a process that, after receiving the adaptation flag λ , has a continuation of type T . Thus the explicit annotation T makes it easy to dynamically check if, after the adaptation, the current process can continue with that type, inside the new monitor. The output flag $c!(\lambda(F), T).P$ contains also the *adaptation function* F . The application of F to the global state will determine the new global type, which provides a new choreography for the system.

Network A process is always controlled by a monitor, which ensures that all performed actions fit the protocol prescribed by the global type. Each monitor controls a single process. So participants correspond to pairs of processes and monitors. We write $\mathcal{M}[P]$ to represent a process P controlled by a monitor $\mathcal{M}$, dubbed *monitored process*. In a reconfiguration phase the monitor controlling the process is changed according to the new global type resulting from the application of the adaptation function to the global state. At this

point the processes whose type does not fit the new monitor must leave the system and new ones can enter it. The data exchange among the participants is done by means of runtime queues (one for each active session). We denote by $s : h$ the *named queue* associated with the session s , where h is a *message queue*. The empty queue is denoted by $\emptyset$. Messages in queues can be either value messages $(p, \Pi, \ell(v))$, indicating that the label ℓ and the value v are sent by participant p to all participants in Π , or adaptation messages $(p, \Pi, \lambda(G))$, indicating that the flag λ and the global type G are sent by participant p to all participants in Π . Queue concatenation, denoted by “.”, has $\emptyset$ as neutral element. A queue is λ -free if it contains no flag.

The sessions are initiated by the “new” constructor applied to a global type (*session initiator*), denoted by $\text{new}(G)$, which generates the monitors and associates them with adequate processes (see Definition 8).

The parallel composition of session initiators, processes with the corresponding monitors and runtime queues form a network. Networks can be restricted on session names.

Definition 4 *Networks* are defined by:

$$N ::= \text{new}(G) \mid \mathcal{M}[P] \mid s : h \mid N \mid N \mid (\nu s)N$$

System A system includes a network, a global state and assumes a collection of processes together with their types (according to the typing rules of Section 3). We use σ to range over global states and we denote by $\mathcal{P}$ the collection of pairs (P, T) . We represent systems as the composition (via “ $\parallel$ ”) of a networks and a global state, without mentioning the process collection, which is considered an implicit parameter.

Definition 5 *Systems* are defined by:

$$\mathcal{S} ::= N \parallel \sigma$$

3 Process types

Process types (called simply *types* where not ambiguous) describe process communication behaviours [27]. They have prefixes corresponding to sending and receiving of labels and flags. In particular an *input type* is a type whose first prefix corresponds to an input action and an *output type* is a type whose first prefix corresponds to an output action, while the *continuation* of a type is the type following its first prefix. A *communication type* is either an input or an output type. The external choice is typed by an intersection type, since an external choice offers both behaviours of the composing processes. Dually, a conditional is an internal choice and so it is typed by a union type. Notice that union and intersection being binary constructor feet well with our binary internal and external choices.

$$\begin{aligned} \text{lin}(\ell(S).T) &= \text{lout}(\ell(S).T) = \{\ell\} \\ \text{lin}(\ell\lambda) &= \text{lout}(\ell\lambda) = \{\lambda\} \\ \text{lin}(\ell(S).T) &= \text{lin}(\ell\lambda) = \text{lout}(\ell(S).T) = \text{lout}(\ell\lambda) = \emptyset \\ \text{lin}(T_1 \wedge T_2) &= \text{lin}(T_1 \vee T_2) = \text{lin}(T_1) \cup \text{lin}(T_2) \\ \text{lout}(T_1 \wedge T_2) &= \text{lout}(T_1 \vee T_2) = \text{lout}(T_1) \cup \text{lout}(T_2) \end{aligned}$$

Table 3 The mappings *lin* and *lout*

To formally define types, we start with the more liberal syntax of pre-types and then we define some restrictions that characterise types of processes.

Definition 6 The set of *pre-types* is inductively defined by:

$$T ::= \ell(S).T \mid \ell(S).T \mid \ell\lambda \mid \ell\lambda \mid T \wedge T \mid T \vee T \mid \text{end}$$

where $\wedge$ and $\vee$ are considered modulo idempotence, commutativity and associativity.

In pre-types and types we assume that $.$ has precedence over $\wedge$ and $\vee$.

In order to define types for processes, we have to avoid intersection between input types with the same first label, which would represent ambiguous external choices: indeed, the types following a same input prefix could be different and this would lead to a communication mismatch, as illustrated in Example 1 of Section 5. For the same reason, process types cannot contain intersections between output types with the same label. Since we have to match types with monitors, where internal choices are always taken by participants sending a label or a flag, we force unions to take output types (possibly combined by intersections or unions) as arguments. Therefore, we formalise the above restrictions by means of two mappings from pre-types to sets of labels and flags (Table 3) and then we define types by using those mappings.

Definition 7 A *type* is a pre-type satisfying the following constraints modulo idempotence, commutativity and associativity of unions and intersections:

- all occurrences of the shape $T_1 \wedge T_2$ are such that

$$\text{lin}(T_1) \cap \text{lin}(T_2) = \text{lout}(T_1) \cap \text{lout}(T_2) = \emptyset$$

- all occurrences of the shape $T_1 \vee T_2$ are such that

$$\text{lin}(T_1) = \text{lin}(T_2) = \text{lout}(T_1) \cap \text{lout}(T_2) = \emptyset.$$

We use T to range over types and $\mathcal{T}$ to denote the set of types. Note that, for example, $(T \wedge T) \vee T$ is a type, whenever T is a type, since types are considered modulo idempotence.

An *environment* Γ is a finite mapping from expression variables to sorts and from process variables to types:

$$\Gamma ::= \emptyset \mid \Gamma, x : S \mid \Gamma, X : T$$

where the notation $\Gamma, x : S$ ($\Gamma, X : T$) means that x (X) does not occur in Γ .

We assume that expressions are typed by sorts, as usual. The typing judgments for expressions are of the shape

$$\Gamma \vdash e : S$$

and the typing rules for expressions are standard.

Only processes with at most one channel can be typed. This choice is justified by the design of monitored processes as session participants and by the absence of delegation. Therefore the typing judgments for processes have the form

$$\Gamma \vdash P \triangleright c : T.$$

Typing rules for processes are given in Table 4. Observe that the type of a process after a reconfiguration is memorised in the (input or output) action in which the adaptation flag is exchanged (see Definition 3). In rules IF and CHOICE we require that the applications of union and intersection on two types form a type (conditions $T_1 \vee T_2 \in \mathcal{T}$ and $T_1 \wedge T_2 \in \mathcal{T}$). Adaptation allows us to avoid recursive types. A recursion variable is always preceded by an adaptation action, i.e. $c!(\lambda, T).X$ (rule RV1) and $c!(\lambda(F), T).X$ (rule RV2). In typing a recursive process $\mu X.P$, rule REC ensures that the type of P is the same as the type associated to X in the environment. Note that $\mu X.P$ is equivalent to $P\{\mu X.P/X\}$ and so, unfolding the process, P will always be associated to all the reconfiguration flags which precede the occurrences of X .

For example, writing the process *Ada* (considered in the Introduction) using the formal syntax, but leaving out labels, $y!(check(F), T_{Ada})$ replaces $y!check(F)$, where the type of the whole process *Ada* is:

$$T_{Ada} = !check \wedge !String. !check \wedge !bye. end$$

The matching between process types and monitors (adequacy) is made rather flexible by using the *subtype* relation on types defined in Table 5. Subtyping is monotone, for input/output prefixes, with respect to continuations and it follows the usual set theoretic inclusion of intersection and union. Notice that we use a weaker definition than standard subtyping on intersection and union types, since it is sufficient to define subtyping on types.

The intuitive meaning of subtyping is that a process with a smaller type has all the behaviours required by a bigger type and possibly more. Therefore *end* is the top type. An input monitor naturally corresponds to an external choice, while an output monitor naturally corresponds to an internal choice. So intersections of input types are adequate for input monitors and unions of output types are adequate for output monitors. Formally, we say that a type is adequate for a monitor if the conditions of the following definition hold.

$\leq$ is the minimal reflexive and transitive relation on $\mathcal{T}$ such that:

$$\begin{aligned} T &\leq \text{end} \quad T_1 \wedge T_2 \leq T_i \quad T_i \leq T_1 \vee T_2 \quad (i = 1, 2) \\ T_1 &\leq T_2 \text{ implies } !\ell(S).T_1 \leq !\ell(S).T_2 \\ T_1 &\leq T_2 \text{ implies } ?\ell(S).T_1 \leq ?\ell(S).T_2 \\ T &\leq T_1 \text{ and } T \leq T_2 \text{ imply } T \leq T_1 \wedge T_2 \\ T_1 &\leq T \text{ and } T_2 \leq T \text{ imply } T_1 \vee T_2 \leq T \\ (T_1 \vee T_2) \wedge T_3 &\leq T \text{ iff } T_1 \wedge T_3 \leq T \text{ and } T_2 \wedge T_3 \leq T \\ T &\leq (T_1 \wedge T_2) \vee T_3 \text{ iff } T \leq T_1 \vee T_3 \text{ and } T \leq T_2 \vee T_3 \end{aligned}$$

Table 5 Subtyping

Definition 8 A type T is *adequate* for a monitor $\mathcal{M}$ (notation $T \propto \mathcal{M}$) if $T \leq |\mathcal{M}|$, where the mapping $|\cdot|$ is defined by:

$$\begin{aligned} |p?\{\ell_i(S_i).\mathcal{M}_i\}_{i \in I}| &= \bigwedge_{i \in I} ?\ell_i(S_i).|\mathcal{M}_i| \\ |\Pi!\{\ell_i(S_i).\mathcal{M}_i\}_{i \in I}| &= \bigvee_{i \in I} !\ell_i(S_i).|\mathcal{M}_i| \\ |p?\{\lambda_i\}_{i \in I}| &= \bigwedge_{i \in I} ?\lambda_i \quad |\Pi!\{\lambda_i\}_{i \in I}| = \bigvee_{i \in I} !\lambda_i \quad |\text{end}| = \text{end} \end{aligned}$$

For instance, the type T_{Ada} defined above is adequate for the monitor of *Ada* obtained by projecting the global type G_2 discussed in the Introduction:

$$\text{Bob!String.}\{\text{iS, iF, aS, Bob}\}.check$$

Decidability of adequacy relies on decidability of subtyping. We show that subtyping is decidable. The proof exploits standard distributivity properties on intersections and unions.

Lemma 1 *Subtyping is decidable.*

Proof A subtyping between two types is equivalent to a set of subtypings, in which no union occurs in the left type and no intersection occurs in the right type. In fact:

- $T_1 \vee T_2 \leq T_3$ iff $T_1 \leq T_3$ and $T_2 \leq T_3$
- $(T_1 \vee T_2) \wedge T_3 \leq T_4$ iff $T_1 \wedge T_3 \leq T_4$ and $T_2 \wedge T_3 \leq T_4$.

