



HAL
open science

Event-based run-time adaptation in communication-centric systems

Cinzia Di Giusto, Jorge A. Pérez

► **To cite this version:**

Cinzia Di Giusto, Jorge A. Pérez. Event-based run-time adaptation in communication-centric systems. Formal Aspects of Computing, 2016, 10.1007/s00165-016-0377-z . hal-01320789

HAL Id: hal-01320789

<https://hal.science/hal-01320789>

Submitted on 24 May 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Event-Based Run-time Adaptation in Communication-Centric Systems

Cinzia Di Giusto¹ and Jorge A. Pérez²

¹ Université de Nice Sophia-Antipolis, CNRS, I3S, UMR 7271, Sophia Antipolis, France

² University of Groningen, Groningen, The Netherlands

Abstract. *Communication-centric systems* are software systems built as assemblies of distributed artifacts that interact following predefined communication protocols. *Session-based concurrency* is a type-based approach to ensure the conformance of communication-centric systems to such protocols. This paper presents a model of session-based concurrency with mechanisms for *run-time adaptation*. Our model allows us to specify communication-centric systems whose session behavior can be dynamically updated at run-time. We improve on previous work by proposing an *event-based* approach: adaptation requests, issued by the system itself or by its context, are assimilated to events which may trigger adaptation routines. These routines exploit type-directed checks to enable the reconfiguration of processes with active protocols. We equip our model with a type system that ensures *communication safety* and *consistency* properties: while safety guarantees absence of run-time communication errors, consistency ensures that update actions do not disrupt already established session protocols. We provide soundness results for binary and multiparty protocols.

Keywords: concurrency, behavioral types, session types, run-time adaptation, process calculi.

1. Introduction

Context and Motivation. Modern software systems are increasingly built as assemblies of distributed artifacts that interact by following predefined communication protocols. Correctness in these *communication-centric* systems, therefore, depends on ensuring that such dialogues conform to their protocols. *Session-based concurrency* is a type-based approach to ensure conformance to prescribed protocols: dialogues are organized into units called *sessions*; interaction patterns are abstracted as *session types* [HVK98], against which specifications may be checked. As these specifications are often given in the π -calculus [MPW92, HVK98], protocol conformance can be analyzed at the level of concurrent processes.

As communication-centric systems typically operate on open and dynamic infrastructures, *run-time adaptation* appears as an indispensable feature to ensure their flexible and uninterrupted operation. Here we understand run-time adaptation as the dynamic reconfiguration of a system's behavior in response to an exceptional event. Examples of exceptional events are, e.g., a varying requirement or a local failure. These events are not necessarily catastrophic but are hard to predict. In this context, conformance to structured

protocols and run-time adaptation appear as intertwined concerns: although the specification of run-time adaptation may not be strictly tied to protocol reconfiguration, steps of run-time adaptation have a direct influence in a system’s communicating behavior and overall correctness.

We are interested in enhancing models of session-based concurrency with forms of run-time adaptation. This entails extending models of session-based concurrency with constructs for representing adaptation, but also developing appropriate analysis and verification techniques for such extended models.

As a first answer to this challenge, in previous work [DP13a, DP15] we extended a typed process framework for binary sessions with constructs from the model of *adaptable processes* [BDPZ12]. In adaptable processes there are two key constructs, namely *located processes* and *update processes*:

- Given a process P , the located process $\text{loc}[P]$ denotes the fact that P resides in *location* loc . Locations are explicit delimiters for process behavior: they denote possibly nested computation sites. Since locations are *transparent*, process P in $\text{loc}[P]$ may interact with processes outside loc .
- An update process $\text{loc}\{U\}$ specifies that the behavior currently enclosed by location loc should be modified by the *adaptation routine* U —intuitively, a function from processes to processes located at loc .

At the operational level, located and update processes are complementary: processes $\text{loc}[P]$ and $\text{loc}\{U\}$ are meant to synchronize in order to realize an adaptation step. As we will see, U is applied to P by “moving” to its enclosing context. The exact operational semantics that justifies this adaptation step will be illustrated shortly. Observe that since processes within locations (such as P in $\text{loc}[P]$) may implement one or several session protocols, the interaction with update processes allow us to specify adaptation routines that concern several running sessions.

We would like to guarantee that adaptation steps resulting from the synchronization of located and update processes preserve overall system correctness. In a session-typed setting this entails ensuring that such synchronizations do not jeopardize the session protocols enclosed by locations. As a simple example of a dangerous adaptation routine, consider the update process $\text{loc}\{\mathbf{0}\}$, which says that the new behavior at loc should be the inactive process $\mathbf{0}$, without any considerations on the current behavior at loc . Indeed, a synchronization between $\text{loc}[P]$ and $\text{loc}\{\mathbf{0}\}$ would result into process $\text{loc}[\mathbf{0}]$, therefore discarding running session protocols in P . Processes depending on the sessions implemented by P would lack complementary partners, thus compromising any conceivable form of protocol conformance.

To rule out adaptation steps that jeopardize established protocols, communication and adaptation actions should be harmonized. In [DP13a, DP15] we defined the notion of *consistency*, according to which a located process can be updated only if it does not contain active (established) session protocols. Adaptation actions are enabled only for locations that do not enclose running sessions; an associated type system ensures protocol conformance but also that session protocols are never disrupted as a result of an adaptation step. This is a simple solution; since adaptation is enabled only when sessions in a location are in a quiescent state, it is an alternative that privileges communication over adaptation. In applications, however, it may be necessary to give communication and adaptation a similar status, enabling adaptation steps also on locations that contain already established sessions. The *main contribution* of this paper is a session-typed framework that enables this kind of adaptation steps, thus overcoming the limitations of the framework introduced in [DP13a, DP15]. Next we elaborate on our approach.

Our Approach. To enable adaptation steps on locations that contain already established sessions, we propose update processes $\text{loc}\{U\}$ in which the adaptation routine U (i) dynamically checks the current state of the protocols running in loc and (ii) determines a new process behavior for loc based on that state.

In their simplest form, our novel update processes concern located processes which implement only one session channel. They are of the form:

$$\text{loc}\{\text{case } x \text{ of } \{(\beta_1 : P_1), (\beta_2 : P_2), \dots, (\beta_m : P_m)\}\} \quad (m \geq 1) \quad (1)$$

where x is a channel variable, β_1, \dots, β_m are session types (cf. § 2.1), and P_1, \dots, P_m are processes in which x occurs free. Intuitively, each pair $(\beta_i : P_i)$ represents an *alternative* for updating location loc : session type β_i denotes the protocol state under which process P_i will be used as the new implementation at loc .

Suppose a process $\text{loc}[Q]$ in which Q implements a session of type α along session channel κ . The interaction between $\text{loc}[Q]$ and the update process in (1) intuitively proceeds as follows: If there is a β_j (with $j \in \{1, \dots, m\}$) that “matches” with α , then the located process is updated to $\text{loc}[P_j[\kappa/x]]$. That is,

the (current) protocol state α guides the selection of the j -th adaptation alternative. Otherwise, if no β_j validates a match, then the located process is kept unchanged as $\text{loc}[Q]$ and the update process is consumed.

In their general form, our novel update processes may act on located processes which enclose more than one running session, thus adding considerable expressiveness. Such an update process is denoted:

$$\text{loc}\{\text{case } x_1, \dots, x_n \text{ of } \{(\beta_i^1; \dots; \beta_i^n) : P_i\}_{i \in I}\} \quad (2)$$

where n is the number of running sessions at location loc . Intuitively, $(\beta_j^1; \dots; \beta_j^n) : P_j$ denotes the j -th alternative for updating loc : session types $\beta_j^1, \dots, \beta_j^n$ denote the protocol states under which process P_j will be used as the new implementation at loc . That is, the update process in (2) is meant to interact with a located process $\text{loc}[Q]$ in which Q implements session types $\alpha_1, \dots, \alpha_n$; to enable an adaptation step, each α_i must “match” with β_j^i for all $i \in \{1, \dots, n\}$.

The semantics for our *typeful* update processes exploits *session monitors*—processes that maintain the (current) protocol state at a given channel—and constructs for *dynamic type inspection*, as put forward in [Kou12, KYHH16].

As just described, our update processes use types to select an alternative for adaptation—that is, types say *how* to adapt. However, update processes by themselves they do not specify *when* adaptation should be triggered. A simple solution is to “embed” update processes within session communication prefixes. Unfortunately, such a specification style would only allow to handle exceptional conditions which can be fully characterized in advance. Other important exceptional conditions, in particular contextual and/or unsolicited run-time conditions, cannot always be anticipated/predicated nor have a clear causal dependency with steps related to structured communications.

To offer a uniform solution to the issue of when to trigger update processes, we propose an *event-based* approach. Our process syntax includes *adaptation signals* that are similar to communication prefixes and that can be used to issue an adaptation request for a given location. Each location is endowed with a *queue* that stores its pending adaptation signals/requests. To detect a request r in the queue of loc and react accordingly, we use the *arrival predicate* $\text{arrive}(\text{loc}, r)$, proposed in [Kou12]. The use of adaptation requests and the arrival predicate allows us to maintain a separation between structured communications (disciplined by session types) and the events that ultimately trigger update processes.

As an example, let upd_E denote an adaptation request. In our model we may specify process

$$\text{if } \text{arrive}(\text{loc}, \text{upd}_E) \text{ then } \text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_i^1; \dots; \beta_i^m) : P_i\}_{i \in I}\} \text{ else } Q$$

which expresses an adaptation policy that checks for the arrival of a request upd_E for location loc . If such a request is detected in the queue of loc then the update process in (2) is triggered; otherwise, process Q is executed. Let us write $\mu\mathcal{X}.P$ to denote a recursive process. The previous update process can be easily modified to express a *persistent* adaptation policy that uses recursion to continuously check for request upd_E :

$$\mu\mathcal{X}.\text{if } \text{arrive}(\text{loc}, \text{upd}_E) \text{ then } \text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_i^1; \dots; \beta_i^m) : P_i\}_{i \in I}\} \text{ else } \mathcal{X}$$

Contributions. We present a process model for session-based concurrency and run-time adaptation based on five constructs: located processes, typeful update processes, session monitors, adaptation signals, and location queues. While located processes are inherited from our previous work [DP15], the use of the other constructs for run-time adaptation is new to this presentation. We equip this model with a session type system that ensures the following key properties:

- **Safety:** Well-typed programs do not exhibit communication errors (e.g., mismatched messages).
- **Consistency:** Well-typed programs do not allow adaptation actions that disrupt already established session protocols.

Safety is the usual guarantee expected from any session type discipline. In contrast, consistency is a guarantee peculiar to our model: it relates the behavior of the adaptation mechanisms with the preservation of prescribed typed interfaces. An example of inconsistent behavior is the interaction between $\text{loc}[P]$ and $\text{loc}\{\mathbf{0}\}$ which, as motivated above, discards the running sessions in P . We prove that well-typed programs are safe and consistent: this ensures that specified session protocols are respected, while forbidding incautious adaptation steps that could disrupt the session behavior of interacting partners.

In the *first part* of the paper we consider a typed process model for *binary sessions*, as our goal is to develop our approach to event-based adaptation in a representative and well-established setting. Nevertheless,

the key ingredients of our approach are sufficiently simple to be expressed in more general scenarios. Hence, in the *second part* of the paper, as a proof of concept, we describe a generalization of our approach to the case of *multiparty, asynchronous* communication. This generalization builds upon the framework developed in [CDV15], which we briefly describe next.

The framework in [CDV15] uses global types (or *choreographies*) and *monitored processes* as main ingredients. A global type abstracts a sequence of communications between two or more participants. By projecting a global type onto each participant, one obtains a series of *local types*. In [CDV15], such local types are used as *monitors* for the processes that implement each partner: a monitor enables the visible actions of the process. The association of a monitor and a process is called a *monitored process*. The whole choreography is implemented as a collection of monitored processes, or *network*. By considering a network together with a *global state* (data/values used to plan adaptation steps), one then obtains a *system*.

The approach in [CDV15] defines a separation between the global type, the monitors (local types), and the processes. Adaptation is specified at the level of global types, which include an explicit signal (called *flag*) that defines a form of *anticipated adaptation* between all participants of the choreography. As a result of this coordinated signal, an adaptation function determines a new choreography for the system, relying on the global state. Such a choreography determines new implementations for all partners, using a centralized repository of typed processes called *collection*.

We extend the framework of [CDV15] with typeful update processes and event-based constructs. As a result, we obtain an alternative framework in which adaptation is not specified at level of global types (unlike [CDV15]); instead, we have events that handle *internal* and *external* adaptation requests:

- Internal adaptation requests are initiated by a participant that announces all other participants that the current global protocol should be abandoned, and that a new global protocol should be adopted instead.
- External adaptation requests are meant to dynamically update or upgrade the local implementation of a single participant.

To express internal adaptation requests, processes may explicitly output a message whose communication object denotes a choreography. This message can be accessed by all protocol participants (using the shared queue) so to determine their updated behavior. Networks in our model are distributed collections of locations, each containing a monitored process, a local collection of processes, and a queue. This queue is useful to model external adaptation requests which, as described above, are targeted to a particular participant. As a result, we obtain a setting in which internal adaptation is no longer anticipated at the level of types (therefore improving [CDV15]) and concerns all participants of the choreography, while external adaptation requests focus on a single participant.

Organization. The rest of the paper is organized as follows:

- §2 illustrates session-based concurrency, our approach and contributions by means of examples.
- §3 presents our process model of binary sessions with constructs for run-time adaptation. Then, §4 presents our session type system, which ensures safety and consistency for processes with adaptation mechanisms (Theorem 4.11).
- In §5 we generalize our approach to multiparty communications, defined as a variant of the framework given in [CDV15]. We show communication safety and consistency of adaptation steps (Theorem 5.19).
- In §6 we discuss related works and in §7 we collect some concluding remarks.

The appendix collects omitted definitions and proofs. This work is an extended version of the paper [DP16]. Here we offer additional explanations and proofs of technical results. Moreover, the approach to adaptation for multiparty sessions (§5), not presented in [DP16], is a new contribution.

2. Our Approach, By Example

In this section we illustrate our approach and contributions by means of an example, the *buyer-seller protocol*.

2.1. Binary Sessions

Background. In §3 and §4, we consider binary session types, as defined in [HVK98]:

$\alpha, \beta ::=$	$?(T).\beta$	input value of type T , continue as β
	$!(T).\beta$	output value of type T , continue as β
	$\&\{n_1:\alpha_1, \dots, n_m:\alpha_m\}$	branching (external choice)
	$\oplus\{n_1:\alpha_1, \dots, n_m:\alpha_m\}$	selection (internal choice)
	$\mu t.\alpha$	recursive session
	t	type variable
	end	terminated session

where T stands for basic types (e.g., booleans, integers) and session types α . Also, n_1, \dots, n_m denote *labels*.

To illustrate session types, consider two participants, a buyer and a seller, which interact as follows. First, buyer sends to seller the name of an item and seller replies back with its price. Then, depending on the amount, buyer may choose to add the item to her shopping cart or to close the transaction. In the latter case the protocol ends. In the former case, buyer must further choose a paying method. From buyer's perspective, this protocol may be described by the session type

$$\alpha_{\text{buy}} = \text{!item}.\text{?amnt}.\alpha_{\text{pay}}$$

where *item* and *amnt* are base types and

$$\alpha_{\text{pay}} = \oplus\{\text{addItem} : \oplus\{\text{ccard} : \alpha_{\text{cc}}, \text{payp} : \alpha_{\text{pp}}\}, \text{cancel} : \mathbf{end}\}.$$

Thus, session type α_{buy} says that protocol α_{pay} may only be enabled after sending a value of type *item* and receiving a value of type *amnt*. Also, *addItem*, *ccard*, *payp*, and *cancel* denote labels in the internal choice. Types α_{cc} and α_{pp} denote the behavior of each payment method. Following the protocol abstracted by α_{buy} , code for buyer may be specified as a π -calculus process. Processes P and R below give two different specifications for buyer:

$$\begin{aligned} P &= \bar{x}(\text{book}).x(a).\text{if } a < 50 \text{ then } x \triangleleft \text{addItem}; x \triangleleft \text{ccard}; P_1 \text{ else } x \triangleleft \text{cancel}; \mathbf{0} \\ R &= \bar{x}(\text{game}).x(b).\text{if } b < 80 \text{ then } x \triangleleft \text{addItem}; x \triangleleft \text{payp}; R_1 \text{ else } x \triangleleft \text{cancel}; \mathbf{0} \end{aligned}$$

Thus, although both P and R implement α_{buy} , their behavior is rather different, for they purchase different items using different payment methods (which are abstracted by processes P_1 and R_1).

Let us now analyze the situation for the seller. To ensure protocol compatibility and absence of communication errors, the session type for seller, denoted β_{sel} , should be *dual* to α_{buy} . Intuitively, duality decrees that every action from buyer must be matched by a complementary action from seller, e.g., every output in α_{buy} is matched by an input in β_{sel} . Formally, this duality is denoted $\alpha_{\text{buy}} \perp_{\mathcal{C}} \beta_{\text{sel}}$ (cf. Definition A.4). Thus, in our example, we let $\beta_{\text{sel}} = \text{?item}.\text{!amnt}.\beta_{\text{pay}}$, where

$$\beta_{\text{pay}} = \&\{\text{addItem} : \&\{\text{ccard} : \beta_{\text{cc}}, \text{payp} : \beta_{\text{pp}}\}, \text{cancel} : \mathbf{end}\}$$

and β_{cc} and β_{pp} denote the duals of α_{cc} and α_{pp} , respectively. A process implementation Q for β_{sel} is the following:

$$Q = y(i).\bar{y}(\text{price}(i)).y \triangleright \{\text{addItem} : y \triangleright \{\text{ccard} : Q_1 \parallel \text{ppal} : Q_2\} \parallel \text{cancel} : \mathbf{0}\}$$

where *price* stands for an auxiliary function.

The interaction of P and Q is defined using *session initialization* constructs: process $\bar{u}(x:\alpha).P$ denotes the *request* of a session of type α ; dually, $u(x:\alpha).P$ denotes the *acceptance* of a session of type α . In both cases, u denotes a (*shared*) *name* used for synchronization. In our example, we may have

$$\text{Sys} = \bar{u}(x : \alpha_{\text{buy}}).P \mid u(y : \beta_{\text{sel}}).Q \longrightarrow (\nu\kappa)(P[\kappa^+/x] \mid Q[\kappa^-/y]) = S_0 \quad (3)$$

Thus, upon synchronization on u , a new session κ is established. Intuitively, in process S_0 session κ is “split” into two *session channels* (or *endpoints*) κ^+ and κ^- : we write $+$ and $-$ to denote their opposing *polarities* [GH05], which make their complementarity manifest. The restriction $(\nu\kappa)$ covers both channels, thus ensuring an interference-free medium for executing the session protocols described by α and β .

Session types may exploit a *subtyping* relation [GH05] to express useful relationships between protocol

specifications. For instance, the type below defines a new payment method for seller:

$$\beta_{\text{gift}} = \&\{\text{addItem} : \&\{\text{giftc} : \beta_{\text{gc}}, \text{ccard} : \beta_{\text{cc}}, \text{payp} : \beta_{\text{pp}}\}, \text{cancel} : \mathbf{end}\} \quad (4)$$

Intuitively, β_{gift} extends β_{pay} with an alternative on label *giftc*. It is safe to use an implementation for β_{gift} wherever an implementation for β_{pay} is required. The *safe substitution* principle that connects β_{gift} and β_{pay} is formalized by a subtyping, denoted \leq_c (cf. Definition A.4). This way, e.g., we have $\beta_{\text{pay}} \leq_c \beta_{\text{gift}}$.

Our Process Model and Approach. We are interested in expressing and reasoning about the run-time modification of session-typed processes such as P and Q above. Such modifications may be desirable as a reaction to exceptional run-time conditions (say, an error) or to implement unanticipated requirements.

Continuing our example, consider again process Sys (cf. (3) above). In our process model, detailed in §3, we have the following reduction step:

$$Sys \longrightarrow (\nu\kappa)(P[\kappa^+/x] \mid \kappa^+[\alpha_{\text{buy}}] \mid Q[\kappa^-/y] \mid \kappa^-[\beta_{\text{buy}}]) = S$$

Session establishment creates *session monitors* $\kappa^+[\alpha_{\text{buy}}]$ and $\kappa^-[\beta_{\text{buy}}]$, processes that maintain the protocol state for the session at κ^+ and κ^- , respectively. These monitors are essential to implement run-time adaptation policies that handle processes with running sessions. Reduction may proceed as follows:

$$S \longrightarrow (\nu\kappa)(\kappa^+(a).\text{if } a < 50 \text{ then } \kappa^+ \triangleleft \text{addItem}; \kappa^+ \triangleleft \text{ccard}; P_1 \text{ else } \kappa^+ \triangleleft \text{cancel}; \mathbf{0} \mid \kappa^+[\text{?amnt. } \alpha_{\text{pay}}] \mid \kappa^-[\text{price(book)}].\kappa^- \triangleright \{\text{addItem} : \kappa^- \triangleright \{\text{ccard} : Q_1 \parallel \text{ppal} : Q_2\} \parallel \text{cancel} : \mathbf{0}\} \mid \kappa^-[\text{!amnt. } \beta_{\text{pay}}]) = S'$$

Observe how the session monitors evolve together with the process, capturing the evolution of the protocol state. In this example, they are useful to record the fact that the buyer part of S' realizes the type $\text{?amnt. } \alpha_{\text{pay}}$, whereas the seller part realizes the type $\text{!amnt. } \beta_{\text{pay}}$. Suppose that we wish to modify at run-time the part of S' realizing the buyer behavior. To preserve protocol correctness, a candidate new implementation must conform, up to \leq_c , to the type $\text{?amnt. } \alpha_{\text{pay}}$; a process realizing any other type will fail to safely interact with the part of S' implementing the seller.