Notice that if $(T_1 \vee T_2) \wedge T_3$ is a type, then both $T_1 \wedge T_3$ and $T_2 \wedge T_3$ are types. A similar argument can be used for erasing the intersections in the right type. Then we have to decide only subtypings of the shape $T \leq T'$, where T is an intersection of communication types and possibly *end*, while T' is either a union of output types or a single input type, both possibly in union with *end* (by Definition 7). Since *end* is the top type:

- If T' contains *end*, then the subtyping holds.
 - If T is *end*, then the subtyping fails, unless T' contains *end*.
- Otherwise, the occurrences of *end* in T can be erased. Thus we are reduced to consider the cases in which T is an intersection of communication types and T' is either a union of output types or a single input type. In the first case,

$$\begin{array}{c}
\Gamma \vdash \mathbf{0} \triangleright c : \text{end} \quad \text{END} \quad \frac{\Gamma \vdash P \triangleright c : T}{\Gamma \vdash \text{op}.P \triangleright c : T} \text{OP} \quad \Gamma, X : T \vdash c?(\lambda, T).X \triangleright c : ?\lambda \quad \text{RV1} \quad \Gamma, X : T \vdash c!(\lambda(F), T).X \triangleright c : !\lambda \quad \text{RV2} \\
\\
\frac{\Gamma, X : T \vdash P \triangleright c : T}{\Gamma \vdash \mu X.P \triangleright c : T} \text{REC} \quad \frac{\Gamma, x : S \vdash P \triangleright c : T}{\Gamma \vdash c? \ell(x).P \triangleright c : ?\ell(S).T} \text{RCV} \quad \frac{\Gamma \vdash P \triangleright c : T \quad \Gamma \vdash e : S}{\Gamma \vdash c!\ell(e).P \triangleright c : !\ell(S).T} \text{SEND} \\
\\
\frac{\Gamma \vdash P \triangleright c : T}{\Gamma \vdash c?(\lambda, T).P \triangleright c : ?\lambda} \text{FRCV} \quad \frac{\Gamma \vdash P \triangleright c : T}{\Gamma \vdash c!(\lambda(F), T).P \triangleright c : !\lambda} \text{FSEND} \\
\\
\frac{\Gamma \vdash e : \text{Bool} \quad \Gamma \vdash P_1 \triangleright c : T_1 \quad \Gamma \vdash P_2 \triangleright c : T_2 \quad T_1 \vee T_2 \in \mathcal{T}}{\Gamma \vdash \text{if } e \text{ then } P_1 \text{ else } P_2 \triangleright c : T_1 \vee T_2} \text{IF} \\
\\
\frac{\Gamma \vdash P_1 \triangleright c : T_1 \quad \Gamma \vdash P_2 \triangleright c : T_2 \quad T_1 \wedge T_2 \in \mathcal{T}}{\Gamma \vdash P_1 + P_2 \triangleright c : T_1 \wedge T_2} \text{CHOICE}
\end{array}$$

Table 4 Typing rules for processes

$$\begin{array}{c}
\text{op}.P \xrightarrow{\text{op}} P \quad \mu X.P \xrightarrow{\tau} P\{\mu X.P/X\} \quad s[p]? \ell(x).P \xrightarrow{s[p]? \ell(v)} P\{v/x\} \quad s[p]!\ell(e).P \xrightarrow{s[p]!\ell(v)} P \quad e \downarrow v \\
\\
s[p]?(\lambda, T).P \xrightarrow{s[p]?(\lambda, T)} P \quad s[p]!(\lambda(F), T).P \xrightarrow{s[p]!(\lambda(F), T)} P \\
\\
\text{if } e \text{ then } P \text{ else } Q \xrightarrow{\tau} P \quad e \downarrow \text{true} \quad \text{if } e \text{ then } P \text{ else } Q \xrightarrow{\tau} Q \quad e \downarrow \text{false} \\
\\
\frac{P \xrightarrow{\alpha} P'}{P + Q \xrightarrow{\alpha} P'} \quad \frac{P \xrightarrow{\delta} P'}{P + Q \xrightarrow{\delta} P' + Q}
\end{array}$$

Table 6 LTS of processes

subtyping $T \leq T'$ holds if and only if at least one of the output prefixes of types in T' is equal to the output prefix of a type in T , and the corresponding continuations are in the subtype relation. In the second case, subtyping $T \leq T'$ holds if and only if T contains a type which has the same input prefix of T' and the corresponding continuations are in the subtype relation. $\square$

4 Semantics

The evolution of a system depends on the evolution of its network and global state. The basic components of networks are the openings of sessions (through the new operator on global types) and the processes associated with monitors. So we start by describing how processes can evolve inside monitors. Monitors guide the communications of processes by choosing the senders/receivers and by allowing only some actions among those offered by the processes. This is formalised by the following LTS for monitors:

$$\begin{array}{c}
p? \{ \ell_i(S_i) .. \mathcal{M}_i \}_{i \in I} \xrightarrow{p? \ell_j} \mathcal{M}_j \quad j \in I \\
\\
\Pi! \{ \ell_i(S_i) .. \mathcal{M}_i \}_{i \in I} \xrightarrow{\Pi! \ell_j} \mathcal{M}_j \quad j \in I \\
\\
p? \{ \lambda_i \}_{i \in I} \xrightarrow{p? \lambda_j} \Pi! \{ \lambda_i \}_{i \in I} \xrightarrow{\Pi! \lambda_j} \quad j \in I
\end{array}$$

Processes can communicate labels and values, flags, adaptation functions and types, or can read/modify the global state through op operations. These behaviours are made explicit by the LTS in Table 6, where the treatment of recursions and conditionals is standard. In the rules for external choice, α ranges over

$$s[p]? \ell(v), s[p]!\ell(v), s[p]?(\lambda, T), s[p]!(\lambda(F), T),$$

and δ ranges over $\{\text{op}, \tau\}$. We omit the symmetric rules.

$$\begin{array}{c}
h \cdot (q, \Pi, \zeta) \cdot (q', \Pi', \zeta') \cdot h' \equiv h \cdot (q', \Pi', \zeta') \cdot (q, \Pi, \zeta) \cdot h' \\
\text{if } \Pi \cap \Pi' = \emptyset \text{ or } q \neq q' \\
\\
h \cdot (q, \Pi, \zeta) \cdot h' \equiv h \cdot (q, \Pi', \zeta) \cdot (q, \Pi'', \zeta) \cdot h' \\
\text{if } \Pi = \Pi' \cup \Pi'' \text{ and } \Pi' \cap \Pi'' = \emptyset
\end{array}$$

where $\zeta ::= \ell(v) \mid \lambda(G)$.

Table 7 Equivalence on message queues

The choices are done by the communication actions, while the operations on the global state are transparent. This is needed since the operations on the memory are recorded neither in the process types nor in the monitors. An operation on the state in an external choice can be performed also if a branch, different from that containing the operation, is executed. The rationale is the independence of the changes in the control data from the communications among session participants. For example,

$$\begin{aligned} & \text{op.s}[p]? \ell_1(x).P_1 + \text{s}[p]? \ell_2(x).P_2 \xrightarrow{\text{op}} \\ & \text{s}[p]? \ell_1(x).P_1 + \text{s}[p]? \ell_2(x).P_2 \xrightarrow{\text{s}[p]? \ell_2(v)} P_2\{v/x\}. \end{aligned}$$

A recursion in an external choice can be unfolded independently from the chosen branch, since unfolding is an internal computation. Similarly, a conditional in an external choice can be evaluated also if a different branch is then executed. For example,

$$\begin{aligned} & (\mu X.\text{if true then s}[p]? \ell_1(x).P_1 \text{ else } X) + \text{s}[p]? \ell_2(x).P_2 \xrightarrow{\tau} \\ & (\text{if true then s}[p]? \ell_1(x).P_1 \text{ else } \mu X.\dots) + \text{s}[p]? \ell_2(x).P_2 \xrightarrow{\tau} \\ & \text{s}[p]? \ell_1(x).P_1 + \text{s}[p]? \ell_2(x).P_2 \xrightarrow{\text{s}[p]? \ell_2(v)} P_2\{v/x\}. \end{aligned}$$

We assume a standard structural equivalence on networks, in which the parallel operator is commutative and associative, and restrictions commute and enlarge their scopes without name clashes. Moreover, any monitored process with end monitor behaves as the neutral element for the parallel and absorbs restrictions, that is:

$$\text{end}[P] \mid N \equiv N \quad \text{and} \quad (\text{vs})\text{end}[P] \equiv \text{end}[P].$$

For message queues, we need an equivalence for commuting independent messages and another one for splitting a message to multiple receivers, see Table 7. The equivalence on message queues induces an equivalence on labelled queues in the obvious way:

$$h \equiv h' \text{ implies } s : h \equiv s : h'.$$

We can distinguish between the transitions which do or do not involve the global state. For simplicity, Table 8 lists the reduction rules of the networks and Table 9 lists the reduction rules of the systems, in which all rules need the global state.

A session starts by reducing a network new G (rule INIT). For each p in the set $\text{pa}(G)$ of the participants in the global type G , we need to find a process P_p in the collection $\mathcal{P}$ associated to the current system. The process P_p must be such that its type T_p is adequate for the monitor which is the projection of G onto p . Then the process (where the channel y has been replaced by $s[p]$) is associated to the corresponding monitor and the empty queue s is created. Lastly, the name s is restricted. In this way we ensure the privacy of the communications in a session (as standard in session calculi [28]). We are interested here in modelling the overall adaptation

strategy, based on decoupling interfaces (i.e. monitors) and implementations (i.e. processes), rather than in the details related to the choice of processes associated to monitors. So we have left this choice arbitrary, the only condition being type adequacy. Note, however, that a natural way of controlling the processes associated to the monitors is given by the choice of the labels and flags which relate them.

The rules IN and OUT define the exchange of messages through queues. The type assignment system ensures that both the type of P is adequate for $\mathcal{M}$ and the type of P' is adequate for $\mathcal{M}'$. Following [11, 3], the agreement between monitors and processes is required; a novelty is that only the monitors define the senders and the receivers of messages.

The rules ADAINCONT and ADAINNEW of Table 8 deal with adaptations, for the session participants which receive the adaptation flag with the new global type. The new global type is needed to compute the new monitor $\mathcal{M}'$ by projection. In the first rule the continuation of the current process inside the monitor has a type which is adequate for $\mathcal{M}'$, so this process will fill $\mathcal{M}'$. In the second rule, instead, it is needed to take from a different process with a type adequate for $\mathcal{M}'$ the collection $\mathcal{P}$. In case such a process does not exist in $\mathcal{P}$, then the system gets stuck.

Evaluation contexts are defined by

$$\mathcal{E} ::= [] \mid \mathcal{E} \mid N \mid (\text{vs})\mathcal{E}$$

The reduction rules for networks can be used for reducing systems thanks to rule SN in Table 9. Rule CXT is a standard contextual rule. Rule OP allows processes to read/modify the global state.

The most interesting rules are ADAOUTCONT and ADAOUTNEW. In both rules participant p sends an adaptation flag and an adaptation function, whose application to the global state gives a new global type G . The global type G may involve new participants (in $\Pi' \setminus (\Pi \cup \{p\})$) which are added to the network by taking processes in $\mathcal{P}$ as in rule INIT. As regards to participant p , the new monitor $G \upharpoonright p$ will be associated or not to the current process according to whether its type is adequate or not for $G \upharpoonright p$, as in rules ADAINCONT and ADAINNEW. In both rules ADAOUTCON and ADAOUTNEW the message with the reconfiguration flag and the global type G will be sent to all participants of the session before the reconfiguration. This is ensured by the well-formedness of global types.

The restriction to λ -free queues deserves some comments. It ensures no new adaptation flag can be thrown until all the receivers of the previous adaptation flag have adapted themselves. A design choice of our framework is to allow a participant to skip an adaptation phase (since it does not appear in the new global type) and then to appear again in the following adaptation. This models a common scenario in which a component is temporarily unavailable and so a new choreography is needed. In the introductory example, the Amer-

$$\begin{array}{c}
\frac{\Pi = \text{pa}(\mathsf{G}) \quad \mathcal{M}_p = \mathsf{G} \upharpoonright p \quad \forall p \in \Pi. (P_p, \mathsf{T}_p) \in \mathcal{P} \ \& \ \mathsf{T}_p \propto \mathcal{M}_p}{\text{new}(\mathsf{G}) \longrightarrow (\nu s) \left(\prod_{p \in \Pi} \mathcal{M}_p[P_p\{s[p]/y\}] \mid s : \emptyset \right)} \text{INIT} \quad \frac{P \xrightarrow{\tau} P'}{\mathcal{M}[P] \longrightarrow \mathcal{M}[P']} \text{TAU} \\
\\
\frac{\mathcal{M} \xrightarrow{q? \ell} \mathcal{M}' \quad P \xrightarrow{s[p]? \ell(v)} P'}{\mathcal{M}[P] \mid s : (q, p, \ell(v)) \cdot h \longrightarrow \mathcal{M}'[P'] \mid s : h} \text{IN} \quad \frac{\mathcal{M} \xrightarrow{\Pi! \ell} \mathcal{M}' \quad P \xrightarrow{s[p]! \ell(v)} P'}{\mathcal{M}[P] \mid s : h \longrightarrow \mathcal{M}'[P'] \mid s : h \cdot (p, \Pi, \ell(v))} \text{OUT} \\
\\
\frac{\mathcal{M} \xrightarrow{q? \lambda} P \xrightarrow{s[p]?(\lambda, \mathsf{T})} P' \quad \mathsf{G} \upharpoonright p = \mathcal{M}' \quad \mathsf{T} \propto \mathcal{M}'}{\mathcal{M}[P] \mid s : (q, p, \lambda(\mathsf{G})) \cdot h \longrightarrow \mathcal{M}'[P'] \mid s : h} \text{ADAInCONT} \\
\\
\frac{\mathcal{M} \xrightarrow{q? \lambda} P \xrightarrow{s[p]?(\lambda, \mathsf{T})} P' \quad \mathsf{G} \upharpoonright p = \mathcal{M}' \quad \mathsf{T} \not\propto \mathcal{M}' \quad (Q, \mathsf{T}') \in \mathcal{P} \quad \mathsf{T}' \propto \mathcal{M}'}{\mathcal{M}[P] \mid s : (q, p, \lambda(\mathsf{G})) \cdot h \longrightarrow \mathcal{M}'[Q\{s[p]/y\}] \mid s : h} \text{ADAInNEW} \\
\\
\frac{N_1 \equiv N'_1 \quad N'_1 \longrightarrow N'_2 \quad N_2 \equiv N'_2}{N_1 \longrightarrow N_2} \text{EQUIV}
\end{array}$$

Table 8 Network reduction

$$\begin{array}{c}
\frac{P \xrightarrow{\text{op}} P'}{\mathcal{M}[P] \parallel \sigma \longrightarrow \mathcal{M}[P'] \parallel \text{op}(\sigma)} \text{OP} \\
\\
\frac{\mathcal{M} \xrightarrow{\Pi! \lambda} P \xrightarrow{s[p]!(\lambda(F), \mathsf{T})} P' \quad F(\sigma) = \mathsf{G} \quad \mathcal{M}_p = \mathsf{G} \upharpoonright p \quad \mathsf{T} \propto \mathcal{M}_p \quad h \text{ } \lambda\text{-free}}{\Pi' = \text{pa}(\mathsf{G}) \quad \forall q \in \Pi'. \mathcal{M}_q = \mathsf{G} \upharpoonright q \quad \forall q \in \Pi' \setminus (\Pi \cup \{p\}). (P_q, \mathsf{T}_q) \in \mathcal{P} \ \& \ \mathsf{T}_q \propto \mathcal{M}_q} \text{ADAOUTCONT} \\
\frac{\mathcal{M}[P] \mid s : h \parallel \sigma \longrightarrow \mathcal{M}_p[P'] \mid \prod_{q \in \Pi' \setminus (\Pi \cup \{p\})} \mathcal{M}_q[P_q\{s[q]/y_q\}] \mid s : h \cdot (p, \Pi, \lambda(\mathsf{G})) \parallel \sigma}{\mathcal{M}[P] \mid s : h \parallel \sigma \longrightarrow \prod_{q \in \Pi' \setminus \Pi} \mathcal{M}_q[P_q\{s[q]/y_q\}] \mid s : h \cdot (p, \Pi, \lambda(\mathsf{G})) \parallel \sigma} \text{ADAOUTNEW} \\
\\
\frac{N \longrightarrow N'}{\mathcal{E}[N] \parallel \sigma \longrightarrow \mathcal{E}[N'] \parallel \sigma} \text{SN} \quad \frac{N \parallel \sigma \longrightarrow N' \parallel \sigma'}{\mathcal{E}[N] \parallel \sigma \longrightarrow \mathcal{E}[N'] \parallel \sigma'} \text{CTX}
\end{array}$$

Table 9 System reduction

ican factory becomes temporarily out of the current choreography. Without the given restriction, when the component becomes available again, we could have two monitored processes with the same session channel, so loosing channel linearity. Observe, however, that by this restriction some participants are allowed to finish their communications before performing an adaptation, while other participants have already self-adapted and then started the new communications.

Note that the λ -freeness of queues can be implemented in several ways without breaking decentralisation, for example by semaphores on queues.

We use $\longrightarrow^*$ with the usual meaning and $\longrightarrow_{\mathcal{P}}, \longrightarrow_{\mathcal{P}}^*$ when we want emphasise the use of $\mathcal{P}$ in rules INIT, ADAInNEW, ADAOUTNEW.

5 Examples

In this section we discuss examples. The first example motivates the restrictions on the definition of process types (Definition 7). The second and the third examples extend the example given in the Introduction, using the syntax of Section 2. Example 2 shows a use of different adaptation flags. Example 3 illustrates the possibility of exchanging the participant who is in charge of sending the adaptation flag. For readability we omit $\{, \}$ in writing global types and monitors.

Example 1 This example shows the necessity of the conditions on types given in Definition 7.

Suppose that any pre-type is a type, by removing the conditions from Definition 7. Then $\mathcal{P}$ could contain (P_1, T_1) and (P_2, T_2) such that:

$$\begin{aligned} P_1 &= y!\ell(3).y?\ell'(x).\mathbf{0} + y!\ell(\text{true}).\mathbf{0} \\ P_2 &= y?\ell(x').y!\ell'(-x').\mathbf{0} \\ T_1 &= !\ell(\text{int}).?\ell'(\text{int}).\text{end} \wedge !\ell(\text{Bool}).\text{end} \\ T_2 &= ?\ell(\text{int}).!\ell'(\text{int}).\text{end} \end{aligned}$$

In particular, we observe that T_1 has an intersection between output types with the same label. Take then the global type $G = 1 \rightarrow 2 : \ell(\text{int}).2 \rightarrow 1 : \ell'(\text{int}).\text{end}$, whose projections on participants 1 and 2 are

$$\mathcal{M}_1 = 2!\ell(\text{int}).2?\ell'(\text{int}).\text{end}, \mathcal{M}_2 = 1?\ell(\text{int}).1!\ell'(\text{int}).\text{end},$$

respectively. According to our definition of adequacy, T_1 and T_2 are adequate for $\mathcal{M}_1$ and $\mathcal{M}_2$, respectively. It is easy to verify that the network $\text{new}(G)$ can reduce to

$$(\nu s)(s[2]1!\ell'(-\text{true}).\mathbf{0} \mid s : \emptyset),$$

which is *stuck*. On the other hand, taking

$$P'_1 = y!\ell(3).y?\ell'(x).\mathbf{0} + y?\ell(x).\mathbf{0}$$

with type $T'_1 = !\ell(\text{int}).?\ell'(\text{int}).\text{end} \wedge ?\ell(\text{Bool}).\text{end}$, we still have that T'_1 and T_2 are adequate for $\mathcal{M}_1$ and $\mathcal{M}_2$, but the network $\text{new}(G)$ smoothly terminates the computation. Note that T'_1 satisfies the conditions of Definition 7, so there is no possible ambiguity on which branch of the external choice must be chosen.

Example 2 Ada wants to consider the possibility to keep the business going, also when both factories are out of use. In this case she can choose either to stop all activities or to start the reconstruction of both factories. In the former case she sends a *stop* adaptation flag and in the latter case a *go* adaptation flag. The global type G_4 of the example in the Introduction must then be replaced by:

$$G_4 = \text{Ada} \rightarrow \{\text{iS}, \text{aS}\} : \{\text{stop}, \text{go}\}$$

If Ada decides to go on with the business when both factories are out, both sellers send their requests to Ada. So we need two new global types:

$$G_g = \begin{cases} \text{Ada} \rightarrow \text{Bob} : \text{String}. \\ \text{iS} \rightarrow \text{Ada} : \text{req}(\text{String}, \text{Int}). \\ \text{aS} \rightarrow \text{Ada} : \text{req}(\text{String}, \text{Int}). \\ \text{Ada} \rightarrow \{\text{iS}, \text{aS}\} : \{\text{check}\} \end{cases} \quad G_s = \text{end}$$

where *req* is the label used to ask the amount of an item, and two adaptation functions:

$$F_g(\text{KO}, \text{KO}) = G_g \quad F_s(\text{KO}, \text{KO}) = G_s$$

A new process for implementing Ada can be:

$$\begin{aligned} \text{Ada}' &= \mu X. y!(\text{check}(F), T_{\text{Ada}'}).X + \\ &\quad y!\text{help}(\text{contract}).(y!(\text{check}(F), T_{\text{Ada}'}).X + \\ &\quad y?\text{req}(x, w).y?\text{req}(x', w').y!(\text{check}(F), T_{\text{Ada}'}).X) + \\ &\quad \text{if } \dots \text{ then } !(stop(F_s), \text{end}).\mathbf{0} \text{ else } !(go(F_g), T_{\text{Ada}'}).X \end{aligned}$$

where

$$\begin{aligned} T_{\text{Ada}'} &= !\text{check} \wedge \\ &\quad !\text{String}.(!\text{check} \wedge ?(\text{String}, \text{Int}).?(\text{String}, \text{Int}).!\text{check}) \\ &\quad \wedge (!\text{stop} \vee !\text{go}) \end{aligned}$$

is the new type of Ada. No modification is required for the processes representing the sellers and the factories.

Note that the process *Ada* of the Introduction could work as well as *Ada'* until both factories are out. In this case *Ada* would no longer agree with the monitor corresponding to the projection of the above G_4 and rule ADAOUTNEW would replace *Ada'* to *Ada*.

Example 3 Assume that Ada has to take a maternity leave. She then decides to transfer the job of monitoring the factories to one of the two sellers, who will play the role of deputy chief in the company. Then Ada writes *iS* (Italian seller) or *aS* (America seller) in the global data. The deputy chief is only in charge to keep the business working, unless both factories are out, in which case he also closes down the business. The deputy chief checks the state of the factories sending as Ada the reconfiguration flag *check*. When Ada is back, the standard management policy is restored.

In this example we use the same global types and adaptation function of the Introduction, but we add eight global types and another adaptation function.

The first four global types consider the case in which the American seller is the deputy chief. The global type G_1^a is the communication protocol when both factories are working:

$$G_1^a = \begin{cases} \text{iS} \rightarrow \text{iF} : \text{req}(\text{String}, \text{Int}). \\ \text{aS} \rightarrow \text{aF} : \text{req}(\text{String}, \text{Int}). \\ \text{aS} \rightarrow \{\text{iS}, \text{iF}, \text{aF}\} : \text{check} \end{cases}$$

If the Italian factory is *OK*, but the American one is *KO* the communication protocol becomes:

$$G_2^a = \begin{cases} \text{iS} \rightarrow \text{iF} : \text{req}(\text{String}, \text{Int}). \\ \text{aS} \rightarrow \text{iF} : \text{req}(\text{String}, \text{Int}). \\ \text{aS} \rightarrow \{\text{iS}, \text{iF}\} : \text{check} \end{cases}$$

If the American factory is *OK*, but the Italian one is *KO*, the global type G_3^a is as expected.

When both factories are *KO* the deputy chief is forced to close the business:

$$G_4^a = \text{aS} \rightarrow \{\text{iS}\} : \text{bye}.\text{end}$$

The other four global types G_1^i - G_4^i prescribe that the adaptation flag is sent by the Italian rather than by the American seller.

The new adaptation function F' , which has a third argument corresponding to the deputy chief in charge, is defined as:

$$\begin{aligned} F'(\text{OK}, \text{OK}, \text{aS}) &= G_1^a & F'(\text{OK}, \text{KO}, \text{aS}) &= G_2^a \\ F'(\text{KO}, \text{OK}, \text{aS}) &= G_3^a & F'(\text{KO}, \text{KO}, \text{aS}) &= G_4^a \\ F'(\text{OK}, \text{OK}, \text{iS}) &= G_1^i & F'(\text{OK}, \text{KO}, \text{iS}) &= G_2^i \\ F'(\text{KO}, \text{OK}, \text{iS}) &= G_3^i & F'(\text{KO}, \text{KO}, \text{iS}) &= G_4^i \end{aligned}$$

A process Ada'' implementing Ada includes a conditional to take into account the maternity leave:

$$\text{Ada}'' = \mu X. P + y!\text{help}(\text{contract}).P + !y.\text{bye}.\mathbf{0}$$

where help is the label for the contract with Bob and

$$P = \text{if } \dots \text{ then } !(check(F), T_{\text{Ada}''}).X \\ \text{else write } \text{dep}.\!(check(F'), \text{end}).\mathbf{0}$$

with $\text{dep} = \text{aS}$ or $\text{dep} = \text{iS}$ and

$$T_{\text{Ada}''} = !check \wedge !\text{help}(\text{String}).!check \wedge !\text{bye}.\text{end}$$

A process implementing sellers uses the adaptation function F' or F according to whether or not Ada is on leave:

$$\begin{aligned} \text{Seller}' &= \mu X. y!\text{req}(\text{item}, \text{amount}).(y?(check, T_{\text{Seller}'}).X + \\ &\quad \text{if } \dots \text{ then } y!(check(F), T_{\text{Seller}'}).X \\ &\quad \text{else } y!(check(F'), T_{\text{Seller}'}).X) \\ &\quad + y?\text{bye}.\mathbf{0} + y!\text{bye}.\mathbf{0} \end{aligned}$$

where:

$$T_{\text{Seller}'} = !\text{req}(\text{String}, \text{Int}).(?check \\ \wedge !check) \wedge ?\text{bye}.\text{end} \wedge !\text{bye}.\text{end}$$

No modification is needed, instead, for the processes representing the other participants, i.e. the two factories and Bob.