We now illustrate how located and update processes, distinctive of our model, could be used in the buyer-seller scenario. As a simple example, process W below

$$W = \mathbf{sys}[\mathbf{buyer}[\bar{u}(x:\alpha_{\text{buy}}).P] \mid \mathbf{seller}[u(y:\beta_{\text{sel}}).Q]] \quad (5)$$

represents an explicitly distributed variant of process Sys given in (3): the two partners now reside in different locations, namely **buyer** and **seller**; location **sys** encloses the whole system. Our process model allows us to define an update process that depends on the current protocol state of the two channels at location **sys**:

$$\mathbf{sys} \left\{ \text{case } x, y \text{ of } \left\{ \begin{array}{l} (\alpha_{\text{buy}}; \beta_{\text{sel}}) : \mathbf{buyer}[R] \mid \mathbf{seller}[Q] \\ (\alpha_{\text{pay}}; \beta_{\text{pay}}) : \mathbf{buyer}[P^*] \mid \mathbf{seller}[Q^*] \end{array} \right\} \right\} \quad (6)$$

This update process defines two alternatives for adaptation:

1. If the current types for the two channels at **sys** are α_{buy} and β_{sel} (i.e., the protocol has just been established) then only the buyer is updated—its new behavior will be given by process R .
2. If the current types for the two channels at **sys** are α_{pay} and β_{pay} (i.e., both item and price information have been already exchanged) then new implementations P^* and Q^* are installed in the respective locations.

Session types are used to guide update processes, but also to ensure that the ensuing process alternatives are sound. The type system we develop in §4 ensures that process $\mathbf{buyer}[R] \mid \mathbf{seller}[Q]$ (the first alternative) has exactly two free session channels x and y , with session types α_{buy} and β_{sel} , respectively. A similar check applies to process $\mathbf{buyer}[P^*] \mid \mathbf{seller}[Q^*]$ and α_{pay} and β_{pay} .

As a simple example of the event-based constructs, we may consider a variant of the process given in §1, which defines a conditional (and persistent) adaptation policy:

$$U = \mu\mathcal{X}.\text{if arrive}(\mathbf{sys}, \text{upd}_E) \text{ then } U_{\text{bs}} \text{ else } \mathcal{X} \quad (7)$$

Above, U_{bs} denotes the update process in (6). Intuitively, the process in (7) persistently checks if an adaptation message upd_E has arrived to the queue of location **sys**, which we denote $\mathbf{sys}[\tilde{r}]$. (We shall write

r, r_0, r_1, \dots to range over adaptation messages.) Notice that such a message may be issued by process W itself or by its context, using an adaptation signal, denoted $\overline{\text{loc}}(\text{upd}_E)$. We therefore envision systems in which communication actions, adaptation policies, and event-based constructs concurrently interact. In our example, we could have the process

$$W \mid U \mid \text{sys}[r_0 \cdot \tilde{r}]$$

which triggers the adaptation of location sys if $r_0 = \text{upd}_E$; otherwise, the communication behavior of W would take place.

2.2. Multiparty Sessions

Background. In § 5 we consider multiparty sessions following the approach in [CDV15], which relies on *global types, monitors, processes, networks, and systems*.

In [CDV15], global types describe communications and explicit adaptation signals; monitors, derived from global types, describe the protocol from the perspective of a single participant. Processes only specify communicating behavior; they are associated to monitors that regulate their behavior and contain information on senders/receivers. Each protocol participant is then defined as a monitored process, a process tied with a monitor. Networks are collections of monitored processes, together with constructs for starting new global protocols and queues for handling asynchronous communication. A system is the composition of a network with a global state, a collection of data that may influence adaptation.

We introduce some notational conventions. Below, $\mathbf{p}, \mathbf{q}, \dots$ denote participants and Π represents a set of participants. Also, l_1, l_2, \dots denote *labels*; S_1, S_2, \dots denote carried types (e.g., Bool, Int); and $\lambda_1, \lambda_2, \dots$ denote *adaptation flags*. Global types abstract sequences of broadcast-like communications from a single participant to a set of participants; such communications include usual labeled messages and also explicit adaptation signals (flags).

$$\begin{array}{l} \mathbf{G}, \mathbf{G}' ::= \mathbf{p} \rightarrow \Pi : \{l_i(S_i).G_i\}_{i \in I} \quad \text{directed communication with labeled alternatives} \\ \quad \quad \quad \mid \mathbf{p} \rightarrow \Pi : \{\lambda_i\}_{i \in I} \quad \text{adaptation flag} \\ \quad \quad \quad \mid \text{end} \quad \quad \quad \text{terminated protocol} \end{array}$$

The set of *monitors* is defined as follows:

$$\begin{array}{l} \mathbf{M} ::= \mathbf{p}^? \{l_i(S_i).M_i\}_{i \in I} \mid \Pi! \{l_i(S_i).M_i\}_{i \in I} \quad \text{communication} \\ \quad \quad \quad \mid \mathbf{p}^? \{\lambda_i\}_{i \in I} \mid \Pi! \{\lambda_i\}_{i \in I} \quad \text{adaptation} \\ \quad \quad \quad \mid \text{end} \quad \quad \quad \text{terminated choreography} \end{array}$$

Formally, global types and monitors are related via a *projection* function that defines how a global type can be decomposed into monitors. Projection also determines valid global types: a global type is said to be *well-formed* if (i) its projections are defined for all participants, and if (ii) exchanges of adaptation flags involve all participants declared in the global type. The syntax of processes is as follows:

$$\begin{array}{l} P ::= c!l(x).P \mid c!l(e).P \mid \text{if } e \text{ then } P \text{ else } Q \mid P + P \mid \mu\mathcal{X}.P \mid \mathcal{X} \mid \mathbf{0} \\ \quad \quad \quad \mid \text{op}.P \mid c?l(\lambda, T).P \mid c!l(\lambda(F), T).P \end{array}$$

Each process owns a unique *channel* c . At run-time, channels specify a session channel and the identity of the participant. Most constructs are standard; processes $\text{op}.P$, $c?l(\lambda, T).P$, and $c!l(\lambda(F), T).P$ deserve additional explanations:

- $\text{op}.P$ denotes an operation on the global state; as such, it influences the behavior of future adaptation actions.
- process $c?l(\lambda, T).P$ receives an adaptation flag λ and has a continuation of type T .
- process $c!l(\lambda(F), T).P$ specifies an *adaptation function* F which, in combination with the global state, will determine a new global type.

Given a monitor \mathbf{M} and a process P , a monitored process is denoted $\mathbf{M}[P]$. Adaptation actions manifest at the level of networks; they are realized directly over monitors but indirectly over processes: the modification of the global type of the system results into a new set of monitors associated to processes; if a current process does not “fit” its new monitor, then it must be replaced with a different process, obtained from a centralized

$e ::= v \mid x, y, z \mid k = k \mid a = a$	expressions	$\text{arrive}(\text{loc}, r)$	arrival predicate
$P ::= \bar{u}(x : \alpha).P$	session request	$\text{loc}[P]$	located process
$u(x : \alpha).P$	session accept	$\text{loc}\{\text{case } x_1, \dots, x_n \text{ of } \{(\beta_i^1; \dots; \beta_i^n) : P_i\}_{i \in I}\}$	update process
$\bar{k}(e).P$	session output	$k[\alpha]$	session monitor
$k(x).P$	session input	$\text{loc}[\tilde{r}]$	location queue
$k \triangleleft n; P$	session selection	$\overline{\text{loc}}(r)$	adaptation signal
$k \triangleright \{n_i : P_i\}_{i \in I}$	session branching		
$\text{close}(k).P$	session closure		
$\mu \mathcal{X}.P$	recursion		
\mathcal{X}	recursion variable		
$P \mid P$	composition		
$(\nu \kappa)P$	channel hiding		
$(\nu u)P$	name hiding		
$\text{if } e \text{ then } P_1 \text{ else } P_2$	conditional		
0	inaction		

Table 1. Syntax of expressions and processes for binary communications. Annotation α denotes a session type. Standard constructs are aligned to the left; novel constructs are aligned to the right.

participant (the bank k) will be added to the conversation. The new global protocol will then be:

$$G' ::= s \rightarrow b : \text{iban}(\text{string}).b \rightarrow k : \text{iban}(\text{string}). \\ b \rightarrow p : \text{ad}(\text{string}).s \rightarrow p : \text{itname}(\text{item})$$

and the idea is that current communications are stopped, and the current active process is substituted with another one that implements the new protocol. The semantics for networks would lead to the following adapted process:

$$\begin{aligned} & (\nu \kappa)(\text{loc}_b[y! \text{card}(\text{number}).y! \text{ad}(\text{address}); \mathcal{P}_b; \text{loc}_b[\epsilon]] \mid \\ & \quad \text{loc}_s[y? \text{card}(x).y! \text{itname}(x_1); \mathcal{P}_s; \text{loc}_s[\epsilon]] \mid \text{loc}_p[y? \text{ad}(x_1).y? \text{itname}(x_2); \mathcal{P}_p; \text{loc}_p[\epsilon]]) \\ \rightarrow & \\ & (\nu \kappa)(\text{loc}_b[y? \text{iban}(x).y! \text{iban}(x).y! \text{ad}(\text{address}); \mathcal{P}_b; \text{loc}_b[\epsilon]] \mid \text{loc}_k[y? \text{iban}(x); \mathcal{P}_k; \text{loc}_k[\epsilon]] \\ & \quad \text{loc}_s[y! \text{iban}(\text{ibstr}).y! \text{itname}(x_1); \mathcal{P}_s; \text{loc}_s[\epsilon]] \mid \text{loc}_p[y? \text{ad}(x_1).y? \text{itname}(x_2); \mathcal{P}_p; \text{loc}_p[\epsilon]]) \end{aligned}$$

We shall revisit this scenario in §5—see Example 5.8.

3. Event-Based Adaptation for Binary Communications

Syntax. We rely on base sets for *names*, ranged over by u, a, b, \dots ; (*session*) *channels*, ranged over by k, κ^p, \dots , with *polarity* $p \in \{+, -\}$; *labels*, ranged over by n, n', \dots ; and *variables*, ranged over by x, y, \dots . *Values*, ranged over by v, v', \dots , may include booleans (written **false** and **true**), integers, names, and channels. We use r to range over *adaptation messages*, such as upd_E in (7) above. We use $\tilde{\cdot}$ to denote finite sequences. Thus, e.g., \tilde{x} is a sequence of variables x_1, \dots, x_n . We use ϵ to denote the empty sequence.

Table 1 reports the syntax of expressions and processes. Processes include usual constructs for input, output, and labeled choice. Common forms of recursion, parallel composition, conditionals, and restriction are also included. Constructs for session establishment are annotated with a session type α , which is useful in derived static analyses. A prefix for closing a session, inherited from [DP15], is convenient to better structure specifications. Variable x is bound in processes $\bar{u}(x:\alpha).P$, $u(x:\alpha).P$, and $k(x).P$. Binding for name and channel restriction is as usual. Also, recursion variable \mathcal{X} is bound in process $\mu \mathcal{X}.P$. Given a process P , its sets of free/bound channels, names, variables, and recursion variables—noted $\text{fc}(P)$, $\text{fn}(P)$, $\text{fv}(P)$, $\text{fpv}(P)$, $\text{bc}(P)$, $\text{bn}(P)$, $\text{bv}(P)$, and $\text{bpv}(P)$, respectively—are as expected. We always rely on usual notions of α -conversion and (capture-avoiding) substitution, denoted $[k/x]$ (for channels) and $[P/\mathcal{X}]$ (for processes). We

write $[k_1, \dots, k_n/x_1, \dots, x_n]$ to stand for an n -ary simultaneous substitution. Processes without free variables or free channels are called *programs*.

Up to here, the language is essentially a synchronous π -calculus with sessions [HVK98]. As motivated in the Introduction, building upon *locations* $\text{loc}, \mathbf{l}_1, \mathbf{l}_2, \dots$, we consider the following novel process constructs for run-time adaptation:

- *located processes*, denoted $\text{loc}[P]$;
- *update processes*, denoted $\text{loc}\{\text{case } x_1, \dots, x_n \text{ of } \{(\beta_i^1; \dots; \beta_i^n) : Q_i\}_{i \in I}\}$;
- *session monitors*, denoted $k[\alpha]$, for a session type α ;
- *location queues*, denoted $\text{loc}[\tilde{r}]$; and
- *adaptation signals*, denoted $\overline{\text{loc}}(r)$.

Moreover, expressions now include the *arrival predicate* $\text{arrive}(\text{loc}, r)$.

We now comment on these elements. *Located processes* and *update processes* have been already discussed. Here we just remark that update processes are assumed to refer to at least one variable x_i and to offer at least one alternative Q_i . Also, variables x_1, \dots, x_n are bound in $\text{loc}\{\text{case } x_1, \dots, x_n \text{ of } \{(\beta_i^1; \dots; \beta_i^n) : Q_i\}_{i \in I}\}$; this process is often abbreviated as $\text{loc}\{\text{case } \tilde{x} \text{ of } \{(\beta_i^1; \dots; \beta_i^n) : Q_i\}_{i \in I}\}$. Update processes generalize the **typecase** introduced in [Kou12], which defines a case-like choice based on a single channel; in contrast, to specify adaptation for locations with multiple open sessions, our update processes define type-directed checks over one or more channels.

Update processes go hand-in-hand with *monitors*, run-time entities which keep the current protocol state at a given channel. We write $\kappa^p[\alpha]$ to denote the monitor that stores the protocol state α for channel κ^p . In [Kou12], a similar construct is used to store in-transit messages in asynchronous communications. For simplicity, here we consider synchronous communication; monitors store only the current protocol state. This choice is aligned with our goal of identifying the elements from the eventful session framework that are essential to run-time adaptation (cf. Remark 4.12).

Location queues, not present in [Kou12], handle adaptation requests, modeled as a possibly empty sequence of messages \tilde{r} . Location queues enable us to give a unified treatment to adaptation requests. Given $\text{loc}[\tilde{r}]$, it is worth observing that messages \tilde{r} are not related to communication as abstracted by session types. This represents the fact that we handle adaptation requests and structured session exchanges as orthogonal issues. An *adaptation signal* $\overline{\text{loc}}(r)$ enqueues request r into the location queue of loc . To this end, as detailed below, the operational semantics defines synchronizations between adaptation signals and location queues. To connect run-time adaptation and communication, our language allows the coupling of update processes with the *arrival predicate on locations*, denoted $\text{arrive}(\text{loc}, r)$. Inspired by the **arrive** predicate in [Kou12], this predicate detects if a message r has been placed in the queue of loc .

Our process language embodies several concerns related to run-time adaptation: using adaptation signals and location queues we may specify *how* an adaptation request is issued; arrival predicates enable us to specify *when* adaptation will be handled; using update processes and monitors we may specify *what* is the goal of an adaptation event.

Semantics. The semantics of our language is given by a *reduction semantics*, the smallest relation generated by the rules in Table 2. We write $P \longrightarrow P'$ for the reduction from P to P' . Reduction relies on a standard notion of structural congruence, denoted \equiv (see Appendix A). It also relies on *evaluation* and *location* contexts:

$$E ::= - \mid \bar{k}(-).P \mid \text{if } - \text{ then } P \text{ else } Q \quad C, D ::= - \mid \text{loc}[C \mid P]$$

Given $C\{-\}$ (resp. $E[-]$), we write $C\{P\}$ (resp. $E[e]$) to denote the process (resp. expression) obtained by filling in occurrences of hole $-$ in C with P (resp. in E with e).

We comment on the reduction rules below. The first four rules formalize session behavior within hierarchies of nested locations. Using duality for session types, denoted \perp_C (cf. Definition A.4), in Rule $\langle R:\text{OPEN} \rangle$ the synchronization on a name u leads to establish a session on fresh channels κ^p and $\kappa^{\bar{p}}$; also, two monitors with the declared session types are created. Duality for polarities p is as expected: $\bar{\bar{p}} = -$ and $\bar{-} = +$. By construction, monitors are *local*, i.e., they are created in the same contexts in which the session is established. Rule $\langle R:\text{COM} \rangle$ represents communication of a value: we require both complementary prefixes and that the monitors support input and output actions. After reduction, prefixes in processes and monitors are consumed. Similarly, Rule $\langle R:\text{SEL} \rangle$ for labeled choice is standard, augmented with monitors. Rule $\langle R:\text{CLO} \rangle$

$\langle \text{R:OPEN} \rangle$	$\frac{\alpha \perp_c \beta}{C\{u(x : \alpha).P\} \mid D\{\bar{u}(y : \beta).Q\} \longrightarrow (\nu\kappa)(C\{P[\kappa^P/x] \mid \kappa^P[\alpha]\} \mid D\{Q[\kappa^{\bar{P}}/y] \mid \kappa^{\bar{P}}[\beta]\})}$
$\langle \text{R:COM} \rangle$	$\frac{C\{\bar{\kappa}^P(v).P \mid \kappa^P[!(T).\alpha]\} \mid D\{\kappa^{\bar{P}}(x).Q \mid \kappa^{\bar{P}}[?(T).\beta]\}}{C\{P \mid \kappa^P[\alpha]\} \mid D\{Q[v/x] \mid \kappa^{\bar{P}}[\beta]\}}$
$\langle \text{R:SEL} \rangle$	$\frac{i \in J}{C\{\kappa^P \triangleright \{n_j : P_j\}_{j \in J} \mid \kappa^P[\&\{n_j : \alpha_j\}_{j \in J}]\} \mid D\{\kappa^{\bar{P}} \triangleleft n_i ; Q \mid \kappa^{\bar{P}}[\oplus\{n_j : \beta_j\}_{j \in J}]\}} \longrightarrow C\{P_i \mid \kappa^P[\alpha_i]\} \mid D\{Q \mid \kappa^{\bar{P}}[\beta_i]\}$
$\langle \text{R:CLO} \rangle$	$\frac{C\{\text{close}(\kappa^P).P \mid \kappa^P[\text{end}]\} \mid D\{\text{close}(\kappa^{\bar{P}}).Q \mid \kappa^{\bar{P}}[\text{end}]\}}{C\{P\} \mid D\{Q\}}$
$\langle \text{R:UREQ} \rangle$	$\frac{C\{\text{loc}[\tilde{r}_1]\} \mid D\{\text{loc}(r)\}}{C\{\text{loc}[\tilde{r}_1 \cdot r]\} \mid D\{\mathbf{0}\}}$
$\langle \text{R:ARR1} \rangle$	$\frac{\tilde{r} = r_1 \cdot \tilde{r}_0}{C\{\text{E}[\text{arrive}(\text{loc}, r_1)]\} \mid D\{\text{loc}[\tilde{r}]\} \longrightarrow C\{\text{E}[\text{true}]\} \mid D\{\text{loc}[\tilde{r}_0]\}}$
$\langle \text{R:ARR2} \rangle$	$\frac{(\tilde{r} = r_2 \cdot \tilde{r}_0 \wedge r_1 \neq r_2) \vee \tilde{r} = \epsilon}{C\{\text{E}[\text{arrive}(\text{loc}, r_1)]\} \mid D\{\text{loc}[\tilde{r}]\} \longrightarrow C\{\text{E}[\text{false}]\} \mid D\{\text{loc}[\tilde{r}]\}}$
$\langle \text{R:UPD1} \rangle$	$\frac{\text{match}_I(\{\alpha_1, \dots, \alpha_m\}, \{\beta_i^1, \dots, \beta_i^m\}_{i \in I}) = \uparrow}{C\{\text{loc}[P]\} \mid D\{\text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_i^1; \dots; \beta_i^m) : Q_i\}_{i \in I}\}\} \longrightarrow C\{\text{loc}[P]\} \mid D\{\mathbf{0}\}}$
$\langle \text{R:UPD2} \rangle$	$\frac{\text{fc}(P) = \{\kappa_1^P, \dots, \kappa_m^P\} \quad \text{match}_I(\{\alpha_1, \dots, \alpha_m\}, \{\beta_i^1, \dots, \beta_i^m\}_{i \in I}) = l \quad \forall j \in [1, \dots, m]. (\kappa_j^P[\alpha_j] \in P) \quad \text{barbs}(P; Q_l[\kappa_1^P, \dots, \kappa_m^P/x_1, \dots, x_m])}{C\{\text{loc}[P]\} \mid D\{\text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_i^1; \dots; \beta_i^m) : Q_i\}_{i \in I}\}\} \longrightarrow C\{\text{loc}[Q_l[\kappa_1^P, \dots, \kappa_m^P/x_1, \dots, x_m]]\} \mid D\{\mathbf{0}\}}$
$\langle \text{R:EVA} \rangle$	$\frac{e \longrightarrow e}{\text{E}[e] \longrightarrow \text{E}[e']}$
$\langle \text{R:PAR} \rangle$	$\frac{P \longrightarrow P'}{P \mid Q \longrightarrow P' \mid Q}$
$\langle \text{R:RESN} \rangle$	$\frac{P \longrightarrow P'}{(\nu a)P \longrightarrow (\nu a)P'}$
$\langle \text{R:RESC} \rangle$	$\frac{P \longrightarrow P'}{(\nu\kappa)P \longrightarrow (\nu\kappa)P'}$
$\langle \text{R:STR} \rangle$	$\frac{P \equiv P' \quad P' \longrightarrow Q' \quad Q' \equiv Q}{P \longrightarrow Q}$
$\langle \text{R:REC} \rangle$	$\frac{\text{rec } \mathcal{X}.P \longrightarrow P[\text{rec } \mathcal{X}.P/\mathcal{X}]$
$\langle \text{R:IFTRUE} \rangle$	$\frac{\text{if true then } P \text{ else } Q \longrightarrow P$
$\langle \text{R:IFFALSE} \rangle$	$\frac{\text{if false then } P \text{ else } Q \longrightarrow Q$

Table 2. Reduction semantics for binary processes. Above, α and β denote session types.

formalizes session termination, discarding involved monitors. The monitors in these three rules allow us to track the evolution of active session protocols. The following five rules in Table 2 define our event-based approach to run-time adaptation. Before commenting on them, we require two auxiliary definitions:

Definition 3.1 (Matching). Given an index set I of alternatives, session types $\alpha_1, \dots, \alpha_m$, and an indexed sequence of session types $\{\beta_i^1, \dots, \beta_i^m\}_{i \in I}$, we define

$$\text{match}_I(\{\alpha_1, \dots, \alpha_m\}, \{\beta_i^1, \dots, \beta_i^m\}_{i \in I}) = \begin{cases} l & \text{if } l \in I \wedge (\forall j \in [1, \dots, m]. \beta_l^j \leq_c \alpha_j) \\ \uparrow & \forall n \in I. (\exists j \in [1, \dots, m]. \beta_n^j \not\leq_c \alpha_j) \end{cases}$$

Intuitively, $\text{match}_I(\{\alpha_1, \dots, \alpha_m\}, \{\beta_i^1, \dots, \beta_i^m\}_{i \in I}) = l$ says that the l -th alternative ensures a match (up to subtyping) between types α_j and β_l^j , for all $j \in [1, \dots, m]$. If there is no alternative satisfying the match (i.e., for every alternative n there is a pair α_j and β_n^j for which match does not hold) then the function is undefined, which is denoted by \uparrow .