Notice that $T_{\text{Ada}''}$ differs from T_{Ada} (see page 6) only for the label help (since in the Introduction we used a simplified syntax). Instead $T_{\text{Seller}'}$ offers more choices than the type of the process Seller in the Introduction. As a consequence, the process Ada of the Introduction is an implementation of Ada also in the present new scenario, which never uses the right of the maternity leave. Instead the process Seller is no longer adequate since it does not implement the chief's behaviour.

6 Safety

In this section, we show subject reduction and progress theorems, following essentially the proof pattern of similar results for multiparty sessions, see e.g. [13]. Indeed, the main innovation of our calculus is that global types, with the corresponding monitors, are reconfigured at each adaptation step. Furthermore, participants of two different global types can coexist inside the same session. This happens when some participants have already performed the adaptation and then they are following the new global type, while other participants are still completing the interactions prescribed by the

old global type. These are the crucial technical difficulties in proving that monitored well-typed processes always behave in a type safe way. Therefore, we need to introduce typing rules for systems, which associate types to session channels.

The type of channel $s[p]$ is formed by taking into account the monitor, which controls the process owning $s[p]$, and the messages in the queue of the session s . The type of a monitored process is the association of the monitor to the session channel owned by the process. For each value message $(p, \Pi, \ell(v))$ in the queue of session s we associate the type $\Pi!\ell(S)$ to $s[p]$, where S is the sort of the value v , preserving the order of messages of queue. So lists of types of this shape form the types of session channels.

A queue can also contain adaptation messages. Note that, thanks to the condition of λ -freeness in rules ADAOUT-CONT and ADAOUT-NEW (Table 9), at most one adaptation message can occur in a queue (modulo structural equivalence, see Table 7). If the queue of session s contains the message $(p, \Pi, \lambda(G))$, then $\Pi!\lambda$ occurs in the type of $s[p]$. A *message type* is then a list of types of the shape $\Pi!\ell(S)$ possibly containing a type of the shape $\Pi!\lambda$. After receiving the message $(p, \Pi, \lambda(G))$, each participant $q \in \Pi$ of session s will behave according to the monitor $G \upharpoonright q$. Therefore the type of $s[q]$ can involve two monitors. One (explicit) monitor (dubbed *active monitor*) is the monitor of the monitored process owing $s[p]$. The other (implicit) monitor (dubbed *virtual monitor*) is the projection onto q of the global type contained in the adaptation message waiting to be received by $s[q]$ (and to become active). A missing virtual monitor is denoted by “—”. In particular, the virtual monitor of the sender of the adaptation message is always missing. So typing rules for queues associate types of the shape $\langle \text{message type}, \text{virtual monitor} \rangle$ (corresponding to the sent messages and the virtual monitor) to session channels.

To sum up, a type of a session channel is either an active monitor, or a pair consisting of a message type and a virtual monitor, or a triple consisting of a message type, an active monitor and a virtual monitor.

Definition 9 *Message types, queue types, virtual monitors, and generalised types* are defined by:

Message Types	$\mathbf{m}$	$::=$	$\varepsilon \mid \Pi!\ell(S) \mid \Pi!\lambda \mid \mathbf{m}; \mathbf{m}$
Virtual Monitors	$\mathcal{V}$	$::=$	$\mathcal{M} \mid -$
Queue Types	$\mathcal{Q}$	$::=$	$\langle \mathbf{m}, \mathcal{V} \rangle$
Generalised Types	χ	$::=$	$\mathcal{M} \mid \mathcal{Q} \mid \langle \mathbf{m}, \mathcal{M}, \mathcal{V} \rangle$

where “ $;$ ” is associative, ε is the type of the empty sequence of messages, such that $\varepsilon; \mathbf{m} = \mathbf{m}$; $\varepsilon = \mathbf{m}$, and “ $-$ ” denotes a missing monitor.

The typing judgements for systems are of the shape

$$\vdash_{\Sigma} \mathcal{S} \triangleright \Delta$$

where Σ is a set of session names (the names of the queues

$$\begin{aligned}
(\Pi! \ell(S); m) \upharpoonright q &= \begin{cases} !\ell(S). (m \upharpoonright q) & \text{if } q \in \Pi \\ m \upharpoonright q & \text{otherwise} \end{cases} & (\Pi! \lambda; m) \upharpoonright q &= \begin{cases} !\lambda. (m \upharpoonright q) & \text{if } q \in \Pi \\ m \upharpoonright q & \text{otherwise} \end{cases} & \varepsilon \upharpoonright q &= \varepsilon \\
p? \{ \ell_i(S_i). \mathcal{M}_i \}_{i \in I} \upharpoonright q &= \begin{cases} ? \{ \ell_i(S_i). (\mathcal{M}_i \upharpoonright q) \}_{i \in I} & \text{if } q = p \\ \mathcal{M}_{i_0} \upharpoonright q & \text{where } i_0 \in I, \text{ if } q \neq p \text{ and } \mathcal{M}_i \upharpoonright q = \mathcal{M}_j \upharpoonright q \forall i, j \in I \end{cases} \\
\Pi! \{ \ell_i(S_i). \mathcal{M}_i \}_{i \in I} \upharpoonright q &= \begin{cases} ! \{ \ell_i(S_i). (\mathcal{M}_i \upharpoonright q) \}_{i \in I} & \text{if } q \in \Pi \\ \mathcal{M}_{i_0} \upharpoonright q & \text{where } i_0 \in I, \text{ if } q \notin \Pi \text{ and } \mathcal{M}_i \upharpoonright q = \mathcal{M}_j \upharpoonright q \forall i, j \in I \end{cases} \\
p? \{ \lambda_i \}_{i \in I} \upharpoonright q &= \begin{cases} ? \{ \lambda_i \}_{i \in I} & \text{if } q = p \\ \varepsilon & \text{otherwise} \end{cases} & \Pi! \{ \lambda_i \}_{i \in I} \upharpoonright q &= \begin{cases} ! \{ \lambda_i \}_{i \in I} & \text{if } q \in \Pi \\ \varepsilon & \text{otherwise} \end{cases} \\
\text{end} \upharpoonright q &= \varepsilon & - \upharpoonright q &= \varepsilon \\
\langle m, \mathcal{V} \rangle \upharpoonright q &= m \upharpoonright q. \mathcal{V} \upharpoonright q & \langle m, \mathcal{M}, \mathcal{V} \rangle \upharpoonright q &= m \upharpoonright q. \mathcal{M} \upharpoonright q. \mathcal{V} \upharpoonright q
\end{aligned}$$

Table 10 Projection of generalised types onto participants

$$\begin{aligned}
\Theta \bowtie \Theta_j \text{ \& } j \in I \text{ imply } !\ell_j(S_j). \Theta \bowtie ? \{ \ell_i(S_i). \Theta_i \}_{i \in I} \\
\forall i \in I \Theta_i \bowtie \Theta'_i \text{ imply } !\{ \ell_i(S_i). \Theta_i \}_{i \in I} \bowtie ? \{ \ell_i(S_i). \Theta'_i \}_{i \in I} \\
j \in I \text{ implies } !\lambda_j \bowtie ? \{ \lambda_i \}_{i \in I} \quad !\{ \lambda_i \}_{i \in I} \bowtie ? \{ \lambda_i \}_{i \in I} \quad \varepsilon \bowtie \varepsilon \\
\Theta_1 \bowtie \Theta_2 \text{ and } \Theta_3 \bowtie \Theta_4 \text{ imply } \Theta_1. \Theta_3 \bowtie \Theta_2. \Theta_4
\end{aligned}$$

Table 11 Duality between projections of generalised types onto participants

which occur free in the network) and Δ is a *session typing*. *Session typings* associate session channels to generalised types:

$$\Delta ::= \emptyset \mid \Delta, s[p] : \chi$$

We apply to the session typings the same conventions used for environments. In particular a session typing Δ_1, Δ_2 is defined only if the domains of Δ_1 and Δ_2 are disjoint.

To ensure type safety it is essential that the communications are performed in a consistent way, i.e. that data are exchanged in the right order and with the right type. Consistency of session typings is defined using *projection of generalised types* and *duality*, given in Tables 10 and 11, respectively. The projection of a generalised type onto a participant q represents the communications offered to q .

The projection of generalised types uses the projection of message types and virtual monitors. We denote these projections by $\chi \upharpoonright q$, $m \upharpoonright q$ and $\mathcal{V} \upharpoonright q$, respectively. The conditions on the equalities of projections correspond to the similar conditions in Table 2. The projection of a generalised type of the shape $\langle m, \mathcal{M}, \mathcal{V} \rangle$ is the concatenation of the projections of m , $\mathcal{M}$ and $\mathcal{V}$. This is meaningful since m represents the message already sent, $\mathcal{M}$ guides the behaviour of the participant before its adaptation and $\mathcal{V}$ will guide its behaviour after the adaptation.

Projection of generalised types (ranged over by Θ) are defined by the following syntax

$$\Theta ::= !\ell(S) \mid !\lambda \mid \Theta. \Theta \mid ? \{ \Theta_i \}_{i \in I} \mid ! \{ \Theta_i \}_{i \in I} \mid \varepsilon$$

We assume $\varepsilon. \Theta = \Theta. \varepsilon = \Theta$, since ε represents no communication.

We write $\Theta \bowtie \Theta'$ when Θ and Θ' are dual according to the definition of Table 11. Note that duality is defined only on Θ which are projections of generalised types.

We can now define consistency as duality of projections.

Definition 10 A session typing Δ is *consistent for the session* s , notation $\text{con}(\Delta, s)$, if $s[p] : \chi \in \Delta$ and $s[q] : \chi' \in \Delta$ with $p \neq q$ imply $\chi \upharpoonright p \bowtie \chi' \upharpoonright p$. A session typing is *consistent* if it is consistent for all sessions which occur in it.

It is easy to check that projections of the same global type are always dual.

Proposition 1 Let G be a global type and $p \neq q$. Then $(G \upharpoonright p) \upharpoonright q \bowtie (G \upharpoonright q) \upharpoonright p$.

This proposition ensures that session typings obtained by projecting global types are consistent.

Table 12 gives the typing rules for systems. A session initiator is typed with the empty set of session names and with the empty session typing (rule NEW). To type a monitored process, we distinguish two cases. If the monitor is end, then the session typing is empty for any process P (rule endP). Otherwise, the channel owned by the process is associated to the monitor, provided that the type of the process (Table 4) is adequate for the monitor, according to Definition 8 (rule MP).

The next three rules type named queues. In these rules the turn-style is decorated by the name of the queue. An empty queue $\emptyset$ is typed with the empty session typing (QINT). Two queue types can be composed only if at most one of them contains a monitor, while the sequence of message types is a message type. Then we define the operator $\sharp$ by:

$$\langle m, \mathcal{V} \rangle \sharp \langle m', - \rangle = \langle m, - \rangle \sharp \langle m', \mathcal{V} \rangle = \langle m; m', \mathcal{V} \rangle$$

$\frac{}{\vdash_{\emptyset} \text{new}(G) \triangleright \emptyset}$ NEW	$\frac{}{\vdash_{\emptyset} \text{end}[P] \triangleright \emptyset}$ endP	$\frac{\vdash P \triangleright s[p] : T \quad \mathcal{M} \neq \text{end} \quad T \propto \mathcal{M}}{\vdash_{\emptyset} \mathcal{M}[P] \triangleright \{s[p] : \mathcal{M}\}}$ MP
$\frac{}{\vdash_{\{s\}} s : \emptyset \triangleright \emptyset}$ QINIT	$\frac{\vdash_{\{s\}} s : h \triangleright \Delta \quad \vdash v : S}{\vdash_{\{s\}} s : h \cdot (p, \Pi, \ell(v)) \triangleright \Delta \# \{s[p] : \langle \Pi! \ell(S), - \rangle\}}$ QSENDV	
	$\frac{\vdash_{\{s\}} s : h \triangleright \Delta}{\vdash_{\{s\}} s : h \cdot (p, \Pi, \lambda(G)) \triangleright \Delta \# (\{s[p] : \langle \Pi! \lambda, - \rangle\} \cup \{s[q] : \langle \varepsilon, G \upharpoonright q \rangle \mid q \in \Pi\})}$ QSENDG	
$\frac{\vdash_{\Sigma_1} N_1 \triangleright \Delta_1 \quad \vdash_{\Sigma_2} N_2 \triangleright \Delta_2 \quad \Sigma_1 \cap \Sigma_2 = \emptyset}{\vdash_{\Sigma_1 \cup \Sigma_2} N_1 \mid N_2 \triangleright \Delta_1 * \Delta_2}$ NPAR	$\frac{\vdash_{\Sigma} N \triangleright \Delta \quad \Delta \approx \Delta'}{\vdash_{\Sigma} N \triangleright \Delta'}$ EQUIV	
$\frac{\vdash_{\Sigma} N \triangleright \Delta \quad \text{con}(\Delta, s)}{\vdash_{\Sigma \setminus \{s\}} (vs)N \triangleright \Delta \setminus s}$ RES	$\frac{\vdash_{\Sigma} N \triangleright \Delta}{\vdash_{\Sigma} N \parallel \sigma \triangleright \Delta}$ SYSTEM	

Table 12 Typing rules for networks and systems

Rules QSENDV and QSENDG use the extension of $\#$ to session typings:

$$\begin{aligned} \Delta \# \Delta' = & \{s[p] : \chi \# \chi' \mid s[p] : \chi \in \Delta \ \& \ s[p] : \chi' \in \Delta'\} \cup \\ & \{s[p] : \chi \mid s[p] : \chi \in \Delta \cup \Delta' \ \& \\ & \quad s[p] \notin \text{dom}(\Delta) \cap \text{dom}(\Delta')\} \end{aligned}$$

Notice that in rules QSENDV and QSENDG the session typings only contain queue types. The queue type $\langle \Pi! \ell(S), - \rangle$ is pushed in the queue type of $s[p]$ for a value message $(p, \Pi, \ell(v))$, where S is the sort of v (rule QSENDV). The queue type $\langle \Pi! \lambda, - \rangle$ is pushed in the queue type of $s[p]$ for an adaptation message $(p, \Pi, \lambda(G))$, while the queue type of $s[q]$ has the projection of G on q as virtual monitor, for all $q \in \Pi$ (rule QSENDG).

For typing the parallel composition of networks, rule NPAR prescribes that each named queue does not occur twice (condition $\Sigma_1 \cap \Sigma_2 = \emptyset$) and composes session typings forming a generalised type out of a queue type and a monitor. We define the composition $*$ between queue types and monitors as:

$$\langle m, \mathcal{V} \rangle * \mathcal{M} = \mathcal{M} * \langle m, \mathcal{V} \rangle = \langle m, \mathcal{M}, \mathcal{V} \rangle$$

We extend $*$ to generalised types and to session typings as expected:

$$\begin{aligned} \Delta * \Delta' = & \{s[p] : \chi * \chi' \mid s[p] : \chi \in \Delta \ \& \ s[p] : \chi' \in \Delta'\} \cup \\ & \{s[p] : \chi \mid s[p] : \chi \in \Delta \cup \Delta' \ \& \\ & \quad s[p] \notin \text{dom}(\Delta) \cap \text{dom}(\Delta')\} \end{aligned}$$

For example, if $\vdash_{\Sigma} N \triangleright \Delta$, then $\vdash_{\Sigma} \text{end}[P] \mid N \triangleright \emptyset * \Delta$ (by rules endP and NPAR) and $\emptyset * \Delta = \Delta$; this fits with the structural equivalence $\text{end}[P] \mid N \equiv N$.

Notice that both $\#$ and $*$ are partial operators on session typings, since they can be undefined when applied to arbitrary generalised types.

In order to take into account the structural congruence between queues (see Table 7), we consider message types modulo the equivalence relation $\approx$ induced by the following rules (where Z stands for either $\ell(S)$ or λ):

$$\begin{aligned} m; \Pi!Z; \Pi'!Z'; m' \approx m; \Pi'!Z'; \Pi!Z; m' \\ \text{if } \Pi \cap \Pi' = \emptyset \\ m; \Pi!Z; m' \approx m; \Pi_1!Z; \Pi_2!Z; m' \\ \text{if } \Pi = \Pi_1 \cup \Pi_2, \Pi_1 \cap \Pi_2 = \emptyset \end{aligned}$$

This equivalence relation on message types extends to generalised types by:

$$\begin{aligned} m \approx m' \text{ implies} \\ \langle m, \mathcal{V} \rangle \approx \langle m', \mathcal{V} \rangle \text{ and } \langle m, \mathcal{M}, \mathcal{V} \rangle \approx \langle m', \mathcal{M}, \mathcal{V} \rangle \end{aligned}$$

We say that two session typings Δ and Δ' are equivalent (notation $\Delta \approx \Delta'$) if

$s[p] : \chi \in \Delta$ implies $s[p] : \chi' \in \Delta'$ with $\chi \approx \chi'$ and vice versa. Rule EQUIV allows to use this equivalence relation.

Rule RES requires the session typing to be consistent for the session s in order to type the restriction on s .

A system can be typed if the network can be typed, while the global state is arbitrary, see rule SYSTEM.

A crucial observation is that virtual monitors occur in generalised types only if queues contain adaptation flags. In other words using the condition of being λ -free (that is a premise of rules ADAOUTCONT and ADAOUTNEW) we get:

If $\vdash_s s : h \triangleright \Delta$ and h is λ -free, then no virtual monitor occurs in Δ .

It is standard to prove an inversion lemma for networks and systems by induction on derivations (Table 12).

Lemma 2 (*Inversion Lemma*)

1. If $\vdash_{\Sigma} \text{new}(G) \triangleright \Delta$, then $\Sigma = \Delta = \emptyset$.
2. If $\vdash_{\Sigma} \text{end}[P] \triangleright \Delta$, then $\Sigma = \Delta = \emptyset$.
3. If $\vdash_{\Sigma} \mathcal{M}[P] \triangleright \Delta$ and $\mathcal{M} \neq \text{end}$, then $\Sigma = \emptyset$ and $\Delta = \{s[p] : \mathcal{M}\}$ and $\vdash P \triangleright s[p] : T$ and $T \propto \mathcal{M}$.
4. If $\vdash_{\Sigma} s : \emptyset \triangleright \Delta$, then $\Sigma = \{s\}$ and $\Delta = \emptyset$.
5. If $\vdash_{\Sigma} s : h \cdot (p, \Pi, \ell(v)) \triangleright \Delta$, then $\Sigma = \{s\}$ and $\Delta \approx \Delta' \# \{s[p] : \langle \Pi! \ell(S), - \rangle\}$ and $\vdash_{\{s\}} s : h \triangleright \Delta'$ and $\vdash v : S$.
6. If $\vdash_{\Sigma} s : h \cdot (p, \Pi, \lambda(G)) \triangleright \Delta$, then $\Sigma = \{s\}$ and $\Delta \approx \Delta' \# (\{s[p] : \langle \Pi! \lambda, - \rangle\} \cup \{s[q] : \langle \varepsilon, G \upharpoonright q \rangle \mid q \in \Pi\})$ and $\vdash_{\{s\}} s : h \triangleright \Delta'$.
7. If $\vdash_{\Sigma} N_1 \mid N_2 \triangleright \Delta$, then $\Sigma = \Sigma_1 \cup \Sigma_2$ and $\Delta = \Delta_1 * \Delta_2$ and $\vdash_{\Sigma_1} N_1 \triangleright \Delta_1$ and $\vdash_{\Sigma_2} N_2 \triangleright \Delta_2$ and $\Sigma_1 \cap \Sigma_2 = \emptyset$.
8. If $\vdash_{\Sigma} (vs)N \triangleright \Delta$, then $\Sigma = \Sigma' \setminus \{s\}$ and $\Delta = \Delta' \setminus s$ and $\vdash_{\Sigma'} N \triangleright \Delta'$ and $\text{con}(\Delta, s)$.
9. If $\vdash_{\Sigma} N \parallel \sigma \triangleright \Delta$, then $\vdash_{\Sigma} N \triangleright \Delta$.

We also need a lemma stating how the typing depends on the first message on the queue. The proof follows immediately from the typing rules of queues.

- Lemma 3**
1. If $\vdash_{\Sigma} s : (p, \Pi, \ell(v)) \cdot h \triangleright \Delta$, then $\Sigma = \{s\}$ and $\Delta \approx \{s[p] : \langle \Pi! \ell(S), - \rangle\} \# \Delta'$ and $\vdash_{\{s\}} s : h \triangleright \Delta'$ and $\vdash v : S$.
 2. If $\vdash_{\Sigma} s : (p, \Pi, \lambda(G)) \cdot h \triangleright \Delta$, then $\Sigma = \{s\}$ and $\Delta \approx (\{s[p] : \langle \Pi! \lambda, - \rangle\} \cup \{s[q] : \langle \varepsilon, G \upharpoonright q \rangle \mid q \in \Pi\}) \# \Delta'$ and $\vdash_{\{s\}} s : h \triangleright \Delta'$.

Monitor LTS transactions reveal the monitor shapes, as detailed in the next lemma, which can be proved by straightforward case analysis.

- Lemma 4**
1. If $\mathcal{M} \xrightarrow{p\ell} \mathcal{M}'$, then $\mathcal{M} = p\{ \ell_i(S_i) .. \mathcal{M}_i \}_{i \in I}$ and $\ell = \ell_j$ and $\mathcal{M}' = \mathcal{M}_j$ for some $j \in I$.
 2. If $\mathcal{M} \xrightarrow{\Pi! \ell} \mathcal{M}'$, then $\mathcal{M} = \Pi! \{ \ell_i(S_i) .. \mathcal{M}_i \}_{i \in I}$ and $\ell = \ell_j$ and $\mathcal{M}' = \mathcal{M}_j$ for some $j \in I$.
 3. If $\mathcal{M} \xrightarrow{p\lambda} \mathcal{M}'$, then $\mathcal{M} = p\{ \lambda_i \}_{i \in I}$ and $\lambda = \lambda_j$ for some $j \in I$.
 4. If $\mathcal{M} \xrightarrow{\Pi! \lambda} \mathcal{M}'$, then $\mathcal{M} = \Pi! \{ \lambda_i \}_{i \in I}$ and $\lambda = \lambda_j$ for some $j \in I$.

The following lemma relates communications offered by processes (as LTS transactions) with their types.

- Lemma 5**
1. If $P \xrightarrow{s[p]?\ell(v)} P'$ and $\vdash P \triangleright s[p] : T$, then either $P = s[p]?\ell(x).P_0$ and $T = ?\ell(S).T'$ or $P = s[p]?\ell(x).P_0 + P''$ and $T = ?\ell(S).T' \wedge T''$, and in both cases $\vdash s[p]?\ell(x).P_0 \triangleright s[p] : ?\ell(S).T'$ and $P' = P_0\{v/x\}$.

2. If $P \xrightarrow{s[p]!\ell(v)} P'$ and $\vdash P \triangleright s[p] : T$, then either $T = !\ell(S).T'$ or $T = !\ell(S).T' \wedge T''$, and $\vdash P' \triangleright s[p] : T'$ and $\vdash v : S$.
3. If $P \xrightarrow{s[p]?\ell(\lambda, T') } P'$ and $\vdash P \triangleright s[p] : T$, then either $T = ?\lambda$ or $T = ?\lambda \wedge T''$, and $\vdash P' \triangleright s[p] : T'$.
4. If $P \xrightarrow{s[p]!\ell(\lambda(F), T') } P'$ and $\vdash P \triangleright s[p] : T$, then either $T = !\lambda$ or $T = !\lambda \wedge T''$, and $\vdash P' \triangleright s[p] : T'$.

Proof The proof is by structural induction on P . We show only Point (1), the proof for the other points being simpler.

If $P \xrightarrow{s[p]?\ell(v)} P'$, then either $P = s[p]?\ell(x).P_0$ or $P = P_1 + P_2$ and $P_i \xrightarrow{s[p]?\ell(v)} P'_i$ for $i = 1$ or $i = 2$. In the first case P is typed by rule RCV and $T = ?\ell(S).T'$. In the second case P is typed by rule CHOICE. Then $\vdash P_i \triangleright s[p] : T_i$ and $T = T_1 \wedge T_2$ for $i = 1, 2$. By induction, either $P_i = s[p]?\ell(x).P_0$ and $T_i = ?\ell(S).T'$ or $P_i = s[p]?\ell(x).P_0 + P'_i$ and $T_i = ?\ell(S).T' \wedge T''$ for $i = 1$ or $i = 2$. In both cases $\vdash s[p]?\ell(x).P_0 \triangleright s[p] : ?\ell(S).T'$ and $P' = P_0\{v/x\}$. $\square$

As usual, session types are not preserved under system reduction: they evolve according to the actions performed by the corresponding participants. This is formalised by the reduction rules given in Table 13, where message types are considered modulo the equivalence relation defined above. The rules in the first line allow us to create monitors and queue types. The rules in the second line get rid of types carrying no information. The subsequent four rules deal with outputs and inputs of labels with sorts and flags. In particular, the rule in the second to last line shows how a virtual monitor becomes the current monitor when a participant adapts itself.