Definition 3.2 (Barbs). Let P be a process. We write $P \downarrow_{\kappa^P}$ if P is structurally congruent to one of the following:

$$\begin{aligned} & (\nu\tilde{\kappa})(C\{\bar{\kappa}^P(v).P_1\} \mid D\{R\}) & (\nu\tilde{\kappa})(C\{\kappa^P(x).P_1\} \mid D\{R\}) \\ & (\nu\tilde{\kappa})(C\{\kappa^P \triangleright \{n_1:P_1 \parallel \dots \parallel n_m:P_m\}\} \mid D\{R\}) & (\nu\tilde{\kappa})(C\{\kappa^P \triangleleft n_i; P'\} \mid D\{R\}) \end{aligned}$$

We write $\text{barbs}(P; Q)$ whenever $P \downarrow_{\kappa^P}$ if and only if $Q \downarrow_{\kappa^P}$.

We may now return to describing the reduction rules:

- Rule $\langle R:\text{UREQ} \rangle$ treats the issue of an adaptation request r as a synchronization between a location queue and an adaptation signal. The queue and the signal may be in different contexts; this enables “remote” requests.
- Rules $\langle R:\text{ARR1} \rangle$ and $\langle R:\text{ARR2} \rangle$ resolve arrival predicates by querying the (possibly remote) queue \tilde{r} .
- Rules $\langle R:\text{UPD1} \rangle$ and $\langle R:\text{UPD2} \rangle$ define the update of the behavior at location loc . Given an index set I over the update process, suitability with respect to the behavior at loc is defined by the function match_I (cf. Definition 3.1). There are two possibilities. If there is no matching alternative then the current protocol state at loc is kept unchanged (Rule $\langle R:\text{UPD1} \rangle$). Otherwise, the predicate holds for an alternative that defines a new protocol state which preserves the barbs of the current state (Rule $\langle R:\text{UPD2} \rangle$). By an abuse of notation, we write $P_1 \in P$ to indicate that P_1 occurs in P , i.e., if $P = (\nu\tilde{\kappa})C\{P_1\}$ for some C and $\tilde{\kappa}$.

In addition, our semantics includes standard and/or self-explanatory treatments for reduction under evaluation contexts, parallel composition, located context, and restriction. Also, it accounts for applications of structural congruence, recursion and conditionals.

Example 3.3. Recall process W given in § 2.1, Equation (5). According to our semantics:

$$\begin{aligned} W & \longrightarrow (\nu\kappa)(\text{sys}[\text{buyer}[P[\kappa^P/x] \mid \kappa^P[\alpha_{\text{buy}}]] \mid \text{seller}[Q[\kappa^{\bar{P}}/y] \mid \kappa^{\bar{P}}[\beta_{\text{sel}}]]]) \\ & \longrightarrow^2 (\nu\kappa)(\text{sys}[\text{buyer}[P' \mid \kappa^P[\alpha_{\text{pay}}]] \mid \text{seller}[Q' \mid \kappa^{\bar{P}}[\beta_{\text{pay}}]]]) \end{aligned}$$

Suppose that following an external request the seller must offer a new payment method (a gift card). Precisely, we would like the seller to act according to the type β_{gift} given in (4). Let α_{gift} be the dual of β_{gift} . We then may define the following update process:

$$R_{xy}^1 = \text{sys}\{\text{case } x, y \text{ of } \{(\alpha_{\text{pay}}; \beta_{\text{pay}}) : \text{buyer}[P' \mid x[\alpha_{\text{gift}}]] \mid \text{seller}[Q'' \mid y[\beta_{\text{gift}}]]\}\}$$

Thus, R_{xy}^1 keeps the expected implementation for the buyer (P'), but updates its associated monitor. For the seller, both the implementation and monitor are updated; above, Q'' is a process that offers the three payment methods. We may then specify the whole system as: $W \mid \mu\mathcal{X}.\text{if arrive}(\text{sys}, \text{upd}_E) \text{ then } R_{xy}^1 \text{ else } \mathcal{X}$. The type system introduced next ensures, among other things, that updates such as R_{xy}^1 consider both a process and its associated monitors, ruling out the possibility of discarding the monitors that enable reduction.

4. Session Types for Eventful Run-time Adaptation

This section introduces a session type system for the process language of § 3. Our main result (Theorem 4.11) is that well-typed programs enjoy both *safety* (absence of run-time communication errors) and *consistency* properties (update actions do not disrupt established sessions). Our development follows the lines of the typed framework in [DP15].

Syntax. The syntax of session types (ranged over by α, β, \dots) has been presented in § 2.1. We consider *basic types* (ranged over by τ, σ, \dots) and write T, S, \dots to range over τ, α . Although our process language copes with run-time adaptation, our type syntax retains the intuitive meaning of standard session types [HVK98], which we now briefly recall.

Type $?(\tau). \alpha$ (resp. $?(\beta). \alpha$) abstracts the behavior of a channel which receives a value of type τ (resp. a channel of type β) and then continues as α . Dually, type $!(\tau). \alpha$ (resp. $!(\beta). \alpha$) represents the behavior of a channel which sends a value of type τ and then continues as α . Type $\&\{n_1 : \alpha_1, \dots, n_m : \alpha_m\}$ describes a branching behavior: it offers m behaviors, and if the j -th alternative is selected then it behaves as described

by type α_j ($1 \leq j \leq m$). In turn, type $\oplus\{n_1 : \alpha_1, \dots, n_m : \alpha_m\}$ describes the behavior of a channel which may select a single behavior among $\alpha_1, \dots, \alpha_m$ and then continues as α_j . We use **end** to type a channel with no communication behavior. Type $\mu t.\alpha$ describes recursive behavior; as usual, we consider recursive types under equi-recursive and contractive assumptions.

Along the paper we have informally appealed to *duality* and *subtyping* over session types (denoted \perp_c and \leq_c , resp.). Since our session type structure is completely standard, we may rely on the (coinductive) definitions given by Gay and Hole [GH05]; see also Definition A.4.

Our typing judgments generalize usual notions with an *interface* \mathcal{I} for each process. Based on the syntactic occurrences of session establishment prefixes $\bar{a}(x:\alpha)$, and $a(x:\alpha)$, the interface of a process describes the services appearing in it. We annotate services with a *qualification* q , which may be ‘lin’ (linear) or ‘un’ (unrestricted). Intuitively, a service is linear if it is offered a finite number of times, otherwise it is persistent or unrestricted. Thus, the interface of a process gives an “upper bound” on the services that it may execute. The typing system uses interfaces to control the behavior contained by locations after an update. We have:

Definition 4.1 (Interfaces). We define an *interface* as the multiset whose underlying set of elements is $I = \{q u:\alpha \mid q \in \{\text{lin}, \text{un}\}\}$ (i.e., a set of assignments from names to qualified session types). We use $\mathcal{I}, \mathcal{I}', \dots$ to range over interfaces. We write $\text{dom}(\mathcal{I})$ to denote the set $\{u \mid q u : \alpha \in \mathcal{I}\}$ and $\#_{\mathcal{I}}(q u : \alpha) = h$ to mean that u occurs h times in \mathcal{I} .

The union of two interfaces is essentially the union of their underlying multisets. We sometimes write $\mathcal{I} \uplus a : \alpha_{\text{lin}}$ and $\mathcal{I} \uplus a : \alpha_{\text{un}}$ to stand for $\mathcal{I} \uplus \{\text{lin } a:\alpha\}$ and $\mathcal{I} \uplus \{\text{un } a:\alpha\}$, respectively. Moreover, we write \mathcal{I}_{lin} (resp. \mathcal{I}_{un}) to denote the subset of \mathcal{I} involving only assignments qualified with lin (resp. un). We now define an ordering relation over interfaces, relying on subtyping:

Definition 4.2 (Interface Ordering). Given interfaces \mathcal{I} and \mathcal{I}' , we write $\mathcal{I} \sqsubseteq \mathcal{I}'$ iff

1. $\forall (\text{lin } a:\alpha)$ such that $\#_{\mathcal{I}_{\text{lin}}}(\text{lin } a:\alpha) = h$ with $h > 0$, then one of the following holds:
 - (a) there exist h distinct elements $(\text{lin } a:\beta_i) \in \mathcal{I}'_{\text{lin}}$ such that $\alpha \leq_c \beta_i$ for $i \in [1..h]$;
 - (b) there exists $(\text{un } a:\beta) \in \mathcal{I}'_{\text{un}}$ such that $\alpha \leq_c \beta$.
2. $\forall (\text{un } a:\alpha) \in \mathcal{I}_{\text{un}}$ then $(\text{un } a:\beta) \in \mathcal{I}'_{\text{un}}$ and $\alpha \leq_c \beta$, for some β .

Typing Environments. We now define our typing environments, following [DP15]. Recall that we write q to range over qualifiers lin and un.

$$\begin{array}{lll}
\Delta & ::= & \emptyset \mid \Delta, k : \alpha \mid \Delta, k : [\alpha] & \text{typing with active sessions} \\
\Gamma & ::= & \emptyset \mid \Gamma, e : \tau \mid \Gamma, u : \langle \alpha_q, \beta_q \rangle & \text{first-order environment (with } \alpha_q \perp_c \beta_q) \\
\Theta & ::= & \emptyset \mid \Theta, \mathcal{X} : \Delta; \mathcal{I} \mid \Theta, \text{loc} : \mathcal{I} & \text{higher-order environment}
\end{array}$$

We consider typings Δ and environments Γ and Θ . Typing Δ collects assignments from channels to session types; it describes currently active sessions. In our system, Δ also includes *bracketed assignments*, denoted $k^p : [\alpha]$, which represent the type for monitors. Subtyping extends to these assignments ($[\alpha] \leq_c [\beta]$ if $\alpha \leq_c \beta$) and thus to typings. We write $\text{dom}(\Delta)$ to denote the set $\{k^p \mid k^p : \alpha \in \Delta \vee k^p : [\alpha] \in \Delta\}$. We write $\Delta, k : \alpha$ where $k \notin \text{dom}(\Delta)$. Furthermore, we write $\Delta, k : \langle \alpha \rangle$ to abbreviate $\Delta, k : \alpha, k : [\alpha]$. That is, $k : \langle \alpha \rangle$ describes both a session and its associated monitor.

Γ is a first-order environment which maps expressions to basic types and names to pairs of qualified session types. The higher-order environment Θ collects assignments of typings to process variables and interfaces to locations. While the former concerns recursive processes, the latter concerns located processes. As we explain next, by relying on the combination of these two pieces of information the type system ensures that run-time adaptation actions preserve the behavioral interfaces of a process. We write $\text{vdom}(\Theta) = \{\mathcal{X} \mid \mathcal{X} : \mathcal{I} \in \Theta\}$ to denote the variables in the domain of Θ . Given these environments, a *type judgment* is of form

$$\Gamma ; \Theta \vdash P \triangleright \Delta; \mathcal{I}$$

meaning that, under environments Γ and Θ , process P has active sessions declared in Δ and interface \mathcal{I} .

Typing rules are shown in Tables 3 and 4. Below we comment on some of the rules in Table 3: the rest are standard and/or self explanatory. Rule $\langle T:\text{ADAPT} \rangle$ types update processes. Notice that the typing rule

$\frac{\langle \text{T:ACC} \rangle}{\alpha \perp_c \beta \quad \Gamma \vdash u \triangleright \langle \alpha_{\text{lin}}, \beta_{\text{lin}} \rangle \quad \gamma \leq_c \alpha \quad \Gamma; \Theta \vdash P \triangleright \Delta, x : \gamma; \mathcal{I}}{\Gamma; \Theta \vdash u(x : \gamma).P \triangleright \Delta; \mathcal{I} \uplus u : \gamma_{\text{lin}}}$	$\frac{\langle \text{T:REQ} \rangle}{\alpha \perp_c \beta \quad \Gamma \vdash u \triangleright \langle \alpha_q, \beta_{\text{lin}} \rangle \quad \gamma \leq_c \beta \quad \Gamma; \Theta \vdash P \triangleright \Delta, x : \gamma; \mathcal{I}}{\Gamma; \Theta \vdash \bar{u}(x : \gamma).P \triangleright \Delta; \mathcal{I} \uplus u : \gamma_{\text{lin}}}$		
$\frac{\langle \text{T:CLO} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I} \quad k \notin \text{dom}(\Delta)}{\Gamma; \Theta \vdash \text{close}(k).P \triangleright \Delta, k : \mathbf{end}; \mathcal{I}}$	$\frac{\langle \text{T:PAR} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta_1; \mathcal{I}_1 \quad \Gamma; \Theta \vdash Q \triangleright \Delta_2; \mathcal{I}_2}{\Gamma; \Theta \vdash P \mid Q \triangleright \Delta_1 \cup \Delta_2; \mathcal{I}_1 \uplus \mathcal{I}_2}$	$\frac{\langle \text{T:REC} \rangle}{\Gamma; \Theta \vdash \mu \mathcal{X}.P \triangleright \Delta; \mathcal{I}}{\Gamma; \Theta, \mathcal{X} : \Delta; \mathcal{I} \vdash P \triangleright \Delta; \mathcal{I}}$	
$\frac{\langle \text{T:LOCENV} \rangle}{\Theta, \text{loc} : \mathcal{I} \vdash \text{loc} \triangleright \mathcal{I}}$	$\frac{\langle \text{T:MSG} \rangle}{\Gamma \vdash r_1 \triangleright \text{msg}}$	$\frac{\langle \text{T:LOCQ} \rangle}{\Gamma \vdash r_1; \tilde{r} \triangleright \text{msg} \quad \Gamma \vdash r_1 \triangleright \text{msg}}$	$\frac{\langle \text{T:ARRIVE} \rangle}{\Theta \vdash \text{loc} \triangleright \mathcal{I} \quad \Gamma \vdash r \triangleright \text{msg}}{\Gamma; \Theta \vdash \text{arrive}(\text{loc}, r) \triangleright \text{bool}}$
$\frac{\langle \text{T:SIG} \rangle}{\Gamma \vdash r \triangleright \text{msg}}{\Gamma; \Theta \vdash \overline{\text{loc}}(r) \triangleright \emptyset; \emptyset}$	$\frac{\langle \text{T:LOC} \rangle}{\Theta \vdash \text{loc} \triangleright \mathcal{I} \quad \Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I}' \quad \mathcal{I}' \sqsubseteq \mathcal{I}}{\Gamma; \Theta \vdash \text{loc}[P] \triangleright \Delta; \mathcal{I}'}$	$\frac{\langle \text{T:QLOC} \rangle}{\Gamma \vdash \tilde{r} \triangleright \text{msg}}{\Gamma; \Theta \vdash \text{loc}[\tilde{r}] \triangleright \emptyset; \emptyset}$	
$\frac{\langle \text{T:CRÉS} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta, \kappa^p : \langle \langle \alpha_1 \rangle \rangle, \kappa^{\bar{p}} : \langle \langle \alpha_2 \rangle \rangle; \mathcal{I} \quad \alpha_1 \perp_c \alpha_2}{\Gamma; \Theta \vdash (\nu \kappa)P \triangleright \Delta; \mathcal{I}}$	$\frac{\langle \text{T:NRES} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I} \cup \mathcal{I}_u \quad u \notin \text{dom}(\mathcal{I})}{\Gamma; \Theta \vdash (\nu u)P \triangleright \Delta; \mathcal{I}}$		
$\frac{\langle \text{T:SUB} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I} \quad \Delta \leq_c \Delta' \quad \mathcal{I} \sqsubseteq \mathcal{I}'}{\Gamma; \Theta \vdash P \triangleright \Delta'; \mathcal{I}'}$	$\frac{\langle \text{T:QUE} \rangle}{\Gamma; \Theta \vdash k[\alpha] \triangleright k : [\alpha]; \emptyset}$		
$\frac{\langle \text{T:ADAPT} \rangle}{\Theta \vdash \text{loc} \triangleright \mathcal{I} \quad \forall j \in J, \text{fv}(Q_j) \setminus \{x_1, \dots, x_m\} = \emptyset \quad \Gamma; \Theta \vdash Q_j \triangleright x_1 : \langle \langle \beta_1^j \rangle \rangle; \dots; x_m : \langle \langle \beta_m^j \rangle \rangle; \mathcal{I}_j \quad \mathcal{I}_j \sqsubseteq \mathcal{I}}{\Gamma; \Theta \vdash \text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_1^j; \dots; \beta_m^j) : Q_j\}_{j \in J}\} \triangleright \emptyset; \emptyset}$			

Table 3. Well-typed processes: Typing rules (Part I).

ensures that each process Q_i has exactly the same active sessions that those declared in the respective case. Also, we require that alternatives contain both processes and monitors. With $\mathcal{I}_j \sqsubseteq \mathcal{I}$ we guarantee that the process behavior does not “exceed” the expected behavior within the location. Rule $\langle \text{T:SUB} \rangle$ takes care of subtyping both for typings Δ and interfaces. Rule $\langle \text{T:CRÉS} \rangle$ types channel restriction; it ensures typing duality among partners of a session and their respective queues. Rule $\langle \text{T:NRES} \rangle$ types hiding of service names, by simply removing their declarations from the interface \mathcal{I} of the process. In the rule, \mathcal{I}_u contains only declarations for u , i.e., $\forall v \neq u, v \notin \text{dom}(\mathcal{I}_u)$. Typing of queues is given by Rule $\langle \text{T:QUE} \rangle$ that simply assigns type $k : [\alpha]$ to queue $k[\alpha]$.

Example 4.3. Consider the process from Example 3.3:

$$(\nu \kappa)(\text{sys}[\text{buyer}[P[\kappa^p/x] \mid \kappa^p[\alpha_{\text{buy}}]] \mid \text{seller}[Q[\kappa^{\bar{p}}/y] \mid \kappa^{\bar{p}}[\beta_{\text{sel}}]]])$$

The type of process R located in sys is $\Gamma; \Theta \vdash R \triangleright \Delta; \mathcal{I}$ with

$$\begin{aligned} \Gamma &:= u : \langle \text{lin} : \alpha_{\text{buy}}, \text{lin} : \beta_{\text{sel}} \rangle \\ \Theta &:= \text{buyer} : \{ \text{lin } \bar{u} : \alpha_{\text{buy}} \}, \text{seller} : \{ \text{lin } u : \beta_{\text{sel}} \}, \text{sys} : \{ \text{lin } \bar{u} : \alpha_{\text{buy}}, \text{lin } u : \beta_{\text{sel}} \} \\ \Delta &:= \kappa^p : \langle \langle \alpha_{\text{buy}} \rangle \rangle, \kappa^{\bar{p}} : \langle \langle \beta_{\text{sel}} \rangle \rangle \\ \mathcal{I} &:= \emptyset \end{aligned}$$

Our type system enjoys the standard *subject reduction* property. We rely on *balanced* typings:

Definition 4.4. We say that typing Δ is *balanced* iff for all $\kappa^p : \alpha \in \Delta$ (resp. $\kappa^p : [\alpha] \in \Delta$) then also $\kappa^{\bar{p}} : \beta \in \Delta$ (resp. $\kappa^{\bar{p}} : [\beta] \in \Delta$), with $\alpha \perp_c \beta$.