Notice that not all the left-hand sides of the reduction rules for networks and systems are typed by consistent session typings. For example,

$$\vdash_{\{s\}} \mathcal{M}[s[1]?\ell(x).s[1]?\ell'(y).\mathbf{0}] \mid s : (2, 1, \ell(\text{true})) \triangleright \Delta$$

where $\mathcal{M} = 2?\ell(\text{Bool}).2?\ell'(\text{Int}).\text{end}$,

$\Delta = \{s[1] : \mathcal{M}, s[2] : \langle 1!\ell(\text{Bool}), - \rangle\}$. Observe that

$$\mathcal{M}[s[1]?\ell(x).s[1]?\ell'(y).\mathbf{0}] \mid s : (2, 1, \ell(\text{true}))$$

matches the left-hand side of the reduction rule IN and Δ is not consistent. The network obtained by putting this network in parallel with $1!\ell'(\text{Int}).\text{end}[s[2]!\ell'(7).\mathbf{0}]$ has a consistent session typing. It is then crucial to show that if the left-hand side of a reduction rule is typed by a session typing, which is consistent when composed with some session typing, then the same property holds for the right-hand side too. It is sufficient to consider the reduction rules which do not contain network and system reductions as premises, i.e. which are the leafs in the reduction trees. This is formalised in the following lemma, which is the key step for proving the Subject Reduction Theorem.

Lemma 6 (*Key Lemma*)

$$\begin{aligned}
& \emptyset \Longrightarrow \{s[p] : \mathcal{M}\} \quad \emptyset \Longrightarrow \{s[p] : \langle \varepsilon, \mathcal{V} \rangle\} \\
& \{s[p] : \langle m, \text{end}, - \rangle\} \Longrightarrow \{s[p] : \langle m, - \rangle\} \quad \{s[p] : \langle \varepsilon, - \rangle\} \Longrightarrow \emptyset \\
& \{s[p] : \langle m, \Pi! \{\ell_i(S_i) \dots \mathcal{M}_i\}_{i \in I}, \mathcal{V} \rangle\} \Longrightarrow \{s[p] : \langle m; \Pi! \ell_j(S_j), \mathcal{M}_j, \mathcal{V} \rangle\} \quad j \in I \\
& \{s[p] : \langle q! \ell_j(S_j); m, \mathcal{V} \rangle, s[q] : \langle m', p? \{\ell_i(S_i) \dots \mathcal{M}_i\}_{i \in I}, \mathcal{V}' \rangle\} \Longrightarrow \{s[p] : \langle m, \mathcal{V} \rangle, s[q] : \langle m', \mathcal{M}_j, \mathcal{V}' \rangle\} \quad j \in I \\
& \{s[p] : \langle m, \Pi! \{\lambda_i\}_{i \in I}, - \rangle\} \Longrightarrow \{s[p] : \langle m; \Pi! \lambda_j, \mathcal{M}, - \rangle\} \quad j \in I \\
& \{s[p] : \langle q! \lambda_j; m, - \rangle, s[q] : \langle m', p? \{\lambda_i\}_{i \in I}, \mathcal{V} \rangle\} \Longrightarrow \{s[p] : \langle m, - \rangle, s[q] : \langle m', \mathcal{V}, - \rangle\} \quad j \in I \\
& \Delta_1 \# \Delta \Longrightarrow \Delta_2 \# \Delta \quad \text{if } \Delta_1 \Longrightarrow \Delta_2 \quad \Delta_1 * \Delta \Longrightarrow \Delta_2 * \Delta \quad \text{if } \Delta_1 \Longrightarrow \Delta_2
\end{aligned}$$

Table 13 Reduction of session typings

1. Let $\vdash_{\Sigma} N \triangleright \Delta$, and $N \longrightarrow N'$ be obtained by any reduction rule different from EQUIV, and $\Delta * \Delta_0$ be consistent for some Δ_0 . Then there is Δ' such that $\vdash_{\Sigma} N' \triangleright \Delta'$ and $\Delta \Longrightarrow^* \Delta'$ and $\Delta' * \Delta_0$ is consistent.
2. Let $\vdash_{\Sigma} \mathcal{S} \triangleright \Delta$, and $\mathcal{S} \longrightarrow \mathcal{S}'$ be obtained by any reduction rule different from SN, CTX, and $\Delta * \Delta_0$ be consistent for some Δ_0 . Then there is Δ' such that $\vdash_{\Sigma} \mathcal{S}' \triangleright \Delta'$ and $\Delta \Longrightarrow^* \Delta'$ and $\Delta' * \Delta_0$ is consistent.

Proof (1). The proof is by cases on network reduction rules. The cases of rule INIT and TAU are trivial, since $\Delta = \Delta'$.

Rule IN. By Lemma 2(7),

$$\vdash_{\Sigma} \mathcal{M}[P] \mid s : (q, p, \ell(v)) \cdot h \triangleright \Delta$$

implies $\Sigma = \Sigma_1 \cup \Sigma_2$

$$\Delta = \Delta_1 * \Delta_2 \quad (1)$$

$$\vdash_{\Sigma_1} \mathcal{M}[P] \triangleright \Delta_1 \quad (2)$$

$$\vdash_{\Sigma_2} s : (q, p, \ell(v)) \cdot h \triangleright \Delta_2 \quad (3)$$

By Lemma 4(1) $\mathcal{M} \xrightarrow{p? \ell} \mathcal{M}'$ implies $\mathcal{M} = q? \{\ell_i(S_i) \dots \mathcal{M}_i\}_{i \in I}$ and $\ell = \ell_j$ and $\mathcal{M}' = \mathcal{M}_j$ for some $j \in I$. By Lemma 2(3), the judgment (2) gives $\Sigma_1 = \emptyset$ and

$$\Delta_1 = \{s[p] : \mathcal{M}\} \quad (4)$$

and $\vdash P \triangleright s[p] : T$ and $T \propto \mathcal{M}$. By Lemma 3(1), the judgment (3) gives $\Sigma_2 = \{s\}$ and

$$\Delta_2 \approx \{s[q] : \langle p! \ell(S), - \rangle\} \# \Delta_3 \quad (5)$$

and

$$\vdash_{\{s\}} s : h \triangleright \Delta_3 \quad (6)$$

and $\vdash v : S$.

By Lemma 5(1) $P \xrightarrow{s[p]? \ell(v)} P'$ and $\vdash P \triangleright s[p] : T$ imply either $P = s[p]? \ell(x). P_0$ and $T = ? \ell(S'). T'$ or $P = s[p]? \ell(x). P_0 + P_1$ and $T = ? \ell(S'). T' \wedge T''$. In both cases

$$\vdash s[p]? \ell(x). P_0 \triangleright s[p] : ? \ell(S'). T' \text{ and } P' = P_0 \{v/x\}.$$

The shapes of T , $\mathcal{M}$ and $T \propto \mathcal{M}$ imply $S' = S_j$ and $T' \propto \mathcal{M}'$.

The consistency of $\Delta * \Delta_0$ implies $S = S'$. The we can derive $\vdash P' \triangleright s[p] : T'$ and

$$\vdash_{\emptyset} \mathcal{M}'[P'] \triangleright \{s[p] : \mathcal{M}'\} \quad (7)$$

Applying NPAR to (6) and (7) we derive

$$\vdash_{\{s\}} \mathcal{M}'[P'] \mid s : h \triangleright \{s[p] : \mathcal{M}'\} * \Delta_3$$

Then $\Delta' = \{s[p] : \mathcal{M}'\} * \Delta_3$. From (1), (4), and (5) we get

$$\Delta \approx \{s[p] : \langle m, q? \{\ell_i(S_i) \dots \mathcal{M}_i\}_{i \in I}, \mathcal{V} \rangle, s[q] : \langle p! \ell(S); m', \mathcal{V}' \rangle\} \cup \Delta'_3$$

where $\ell = \ell_j$ and $S = S_j$ and $j \in I$

for some $m, \mathcal{V}, m', \mathcal{V}', \Delta'_3$ such that

$$\Delta_3 = \{s[p] : \langle m, \mathcal{V} \rangle, s[q] : \langle m', \mathcal{V}' \rangle\} \cup \Delta'_3.$$

Since

$$\{s[q] : \langle p! \ell(S); m', \mathcal{V}' \rangle, s[p] : \langle m, q? \{\ell_i(S_i) \dots \mathcal{M}_i\}_{i \in I}, \mathcal{V} \rangle\} \Longrightarrow \{s[q] : \langle m', \mathcal{V}' \rangle, s[p] : \langle m, \mathcal{M}', \mathcal{V} \rangle\}$$

then we get $\Delta \Longrightarrow^* \Delta'$. The only differences between Δ and Δ' are:

- the erasure of the message $p! \ell(S)$ in the type of $s[q]$;
- the replacement of the monitor $\mathcal{M}'$ to the monitor $q? \{\ell_i(S_i) \dots \mathcal{M}_i\}_{i \in I}$ in the type of $s[p]$.

It is then easy to check that the consistency of $\Delta * \Delta_0$ implies the consistency of $\Delta' * \Delta_0$.

Rule ADAINNEW. By Lemma 2(7),

$$\vdash_{\Sigma} \mathcal{M}[P] \mid s : (q, p, \lambda(G)) \cdot h \triangleright \Delta$$

implies $\Sigma = \Sigma_1 \cup \Sigma_2$

$$\Delta = \Delta_1 * \Delta_2 \quad (8)$$

$$\vdash_{\Sigma_1} \mathcal{M}[P] \triangleright \Delta_1 \quad (9)$$

$$\vdash_{\Sigma_2} s : (q, p, \lambda(G)) \cdot h \triangleright \Delta_2 \quad (10)$$

By Lemma 4(3) $\mathcal{M} \xrightarrow{q? \lambda} \mathcal{M}'$ implies

$$\mathcal{M} = q? \{\lambda_i\}_{i \in I} \quad (11)$$

and $\lambda = \lambda_j$ for some $j \in I$. By Lemma 2(3), the judgment (9) gives $\Sigma_1 = \emptyset$ and

$$\Delta_1 = \{s[p] : \mathcal{M}\} \quad (12)$$

and $\vdash P \triangleright s[p] : T$ and $T \propto \mathcal{M}$. By Lemma 3(2), the judgment (10) gives $\Sigma_2 = \{s\}$ and

$$\Delta_2 \approx \{s[q] : \langle p! \lambda, - \rangle, s[p] : \langle \varepsilon, \mathcal{M}' \rangle\} \# \Delta_3 \quad (13)$$

(taking into account that $G \vdash q = \mathcal{M}'$) and

$$\vdash_{\{s\}} s : h \triangleright \Delta_3 \quad (14)$$

We can obtain

$$\vdash_{\emptyset} \mathcal{M}'[Q\{s[p]/y\}] \triangleright \{s[p] : \mathcal{M}'\} \quad (15)$$

by using rule MP, since $(Q, T') \in \mathcal{P}$, $T' \propto \mathcal{M}'$ and $\mathcal{M}' \neq \text{end}$ (which is implied by the premise $T \not\propto \mathcal{M}'$ of the rule ADAINNEW). Hence, we apply rule NPAR to the judgments (14) and (15) and we derive:

$$\vdash_{\{s\}} \mathcal{M}'[Q\{s[p]/y\}] \mid s : h \triangleright \{s[p] : \mathcal{M}'\} * \Delta_3$$

Then $\Delta' = \{s[p] : \mathcal{M}'\} * \Delta_3$. Notice that (8), (12), (11) and (13) imply

$$\Delta \approx \{s[p] : \langle m, q? \{ \lambda_i \}_{i \in I}, \mathcal{M}' \rangle \} \cup \{s[q] : \langle p! \lambda; m', - \rangle \} \# \Delta'_3$$

for some m, m', Δ'_3 such that

$$\Delta_3 \approx \{s[p] : \langle m, - \rangle, s[q] : \langle m', - \rangle \} \cup \Delta'_3.$$

Note that $s[p]$ has $\mathcal{M}'$ as virtual monitor in Δ_2 and then no virtual monitor in Δ_3 . Instead $s[q]$ has no virtual monitor being the sender of the adaptation. Since

$$\begin{aligned} & \{s[q] : \langle p! \lambda; m', - \rangle, s[p] : \langle m, q? \{ \lambda_i \}_{i \in I}, \mathcal{M}' \rangle \} \implies \\ & \{s[q] : \langle m', - \rangle, s[p] : \langle m, \mathcal{M}', - \rangle \} \end{aligned}$$

we get $\Delta \implies^* \Delta'$. The only differences between Δ and Δ' are:

- the erasure of the message $p! \lambda$ in the type of $s[q]$;
- the erasure of the monitor $q? \{ \lambda_i \}_{i \in I}$ in the type of $s[p]$;
- the monitor $\mathcal{M}'$ is virtual in the type of $s[p]$ in Δ and it is active in the type of $s[p]$ in Δ' .

It is then easy to check that the consistency of $\Delta * \Delta_0$ implies the consistency of $\Delta' * \Delta_0$.

The proof for rules OUT and ADAINCONT are similar and simpler than those for rules IN and ADAINNEW, respectively.

(2). The proof is by cases on system reduction rules. The case of rule OP is trivial, since $\Delta = \Delta'$.

Rule ADAOUTCONT. Being h λ -free, Δ does not contain virtual monitors. By Lemma 2(9) and (7),

$$\vdash_{\Sigma} \mathcal{M}[P] \mid s : h \parallel \sigma \triangleright \Delta \text{ implies } \Sigma = \Sigma_1 \cup \Sigma_2,$$

$$\Delta = \Delta_1 * \Delta_2 \quad (16)$$

$$\vdash_{\Sigma_1} \mathcal{M}[P] \triangleright \Delta_1 \quad (17)$$

$$\vdash_{\Sigma_2} s : h \triangleright \Delta_2 \quad (18)$$

By Lemma 4(4), $\mathcal{M} \xrightarrow{\Pi! \lambda}$ implies

$$\mathcal{M} = \Pi! \{ \lambda_i \}_{i \in I} \quad (19)$$

and $\lambda = \lambda_j$ for some $j \in I$. By Lemma 2(3), the judgment (17) gives $\Sigma_1 = \emptyset$ and

$$\Delta_1 = \{s[p] : \mathcal{M}\} \quad (20)$$

$$\vdash P \triangleright s[p] : T \quad (21)$$

and $T \propto \mathcal{M}$.

By Lemma 5(4), the judgment (21) and $P \xrightarrow{s[p]!(\lambda(F), T')} P'$ imply $\vdash P' \triangleright s[p] : T'$. We consider only the case $\mathcal{M}_p \neq \text{end}$, the proof for the case $\mathcal{M}_p = \text{end}$ being similar but for the use of rule $\{s[p] : \langle m, \text{end}, - \rangle\} \implies \{s[p] : \langle m, - \rangle\}$. Since $T' \propto \mathcal{M}_p$, then we derive

$$\vdash_{\emptyset} \mathcal{M}_p[P'] \triangleright \{s[p] : \mathcal{M}_p\} \quad (22)$$

by rule MP. Similarly, for all $q \in \Pi' \setminus (\Pi \cup \{p\})$ we can derive

$$\vdash_{\emptyset} \mathcal{M}_q[P_q\{s[q]/y_q\}] \triangleright \{s[q] : \mathcal{M}_q\} \quad (23)$$

by rule MP, since $(P_q, T_q) \in \mathcal{P}$ and $T_q \propto \mathcal{M}_q$.

By Lemma 2(4), (5) and (6), the judgment (18) gives $\Sigma_2 = \{s\}$. Rule QSENDG applied to the judgment (18) derives

$$\vdash_{\{s\}} s : h \cdot (p, \Pi, \lambda(G)) \triangleright \Delta_2 \# \Delta_3 \quad (24)$$

where $\Delta_3 = \{s[p] : \langle \Pi! \lambda, - \rangle\} \cup \{s[q] : \langle \varepsilon, G \vdash q \rangle \mid q \in \Pi\}$. Notice that $\Delta_2 \# \Delta_3$ is defined, since Δ_2 does not contain virtual monitors. Applying rule NPAR to judgments (22), (23) and (24) we conclude

$$\vdash_{\{s\}} \mathcal{M}_p[P'] \mid N \mid s : h \cdot (p, \Pi, \lambda(G)) \triangleright \Delta'$$

where

$$N = \prod_{q \in \Pi' \setminus (\Pi \cup \{p\})} \mathcal{M}_q[P_q\{s[q]/y_q\}] \text{ and}$$

$$\Delta' = \{s[p] : \mathcal{M}_p\} * \{s[q] : \mathcal{M}_q \mid q \in \Pi' \setminus (\Pi \cup \{p\})\} * (\Delta_2 \# \Delta_3).$$

Notice that (16), (20) and (19) imply

$$\Delta = \{s[p] : \langle m, \Pi! \{ \lambda_i \}_{i \in I}, - \rangle\} * \Delta'_2$$

for some m such that $\Delta_2 \approx \{s[p] : \langle m, - \rangle\} \# \Delta'_2$. Since

$$\{s[p] : \langle m, \Pi! \{ \lambda_i \}_{i \in I}, - \rangle\} \implies \{s[p] : \langle m; \Pi! \lambda, \mathcal{M}_p, - \rangle\} \text{ for } \lambda = \lambda_j \text{ with } j \in I$$

$$\emptyset \implies \{s[q] : \langle \varepsilon, G \vdash q \rangle\} \text{ for } q \in \Pi$$

$$\emptyset \implies \{s[q] : \mathcal{M}_q\} \text{ for } q \in \Pi' \setminus \Pi \cup \{p\}$$

we get $\Delta \implies^* \Delta'$. The session typing Δ' contains only monitors which are projections of the global type G . Therefore the consistency of $\Delta * \Delta_0$ implies the consistency of $\Delta' * \Delta_0$.

The proof for rule ADAOUTNEW proceeds as in the previous case. $\square$

The next lemma shows that typings for systems are invariant under structural equivalence of networks, as expected.

Lemma 7 *If $\vdash_{\Sigma} N \parallel \sigma \triangleright \Delta$ and $N \equiv N'$, then $\vdash_{\Sigma} N' \parallel \sigma \triangleright \Delta$.*

Proof By Lemma 2(9) and rule SYSTEM it is enough to show that $\vdash_{\Sigma} N \triangleright \Delta$ and $N \equiv N'$ imply $\vdash_{\Sigma} N' \triangleright \Delta$. The proof is by induction on the definition of structural equivalence, observing that $\vdash_{\emptyset} \text{end}[P] \triangleright \emptyset$ and using typing rule EQUIV. $\square$

Theorem 1 (Subject Reduction) *If $\vdash_{\Sigma} \mathcal{S} \triangleright \Delta$ with Δ consistent and $\mathcal{S} \longrightarrow^* \mathcal{S}'$, then $\vdash_{\Sigma} \mathcal{S}' \triangleright \Delta'$ for some consistent Δ' such that $\Delta \implies^* \Delta'$.*

Proof It is enough to show the statement for the case $\mathcal{S} \equiv \mathcal{E}[N] \parallel \sigma$ and $\mathcal{S}' \equiv \mathcal{E}[N'] \parallel \sigma'$, where either $N \longrightarrow N'$ or $N \parallel \sigma \longrightarrow N' \parallel \sigma'$ by one of the rules considered in Lemma 6. By the structural equivalence on networks we can assume $\mathcal{E} = (\vec{v}\vec{s})([\] \mid N_0)$ without loss of generality. Lemma 7 and Lemma 2(9), (8) and (7) applied to $\vdash_{\Sigma} \mathcal{S} \triangleright \Delta$ give $\vdash_{\Sigma_1} N \triangleright \Delta_1$ and $\vdash_{\Sigma_0} N_0 \triangleright \Delta_0$, where $\Sigma = (\Sigma_0 \cup \Sigma_1) \setminus \vec{s}$ and $\Delta = (\Delta_0 * \Delta_1) \setminus \vec{s}$. The consistency of Δ implies the consistency of $\Delta_0 * \Delta_1$ by Lemma 2(8). In the case $N \longrightarrow N'$, by Lemma 6(1) there is Δ'_1 such that $\vdash_{\Sigma_1} N' \triangleright \Delta'_1$ and $\Delta_1 \Longrightarrow^* \Delta'_1$ and $\Delta_0 * \Delta'_1$ is consistent. In the case $N \parallel \sigma \longrightarrow N' \parallel \sigma'$, by Lemma 6(2) there is Δ'_1 such that $\vdash_{\Sigma_1} N' \parallel \sigma' \triangleright \Delta'_1$ (that is $\vdash_{\Sigma_1} N' \triangleright \Delta'_1$ by Lemma 2(9)) and $\Delta_1 \Longrightarrow^* \Delta'_1$ and $\Delta_0 * \Delta'_1$ is consistent. Therefore, we derive $\vdash_{\Sigma} \mathcal{S}' \triangleright \Delta'$, where $\Delta' = (\Delta_0 * \Delta'_1) \setminus \vec{s}$ by applying typing rules NPAR, RES, and SYSTEM. Observe that $\Delta \Longrightarrow^* \Delta'$ and Δ' is consistent. $\square$

We say that a system is *initial* when its network is a parallel composition of session initiators, which is always typeable. The type system can guarantee progress, proviso that the collection of processes and types contains at least one process for each monitor which is created at run time in the adaptations. This can also be statically checked when the domains of the adaptation functions which occur in processes are finite. We say that a collection $\mathcal{P}$ is *complete* if, for every global type G in the domain of an adaptation function which occurs in a process belonging to $\mathcal{P}$, there are processes in $\mathcal{P}$ whose types are adequate for the monitors obtained by projecting G onto its participants.

Theorem 2 (Progress) *If $\mathcal{P}$ is complete, $\mathcal{S}$ is an initial system and $\mathcal{S} \longrightarrow_{\mathcal{P}}^* \mathcal{S}'$, then $\mathcal{S}'$ has progress, i.e.*

1. *every input monitored process will eventually receive a message, and*
2. *every message in a queue will eventually be received by an input monitored process.*

Proof Coppo et al. [13] show that an initial system without adaptation flags has progress (by rewriting the result of [13] in our setting).

With respect to the framework of [13], each adaptation step, in our model, can be seen as the starting of a new interaction protocol where all participants can be implemented thanks to the completeness of $\mathcal{P}$. Therefore, the following key features ensure the progress property in our case:

1. all the interaction protocols prescribed by a global type are terminating, either by exchanging adaptation flags or reaching end;
2. a participant is a process, monitored by a projection of a global type, with at most one channel, so there is no communications among participants monitored by different global types, even when they are in the same session;

3. by Subject Reduction, systems are well typed and reduction preserves the consistency of session typings, thus all communications take place in the order prescribed by the global types.

Notice that Point (2) holds thanks to the condition of λ -freeness in rules ADAOUTCONT and ADAOUTNEW.

Concluding, the computation can be seen as a succession of independent terminating communication protocols, each of which has the progress property. So the whole computation has it. $\square$

7 Related Work

The literature includes several works aimed at studying adaptive systems in different application contexts and by different perspectives on the conceptual notion of adaptation. The paper [7] provides a valuable discussion on this issue and an interesting classification of various approaches. The state-of-the-art in service choreography adaptation is analysed in [31]. We focus here on the papers which are more related to the distinguishing features of our approach.

Adaptable processes In [5] Bravetti et al. present a calculus in which adaptable processes can be modified by “update patterns”. Run-time adaptation of structured communications is approached in [20] by combining the constructors for adaptable processes of [5] with the session type system of [22] for the Boxed Ambient calculus [8]. Session behaviours are never disrupted by adaptation actions, since processes engaged in active sessions cannot be updated. This calculus deals with adaptations of single processes, not with adaptations of the choreography of communicating processes, and only considers dyadic sessions and synchronous communications.

Adaptable choreographies The paper most similar to ours is [1], where global and session types are used to guarantee deadlock-freeness in a calculus of multiparty sessions with asynchronous communications. Only part of the running code is updated. Two different conditions are given for ensuring liveness. The first condition requires that all channel queues are empty before updating. The second condition requires a partial order between the session participants with a unique minimal element. The participants are then updated following this order. Our adaptation framework allows the progress property to be guaranteed without assuming such conditions.