Similarly to what has been proposed in [HYC08], here we introduce a reduction over session typings:

Definition 4.5 (Reduction for Typings). Reduction for typings, denoted $\Delta \mapsto \Delta'$, is defined by the following rules:

1. $\Delta, \kappa^p : \langle \langle (T).\alpha \rangle \rangle, \kappa^{\bar{p}} : \langle \langle (T).\beta \rangle \rangle \mapsto \Delta, \kappa^p : \langle \langle \alpha \rangle \rangle, \kappa^{\bar{p}} : \langle \langle \beta \rangle \rangle$.
2. $\Delta, \kappa^p : \langle \langle \&\{n_1:\alpha_1, \dots, n_m:\alpha_m\} \rangle \rangle, \kappa^{\bar{p}} : \langle \langle \oplus\{n_1:\beta_1, \dots, n_m:\beta_m\} \rangle \rangle \mapsto \Delta, \kappa^p : \langle \langle \alpha_i \rangle \rangle, \kappa^{\bar{p}} : \langle \langle \beta_i \rangle \rangle$, with $i \in \{1, \dots, m\}$.

$\frac{\langle \text{T:BOOL} \rangle}{\Gamma \vdash \text{true, false} \triangleright \text{bool}}$	$\frac{\langle \text{T:NAME} \rangle}{\Gamma \vdash u \triangleright \text{name}}$	$\frac{\langle \text{T:BVAR} \rangle}{\Gamma, x : \text{bool} \vdash x \triangleright \text{bool}}$	$\frac{\langle \text{T:NVAR} \rangle}{\Gamma, x : \text{name} \vdash x \triangleright \text{name}}$
$\frac{\langle \text{T:EQ} \rangle}{d = u \vee d = \kappa^p \vee d = x}$	$\frac{\langle \text{T:SER} \rangle}{\alpha \perp_c \beta}$	$\frac{\langle \text{T:RVAR} \rangle}{\Gamma; \Theta, \mathcal{X} : \Delta, \mathcal{I} \vdash \mathcal{X} : \Delta; \mathcal{I}}$	$\frac{\langle \text{T:NIL} \rangle}{\Gamma; \Theta \vdash \mathbf{0} \triangleright \emptyset; \emptyset}$
$\frac{\langle \text{T:THR} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta, k : \beta; \mathcal{I}}$	$\frac{\langle \text{T:CAT} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta, k : \beta, x : \alpha; \mathcal{I}}$		$\frac{\langle \text{T:CAT} \rangle}{\Gamma; \Theta \vdash k(x).P \triangleright \Delta, k : ?(\alpha).\beta; \mathcal{I}}$
$\frac{\langle \text{T:IN} \rangle}{\Gamma, x : \tau; \Theta \vdash P \triangleright \Delta, k : \alpha; \mathcal{I}}$	$\frac{\langle \text{T:OUT} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta, k : \alpha; \mathcal{I} \quad \Gamma \vdash e \triangleright \tau}$		$\frac{\langle \text{T:OUT} \rangle}{\Gamma; \Theta \vdash \bar{k}(e).P \triangleright \Delta, k : !(\tau).\alpha; \mathcal{I}}$
$\frac{\langle \text{T:WEAKC} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I} \quad \kappa^+, \kappa^- \notin \text{dom}(\Delta)}$	$\frac{\langle \text{T:WEAKN} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I} \quad u \notin \text{dom}(\mathcal{I})}$		$\frac{\langle \text{T:WEAKN} \rangle}{\Gamma; \Theta \vdash (\nu u)P \triangleright \Delta; \mathcal{I}}$
$\frac{\langle \text{T:IF} \rangle}{\Gamma; \Theta \vdash e \triangleright \text{bool} \quad \Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I} \quad \Gamma; \Theta \vdash Q \triangleright \Delta; \mathcal{I}}$	$\frac{\langle \text{T:SEL} \rangle}{\Gamma; \Theta \vdash P \triangleright \Delta, k : \alpha_i; \mathcal{I} \quad 1 \leq i \leq m}$		$\frac{\langle \text{T:SEL} \rangle}{\Gamma; \Theta \vdash k \triangleleft n_i; P \triangleright \Delta, k : \oplus\{n_1 : \alpha_1, \dots, n_m : \alpha_m\}; \mathcal{I}}$
$\frac{\langle \text{T:IF} \rangle}{\Gamma; \Theta \vdash \text{if } e \text{ then } P \text{ else } Q \triangleright \Delta; \mathcal{I}}$			
$\frac{\langle \text{T:SEL} \rangle}{\Gamma; \Theta \vdash P_1 \triangleright \Delta, k : \alpha_1; \mathcal{I}_1 \quad \dots \quad \Gamma; \Theta \vdash P_m \triangleright \Delta, k : \alpha_m; \mathcal{I}_m \quad \mathcal{I} = \mathcal{I}_1 \uplus \dots \uplus \mathcal{I}_m}$			
$\frac{\langle \text{T:SEL} \rangle}{\Gamma; \Theta \vdash k \triangleright \{n_1 : P_1 \parallel \dots \parallel n_m : P_m\} \triangleright \Delta, k : \&\{n_1 : \alpha_1, \dots, n_m : \alpha_m\}; \mathcal{I}}$			

Table 4. Well-typed processes: Typing rules (Part II).

Theorem 4.6 (Subject Reduction). If $\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I}$ with Δ balanced and $P \longrightarrow Q$ then $\Gamma; \Theta \vdash Q \triangleright \Delta'; \mathcal{I}'$, for some $\mathcal{I}' \subseteq \mathcal{I}$ and balanced Δ' such that either $\Delta' \leq_c \Delta$ or $\Delta \longmapsto \Delta'$.

Proof. By induction on the last rule applied in the reduction. See Appendix B for details. \square

We now define and state *safety* and *consistency* properties. While safety guarantees adherence to prescribed session types and absence of run-time *errors*, consistency ensures that sessions are not jeopardized by careless run-time adaptation actions. Defining both properties requires the following notions of κ -processes and (located) κ -redexes:

Definition 4.7 (κ -processes, κ -redexes, errors). A process P is a κ -process if it is a prefixed process with subject κ^p , i.e., P is one of the following:

$$\kappa^p(x).P' \quad \bar{\kappa}^p(v).P' \quad \text{close}(\kappa^p).P' \quad \kappa^p \triangleright \{n_i : P_i\}_{i \in I} \quad \kappa^p \triangleleft n.P'$$

Process P is a κ -redex if it contains the composition of exactly two κ -processes with *opposing polarities*, i.e., for some contexts C, D , and E , and processes P_1, P_2, \dots, P_m , and P' , process P is structurally congruent to one of the following:

$$\begin{array}{l|l} (\nu \tilde{\kappa})(C\{\bar{\kappa}^p(v).P_1\} & D\{\kappa^{\bar{p}}(x).P_2\}) \\ (\nu \tilde{\kappa})(C\{\kappa^p \triangleright \{n_1 : P_1 \parallel \dots \parallel n_m : P_m\}\} & D\{\kappa^{\bar{p}} \triangleleft n_i; P'\}) \\ (\nu \tilde{\kappa})(C\{\text{close}(\kappa^p).P_1\} & D\{\text{close}(\kappa^{\bar{p}}).P_2\}) \end{array}$$

A *located κ -redex* is a κ -redex in which one or both of its constituent κ -processes are contained by least one located process. P is an *error* if $P \equiv (\nu \tilde{\kappa})(Q \mid R)$ where, for some κ , Q contains either exactly two κ -processes that do not form a κ -redex or three or more κ -processes.

Example 4.8. The following processes

$$l_2[l_1[\kappa^p(\tilde{x}).P_1 \mid \bar{\kappa}^{\bar{p}}(v).P_2] \quad l_1[\kappa^p(\tilde{x}).P_1 \mid l_2[\bar{\kappa}^{\bar{p}}(v).P_2] \quad l_1[\kappa^p(\tilde{x}).P_1 \mid \bar{\kappa}^{\bar{p}}(v).P_2]$$

are located κ -redexes, whereas process $\kappa^p(\tilde{x}).P_1 \mid \bar{\kappa}^{\bar{p}}(v).P_2$ is not.

Informally, a process P is called *consistent* if whenever it has a located κ -redex then update actions do not destroy such a redex. Below, we formalize this intuition. Let us write $P \longrightarrow_{\text{upd}} P'$ for any reduction inferred using Rule $\langle \text{R:UPD2} \rangle$ (cf. Table 2). We then define:

Definition 4.9 (Safety and Consistency). Let P be a process.

- We say P is *safe* if it never reduces into an error.
- We say P is *update-consistent* if and only if, for all P' and κ such that $P \longrightarrow^* P'$ and P' contains a κ -redex, if $P' \longrightarrow_{\text{upd}} P''$ then P'' contains a κ -redex.

Lemma 4.10. If $\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I}$ with Δ balanced then P is not an error process.

Proof. Follows from session type duality, Definition 4.4 (balanced typing), and Definition 4.7 (error). \square

We now state our main result; it follows as a consequence of Theorem 4.6.

Theorem 4.11 (Typing Ensures Safety and Consistency). If $\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I}$ with Δ balanced then program P is safe and update-consistent.

Proof. Safety means that P never reduces into an error. Since Δ is balanced, by Lemma 4.10 we know that P is not an error. Theorem 4.6 ensures that reduction preserves balanced typings; therefore, P never reduces to an error.

Proving update-consistency entails showing that for all P_0 and κ such that $P \longrightarrow^* P_0$ and P_0 contains a κ -redex, if $P_0 \longrightarrow_{\text{upd}} P_1$ then P_1 contains a κ -redex. That is, P_0 has both the κ -redex and an update action which preserves it. By Theorem 4.6 we know that P_0 is well-typed under a balanced typing. To show consistency we distinguish two cases. The first case is when the κ -redex is *not* located in P_0 : then, the update action cannot affect it (i.e., the update concerns located process in P_0 which do not enclose the κ -processes that constitute the κ -redex) and the thesis follows. The second case is when the κ -redex is located in P_0 : therefore, κ -processes that constitute the κ -redex could be affected by the reduction $P_0 \longrightarrow_{\text{upd}} P_1$. There are several sub-cases, depending on which κ -processes are located. We consider the sub-case in which one κ -process (say, $R_1 = \kappa^p(x).R'$) is located whereas the other (say, $R_2 = \bar{\kappa}^{\bar{p}}(v).R''$) is not; other cases are similar. The proof follows directly from Rule $\langle \text{R:UPD2} \rangle$ (cf. Table 2) which replaces R_1 with an alternative Q_l . The rule not only ensures that R_1 and Q_l have matching session types (up to subtyping) but also that they have the same barbs, i.e., the same top-level actions. This ensures that the κ -redex is preserved under the reduction $\longrightarrow_{\text{upd}}$, and so the thesis follows. \square

Remark 4.12 (Asynchronous Communication). In this section, we have focused on *synchronous* communication: this allows us to give a compact semantics, relying on a standard type structure. To account for asynchrony, we would require a run-time syntax for programs with queues for in-transit messages (values, sessions, labels). The type system should be extended to accommodate these new run-time processes. In our case, an extension with asynchrony would rely on the machinery defined in [Kou12]; we present such a machinery (for multiparty sessions) in the following section.

Remark 4.13 (Incremental Adaptation). Adaptation in our framework is “incremental” in that modifications always preserve/extend active session protocols, exploiting subtyping. Our framework can be modified so that arbitrary protocols are installed as a result of an update. One would need to ensure that the two endpoints of a session are present in the same location: arbitrary updates are safe as long as both endpoints are simultaneously updated with dual protocols. This alternative requires modifying definitions for matching (Definition 3.1) and interface ordering (Definition 4.2).

5. Event-Based Adaptation for Multiparty Communications

Here we generalize our approach to the case of multiparty communications. As already explained, we present a proof of concept that integrates type-directed constructs for update and eventful constructs for adaptation requests into the framework put forward by Coppo et al. [CDV15].

As described in § 2.2, in the framework of [CDV15] communication actions from adaptation mechanisms are tied together, as adaptation flags occur within global type specifications at the same level of protocol exchanges. Our event-based approach treats adaptation at the level of processes, defining a separation between communication and adaptation concerns, which in our opinion is beneficial for specification and analysis.

Integrating our event-based approach entails novelties with respect to the notions of global types and networks in [CDV15]. Monitored processes will be now embedded into distributed locations which contain local collections and queues that govern local and external adaptation. Furthermore, at the level of processes we replace processes $c?l(\lambda, \mathbb{T}).P$, and $c!l(\lambda(F), \mathbb{T}).P$ with a construct $c!(\mathbb{G}).P$ which enforces a locally-motivated adaptation based on global type \mathbb{G} . Overall, we have a two-level framework of processes and systems (instead of processes, networks, and systems, as in [CDV15]).

5.1. Syntax of Types

5.1.1. Global Types

Global types specify structured interactions between two or more participants. We distinguish between communication and adaptation concerns; for this reason, unlike [CDV15], our global types do not mention synchronizations related to adaptation. We write $\mathfrak{p}, \mathfrak{q}, \dots$ and Π, Π', \dots to denote participants and sets of participants, respectively. Exchange of values take place on *labels*, denoted λ, λ', \dots in the sequel, which are useful to express choices. We write $\text{loc}, \text{loc}', \dots$ to range over locations. As in [CDV15], we assume that communicated values are extracted from a set of sorts S . Formally, we have:

Definition 5.1 (Global Types). The set of *global types* is defined by

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

As we follow [CDV15], our syntax of global types is admittedly simple: a global type $\mathfrak{p} \rightarrow \mathfrak{q} : \{\lambda_i(S_i).G_i\}_{i \in I}$ represents a labeled communication from participant \mathfrak{p} to participant \mathfrak{q} . Global type end denotes the terminated protocol. We write $\text{part}(\mathbb{G})$ to denote the set of participants declared in \mathbb{G} .

5.1.2. Monitors

Monitors capture the local view that participants have of a global type. They are used to enable actions of the process implementations; they also contain information on the participants and labels involved in the directed exchanges.

Definition 5.2 (Monitors). The set of *monitors* is defined by:

$$\mathcal{M} ::= \mathfrak{p}?\{\lambda_i(S_i).\mathcal{M}_i\}_{i \in I} \mid \mathfrak{p}!\{\lambda_i(S_i).\mathcal{M}_i\}_{i \in I} \mid \text{end}$$

An input monitor $\mathfrak{p}?\{\lambda_i(S_i).\mathcal{M}_i\}_{i \in I}$ can be associated to a process which can receive, for each $i \in I$, a value of sort S_i with label λ_i , with a continuation that is associated to \mathcal{M}_i (external choice). The output monitor $\mathfrak{p}!\{\lambda_i(S_i).\mathcal{M}_i\}_{i \in I}$ is dual: it can be associated to a process which sends a value of sort S_i , distinguished by label λ_i , for each $i \in I$, with a continuation that is associated to \mathcal{M}_i (internal choice).

Following the approach of multiparty session types [HYC08], the monitor of a participant is obtained via a *projection function*, defined below.

Definition 5.3 (Projection of a Global Type). The projection of a global type \mathbb{G} onto a monitor for a participant \mathfrak{p} , denoted $\mathbb{G} \downarrow_{\mathfrak{p}}$, is defined as follows:

- $(\mathfrak{p} \rightarrow \mathfrak{q} : \{\lambda_i(S_i).G_i\}_{i \in I}) \downarrow_{\mathfrak{r}} = \begin{cases} \mathfrak{p}?\{\lambda_i(S_i).G_i \downarrow_{\mathfrak{r}}\}_{i \in I} & \text{if } \mathfrak{r} = \mathfrak{q} \\ \mathfrak{q}!\{\lambda_i(S_i).G_i \downarrow_{\mathfrak{r}}\}_{i \in I} & \text{if } \mathfrak{r} = \mathfrak{p} \\ G_1 \downarrow_{\mathfrak{r}} & \text{if } \mathfrak{r} \neq \mathfrak{p} \text{ and } \mathfrak{r} \neq \mathfrak{q} \text{ and } G_i \downarrow_{\mathfrak{r}} = G_j \downarrow_{\mathfrak{r}} \text{ for all } i, j \in I \end{cases}$
- $\text{end} \downarrow_{\mathfrak{r}} = \text{end}$

We will say that a global type is *well-formed* if its projections are defined for all participants. We shall work only with well-formed global types. We assume that each well-formed \mathbb{G} is associated with a mapping $\mathcal{L}_{\mathbb{G}}$ from participants to locations such that $\mathfrak{p} \neq \mathfrak{p}'$ implies $\mathcal{L}_{\mathbb{G}}(\mathfrak{p}) \neq \mathcal{L}_{\mathbb{G}}(\mathfrak{p}')$, for all $\mathfrak{p}, \mathfrak{p}' \in \text{part}(\mathbb{G})$.

5.1.3. Process Types

We now describe types for the process language that will be introduced in Definition 5.6. While in multiparty session types [HYC08], the process types correspond to *local types* (the projection of the global type

$\frac{\langle \text{M:EXP} \rangle}{\Gamma, e : S \vdash e : S}$	$\frac{\langle \text{M:END} \rangle}{\Gamma \vdash \mathbf{0} \triangleright c : \text{end}}$	$\frac{\langle \text{M:ADAPT} \rangle}{\Gamma \vdash P \triangleright c : T}$	$\frac{\langle \text{M:REC} \rangle}{\Gamma, X : T \vdash P \triangleright c : T}$	$\frac{\langle \text{M:IF} \rangle}{\Gamma \vdash e : S \quad \Gamma \vdash P_1 \triangleright c : T \quad \Gamma \vdash P_2 \triangleright c : T}$
$\frac{\langle \text{M:RCV} \rangle}{\Gamma, x_1 : S_1 \vdash P_1 \triangleright c : T_1 \quad \dots \quad \Gamma, x_n : S_n \vdash P_n \triangleright c : T_n}$		$\frac{\langle \text{M:SEND} \rangle}{\Gamma \vdash P \triangleright c : T_k \quad \Gamma \vdash e : S \quad k \in I}$		
$\Gamma \vdash c? \{ \lambda_i(x_i).P_i \}_{i \in I} \triangleright c : \& \{ \lambda_i(S_i).T_i \}_{i \in I}$		$\Gamma \vdash c! \lambda(e).P \triangleright c : \oplus \{ \lambda_i(S_i).T_i \}_{i \in I}$		

Table 5. Typing rules for multiparty processes.

onto every participant), the approach in [CDV15] uses a discipline of process types which is different from local types (i.e., monitors). This adds flexibility, and is useful to establish an independence between local implementations (as formalized by processes) and choreographic coordination (as formalized by global types and monitors). Clearly, process types cannot be completely independent from monitors; a basic coherence between the two should exist. This coherence is defined in terms of *adequacy* (cf. Definition 5.5), a relation that uses subtyping to determine which process types “fit” a given monitor.

Following [Pad10], the process types defined on [CDV15] exploit intersection and union types; this ensures an intuitive subtyping relation. Here we consider a syntax for process types that is certainly closer to the local types of [HYC08] than the syntax in [CDV15]. This is for the sake of presentation, as we wish to keep a simple conceptual relation with the binary session types introduced in §4. We then have:

Definition 5.4 (Process Types). The syntax of process types is defined as follows:

$$T ::= \& \{ \lambda_i(S_i).T_i \}_{i \in I} \mid \oplus \{ \lambda_i(S_i).T_i \}_{i \in I} \mid \mu t.T \mid t \mid \text{end}$$

We write \mathcal{T} to denote the set of all types, and use T, T', \dots to range over process types.

Thus, intuitively, process types $\& \{ \lambda_i(S_i).T_i \}_{i \in I}$ and $\oplus \{ \lambda_i(S_i).T_i \}_{i \in I}$ are used to type labelled inputs and outputs, respectively. Recursive types are as expected, following usual considerations for type equality.

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

We now introduce *typing judgments* for expressions e and processes P :

$$\Gamma \vdash e : S \quad \Gamma \vdash P \triangleright c : T$$

Expressions are typed by sorts; we assume standard typing rules. The judgments for processes which makes it explicit that processes with at most one channel can be typed.

Typing rules for processes are given in Table 5. Due to the absence of explicit flags for adaptation, they are a subset of the typing rules given in [CDV15]. We also consider Rule $\langle \text{M:ADAPT} \rangle$ which represents the fact that internal adaptation requests are defined independently from communication protocols abstracted by types.

The relation between process types and monitors is formalized via *adequacy*. Adequacy relies on a subtyping relation, here denoted \leq ; it corresponds to a finite variant of the subtyping \leq_c (formally given in Definition A.4). Our notion of adequacy is simple, and highlights how the difference between monitors and process types lies in the information about participant identities:

Definition 5.5 (Adequacy). Let $|\cdot|$ be a mapping from monitors to types defined as follows:

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

We then say that type T is *adequate* to monitor \mathcal{M} (notation $T \propto \mathcal{M}$) if $T \leq |\mathcal{M}|$.

5.2. Processes and Networks

5.2.1. Syntax

Processes abstract the protocol implementations for each participant. The association of a process P with a monitor \mathcal{M} results into a *monitored process* $\mathcal{M}[P]$, which realizes the behavior of a single participant. As hinted at above, the main differences of our framework with respect to [CDV15] appear at the level of monitored processes and networks (collections of monitored processes).

The definition of processes implementing multiparty services, given next, is rather standard. We have operators for sending and receiving messages, conditionals, and recursion. On top of these operators we add a primitive for broadcasting an *internal adaptation request*, denoted $c!(G).P$. Intuitively, this primitive is an output action that mentions a new global type to be implemented by protocol participants.