The paper [14], building on [30, 16, 17] and [15], proposes a rule-based approach in which all interactions, under all possible changes produced by the adaptation rules, proceed as prescribed by an abstract model. In particular, the system is deadlock-free by construction. The adaptive system is composed by interacting participants deployed on different locations, each executing its own code. Adaptation

is performed by distributed adaptation servers, which are repositories of adaptation rules. Rules can be added or removed at any moment, while the system is running. Applicability depends on execution environment and properties of the code region to be replaced. If a rule is applied, it replaces part of the code of (some of) the participants with a newer version, able to better meet the requirements. Adaptations of different participants are coordinated ensuring coherent behaviour. Data and control flow statements are done in a Java-style syntax. Central to the technical development are the notions of adaptive interaction oriented choreography and adaptive process oriented choreography, which resemble our global types and monitors. Although there are many analogies between this and our paper, there are important differences. In [14] auxiliary communications are needed to ensure that all participants take the same branch in conditionals, and new participants cannot be added by an adaptation. Moreover in [14] adaptation involves only a part of the choreography and can be applied in any moment, while in our calculus the interaction protocols contain the adaptation points and the reconfiguration step applies to the whole system.

The language of service choreographies defined in [6] is extended with two operators in [4] to take into account adaptation. One operator allows for the specification of adaptable scopes that can be dynamically modified, while the second may dynamically update code in one of such scopes. A similar extension is given for the service contract language of [6], so that projection can relate the two adaptable languages. The main difference with our proposal is that the set of possible participants to the modified choreographies is fixed once for all, in the adaptable scopes.

Monitors In the literature there are many calculi in which the process behaviour is statically and/or dynamically controlled by means of monitors, for example [21, 26]. The works that most influenced the present paper are [11, 3]. The calculus in those papers is a multiparty session calculus with assertions, and therefore it is much more expressive than our calculus. In fact the monitors in [11, 3] prescribe not only the types of the exchanged data, but also that the values of these data satisfy some predicates. Another main difference is that those monitors contain information on the behaviours of all session participants, while our monitors represent the behaviour of single participants.

Intersection and union types In the present paper we type processes with intersections and unions taking inspiration from [33]. The type syntax in that paper is more liberal than ours, for example not requiring labels in an intersection and in a union be different, so more processes can be typed. Also in [33], the most interesting processes correspond to external choices between inputs and internal choices between outputs.

Subtyping for intersections and unions is naturally inspired by their set-theoretical interpretation. Considering the mapping between monitors and types of Definition 8, in this paper we give a subtyping which is the opposite of that considered in [33]. Both subtypings have been largely used [23, 24, 37, 34, 35, 18, 9, 29, 32]. The main reason of this difference is that in typing processes one can either *assume* or *derive* the types of channels. In the simple case of a process P with only one channel y the typing judgments have the shapes $y : T \vdash P$ and $\vdash P \triangleright \{y : T\}$, respectively. This is the reason why subtyping in [23] and in [32] is defined in opposite ways. Choices between fewer inputs are smaller in the subtyping of [23] and bigger in the subtyping of [32]. Choices between outputs behave dually.

8 Conclusion

We have presented a formal model of self-adaptation in multiparty sessions. The framework is based on self-adaptive monitors and global types. In the service-oriented context, global types can be exploited as choreographies of service interactions and monitors as local protocol specifications of services. From this perspective, the present paper provides a formal model for assessing the impact of highly evolving environments, which demand for dynamic self-reconfigurations of the whole system. Interactions among services may be added or removed as well as new services may be required. We use a type discipline to ensure that service collaborations will behave in a safe way after dynamic adaptations.

Differently from approaches focusing on adaptation as code modification in software systems, our approach is choreography centred (similar to [14] for this aspect). When dynamic conditions demand a change, the global choreography updates itself together with the new monitors which prescribe the new behaviours to the participants. A process fills (implements) a given monitor if its type is adequate for that monitor, otherwise a different implementation (process) need to be found. Notably, we proved that all monitored processes behave correctly and interact with each other in a safe way, once the adaptation has been performed.

As a main feature, we achieve a decentralised control of the adaptation and a notable flexibility in the dynamics of self-reconfiguration. According to its monitor, any participant can be in charge of checking global data and sending the adaptation request, instead of devoting a centralised mechanism to this task. Furthermore, the dynamic reconfiguration can add new participants, while some of the old participants are not longer involved. Finally, processes, that are simply implementation code, can follow different incompatible computational paths, thus each participant can be differently implemented in the various adaptation steps.

One apparent limitation of our calculus is that processes can only operate on a single channel. This limitation can

be addressed by extending the process language and its typing rules, without major consequences on the rest of the development. Instead, adding session delegation would require further investigation.

We plan to experiment with implementations of our approach, to evaluate its feasibility.

We are working toward a quantitative version of our model, where the global state also contains dynamically evolving semantic information about processes, such as reputation or performance rates. Using this information, adaptation functions will be able to choose a single process among all the processes matching a monitor, as one of the best implementations for that participant. In the present calculus, this issue results in an arbitrary choice, since processes can be taken solely on the basis of their compatibility with monitors from the point of view of safe adaptations. In a realistic application, instead, it would be interesting to involve other requests concerning quantitative aspects.

Acknowledgements The authors gratefully thank the anonymous referees for their accurate and enlightening remarks, that strongly improved both the presentation and the technical development.

This work was partially supported by EU Collaborative project ASCENS 257414, ICT COST Action IC1201 BETTY, MIUR PRIN Project CINA Prot. 2010LHT4KM and Torino University/Compagnia San Paolo Project SALT.

References

1. Anderson, G., Rathke, J.: Dynamic Software Update for Message Passing Programs. In: APLAS'12, *LNCS*, vol. 7705, pp. 207–222. Springer (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: CONCUR'08, *LNCS*, vol. 5201, pp. 418–433. Springer (2008)
3. Bocchi, L., Chen, T.C., Demangeon, R., Honda, K., Yoshida, N.: Monitoring Networks through Multiparty Session Types. In: FMOODS/FORTE'13, *LNCS*, vol. 7892, pp. 50–65. Springer (2013)
4. Bravetti, M., Carbone, M., Hildebrandt, T., Lanese, I., Mauro, J., Pérez, J.A., Zavattaro, G.: Towards Global and Local Types for Adaptation. In: SEFM'13, *LNCS*, vol. 8368, pp. 3–14. Springer (2014)
5. Bravetti, M., Di Giusto, C., Pérez, J.A., Zavattaro, G.: Adaptable Processes. *Logical Methods in Computer Science* **8**(4) (2012)
6. Bravetti, M., Zavattaro, G.: Towards a Unifying Theory for Choreography Conformance and Contract Compliance. In: SC'07, *LNCS*, vol. 4829, pp. 34–50. Springer (2007)
7. Bruni, R., Corradini, A., Gadducci, F., Lluch-Lafuente, A., Vandin, A.: A Conceptual Framework for Adaptation. In: FASE'12, *LNCS*, vol. 7212, pp. 240–254. Springer (2012)
8. Bugliesi, M., Castagna, G., Crafa, S.: Access Control for Mobile Agents: The Calculus of Boxed Ambients. *ACM Transactions on Programming Languages and Systems* **26**(1), 57–124 (2004)
9. Carbone, M., Honda, K., Yoshida, N.: Structured Communication-Centered Programming for Web Services. *ACM Transactions on Programming Languages and Systems* **34**(2), 8:1–8:78 (2012)
10. Castagna, G., Dezani-Ciancaglini, M., Padovani, L.: On Global Types and Multi-Party Sessions. *Logical Methods in Computer Science* **8**, 1–45 (2012)
11. Chen, T.C., Bocchi, L., Deniérou, P.M., Honda, K., Yoshida, N.: Asynchronous Distributed Monitoring for Multiparty Session Enforcement. In: TGC'11, *LNCS*, vol. 7173, pp. 25–45. Springer (2012)
12. Coppo, M., Dezani-Ciancaglini, M., Venneri, B.: Self-Adaptive Monitors for Multiparty Sessions. In: PDP'14, pp. 688–696. IEEE (2014)
13. Coppo, M., Dezani-Ciancaglini, M., Yoshida, N., Padovani, L.: Global Progress for Dynamically Interleaved Multiparty Sessions. *Mathematical Structures in Computer Science* (2013). To appear
14. Dalla Preda, M., Gabbriellini, M., Giallorenzo, S., Lanese, I., Mauro, J.: Deadlock Freedom by Construction for Distributed Adaptive Applications. *CoRR* (2014). URL <http://arxiv.org/abs/1407.0970>
15. Dalla Preda, M., Giallorenzo, S., Lanese, I., Mauro, J., Gabbriellini, M.: AIOCJ: A Choreographic Framework for Safe Adaptive Distributed Applications. In: SLE'14, *LNCS*, vol. 8706, pp. 161–170. Springer (2014)
16. Dalla Preda, M., Lanese, I., Mauro, J., Gabbriellini, M.: Adaptive Choreographies (2013). URL <http://www.cs.unibo.it/~lanese/publications/adaptchor.pdf.gz>
17. Dalla Preda, M., Lanese, I., Mauro, J., Gabbriellini, M., Giallorenzo, S.: Safe Run-time Adaptation of Distributed Systems (2013). URL <http://www.cs.unibo.it/~lanese/publications/fulltext/safeadapt.pdf.gz>
18. Demangeon, R., Honda, K.: Full Abstraction in a Subtyped pi-Calculus with Linear Types. In: CONCUR'11, *LNCS*, vol. 6901, pp. 280–296. Springer (2011)
19. Deniérou, P.M., Yoshida, N.: Dynamic Multirole Session Types. In: POPL'11, pp. 435–446. ACM Press (2011)
20. Di Giusto, C., Pérez, J.A.: Disciplined Structured Communications with Consistent Runtime Adaptation. In: SAC'13, pp. 1913–1918. ACM Press (2013)
21. Ferrari, G.L., Moggi, E., Pugliese, R.: Guardians for Ambient-based Monitoring. *ENTCS* **66**(3), 52–75 (2002)
22. Garralda, P., Compagnoni, A.B., Dezani-Ciancaglini, M.: BASS: Boxed Ambients with Safe Sessions. In: PPDP'06, pp. 61–72. ACM Press (2006)
23. Gay, S., Hole, M.: Subtyping for Session Types in the Pi Calculus. *Acta Informatica* **42**(2/3), 191–225 (2005)
24. Gay, S.J.: Bounded Polymorphism in Session Types. *Mathematical Structures in Computer Science* **18**(5), 895–930 (2008)
25. Ghezzi, C., Pradella, M., Salvaneschi, G.: An Evaluation of the Adaptation Capabilities in Programming Languages. In: SEAMS'11, pp. 50–59. ACM Press (2011)
26. Gorla, D., Hennessy, M., Sassone, V.: Security Policies as Membranes in Systems for Global Computing. *ENTCS* **138**(1), 23–42 (2005)
27. Honda, K., Vasconcelos, V.T., Kubo, M.: Language Primitives and Type Disciplines for Structured Communication-based Programming. In: ESOP'98, *LNCS*, vol. 1381, pp. 22–138. Springer (1998)
28. Honda, K., Yoshida, N., Carbone, M.: Multiparty Asynchronous Session Types. In: POPL'08, pp. 273–284. ACM Press (2008)
29. Kouzapas, D., Yoshida, N., Honda, K.: On Asynchronous Session Semantics. In: FMOODS/FORTE'11, *LNCS*, vol. 6722, pp. 228–243. Springer (2011)
30. Lanese, I., Bucciarone, A., Montesi, F.: A Framework for Rule-Based Dynamic Adaptation. In: TGC'10, *LNCS*, vol. 6084, pp. 284–300. Springer (2010)
31. Leite, L.A.F., Oliva, G.A., Nogueira, G.M., Gerosa, M.A., Kon, F., Milojicic, D.S.: A Systematic Literature Review of Service Choreography Adaptation. *Service Oriented Computing and Applications* **7**(3), 199–216 (2013)

32. Mostrous, D., Yoshida, N., Honda, K.: Global Principal Typing in Partially Commutative Asynchronous Sessions. In: ESOP'09, *LNCS*, vol. 5502, pp. 316–332. Springer (2009)
33. Padovani, L.: Session Types = Intersection Types + Union Types. In: ITRS'10, *EPTCS*, vol. 45, pp. 71–89 (2010)
34. Padovani, L.: Fair Subtyping for Multi-party Session Types. In: COORDINATION'11, *LNCS*, vol. 6721, pp. 127–141. Springer (2011)
35. Padovani, L.: Fair Subtyping for Open Session Types. In: ICALP'13, *LNCS*, vol. 7966, pp. 373–384. Springer (2013)
36. Psailer, H., Juszczak, L., Skopik, F., Schall, D., Dustdar, S.: Runtime Behavior Monitoring and Self-Adaptation in Service-Oriented Systems. In: SASO'10, pp. 164–173. IEEE Computer Society (2010)
37. Vasconcelos, V.T.: Fundamentals of Session Types. In: SFM'09, *LNCS*, vol. 5569, pp. 158–186. Springer (2009)