Definition 5.6 (Processes). The syntax of processes is defined as:

$$P ::= c?\{\lambda_i(x_i).P_i\}_{i \in I} \mid c!\lambda(e).P \mid \mu\mathcal{X}.P \mid \mathcal{X} \mid \text{if } e \text{ then } P \text{ else } Q \mid \mathbf{0} \\ \mid c!(G).P$$

Each process owns a unique channel, denoted y . Given a session κ and a participant \mathfrak{p} , channel y in the user code is substituted at run-time by a session channel $\kappa[\mathfrak{p}]$. In the following we write c to denote either y (the user channel) or $\kappa[\mathfrak{p}]$ (the session channel). We assume that processes pass around expressions, defined as expected. Every expression e reduces to a value v ; this is denoted $e \rightsquigarrow v$.

As an example of internal adaptation requests, suppose participants $\mathfrak{p}_1, \dots, \mathfrak{p}_j, \dots, \mathfrak{p}_n$ establish session κ to communicate according to a specific protocol (say, **Skype**). To request a protocol change, participant \mathfrak{p}_j may issue an adaptation request to all its partners by including the process $\kappa[\mathfrak{p}_j]!(\text{GTalk}).P$ in its local implementation. As we will formalize later on, this request will arrive to the event queue of session κ , and then all participants will proceed to adapt their local behavior to the new global type **GTalk**.

Our notion of networks departs significantly from that in [CDV15] by considering an explicitly distributed setting based on (inactive) locations and event queues:

Definition 5.7 (Networks). The syntax of *networks* N and *adaptation requests* r is defined as

N	$::=$	$\text{new}(G)$	initialize global type
		$\text{loc}[\mathcal{M}[P]; \mathcal{P}; \text{loc}[r]]$	active location
		$\text{loc}[\mathcal{P}; \text{loc}[r]]$	inactive location
		$\kappa : h$	message queue
		$\kappa^e[h]$	internal adaptation: event queue
		$(\nu\kappa)N$	restriction
		$N \mid N$	parallel composition
r	$::=$	ϵ	empty request
		$(\text{add} : \langle Q, T \rangle)$	external adaptation: upgrade request
		$(\text{upd} : \text{case } x \text{ of } \{(T_i) : Q_i\}_{i \in I})$	external adaptation: update request

A location represents a place where several programs can be executed. We assume a preexisting set of locations which are “filled in” with appropriate monitored processes upon initialization of global types (see below). An active location $\text{loc}[\mathcal{M}[P]; \mathcal{P}; \text{loc}[r]]$ is composed of a monitored process $\mathcal{M}[P]$, a local collection of typed processes \mathcal{P} (i.e., a set of pairs (Q, T) where Q is a process and T is its corresponding type), and a queue $\text{loc}[r]$ (where r stands for an *external* adaptation request). For simplicity, we assume that in each location only one session can be active. An inactive location is simply a location without a monitored process. The queue of a location connects it to an unspecified environment that can (i) add new processes to the collection (or upgrade the an existing process), or (ii) update the behavior of a monitored process. This is similar to the update processes for the binary framework.

As we shall see, initializing a global type G entails setting up monitored processes in appropriate locations, establishing a session channel κ , and creating a message queue $\kappa : h$ for supporting asynchronous communication. We write \emptyset to denote the empty message queue; messages are of the form $(\mathfrak{p}, \mathfrak{q}, \lambda(v))$ thus denoting the fact that participant \mathfrak{p} sends value v and label λ to participant \mathfrak{q} . Message concatenation is denoted with ‘.’.

$$\begin{array}{c}
\frac{\langle \text{M:P:IN} \rangle}{\kappa[\mathbf{p}]?\{\lambda_i(x_i).P_i\}_{i \in I} \xrightarrow{\kappa[\mathbf{p}]?\lambda_j(v_j)} P_j[v/x]} \quad \frac{\langle \text{M:P:OUT} \rangle}{e \sim v} \quad \frac{\langle \text{M:P:ADAPOUT} \rangle}{\kappa[\mathbf{p}]\lambda(G).P \xrightarrow{\kappa[\mathbf{p}]\lambda(G)} P} \\
\frac{\langle \text{M:P:IFT} \rangle}{e \sim \mathbf{true}} \quad \frac{\langle \text{M:P:IFF} \rangle}{e \sim \mathbf{false}} \quad \frac{\langle \text{M:P:REC} \rangle}{\mu\mathcal{X}.P \xrightarrow{\tau} P[\mu\mathcal{X}.P/X]} \\
\text{if } e \text{ then } P_1 \text{ else } P_2 \xrightarrow{\tau} P_1 \quad \text{if } e \text{ then } P_1 \text{ else } P_2 \xrightarrow{\tau} P_2
\end{array}$$

Table 6. LTS for multiparty processes.

The event queue $\kappa^e[h]$ handles internal adaptation requests, associated to process $c!(G).P$ introduced earlier.

5.2.2. Semantics

The semantics of networks is given in terms of a reduction relation, denoted $N \longrightarrow N'$. We require some auxiliary definitions, in order to connect the behavior of monitors and processes to that of whole networks; we also require a structural congruence to capture, among other things, asynchronous message communication as handled via queues. As in [CDV15], we assume that all local collections \mathcal{P} are *complete*, which intuitively means that such collections contain all conceivable pairs of processes and types.

LTS for Monitors and Processes. Reduction for networks relies on LTSs for monitors and processes. The LTS for monitors is rather simple; it is used to enable the actions to be performed by processes:

$$\mathbf{p}?\{\lambda_i(S_i).\mathcal{M}_i\}_{i \in I} \xrightarrow{\mathbf{p}?\lambda_j} \mathcal{M}_j \quad (j \in I) \quad \mathbf{q}!\{\lambda_i(S_i).\mathcal{M}_i\}_{i \in I} \xrightarrow{\mathbf{q}!\lambda_j} \mathcal{M}_j \quad (j \in I)$$

The LTS for processes relies on the following set of labels, with associated rules as in Table 6.

$$\alpha ::= \tau \mid \kappa[\mathbf{p}]?\lambda(v) \mid \kappa[\mathbf{p}]\lambda(v) \mid \kappa[\mathbf{p}]\lambda(G)$$

Structural Congruence. Reduction for networks relies also on a *structural equivalence* and on *evaluation contexts*. The structural equivalence on networks, denoted \equiv , defines parallel composition as a commutative, associative operator, and which allows restriction to reduce and enlarge their scope without name clashes. Also, it decrees that any monitored process with the end monitor corresponds to the neutral element for parallel composition and absorbs restriction. More formally:

$$\begin{array}{l}
\text{end}[P] \mid N \equiv N \quad (\nu\kappa)\text{end}[P] \equiv \text{end}[P] \quad (\nu\kappa)N_1 \mid N_2 \equiv (\nu\kappa)N_1 \mid N_2 \\
(\nu\kappa)(\kappa : h) \mid N \equiv N \quad (\nu\kappa)(\nu\kappa')N \equiv (\nu\kappa')(\nu\kappa)N
\end{array}$$

We consider also a structural equivalence for message queues which allows us to commute messages involving different receivers or senders:

$$h \cdot (\mathbf{p}, \mathbf{q}, \lambda_i(v_i)) \cdot (\mathbf{p}', \mathbf{q}', \lambda_j(v_j)) \cdot h' \equiv h \cdot (\mathbf{p}', \mathbf{q}', \lambda_j(v_j)) \cdot (\mathbf{p}, \mathbf{q}, \lambda_i(v_i)) \cdot h' \quad \text{if } \mathbf{p} \neq \mathbf{p}' \text{ or } \mathbf{q} \neq \mathbf{q}'$$

This congruence is extended to networks in the expected way by letting: $h \equiv h'$ implies $\kappa : h \equiv \kappa : h'$.

Finally, we have the following definition for evaluation contexts, denoted \mathcal{E} :

$$\mathcal{E} ::= [] \mid \mathcal{E} \mid N \mid (\nu\kappa)\mathcal{E}$$

Reduction Rules for Networks. The rules of the reduction semantics for networks are given in Table 7; we now describe them:

- Rule $\langle \text{N:OPEN} \rangle$ initializes a global type, denoted G in the rule. To that end, it considers all the inactive locations associated to the declared participants of G , denoted Π . For each $\mathbf{p} \in \Pi$, a process implementation is picked up from the local collection $\mathcal{P}_{\mathbf{p}}$. The process type of this implementation is expected to be adequate to the monitor $G \downarrow_{\mathbf{p}}$. Observe that mapping \mathcal{L}_G ensures that each participant is assigned to a different (inactive) location. A fresh session name κ is created for protocol participants; associated message and event queues are also created.

$\langle \text{N:OPEN} \rangle$	$\frac{\Pi = \text{part}(\mathbf{G}) \quad \forall p \in \Pi. (\mathcal{M}_p = \mathbf{G} \downarrow_p \wedge (P_p, T_p) \in \mathcal{P}_p \wedge T_p \propto \mathcal{M}_p \wedge \mathcal{L}_{\mathbf{G}}(p) = \text{loc}_p)}{\text{new}(\mathbf{G}) \mid \prod_{p \in \Pi} \text{loc}_p[\mathcal{P}_p; \text{loc}_p[\epsilon]] \longrightarrow (\nu \kappa) \left(\prod_{p \in \Pi} \text{loc}_p[\mathcal{M}_p[P_p[\kappa[p]/y]]; \mathcal{P}_p; \text{loc}_p[\epsilon]] \mid \kappa : \emptyset \mid \kappa^e[\epsilon] \right)}$
$\langle \text{N:SEND} \rangle$	$\frac{\mathcal{M} \xrightarrow{q! \lambda} \mathcal{M}' \quad P \xrightarrow{\kappa[p]! \lambda(v)} P'}{\text{loc}_p[\mathcal{M}[P]; \mathcal{P}; \text{loc}_p[\epsilon]] \mid \kappa : h \mid \kappa^e[\epsilon] \longrightarrow \text{loc}_p[\mathcal{M}'[P']; \mathcal{P}; \text{loc}_p[\epsilon]] \mid \kappa : (p, q, \lambda(v)) \cdot h \mid \kappa^e[\epsilon]}$
$\langle \text{N:RECV} \rangle$	$\frac{\mathcal{M} \xrightarrow{q? \lambda} \mathcal{M}' \quad P \xrightarrow{\kappa[p]? \lambda(v)} P'}{\text{loc}_p[\mathcal{M}[P]; \mathcal{P}; \text{loc}_p[\epsilon]] \mid \kappa : (q, p, \lambda(v)) \cdot h \mid \kappa^e[\epsilon] \longrightarrow \text{loc}_p[\mathcal{M}'[P']; \mathcal{P}; \text{loc}_p[\epsilon]] \mid \kappa : h \mid \kappa^e[\epsilon]}$
$\langle \text{N:TAU} \rangle$	$\frac{P \xrightarrow{\tau} P'}{\text{loc}[\mathcal{M}[P]; \mathcal{P}; \text{loc}[\epsilon]] \mid \kappa : h \mid \kappa^e[\epsilon] \longrightarrow \text{loc}[\mathcal{M}[P']; \mathcal{P}; \text{loc}[\epsilon]] \mid \kappa : h \mid \kappa^e[\epsilon]}$
$\langle \text{N:INMSG} \rangle$	$\frac{P \xrightarrow{\kappa!(\mathbf{G})} P'}{\text{loc}[\mathcal{M}[P]; \mathcal{P}; \text{loc}[\epsilon]] \mid \kappa^e[\epsilon] \mid \kappa : h \longrightarrow \text{loc}[\mathcal{M}[P']; \mathcal{P}; \text{loc}[\epsilon]] \mid \kappa^e[\mathbf{G}] \mid \kappa : h}$
$\langle \text{N:INUPD} \rangle$	$\frac{\forall p_i \in \text{part}(\mathbf{G}'). \text{InUpd}(p_i, \Pi, \mathbf{G}', \mathcal{P}_{p_i}, \kappa, \mathcal{M}_{p_i}, \mathcal{M}'_{p_i}, P_{p_i}, P'_{p_i})}{(\nu \kappa) \left(\prod_{p \in \Pi} \text{loc}_p[\mathcal{M}_p[P_p]; \mathcal{P}_p; \text{loc}_p[\epsilon]] \mid \kappa : h \mid \kappa^e[\mathbf{G}'] \right) \mid \prod_{q \in \text{part}(\mathbf{G}') \setminus \Pi} \text{loc}_q[\mathcal{P}_q; \text{loc}_q[\epsilon]] \longrightarrow (\nu \kappa) \left(\prod_{p \in \text{part}(\mathbf{G}')} \text{loc}_p[\mathcal{M}'_p[P'_p]; \mathcal{P}_p; \text{loc}_p[\epsilon]] \mid \kappa : \epsilon \mid \kappa^e[\epsilon] \right) \mid \prod_{q \in \Pi \setminus \text{part}(\mathbf{G}')} \text{loc}_q[\mathcal{P}_q; \text{loc}_q[\epsilon]]}$
$\langle \text{N:ADDLOC1} \rangle$	$\frac{\exists (R, T_R) \in \mathcal{P} \quad T_R \leq T_Q \quad \mathcal{P}' = \mathcal{P} \cup \{(Q, T_Q)\} \setminus \{(R, T_R)\}}{\text{loc}[\mathcal{M}[P]; \mathcal{P}; \text{loc}[(\text{add} : \langle Q, T_Q \rangle)]] \longrightarrow \text{loc}[\mathcal{M}[P]; \mathcal{P}'; \text{loc}[\epsilon]]}$
$\langle \text{N:ADDLOC2} \rangle$	$\frac{\exists (R, T_R) \in \mathcal{P} \quad T_R \leq T_Q \quad \mathcal{P}' = \mathcal{P} \cup \{(Q, T_Q)\} \setminus \{(R, T_R)\}}{\text{loc}[\mathcal{P}; \text{loc}[(\text{add} : \langle Q, T_Q \rangle)]] \longrightarrow \text{loc}[\mathcal{P}'; \text{loc}[\epsilon]]}$
$\langle \text{N:UPDLOC1} \rangle$	$\frac{\text{match}_I^{\times}(\{\mathcal{M}\}, \{T_i\}_{i \in I}) = l \wedge R = Q_l[\kappa[p]/y]}{\text{loc}_p[(\mathcal{M}[P]; \mathcal{P}; \text{loc}_p[(\text{upd} : \text{case } x \text{ of } \{(T_i) : Q_i\}_{i \in I})])] \longrightarrow \text{loc}_p[\mathcal{M}[R]; \mathcal{P}; \text{loc}_p[\epsilon]]}$
$\langle \text{N:UPDLOC2} \rangle$	$\frac{\text{match}_I^{\times}(\{\mathcal{M}\}, \{T_i\}_{i \in I}) = \uparrow}{\text{loc}_p[(\mathcal{M}[P]; \mathcal{P}; \text{loc}_p[(\text{upd} : \text{case } x \text{ of } \{(T_i) : Q_i\}_{i \in I})])] \longrightarrow \text{loc}_p[\mathcal{M}[P]; \mathcal{P}; \text{loc}_p[\epsilon]]}$
$\langle \text{N:EQUIV} \rangle$	$\frac{N_1 \equiv N'_1 \quad N_1 \longrightarrow N_2 \quad N_2 \equiv N'_2}{N'_1 \longrightarrow N'_2}$
$\langle \text{N: EVAL} \rangle$	$\frac{N \longrightarrow N'}{\mathcal{E}[N] \longrightarrow \mathcal{E}[N']}$

Table 7. Reduction semantics for multiparty processes.

- Rules $\langle \text{N:SEND} \rangle$ and $\langle \text{N:RECV} \rangle$ regulate output and input actions from/to the message queue of an established session. The action must be enabled by the label of the monitor, which also describes information on the receiver/sender of the exchanged value. Both monitors and processes evolve as a result of the interaction with the queue. In both rules, a reduction is enabled provided the event queue is empty.
- Rule $\langle \text{N:TAU} \rangle$ formalizes the fact that evaluation of conditional expressions and recursion unfolding within a (monitored) process can occur without affecting its monitor.
- Rules $\langle \text{N:INMSG} \rangle$ and $\langle \text{N:INUPD} \rangle$ formalize *internal adaptation*. Similarly as Rule $\langle \text{N:TAU} \rangle$, issuing an internal adaptation request is a behavior specified by the process but independent from its monitor.

Rule $\langle N:INMSG \rangle$ is enabled if the event queue is empty; it adds a global type to the event queue of the session.

Once a global type is added to the event queue, Rule $\langle N:INUPD \rangle$ can be triggered. This rule assumes that a fixed set of participants (denoted Π) is already executing a global protocol and should adapt to execute a new protocol G' . The rule sets up local implementations for all participants in $\mathbf{part}(G')$, also considering the existing behaviors in the set of participants Π . There are three possibilities:

- (a) Participant p belongs to both Π and $\mathbf{part}(G')$ and location \mathbf{loc}_p encloses a monitored process that “fits” with the projection of G' onto p . In this case, the monitored process is kept as it is.
- (b) Participant p belongs to both Π and $\mathbf{part}(G')$, but the current monitored process enclosed in location \mathbf{loc}_p is not compatible with the projection of G' onto p . In this case, a new process implementation is selected from the local collection \mathcal{P}_p and instantiated with an appropriate session name.
- (c) Participant p belongs to $\mathbf{part}(G')$ but not to Π . In this case, the respective location is made active by setting up an appropriate monitored process extracted from the collection.

Consequently, the premise of Rule $\langle N:INUPD \rangle$ uses the following auxiliary predicate, which formalizes these three possibilities as a disjunction:

$$\begin{aligned} \text{InUpd}(p, \Pi, G', \mathcal{P}_p, \kappa, \mathcal{M}_p, \mathcal{M}'_p, P_p, P'_p) = \\ (\mathbf{p} \in \Pi \cap \mathbf{part}(G') \wedge G' \downarrow_p = \mathcal{M}_p = \mathcal{M}'_p \wedge P'_p = P_p) \vee \\ (\mathbf{p} \in \Pi \cap \mathbf{part}(G') \wedge G' \downarrow_p = \mathcal{M}'_p \neq \mathcal{M}_p \wedge (Q_p, T_p) \in \mathcal{P}_p \wedge T_p \propto \mathcal{M}'_p \wedge P'_p = Q_p[\kappa[p]/y]) \vee \\ (\mathbf{p} \in \mathbf{part}(G') \setminus \Pi \wedge (Q_p, T_p) \in \mathcal{P}_p \wedge G' \downarrow_p = \mathcal{M}'_p \wedge T_p \propto \mathcal{M}'_p \wedge P'_p = Q_p[\kappa[p]/y]) \end{aligned}$$

Notice that the left-hand side of the conclusion of Rule $\langle N:INUPD \rangle$ considers a series of active locations (for participants in Π) together with a series of inactive locations, corresponding to participants in $\mathbf{part}(G')$ but not in Π , i.e., the locations that will become active as a result of the reduction step (cf. item (c) above). Similarly, the right-hand side of the rule features a set of inactive locations, corresponding to participants present in Π but not in $\mathbf{part}(G')$; these correspond to participants excluded as a result of adaptation. Also, observe that both message and event queues are emptied after reduction.

- Rules $\langle N:ADDLOC1 \rangle$, $\langle N:ADDLOC2 \rangle$, $\langle N:UPDLOC1 \rangle$, and $\langle N:UPDLOC2 \rangle$ formalize *external adaptation*. These rules are enabled by the addition of an adaptation request in the location queue. Such a request is meant to come from the context of a location, and contains typed processes.
 - Rule $\langle N:ADDLOC1 \rangle$ upgrades a process pair present in the local collection of an active location; this collection is denoted \mathcal{P} in the rule. Rule $\langle N:ADDLOC2 \rangle$ is similar but applies to inactive locations. Since we assume that the collection is complete, these rules implement a form of upgrade on the local collection.
 - Rules $\langle N:UPDLOC1 \rangle$ and $\langle N:UPDLOC2 \rangle$ formalize type-directed adaptation for monitored processes, in the spirit of the update processes for binary sessions. The external adaptation message contains a finite series of alternatives to the current behavior of the monitored process. The rule tries to find a compatible update (Rule $\langle N:UPDLOC1 \rangle$); if no suitable update is present then the external request is dismissed, and the monitored process is kept unchanged (Rule $\langle N:UPDLOC2 \rangle$). Notice that compatibility depends on a match predicate, denoted $\text{match}_I^\times(\{\mathcal{M}\}, \{T_i\}_{i \in I})$, defined as in Definition 3.1 but considering monitors and process types instead of binary session types, and adequacy rather than subtyping.
- Rules $\langle N:EQUIV \rangle$ and $\langle N:EVAL \rangle$ define self-explanatory treatments for incorporating structural equivalence and evaluation contexts into reduction.

Example 5.8. Consider the example of § 2.2. Using Rule $\langle N:OPEN \rangle$ for session establishment we have:

$$\begin{aligned} \text{new}(G) \mid \mathbf{loc}_b[\mathcal{P}_b; \mathbf{loc}_b[\epsilon]] \mid \mathbf{loc}_s[\mathcal{P}_s; \mathbf{loc}_s[\epsilon]] \mid \mathbf{loc}_p[\mathcal{P}_p; \mathbf{loc}_p[\epsilon]] \\ \longrightarrow \\ (\nu \kappa)(\mathbf{loc}_b[\mathcal{M}_b[P_b]; \mathcal{P}_b; \mathbf{loc}_b[\epsilon]] \mid \mathbf{loc}_s[\mathcal{M}_s[P_s]; \mathcal{P}_s; \mathbf{loc}_s[\epsilon]] \mid \mathbf{loc}_p[\mathcal{M}_p[P_p]; \mathcal{P}_p; \mathbf{loc}_p[\epsilon]] \mid \kappa : \emptyset \mid \kappa^\epsilon[\epsilon]) \end{aligned}$$

with

$$\begin{aligned} \mathcal{M}_b &:= \text{slit}(item).s?price(int).s!ok(bool).s!card(int).p!ad(string) \\ \mathcal{M}_s &:= \text{b?it}(item).b!price(int).b?ok(bool).b?card(int).p!itname(item) \\ \mathcal{M}_p &:= \text{b?ad}(string).s?itname(item) \end{aligned}$$

$$\begin{aligned}
(p!\lambda(S); m) \downarrow_q &= \begin{cases} !\lambda(S).m \downarrow_q & \text{if } p = q \\ m \downarrow_q & \text{otherwise} \end{cases} \quad \epsilon \downarrow_q = \epsilon \quad \text{end } \downarrow_q = \epsilon \quad \langle m, \mathcal{M} \rangle \downarrow_q = (m \downarrow_q).(\mathcal{M} \downarrow_q) \\
p?\{\lambda_i(S_i).\mathcal{M}_i\}_{i \in I} \downarrow_q &= \begin{cases} ?\{\lambda_i(S_i).\mathcal{M}_i \downarrow_q\}_{i \in I} & \text{if } p = q \\ \mathcal{M}_k \downarrow_q & \text{where } k \in I, p \neq q \text{ and } \forall i, j \in I \mathcal{M}_i = \mathcal{M}_j \end{cases} \\
p!\{\lambda_i(S_i).\mathcal{M}_i\}_{i \in I} \downarrow_q &= \begin{cases} !\{\lambda_i(S_i).\mathcal{M}_i \downarrow_q\}_{i \in I} & \text{if } p = q \\ \mathcal{M}_k \downarrow_q & \text{where } k \in I, p \neq q \text{ and } \forall i, j \in I \mathcal{M}_i = \mathcal{M}_j \end{cases}
\end{aligned}$$

Table 8. Projection of generalized types onto participants.

Remark 5.9 (Adaptation Functions). We have proposed using processes communicating global types (here denoted $c!\lambda(G).P$), rather than adaptation flags, to specify internal adaptation, i.e., adaptation requests issued by the system itself. This is a simple way of formalizing this important class of adaptation routines. As an alternative, a less direct mechanism would consist in communicating a value on which a future adaptation should depend on, rather than the expected new choreography. Such a value could be used as a parameter for the adaptation functions used in [CDV15]. This would add flexibility to the specification of internal adaptation requests.

Having introduced networks and their semantics, we now proceed to investigate type-based analysis techniques for them.

5.3. Safety and Consistency

This section generalizes the results of §4 to the multiparty asynchronous case. In order to lighten the syntax, the types for networks do not include any information on interfaces and only describe an active session. Indeed, as interfaces provide information about the services the location may execute and as each location contains a complete collection of processes \mathcal{P} then the interface type will be the same for each location. We have already described how to type a process, we now show how to type its monitor and relative queues and how to combine them into networks.

As communication is asynchronous the information on the type of a session participant p ($\kappa[p]$) can be split between the process implementing the session and the associated session queue. We thus have *generalized types*, which can be either a message type, a monitor type or the combination of the two. More precisely:

Definition 5.10. Session types, generalized types, message and queues types are defined by:

$$\begin{aligned}
\text{Session types} \quad \Delta &::= \emptyset \mid \Delta, \kappa[p] : \chi \\
\text{Generalized types} \quad \chi &::= \mathcal{M} \mid \mathcal{Q} \mid m \\
\text{Queue types} \quad \mathcal{Q} &::= \langle m, \mathcal{M} \rangle \\
\text{Message types} \quad m &::= \epsilon \mid q!\lambda(S) \mid m; m
\end{aligned}$$

The typing judgments for networks are of shape

$$\Gamma \vdash_{\Sigma} N \triangleright \Delta$$

where Σ lists the free session names in N . As for the binary case, type safety relies on a notion of balanced typing: communications are performed in the right order and with the proper type. The definition of duality \bowtie is the same as for the binary case (see Definition A.2).

Definition 5.11. A session typing Δ is balanced for the session κ (denoted $\text{bal}(\Delta, \kappa)$) if $\kappa[p] : \chi \in \Delta$ and $\kappa[q] : \chi' \in \Delta$ with $p \neq q$ imply $\chi \downarrow_q \bowtie \chi' \downarrow_p$ where the projection of generalized types onto participants is given in Table 8.

It is easy to see that session types obtained from the projection of global types are balanced.

Proposition 5.12. Let G be a global type and $p \neq q$, then $(G \downarrow_p) \downarrow_q \bowtie (G \downarrow_q) \downarrow_p$.

We are now ready to complete the typing system with typing rules for networks and queues, given in Table 9. We briefly comment on the rules. A session initiator is typed with the empty session typing

$\langle \text{M:NEW} \rangle$ $\frac{}{\Gamma \vdash_{\emptyset} \mathbf{new}(G) \triangleright \emptyset}$	$\langle \text{M:END} \rangle$ $\frac{}{\Gamma \vdash_{\emptyset} \mathbf{end}[P] \triangleright \emptyset}$	$\langle \text{M:MP} \rangle$ $\frac{\Gamma \vdash P \triangleright \kappa[\mathbf{p}] : T \quad \mathcal{M} \neq \mathbf{end} \quad T \propto \mathcal{M}}{\Gamma \vdash_{\emptyset} \mathcal{M}[P] \triangleright \kappa[\mathbf{p}] : \mathcal{M}}$	
$\langle \text{M:CQUEUE} \rangle$ $\frac{}{\Gamma \vdash_{\emptyset} \kappa^e[h] \triangleright \emptyset}$	$\langle \text{M:QINIT} \rangle$ $\frac{}{\Gamma \vdash_{\{\kappa\}} \kappa : \emptyset \triangleright \emptyset}$	$\langle \text{M:QSEND} \rangle$ $\frac{\Gamma \vdash_{\{\kappa\}} \kappa : h \triangleright \Delta \quad \Gamma \vdash v : S}{\Gamma \vdash_{\{\kappa\}} \kappa : (\mathbf{p}, \mathbf{q}, \lambda(v)) \cdot h \triangleright \Delta \# \{\kappa[\mathbf{p}] : \mathbf{q}!\lambda(S)\}}$	
$\langle \text{M:PAR} \rangle$ $\frac{\Gamma \vdash_{\Sigma_1} N_1 \triangleright \Delta_1 \quad \Gamma \vdash_{\Sigma_2} N_2 \triangleright \Delta_2}{\Gamma \vdash_{\Sigma_1 \cup \Sigma_2} N_1 \mid N_2 \triangleright \Delta_1 \star \Delta_2}$		$\langle \text{M:EQUIV} \rangle$ $\frac{\Gamma \vdash_{\Sigma} N \triangleright \Delta \quad \Delta \approx \Delta'}{\Gamma \vdash_{\Sigma} N \triangleright \Delta'}$	$\langle \text{M:RES} \rangle$ $\frac{\Gamma \vdash_{\Sigma} N \triangleright \Delta \quad \mathbf{bal}(\Delta, \kappa)}{\Gamma \vdash_{\Sigma \setminus \{\kappa\}} (\nu \kappa) N \triangleright \Delta \setminus \{\kappa\}}$
$\langle \text{M:LOC1} \rangle$ $\frac{\forall (P, T_P) \in \mathcal{P} \quad \Gamma \vdash P \triangleright T_P}{\Gamma \vdash_{\emptyset} \mathbf{loc}[P; \mathbf{loc}[r]] \triangleright \emptyset}$		$\langle \text{M:LOC2} \rangle$ $\frac{\forall (P, T_P) \in \mathcal{P} \quad \Gamma \vdash P \triangleright T_P \quad \Gamma \vdash_{\emptyset} \mathcal{M}[P] \triangleright \kappa[\mathbf{p}] : \mathcal{M}}{\Gamma \vdash_{\emptyset} \mathbf{loc}[\mathcal{M}[P]; \mathcal{P}; \mathbf{loc}[r]] \triangleright \kappa[\mathbf{p}] : \mathcal{M}}$	

Table 9. Typing rules for networks and queues.

(Rule $\langle \text{M:NEW} \rangle$). If $\mathcal{M} = \mathbf{end}$ then the monitored process $\mathcal{M}[P]$ is typed with the empty session type (Rule $\langle \text{M:END} \rangle$); otherwise, if the type of P is adequate with respect to the monitor then the type of the monitored process is $\kappa[\mathbf{p}] : \mathcal{M}$ (Rule $\langle \text{M:MP} \rangle$). The event queue κ^e is typed with the completed session type (Rule $\langle \text{M:CQUEUE} \rangle$). Rules $\langle \text{M:QINIT} \rangle$ and $\langle \text{QSEND} \rangle$ type the session message queue. In particular, Rule $\langle \text{M:QSEND} \rangle$ makes use of the operator $\#$ that dispatches message types to proper session channels:

$$\begin{aligned} m \# m' &::= m; m' \\ \Delta \# \Delta' &::= \{\kappa[\mathbf{p}] : \chi \# \chi' \mid \kappa[\mathbf{p}] : \chi \in \Delta \wedge \kappa[\mathbf{p}] : \chi' \in \Delta'\} \cup \{\kappa[\mathbf{p}] : \chi \mid \kappa[\mathbf{p}] : \chi \in \Delta \cup \Delta' \wedge \kappa[\mathbf{p}] : \chi \notin \Delta \cap \Delta'\} \end{aligned}$$

Rule $\langle \text{M:PAR} \rangle$ combines networks; it relies on an operator \star , which combines the typing information coming from queues and monitored processes:

$$\begin{aligned} m \star \mathcal{M} &::= \langle m, \mathcal{M} \rangle \\ \mathcal{M} \star m &::= \langle m, \mathcal{M} \rangle \\ \Delta \star \Delta' &::= \{\kappa[\mathbf{p}] : \chi \star \chi' \mid \kappa[\mathbf{p}] : \chi \in \Delta \wedge \kappa[\mathbf{p}] : \chi' \in \Delta'\} \cup \{\kappa[\mathbf{p}] : \chi \mid \kappa[\mathbf{p}] : \chi \in \Delta \cup \Delta' \wedge \kappa[\mathbf{p}] : \chi \notin \Delta \cap \Delta'\} \end{aligned}$$

In order to take into account the structural congruence on queues (Rule $\langle \text{M:EQUIV} \rangle$) we build an equivalence relation (\approx) on types that is induced by the following equivalence rule on message types:

$$m; \mathbf{q}!\lambda(S); \mathbf{q}'!\lambda(S'); m' \approx m; \mathbf{q}'!\lambda(S'); \mathbf{q}!\lambda(S); m'$$

with $\mathbf{q} \neq \mathbf{q}'$, and is extended to generalized types by letting if $m \approx m'$ then $\langle m, \mathcal{M} \rangle \approx \langle m', \mathcal{M} \rangle$.

As expected, the type of a restricted network depends on the consistency of the session κ (Rule $\langle \text{M:RES} \rangle$). Finally, Rules $\langle \text{M:LOC1} \rangle$ and $\langle \text{M:LOC2} \rangle$ are used to type locations: as in the binary case, locations do not add any information to the session type.

Example 5.13. Consider Example 5.8, and the process

$$R = \mathbf{loc}_b[\mathcal{M}_b[P_b]; \mathcal{P}_b; \mathbf{loc}_b[\epsilon]] \mid \mathbf{loc}_s[\mathcal{M}_s[P_s]; \mathcal{P}_s; \mathbf{loc}_s[\epsilon]] \mid \mathbf{loc}_p[\mathcal{M}_p[P_p]; \mathcal{P}_p; \mathbf{loc}_p[\epsilon]] \mid \kappa : \emptyset \mid \kappa^e[\epsilon]$$

obtained after the session establishment. According to the typing rules, we obtain the following typing judgment

$$\emptyset \vdash_{\{\kappa\}} R \triangleright \kappa[\mathbf{b}] : \langle \emptyset, \mathcal{M}_b \rangle, \kappa[\mathbf{s}] : \langle \emptyset, \mathcal{M}_s \rangle, \kappa[\mathbf{p}] : \langle \emptyset, \mathcal{M}_p \rangle.$$

We now state the *subject reduction* theorem, which guarantees that communications are performed by (located) monitored process in a safe manner, following the protocols prescribed by global types. Similarly as for the case of binary sessions, we introduce a reduction over session types.

Definition 5.14 (Reduction for Typings). Reduction for typings, denoted $\Delta \mapsto \Delta'$, is defined by the following rules:

1. $\Delta, \kappa[\mathbf{p}] : \langle h, \mathbf{q}!\lambda(S).\mathcal{M} \rangle \mapsto \Delta, \kappa[\mathbf{p}] : \langle \mathbf{q}!\lambda(S); h, \mathcal{M} \rangle$
2. $\Delta, \kappa[\mathbf{p}] : \langle \mathbf{q}!\lambda(S); h, \mathcal{M} \rangle, \kappa[\mathbf{q}] : \langle h', \mathbf{p}?\lambda(S).\mathcal{M}' \rangle \mapsto \Delta, \kappa[\mathbf{p}] : \langle h, \mathcal{M} \rangle, \kappa[\mathbf{q}] : \langle h', \mathcal{M}' \rangle$

Theorem 5.15 (Subject Reduction). If $\Gamma \vdash_{\Sigma} N \triangleright \Delta$ with Δ balanced and $N \longrightarrow N'$ then $\Gamma \vdash_{\Sigma} N' \triangleright \Delta'$, for some balanced Δ' such that either $\Delta = \Delta'$ or $\Delta \mapsto \Delta'$.

Proof. By induction on the last rule applied in the reduction. See Appendix B for details. \square

Notice that we can also prove a form of *progress*: i.e., in the absence of internal adaptation events (only external adaptation), every input monitored process will eventually receive a message and conversely every message in a queue will eventually be received by an input monitored process.¹ The proof follows the one presented in [CDV15, CDYP16], which crucially depends on the assumption that a process owns only one session channel.

We now extend *safety* and *consistency* properties to the multiparty case. Similarly as before we define κ -redexes, and *error processes*. Major changes are due to the fact that here we consider asynchronous communication.

Definition 5.16 (κ -redexes, errors). A network N contains a κ -redex if it is of one of the following forms:

- (a) $\mathcal{E}[\text{loc}_{\mathfrak{p}}[\mathcal{M}_{in}[\kappa[\mathfrak{p}]?\{\lambda_i(x_i).P_i\}_{i \in I}; \mathcal{P}_{\mathfrak{p}}; \text{loc}_{\mathfrak{p}}[r]] \mid \text{loc}_{\mathfrak{q}}[\mathcal{M}_o[\kappa[\mathfrak{q}]!\lambda(v).P]; \mathcal{P}_{\mathfrak{q}}; \text{loc}_{\mathfrak{q}}[r']]]]$
- (b) $\mathcal{E}[\text{loc}[\mathcal{M}_{in}[\kappa[\mathfrak{p}]?\{\lambda_i(x_i).P_i\}_{i \in I}; \mathcal{P}; \text{loc}[r]] \mid \kappa : (\mathfrak{q}, \mathfrak{p}, \lambda(v)) \cdot h]$

where $\mathcal{M}_{in} = p?\{\lambda_i(S_i).\mathcal{M}_i\}_{i \in I}$, $\mathcal{M}_o = p!\lambda(S).\mathcal{M}'$ and there exists $j \in I$ such that $\lambda_j = \lambda$ and $S_j = S$.

A network $N \equiv (\nu \tilde{\kappa})(N')$ is an *error* if N' does not contain a κ -redex.

The introduction of internal updates breaks the update-consistency property as stated in §4. This is expected, as internal updates are meant to globally change the communication protocol agreed upon session establishment. Nonetheless, as external updates behave exactly as in the binary case, the property is preserved for them. Below we write $N \longrightarrow_{\text{upd}} N'$ for any reduction inferred using Rule $\langle \text{N:UPDLOC1} \rangle$ (cf. Table 7). We may then define:

Definition 5.17 (Safety and Consistency). Let N be a network.

- We say N is *safe* if it never reduces into an error.
- We say N is *external update-consistent* if and only if, for all N_0 and κ such that $N \longrightarrow^* N_0$ and N_0 contains a κ -redex, if $N_0 \longrightarrow_{\text{upd}} N_1$ then N_1 contains the same κ -redex.

As before our main result follows as a consequence of Theorem 5.15.

Lemma 5.18. If $\Gamma \vdash_{\Sigma} N \triangleright \Delta$ with Δ balanced then network N contains no error.

Proof. Follows directly from Definition 5.11. \square

Theorem 5.19 (Typing Ensures Safety and Consistency). If $\Gamma \vdash_{\Sigma} N \triangleright \Delta$ with Δ balanced then network N is safe and external update consistent.

Proof. Safety means that N never reduces into an error. Since Δ is balanced by Lemma 5.18 we know that N is not an error. Theorem 5.15 ensures that reduction preserves balanced typings; therefore, N never reduces to an error.

Proving update-consistency entails showing that for all κ such that N_0 contains a κ -redex, and $N \longrightarrow^* N_0$, if $N_0 \longrightarrow_{\text{upd}} N_1$ then N_1 contains the same κ -redex. By Definition 5.16 N_0 , there are two cases. We consider only the case (a), i.e., the redex is of the form

$$\mathcal{E}[\text{loc}_{\mathfrak{p}}[\mathcal{M}_{in}[\kappa[\mathfrak{p}]?\{\lambda_i(x_i).P_i\}_{i \in I}; \mathcal{P}_{\mathfrak{p}}; \text{loc}_{\mathfrak{p}}[r]] \mid \text{loc}_{\mathfrak{q}}[\mathcal{M}_o[\kappa[\mathfrak{q}]!\lambda(v).P]; \mathcal{P}_{\mathfrak{q}}; \text{loc}_{\mathfrak{q}}[r']]]]$$

as case (b) is similar. If Rule $\langle \text{N:UPDLOC1} \rangle$ is applied either $\text{loc}_{\mathfrak{p}}$ or $\text{loc}_{\mathfrak{q}}$ are updated. Without loss of generality suppose that $\text{loc}_{\mathfrak{p}}$ is updated. Now function match guarantees that the type of the updated process P located in $\text{loc}_{\mathfrak{p}}$ is kept unchanged. As P only implements the behavior of participant \mathfrak{p} , by inversion on typing we know that the κ -process needs to be present also in the updated process. Thus we can conclude that the κ -redex is still present. \square

¹ The restriction to external adaptation is expected, as internal updates are meant to modify the protocol for current participants, and so sent messages could be left undelivered.

6. Related Work

Many previous works have investigated formal approaches to dynamic/run-time adaptation in different settings; below we describe some of these efforts.

An early work devoted to software process evolution is [JC93], which describes the EPOS approach to evolution and customization. Correctness criteria for adaptive workflows are surveyed in [RRD04]. In the setting of workflows, modifications may occur at the level of schemas or at the level of individual workflow instances. It is natural to represent workflows as (Petri) nets with special characteristics; as such, the focus of the comparison in [RRD04] is on approaches based on Petri nets with different semantics. Two correctness criteria compared are graph equivalence (focused on changes in schemas) and trace equivalence (focused on changes in instances). The work [MGR04] describes a workflow management system that supports rule-based adaptation; it provides automated support for predictive and reactive adaptation strategies. The modeling approach and checked properties in these works are considerably different from those in our work.

The paper [Fox11] develops a model for dynamically adaptive services using COWS (the Calculus for Orchestration of Web Services). In particular, a model for adaptation managers is proposed; it exploits the constructs that COWS provides for timed behaviors, and the associated model checker CMC. The proposed model enables structural and functional modifications: while the former concerns the system configuration graph, the latter involves component replacement. Services are assessed in terms of responsiveness, availability, and reliability. This work does not appeal to static verification based on behavioral types. Also, properties related to protocol conformance, such as safety or consistency, are not explicitly considered in [Fox11].

To the best of our knowledge our work develops the first application of constructs for type-directed checks and event-based communication for specifying and analyzing run-time adaptation for session-based concurrent systems. In fact, although such constructs have been proposed in previous works, their application for run-time adaptation seems to be a new contribution of our work. As such, our work develops an unexplored perspective for existing elements in session-based languages with the aim of enhancing models and reasoning techniques for communication-centric systems which may be updated at run-time.

The combination of static typing and type-directed tests for dynamic reconfiguration is not new. For instance, the work [SC06] studies this combination in the very different setting of object-oriented component programming. As already discussed, we build upon constructs proposed in [HKP⁺10, KYH11, KYHH16, Kou12] for defining type-directed checks and expressing eventful sessions. The earliest works on eventful sessions, covering theory and implementation issues, are [HKP⁺10, KYH11]. Kouzapas's PhD thesis [Kou12] provides a unified presentation of the eventful session framework, with case studies including event selectors (a building block in event-driven systems) and transformations between multithreaded and event-driven programs. At the level of types, the work in [Kou12] introduces session set types to support the `typecase` construct. In contrast, we use dynamic type inspection only for run-time adaptation; in [Kou12] `typecase` is part of the process syntax. This choice enables us to retain a standard session type syntax.

Previous works [DP15, AR12, BCH⁺13, CDV15, CDP14, DGL⁺14] have addressed forms of run-time adaptation for models of communicating systems based on binary and multiparty session types. As we elaborate next, none of them features the distinctive constructs of our framework: adaptation routines based on type-directed checks on session protocols and eventful constructs that handle and trigger such routines.

We have already mentioned how the present development was motivated as an enhancement to our previous work [DP15], in which update processes do not consider type-directed checks and adaptation of located processes with running sessions is disallowed by typing. While [DP15] addresses binary session types, all of [AR12, BCH⁺13, CDV15, CDP14, DGL⁺14] concern models based on multiparty sessions and/or choreographies. The work [AR12] studies dynamic update for message passing programs; a form of consistency for updates over threads is ensured using multiparty session types, following an asynchronous communication discipline. The paper [BCH⁺13] develops a model of choreographies with interleaved sessions and adaptation as well as an associated endpoint language; however, typing is not addressed. The key ideas of the self-adaptable multiparty sessions in [CDV15] have been summarized in § 2. Our proposal for event-based adaptation for multiparty protocols, described in § 5, departs from the approach in [CDV15] and extends it to cover forms of unanticipated adaptation, including internal and external adaptation requests. Another work derived from [CDV15] is [CDP14] in which adaptation of multiparty protocols is coupled with mechanisms for enforcing secure information flow and access control. As such, it is both technically and conceptually different from the framework that we have developed in § 5. Finally, the paper [DGL⁺14] describes a framework for programming distributed adaptive applications. The framework relies on a language for choreographies in which adaptation is specified following a rule-based approach.

The work [BCD⁺13] develops a *monitored* variant of the π -calculus and equips it with a combination of static and dynamic verification techniques based on multiparty session types. The goal is to ensure that a collection of decentralized and possibly monitored local processes properly implement a protocol given as a global type. As in our work, in [BCD⁺13] we find located processes which represent *principals*; a *network* is a set of principals with a *global queue*—a global transport. The semantics of global specifications induces *assertion environments* which in turn are essential to define the behavior of monitored networks: a network may performed actions allowed by its monitor. Key properties include transparency and safety: while transparency says that a monitored process/network behaves as an unmonitored (but well-behaved) process/network, safety says that a protocol (local and global) is respected.

7. Concluding Remarks

In this paper, we have introduced a novel approach to run-time adaptation of communicating processes whose structured protocols adhere to a session type discipline. Our study has addressed the case of protocols abstracted as binary sessions; as a proof of concept, we have also explored the case of adaptation in multiparty sessions with asynchronous communication.

Our approach is based on two kinds of constructs: (i) an update process construct based on type-directed checks and (ii) on constructs for eventful sessions, namely queues equipped with an arrival predicate. This combination of constructs allows us to specify dynamic modifications on the behavior of session-typed processes, therefore providing a basis for reasoning about expressive models of communication-centric systems with run-time adaptation. We may specify *what* should be the content of an adaptation routine, but also *when* it should be triggered. The constructs on which our approach relies are not new to this paper, but have been introduced and thoroughly studied by Kouzapas et al. in a series of works [HKP⁺10, KYH11, KYHH16, Kou12]. Such works investigate foundational and practical issues for eventful binary sessions; in our view, this adds significant value to our proposal, as the key syntactic elements on which it stands have been already validated (in theory and in practice) by previous independent studies.

The case of binary sessions represents a well-studied and representative setting, in which the distinguishing aspects of our approach can be cleanly introduced. We identified the strictly necessary eventful process constructs that enhance and refine the mechanisms for run-time adaptation given in [DP13a, DP15]. Adaptation requests are handled via event detectors and queues associated to locations. Our approach enables us to specify rich forms of updates on locations with running sessions; this is a concrete improvement with respect to [DP15], in which updates are only allowed when locations do not enclose running sessions. To rule out update steps that jeopardize running session protocols, we also introduced a type system that ensures communication safety and update consistency for session programs.

Concerning our proof of concept for multiparty sessions, we described how our approach can be integrated on top of the framework for self-adaptable multiparty, asynchronous sessions by Coppo et al. [CDV15]. In this more general setting we may specify internal and external adaptation requests; we rely on a process model which organizes communication and adaptation components in terms of on distributed locations. We notice that expressing both internal and external exceptional events is useful in practice; for instance, both kinds of events coexist in BPMN 2.0 (see, e.g., [DRMR13, Chap.4]).

Directions for future work include validating our approach in concrete case studies and different settings. We intend to revisit the model of *supervision trees* (a mechanism for fault-tolerance in Erlang) that we gave in [DP13b]. Also, it would be interesting to explore if our approach based on events can be harmonized with the model of security-driven adaptation that have been recently developed in [CDP14], in which read and write violations to security policies determine adaptation events. Finally, it would be worth accommodating existing static techniques for progress (deadlock-freedom) into our typing disciplines for run-time adaptation. We see consistency and progress as orthogonal guarantees; we do not find major obstacles preventing an extension of our typing systems with the existing techniques such as those developed in, e.g., [VV13, PVV14].

Acknowledgments. We are grateful to Ilaria Castellani, Mariangiola Dezani-Ciancaglini, and to the anonymous reviewers, whose remarks and suggestions helped us much to improve the paper. This research was partially supported by a Short-Term Scientific Mission grant to Di Giusto from COST Action IC1201: Behavioural Types for Reliable Large-Scale Software Systems. Pérez is also affiliated to the NOVA Laboratory for Computer Science and Informatics (NOVA LINCS Ref. UID/CEC/04516/2013), Universidade Nova de Lisboa, Portugal.

References

- [AR12] Gabrielle Anderson and Julian Rathke. Dynamic software update for message passing programs. In Ranjit Jhala and Atsushi Igarashi, editors, *Programming Languages and Systems - 10th Asian Symposium, APLAS 2012, Kyoto, Japan, December 11-13, 2012. Proceedings*, volume 7705 of *Lecture Notes in Computer Science*, pages 207–222. Springer, 2012.
- [BCD⁺13] Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. In Dirk Beyer and Michele Boreale, editors, *Formal Techniques for Distributed Systems - Joint IFIP WG 6.1 International Conference, FMOODS/FORTE 2013, Held as Part of the 8th International Federated Conference on Distributed Computing Techniques, DisCoTec 2013, Florence, Italy, June 3-5, 2013. Proceedings*, volume 7892 of *Lecture Notes in Computer Science*, pages 50–65. Springer, 2013.
- [BCH⁺13] Mario Bravetti, Marco Carbone, Thomas T. Hildebrandt, Ivan Lanese, Jacopo Mauro, Jorge A. Pérez, and Gianluigi Zavattaro. Towards global and local types for adaptation. In Steve Counsell and Manuel Núñez, editors, *Software Engineering and Formal Methods - SEFM 2013 Collocated Workshops*, volume 8368 of *Lecture Notes in Computer Science*, pages 3–14. Springer, 2013.
- [BDPZ12] Mario Bravetti, Cinzia Di Giusto, Jorge A. Pérez, and Gianluigi Zavattaro. Adaptable Processes. *Logical Methods in Computer Science*, 8(4), 2012.
- [CDP14] Ilaria Castellani, Mariangiola Dezani-Ciancaglini, and Jorge A. Pérez. Self-adaptation and secure information flow in multiparty structured communications: A unified perspective. In Marco Carbone, editor, *Proceedings Third Workshop on Behavioural Types, BEAT 2014, Rome, Italy, 1st September 2014.*, volume 162 of *EPTCS*, pages 9–18, 2014.
- [CDV15] Mario Coppo, Mariangiola Dezani-Ciancaglini, and Betti Venneri. Self-adaptive multiparty sessions. *Service Oriented Computing and Applications*, 9(3-4):249–268, 2015.
- [CDYP16] Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. Global progress for dynamically interleaved multiparty sessions. *Mathematical Structures in Computer Science*, 26(2):238–302, 2016.
- [DGL⁺14] Mila Dalla Preda, Saverio Giallorenzo, Ivan Lanese, Jacopo Mauro, and Maurizio Gabbriellini. AIOJ: A choreographic framework for safe adaptive distributed applications. In Benoît Combemale, David J. Pearce, Olivier Barais, and Jurgen J. Vinju, editors, *Software Language Engineering - 7th International Conference, SLE 2014*, volume 8706 of *Lecture Notes in Computer Science*, pages 161–170. Springer, 2014.
- [DP13a] Cinzia Di Giusto and Jorge A. Pérez. Disciplined structured communications with consistent runtime adaptation. In Sung Y. Shin and José Carlos Maldonado, editors, *Proceedings of the 28th Annual ACM Symposium on Applied Computing, SAC '13, Coimbra, Portugal, March 18-22, 2013*, pages 1913–1918. ACM, 2013.
- [DP13b] Cinzia Di Giusto and Jorge A. Pérez. Session types with runtime adaptation: Overview and examples. In *PLACES*, volume 137 of *EPTCS*, pages 21–32, 2013.
- [DP15] Cinzia Di Giusto and Jorge A. Pérez. Disciplined structured communications with disciplined runtime adaptation. *Sci. Comput. Program.*, 97:235–265, 2015.
- [DP16] Cinzia Di Giusto and Jorge A. Pérez. An Event-Based Approach to Runtime Adaptation in Communication-Centric Systems. In *Post-proceedings of WS-FM:FASOCC 2014 – 11th International Workshop on Web Services and Formal Methods: Formal Aspects of Service-Oriented and Cloud Computing*, Lecture Notes in Computer Science. Springer, 2016. To appear.
- [DRMR13] Marlon Dumas, Marcello La Rosa, Jan Mendling, and Hajo A. Reijers. *Fundamentals of Business Process Management*. Springer, 2013.
- [Fox11] Jorge Fox. A formal orchestration model for dynamically adaptable services with COWS. In *Proceedings International Conference on Adaptive and Self-Adaptive Systems and Applications (ADAPTIVE 2011)*, IARIA, pages 67–72, 2011.
- [GH05] Simon J. Gay and Malcolm Hole. Subtyping for session types in the pi calculus. *Acta Inf.*, 42(2-3):191–225, 2005.
- [HKP⁺10] Raymond Hu, Dimitrios Kouzapas, Olivier Pernet, Nobuko Yoshida, and Kohei Honda. Type-safe eventful sessions in java. In Theo D'Hondt, editor, *ECOOP 2010 - Object-Oriented Programming, 24th European Conference, Maribor, Slovenia, June 21-25, 2010. Proceedings*, volume 6183 of *Lecture Notes in Computer Science*, pages 329–353. Springer, 2010.
- [HVK98] Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In Chris Hankin, editor, *Programming Languages and Systems - ESOP'98, 7th European Symposium on Programming, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings*, volume 1381 of *Lecture Notes in Computer Science*, pages 122–138. Springer, 1998.
- [HYC08] Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In George C. Necula and Philip Wadler, editors, *POPL*, pages 273–284. ACM, 2008.
- [JC93] Maria Letizia Jaccheri and Reidar Conradi. Techniques for process model evolution in EPOS. *IEEE Trans. Software Eng.*, 19(12):1145–1156, 1993.
- [Kou12] Dimitrios Kouzapas. *A Study of Bisimulation Theory for Session Types*. PhD thesis, Imperial College London, 2012.
- [KYH11] Dimitrios Kouzapas, Nobuko Yoshida, and Kohei Honda. On asynchronous session semantics. In Roberto Bruni and Jürgen Dingel, editors, *Formal Techniques for Distributed Systems - Joint 13th IFIP WG 6.1 International Conference, FMOODS 2011, and 31st IFIP WG 6.1 International Conference, FORTE 2011, Reykjavik, Iceland, June 6-9, 2011. Proceedings*, volume 6722 of *Lecture Notes in Computer Science*, pages 228–243. Springer, 2011.
- [KYHH16] Dimitrios Kouzapas, Nobuko Yoshida, Raymond Hu, and Kohei Honda. On asynchronous eventful session semantics. *Mathematical Structures in Computer Science*, 26(2):303–364, 2016.

- [MGR04] Robert Müller, Ulrike Greiner, and Erhard Rahm. `AgentWork`: a workflow system supporting rule-based workflow adaptation. *Data Knowl. Eng.*, 51(2):223–256, 2004.
- [MPW92] Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, i. *Inf. Comput.*, 100(1):1–40, 1992.
- [Pad10] Luca Padovani. Session types = intersection types + union types. In Elaine Pimentel, Betti Venneri, and Joe B. Wells, editors, *Proceedings Fifth Workshop on Intersection Types and Related Systems, ITRS 2010, Edinburgh, U.K., 9th July 2010.*, volume 45 of *EPTCS*, pages 71–89, 2010.
- [PVV14] Luca Padovani, Vasco Thudichum Vasconcelos, and Hugo Torres Vieira. Typing liveness in multiparty communicating systems. In eva Kühn and Rosario Pugliese, editors, *Coordination Models and Languages - 16th IFIP WG 6.1 International Conference, COORDINATION 2014, Held as Part of the 9th International Federated Conferences on Distributed Computing Techniques, DisCoTec 2014, Berlin, Germany, June 3-5, 2014, Proceedings*, volume 8459 of *Lecture Notes in Computer Science*, pages 147–162. Springer, 2014.
- [RRD04] Stefanie Rinderle, Manfred Reichert, and Peter Dadam. Correctness criteria for dynamic changes in workflow systems - a survey. *Data Knowl. Eng.*, 50(1):9–34, 2004.
- [SC06] João Costa Seco and Luís Caires. Types for dynamic reconfiguration. In Peter Sestoft, editor, *Programming Languages and Systems, 15th European Symposium on Programming, ESOP 2006, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2006, Vienna, Austria, March 27-28, 2006, Proceedings*, volume 3924 of *Lecture Notes in Computer Science*, pages 214–229. Springer, 2006.
- [VV13] Hugo Torres Vieira and Vasco Thudichum Vasconcelos. Typing progress in communication-centred systems. In Rocco De Nicola and Christine Julien, editors, *Coordination Models and Languages, 15th International Conference, COORDINATION 2013, Held as Part of the 8th International Federated Conference on Distributed Computing Techniques, DisCoTec 2013, Florence, Italy, June 3-5, 2013. Proceedings*, volume 7890 of *Lecture Notes in Computer Science*, pages 236–250. Springer, 2013.

A. Supplementary Definitions

A.1. Structural Congruence

Definition A.1 (Structural Congruence). Structural congruence is the smallest congruence relation on processes that is generated by the following laws:

$$\begin{array}{ll}
P|Q \equiv Q|P & (P|Q)|R \equiv P|(Q|R) \\
P|\mathbf{0} \equiv P & P \equiv Q \text{ if } P \equiv_{\alpha} Q \\
(\nu s)\mathbf{0} \equiv \mathbf{0} & (\nu s)(\nu s')P \equiv (\nu s')(\nu s)P \\
(\nu s)P|Q \equiv (\nu s)(P|Q) \text{ (if } s \notin \text{fc}(Q) \cup \text{fn}(Q)) & (\nu s)\text{loc}[P] \equiv \text{loc}[(\nu s)P]
\end{array}$$

with s, s', \dots ranges over both names and session channels. The extension of \equiv to contexts is as expected.

A.2. Coinductive Subtyping and Duality

For all types, define $\mathbf{unfold}(T)$ by recursion on the structure of T :

$$\mathbf{unfold}(\mu t.T) = \mathbf{unfold}(T[\mu t.T/t])$$

and $\mathbf{unfold}(T) = T$ otherwise.

The following two definitions (*duality relations* and *type simulation*) are used by Definition A.4. Given an index set $I = \{1, \dots, m\}$, we use $\&\{n_i : T_i\}_{i \in I}$ and $\oplus\{n_i : T_i\}_{i \in I}$ to abbreviate $\&\{n_1 : T_1, \dots, n_m : T_m\}$ and $\oplus\{n_1 : T_1, \dots, n_m : T_m\}$, respectively.

Definition A.2 (Duality Relation). A relation $\mathcal{R} \subseteq \mathcal{T} \times \mathcal{T}$ is a *duality relation* if $(T, S) \in \mathcal{R}$ implies the following conditions:

1. If $\mathbf{unfold}(T) = \tau$ then $\mathbf{unfold}(S) = \sigma$ and $\tau \leq_c \sigma$ and $\sigma \leq_c \tau$.
2. If $\mathbf{unfold}(T) = \mathbf{end}$ then $\mathbf{unfold}(S) = \mathbf{end}$.
3. If $\mathbf{unfold}(T) = ?(T_2).T_1$ then $\mathbf{unfold}(S) = !(S_2).S_1$ and $(T_1, S_1) \in \mathcal{R}$ and $T_2 \leq_c S_2$ and $S_2 \leq_c T_2$.
4. If $\mathbf{unfold}(T) = !(T_2).T_1$ then $\mathbf{unfold}(S) = ?(S_2).S_1$ and $(T_1, S_1) \in \mathcal{R}$ and $T_2 \leq_c S_2$ and $S_2 \leq_c T_2$.
5. If $\mathbf{unfold}(T) = ?(\tau_1, \dots, \tau_n).T_1$ then $\mathbf{unfold}(S) = ?(\sigma_1, \dots, \sigma_n).S_1$ then for all $i \in [1..n]$, we have that $(T_1, S_1) \in \mathcal{R}$ and $\tau_i \leq_c \sigma_i$ and $\sigma_i \leq_c \tau_i$.
6. If $\mathbf{unfold}(T) = !(\tau_1, \dots, \tau_n).T_1$ then $\mathbf{unfold}(S) = ?(\sigma_1, \dots, \sigma_n).S_1$ then for all $i \in [1..n]$, we have that $(T_1, S_1) \in \mathcal{R}$ and $\tau_i \leq_c \sigma_i$ and $\sigma_i \leq_c \tau_i$.
7. If $\mathbf{unfold}(T) = \&\{n_1 : T_1 \dots n_m : T_m\}$ then $\mathbf{unfold}(S) = \oplus\{n_1 : S_1 \dots n_m : S_m\}$ and for all $i \in [1..m]$, we have that $(T_i, S_i) \in \mathcal{R}$.
8. If $\mathbf{unfold}(T) = \oplus\{n_1 : T_1 \dots n_m : T_m\}$ then $\mathbf{unfold}(S) = \&\{n_1 : S_1 \dots n_m : S_m\}$ and for all $i \in [1..m]$, we have that $(T_i, S_i) \in \mathcal{R}$.

Definition A.3 (Type Simulation). A relation $\mathcal{R} \subseteq \mathcal{T} \times \mathcal{T}$ is a *type simulation* if $(T, S) \in \mathcal{R}$ implies the following conditions:

1. If $\mathbf{unfold}(T) = \tau$ then $\mathbf{unfold}(S) = \sigma$ and $\tau \leq_B \sigma$.
2. If $\mathbf{unfold}(T) = \mathbf{end}$ then $\mathbf{unfold}(S) = \mathbf{end}$.
3. If $\mathbf{unfold}(T) = ?(T_2).T_1$ then $\mathbf{unfold}(S) = ?(S_2).S_1$ and $(T_1, S_1) \in \mathcal{R}$ and $(T_2, S_2) \in \mathcal{R}$.
4. If $\mathbf{unfold}(T) = !(T_2).T_1$ then $\mathbf{unfold}(S) = !(S_2).S_1$ and $(T_1, S_1) \in \mathcal{R}$ and $(S_2, T_2) \in \mathcal{R}$.
5. If $\mathbf{unfold}(T) = ?(\tau_1, \dots, \tau_n).T_1$ then $\mathbf{unfold}(S) = ?(\sigma_1, \dots, \sigma_n).S_1$ then for all $i \in [1..n]$, we have that $(\tau_i, \sigma_i) \in \mathcal{R}$ and $(T_1, S_1) \in \mathcal{R}$.
6. If $\mathbf{unfold}(T) = !(\tau_1, \dots, \tau_n).T_1$ then $\mathbf{unfold}(S) = !(\sigma_1, \dots, \sigma_n).S_1$ then for all $i \in [1..n]$, we have that $(\sigma_i, \tau_i) \in \mathcal{R}$ and $(T_1, S_1) \in \mathcal{R}$.
7. If $\mathbf{unfold}(T) = \oplus\{n_i : T_i\}_{i \in I}$ then $\mathbf{unfold}(S) = \oplus\{n_j : S_j\}_{j \in J}$ and $I \subseteq J$ for all $i \in I$, we have that $(T_i, S_i) \in \mathcal{R}$.
8. If $\mathbf{unfold}(T) = \&\{n_i : T_i\}_{i \in I}$ then $\mathbf{unfold}(S) = \&\{n_j : S_k\}_{j \in J}$ and $J \subseteq I$ for all $j \in J$, we have that $(T_j, S_j) \in \mathcal{R}$.

Based on the above definitions, we can now define:

Definition A.4 (Coinductive Duality and Subtyping). Let T, S be types.

- The *coinductive duality relation*, denoted \perp_c , is defined by $T \perp_c S$ if and only if there exists a duality relation \mathcal{R} such that $(T, S) \in \mathcal{R}$.
- The *coinductive subtyping relation*, denoted \leq_c , is defined by $T \leq_c S$ if and only if there exists a type simulation \mathcal{S} such that $(T, S) \in \mathcal{S}$. The extension of \leq_c to typings, written $\Delta \leq_c \Delta'$, arises as expected.

B. Omitted Proofs

B.1. Binary Case

The following auxiliary result concerns substitutions for channels, expressions, and process variables. Observe how the case of process variables has been relaxed so as to allow substitution with a process with “smaller” interface (in the sense of \sqsubseteq , cf. Definition 4.2). This extra flexibility is in line with the typing rule for located processes (cf. Rule $\langle T:\text{LOC} \rangle$, Table 3), and will be useful later on in proofs.

Lemma B.1 (Substitution Lemma).

1. If $\Gamma; \Theta \vdash P \triangleright \Delta, x : \alpha; \mathcal{I}$ then $\Gamma; \Theta \vdash P[\kappa^p/x] \triangleright \Delta, \kappa^p : \alpha; \mathcal{I}$.
2. If $\Gamma, x : \tau; \Theta \vdash P \triangleright \Delta; \mathcal{I}$ and $\Gamma \vdash e \triangleright \tau$ then $\Gamma; \Theta \vdash P[e/x] \triangleright \Delta; \mathcal{I}$.

Proof. Easily shown by induction on the structure of P . \square

As reduction may occur inside contexts, in proofs it is useful to have *typed contexts*. These are contexts in which the hole has associated typing information—concretely, the typing for processes which may fill in the hole. Defining context requires a simple extension of judgments, in the following way:

$$\mathcal{H}; \Gamma; \Theta \vdash C \triangleright \Delta; \mathcal{I}$$

Intuitively, \mathcal{H} contains the description of the type associated to the hole in C . Typing rules are extended in the expected way. Because contexts have a single hole, \mathcal{H} is either empty or has exactly one element. When \mathcal{H} is empty, we write $\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I}$ instead of $\cdot; \Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I}$. Two additional typing rules are required:

$$\langle T:\text{HOLE} \rangle \frac{}{\bullet_{\Gamma; \Theta \vdash \Delta; \mathcal{I}}; \Gamma; \Theta \vdash \bullet \triangleright \Delta; \mathcal{I}} \quad \langle T:\text{FILL} \rangle \frac{\bullet_{\Gamma; \Theta \vdash \Delta; \mathcal{I}}; \Gamma; \Theta \vdash C \triangleright \Delta_1; \mathcal{I}_1 \quad \Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I}}{\Gamma; \Theta \vdash C\{P\} \triangleright \Delta_1; \mathcal{I}_1}$$

Axiom $\langle T:\text{HOLE} \rangle$ allows us to introduce typed holes into contexts. In Rule $\langle T:\text{FILL} \rangle$, P is a process (it does not have any holes), and C is a context with a hole of type $\Gamma; \Theta \vdash \Delta; \mathcal{I}$. The substitution of occurrences of \bullet in C with P , noted $C\{P\}$, is sound as long as the typings of P coincide with those declared in \mathcal{H} for C . We introduce some convenient notation for typed holes.

Notation B.2. Let us use $\mathcal{S}, \mathcal{S}', \dots$ to range over judgments attached to typed holes. This way, $\bullet_{\mathcal{S}}$ denotes the valid typed hole associated to $\mathcal{S} = \Gamma; \Theta \vdash \Delta; \mathcal{I}$.

Lemma B.3. Let P and C be a process and a typed context such that

$$\Gamma; \Theta \vdash C\{P\} \triangleright \Delta; \mathcal{I}$$

is a derivable judgment. There exist Δ_1, \mathcal{I}_1 such that (i) $\Gamma; \Theta \vdash P \triangleright \Delta_1; \mathcal{I}_1$ is a well-typed process, and (ii) $\Delta_1 \sqsubseteq \Delta$ and $\mathcal{I}_1 \sqsubseteq \mathcal{I}$.

Lemma B.4. Let C be a context. Suppose $\bullet_{\mathcal{S}}; \Gamma; \Theta \vdash C \triangleright \Delta_C \cup \Delta_{\mathcal{S}}; \mathcal{I}_C \uplus \mathcal{I}_{\mathcal{S}}$ with $\mathcal{S} = \Gamma; \Theta \vdash \Delta_{\mathcal{S}}; \mathcal{I}_{\mathcal{S}}$ is well-typed. Let $\mathcal{S}' = \Gamma; \Theta \vdash \Delta_{\mathcal{S}'}; \mathcal{I}_{\mathcal{S}'}$. Then

$$\bullet_{\mathcal{S}'}; \Gamma; \Theta \vdash C \triangleright \Delta_C \cup \Delta_{\mathcal{S}'}; \mathcal{I}_C \uplus \mathcal{I}_{\mathcal{S}'}$$

is a derivable judgment.

Theorem B.5 (Subject Congruence). If $\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I}$ and $P \equiv Q$ then $\Gamma; \Theta \vdash Q \triangleright \Delta; \mathcal{I}$.

Proof. By induction on the derivation of $P \equiv Q$, with a case analysis on the last applied rule. \square

Theorem B.6 (Subject Reduction - 4.6). If $\Gamma; \Theta \vdash P \triangleright \Delta; \mathcal{I}$ with Δ balanced and $P \longrightarrow Q$ then $\Gamma; \Theta \vdash Q \triangleright \Delta'; \mathcal{I}'$, for some $\mathcal{I}' \subseteq \mathcal{I}$ and balanced Δ' such that either $\Delta' \leq_c \Delta$ or $\Delta \mapsto \Delta'$.

Proof. By induction on the last rule applied in the reduction. We assume that $e \downarrow c$ is a type preserving operation, for every e . We examine only a few interesting cases, namely those for session establishment, runtime update, and intra-session communication; remaining cases are similar or simpler.

1. Case $\langle R:OPEN \rangle$: From Table 2 we have:

$$C\{u(x : \alpha).P_1\} \mid D\{\bar{u}(y : \beta).P_2\} \longrightarrow (\nu\kappa)(C\{P_1[\kappa^+/x] \mid \kappa^+[\alpha]\} \mid D\{P_2[\kappa^-/y] \mid \kappa^-[\beta]\})$$

with $\alpha \perp_c \beta$. By assumption $\Gamma; \Theta \vdash C\{u(x : \alpha).P_1\} \mid D\{\bar{u}(y : \beta).P_2\} \triangleright \Delta; \mathcal{I}$ with balanced Δ . Then, by inversion on typing, using Rules $\langle T:ACCEPT \rangle$, $\langle T:REQUEST \rangle$, and $\langle T:PAR \rangle$ we infer there exist Δ', \mathcal{I}' such that

$$\frac{(8) \quad (9)}{\Gamma; \Theta \vdash C\{u(x : \alpha).P_1\} \mid D\{\bar{u}(y : \beta).P_2\} \triangleright \Delta; \mathcal{I}}$$

where, letting $\Delta = \Delta'_1 \cup \Delta'_2$, subtree (8) is as follows:

$$\frac{\bullet_{S_1}; \Gamma; \Theta \vdash C \triangleright \Delta'_1; \mathcal{I}'_1 \uplus u : \alpha_{\text{lin}} \quad \frac{\alpha \leq_c \alpha' \quad \alpha' \perp_c \beta' \quad \Gamma \vdash u \triangleright \langle \alpha'_{\text{lin}}, \beta'_{\text{lin}} \rangle \quad \Gamma; \Theta \vdash P_1 \triangleright \Delta_1, x : \alpha; \mathcal{I}_1}{\Gamma; \Theta \vdash u(x : \alpha).P_1 \triangleright \Delta_1; \mathcal{I}_1 \uplus u : \alpha_{\text{lin}}}}{\Gamma; \Theta \vdash C\{u(x : \alpha).P_1\} \triangleright \Delta'_1; \mathcal{I}'_1 \uplus u : \alpha_{\text{lin}}}$$

with

$$S_1 = \Gamma; \Theta \vdash \Delta_1; \mathcal{I}_1 \uplus u : \alpha_{\text{lin}}$$

Then, subtree (9) is as follows:

$$\frac{\bullet_{S_2}; \Gamma; \Theta \vdash D \triangleright \Delta'_2; \mathcal{I}'_2 \uplus u : \beta_{\text{lin}} \quad \frac{\beta \leq_c \beta' \quad \alpha' \perp_c \beta' \quad \Gamma \vdash u \triangleright \langle \alpha_{\text{lin}}, \beta_{\text{lin}} \rangle \quad \Gamma; \Theta \vdash P_2 \triangleright \Delta_2, y : \beta; \mathcal{I}_2}{\Gamma; \Theta \vdash \bar{u}(y : \beta).P_2 \triangleright \Delta_2; \mathcal{I}_2 \uplus u : \beta_{\text{lin}}}}{\Gamma; \Theta \vdash D\{\bar{u}(y : \beta).P_2\} \triangleright \Delta'_2; \mathcal{I}'_2 \uplus u : \beta_{\text{lin}}}$$

with

$$S_2 = \Gamma; \Theta \vdash \Delta_2; \mathcal{I}_2 \uplus u : \beta_{\text{lin}}$$

By Lemma B.3 we have that $\Delta_1 \subseteq \Delta'_1$ and $\Delta_2 \subseteq \Delta'_2$. We also infer $\mathcal{I}_1 \sqsubseteq \mathcal{I}'_1$, $\mathcal{I}_2 \sqsubseteq \mathcal{I}'_2$, and $\mathcal{I}' \sqsubseteq \mathcal{I}$. Now, using Lemma B.1(1) on judgments for P_1 and P_2 , we obtain:

- (a) $\Gamma; \Theta \vdash P_1[\kappa^+/x] \triangleright \Delta_1, \kappa^+ : \alpha; \mathcal{I}_1$.
- (b) $\Gamma; \Theta \vdash P_2[\kappa^-/y] \triangleright \Delta_2, \kappa^- : \beta; \mathcal{I}_2$.

We may now reconstruct the derivation given in (8) using Lemma B.4 and Rule $\langle T:PAR \rangle$:

$$(11) \quad \frac{\frac{\Gamma; \Theta \vdash P_1[\kappa^+/x] \triangleright \Delta_1, \kappa^+ : \alpha; \mathcal{I}_1 \quad \Gamma; \Theta \vdash \kappa^+[\alpha] \triangleright \kappa^+ : [\alpha]; \emptyset}{\Gamma; \Theta \vdash P_1[\kappa^+/x] \mid \kappa^+[\alpha] \triangleright \Delta_1, \kappa^+ : \alpha, \kappa^+ : [\alpha]; \mathcal{I}_1}}{\Gamma; \Theta \vdash C\{P_1[\kappa^+/x] \mid \kappa^+[\alpha]\} \triangleright \Delta'_1, \kappa^+ : \alpha, \kappa^+ : [\alpha]; \mathcal{I}'_1}}$$

with

$$\bullet_{S_3}; \Gamma; \Theta \vdash C \triangleright \Delta'_1, \kappa^+ : \alpha, \kappa^+ : [\alpha]; \mathcal{I}'_1$$

and

$$S_3 = \Gamma; \Theta \vdash \Delta_1, \kappa^+ : \alpha, \kappa^+ : [\alpha]; \mathcal{I}_1$$

For D , we proceed analogously from (9) and infer:

$$(12) \quad \frac{\bullet_{S_4}; \Gamma; \Theta \vdash D \triangleright \Delta'_2, \kappa^- : \beta, \kappa^- : [\beta]; \mathcal{I}'_2 \quad \Gamma; \Theta \vdash P_2[\kappa^-/y] \mid \kappa^-[\beta] \triangleright \Delta_2, \kappa^- : \beta, \kappa^- : [\beta]; \mathcal{I}_2}{\Gamma; \Theta \vdash D\{P_2[\kappa^-/y] \mid \kappa^-[\beta]\} \triangleright \Delta'_2, \kappa^- : \beta, \kappa^- : [\beta]; \mathcal{I}'_2}}$$

with

$$\mathcal{S}_4 = \Gamma; \Theta \vdash \Delta_2, \kappa^- : \beta, \kappa^- : [\beta]; \mathcal{I}_2.$$

We may finally derive the type for the result of the reduction: using Rules $\langle \text{T:PAR} \rangle$ and $\langle \text{T:CREs} \rangle$ we obtain:

$$\frac{\frac{\Gamma; \Theta \vdash C\{P_1[\kappa^+/x] \mid \kappa^+[\alpha]\} \mid D\{P_2[\kappa^-/y] \mid \kappa^-[\beta]\} \triangleright \Delta, \kappa^+ : \alpha, \kappa^- : \beta, \kappa^+ : [\alpha], \kappa^- : [\beta]; \mathcal{I}'_1 \uplus \mathcal{I}'_2}{\Gamma; \Theta \vdash (\nu\kappa)C\{P_1[\kappa^+/x] \mid \kappa^+[\alpha]\} \mid D\{P_2[\kappa^-/y] \mid \kappa^-[\beta]\} \triangleright \Delta; \mathcal{I}'_1 \uplus \mathcal{I}'_2}}{\Gamma; \Theta \vdash C\{P_1[\kappa^+/x] \mid \kappa^+[\alpha]\} \mid D\{P_2[\kappa^-/y] \mid \kappa^-[\beta]\} \triangleright \Delta; \mathcal{I}'_1 \uplus \mathcal{I}'_2}} \quad (10) \quad (12)$$

Notice that Δ is unchanged. This concludes this case.

2. Case $\langle \text{R:UPD2} \rangle$: From Table 2 we have:

$$C\{\text{loc}[P]\} \mid D\{\text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_1^i; \dots; \beta_m^i) : Q_i\}_{i \in I}\}\} \longrightarrow C\{\text{loc}[V]\} \mid D\{\mathbf{0}\}$$

with $V = Q_l[\kappa_1^p, \dots, \kappa_m^p / x_1, \dots, x_m]$. By assumption we have

$$\Gamma; \Theta \vdash C\{\text{loc}[P]\} \mid D\{\text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_1^i; \dots; \beta_m^i) : Q_i\}_{i \in I}\}\} \triangleright \Delta; \mathcal{I}$$

with Δ balanced. Then, by inversion on typing, using Rules $\langle \text{T:FILL} \rangle$, $\langle \text{T:PAR} \rangle$, $\langle \text{T:ADAPT} \rangle$, and $\langle \text{T:LOC} \rangle$ we infer:

$$\frac{\Gamma; \Theta \vdash C\{\text{loc}[P]\} \mid D\{\text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_1^i; \dots; \beta_m^i) : Q_i\}_{i \in I}\}\} \triangleright \Delta; \mathcal{I}}{\Gamma; \Theta \vdash C\{\text{loc}[P]\} \mid D\{\text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_1^i; \dots; \beta_m^i) : Q_i\}_{i \in I}\}\} \triangleright \Delta; \mathcal{I}} \quad (13) \quad (14)$$

Let $\Delta = \Delta'_1 \cup \Delta'_2$ and $\mathcal{I} = \mathcal{I}'_1 \uplus \mathcal{I}'_2$, subtree (13) is as follows:

$$\frac{\frac{\frac{\bullet_{\mathcal{S}_1}; \Gamma; \Theta \vdash C \triangleright \Delta'_1; \mathcal{I}'_1}{\Gamma; \Theta \vdash C \triangleright \Delta'_1; \mathcal{I}'_1} \quad \frac{\mathcal{I}_1 \sqsubseteq \mathcal{I}'_1 \quad \Theta \vdash \text{loc} \triangleright \mathcal{I}'_1 \quad \Gamma; \Theta \vdash P_1 \triangleright \Delta_1; \mathcal{I}_1}{\Gamma; \Theta \vdash \text{loc}[P] \triangleright \Delta_1; \mathcal{I}_1}}{\Gamma; \Theta \vdash C\{\text{loc}[P]\} \triangleright \Delta'_1; \mathcal{I}'_1}}{\Gamma; \Theta \vdash C\{\text{loc}[P]\} \triangleright \Delta'_1; \mathcal{I}'_1}} \quad (13)$$

with $\mathcal{S}_1 = \Gamma; \Theta \vdash \Delta_1; \mathcal{I}_1$, and $\mathcal{I}_1 \sqsubseteq \mathcal{I}'_1$ (by Lemma B.3). Subtree (14) is as follows:

$$\frac{\frac{\bullet_{\mathcal{S}_2}; \Gamma; \Theta \vdash D \triangleright \Delta'_2; \mathcal{I}'_2 \quad \frac{\Theta \vdash \text{loc} \triangleright \mathcal{I} \quad \forall j \in J, \mathcal{I}_j \sqsubseteq \mathcal{I} \quad \Gamma; \Theta \vdash Q_i \triangleright \langle x_1 : \beta_1^j \rangle; \dots; \langle x_m : \beta_m^j \rangle; \mathcal{I}_j}{\Gamma; \Theta \vdash \text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_1^i; \dots; \beta_m^i) : Q_i\}_{i \in I}\} \triangleright \emptyset; \emptyset}}{\Gamma; \Theta \vdash D\{\text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_1^i; \dots; \beta_m^i) : Q_i\}_{i \in I}\}\} \triangleright \Delta'_2; \mathcal{I}'_2}}{\Gamma; \Theta \vdash D\{\text{loc}\{\text{case } x_1, \dots, x_m \text{ of } \{(\beta_1^i; \dots; \beta_m^i) : Q_i\}_{i \in I}\}\} \triangleright \Delta'_2; \mathcal{I}'_2}} \quad (14)$$

with $\mathcal{S}_2 = \Gamma; \Theta \vdash \emptyset; \emptyset$. We now reconstruct the derivation after the reduction, using Rules $\langle \text{T:PAR} \rangle$, $\langle \text{T:FILL} \rangle$, and Lemma B.4. By derivation (14) we know that

$$\Gamma; \Theta \vdash Q_l \triangleright \langle x_1 : \beta_1^l \rangle; \dots; \langle x_m : \beta_m^l \rangle; \mathcal{I}_l$$

thus applying Lemma B.1(1) we have:

$$\frac{\frac{\bullet_{\mathcal{S}_4}; \Gamma; \Theta \vdash D \triangleright \Delta'_2; \mathcal{I}'_2 \quad \Gamma; \Theta \vdash \mathbf{0} \triangleright \emptyset; \emptyset}{\Gamma; \Theta \vdash D\{\mathbf{0}\} \triangleright \Delta'_2; \mathcal{I}'_2}}{\Gamma; \Theta \vdash C\{V\} \mid D\{\mathbf{0}\} \triangleright \Delta'_1 \cup \Delta'_2; \mathcal{I}'_3 \uplus \mathcal{I}'_2}} \quad (16)$$

$$\frac{\bullet_{\mathcal{S}_5}; \Gamma; \Theta \vdash C \triangleright \Delta'_1; \mathcal{I}'_1 \quad \Gamma; \Theta \vdash V \triangleright \langle x_1 : \beta_1^l \rangle; \dots; \langle x_m : \beta_m^l \rangle; \mathcal{I}_l}{\Gamma; \Theta \vdash C\{V\} \triangleright \Delta'_1; \mathcal{I}'_1}} \quad (16)$$

with $\mathcal{S}_5 = \Gamma; \Theta \vdash \langle x_1 : \beta_1^l \rangle; \dots; \langle x_m : \beta_m^l \rangle; \mathcal{I}_l$. By Lemma B.3 we know $\mathcal{I}_l \sqsubseteq \mathcal{I}'_1$. Moreover by Lemma B.4, and by application of Rule $\langle \text{R:UPD2} \rangle$ we have $\Delta'_1 \leq_c \Delta'_1$. Thus $\Delta'_1 \cup \Delta'_2 \leq_c \Delta$. This concludes the analysis for this case.

3. Case $\langle \text{R:COM} \rangle$: From Table 2 we have:

$$C\{\overline{\kappa}^p(v).P_1 \mid \kappa^p[!(\tau).\alpha]\} \mid D\{\overline{\kappa}^p(x).P_2 \mid \kappa^p[?(\tau).\beta]\} \longrightarrow C\{P_1 \mid \kappa^p[\alpha]\} \mid D\{P_2[v/x] \mid \kappa^p[\beta]\}$$

By assumption, we have $\Gamma; \Theta \vdash C\{\overline{\kappa}^p(v).P_1 \mid \kappa^p[!(\tau).\alpha]\} \mid D\{\overline{\kappa}^p(x).P_2 \mid \kappa^p[?(\tau).\beta]\} \triangleright \Delta; \mathcal{I}$, with Δ balanced and $(\alpha \perp_c \beta)$. By inversion on typing, using Rules $\langle \text{T:FILL} \rangle$, $\langle \text{T:PAR} \rangle$, $\langle \text{T:IN} \rangle$, and $\langle \text{T:OUT} \rangle$, we infer:

$$\frac{(17) \quad (19)}{\Gamma; \Theta \vdash C\{\overline{\kappa}^p(v).P_1 \mid \kappa^p[!(\tau).\alpha]\} \mid D\{\overline{\kappa}^p(x).P_2 \mid \kappa^p[?(\tau).\beta]\} \triangleright \Delta'; \mathcal{I}'_1 \uplus \mathcal{I}'_2}$$

where:

$$\begin{aligned} \Delta &= \Delta'_1 \cup \Delta'_2, \kappa^p : !(\tau).\alpha, \kappa^p : [!(\tau).\alpha], \kappa^p : ?(\tau).\beta, \kappa^p : [?(\tau).\beta] \\ \mathcal{I} &= \mathcal{I}'_1 \uplus \mathcal{I}'_2 \end{aligned}$$

We have that subtree (17) is as follows:

$$(18) \quad \frac{\frac{\Gamma; \Theta \vdash P_1 \triangleright \Delta_1, \kappa^p : \alpha; \mathcal{I}_1 \quad \Gamma \vdash v : \tau}{\Gamma; \Theta \vdash \overline{\kappa}^p(v).P_1 \triangleright \Delta_1, \kappa^p : !(\tau).\alpha; \mathcal{I}_1} \quad \Gamma; \Theta \vdash \kappa^p[!(\tau).\alpha] \triangleright \kappa^p : [!(\tau).\alpha]; \emptyset}{\Gamma; \Theta \vdash C\{\overline{\kappa}^p(v).P_1 \mid \kappa^p[!(\tau).\alpha]\} \triangleright \Delta'_1, \kappa^p : !(\tau).\alpha, \kappa^p : [!(\tau).\alpha]; \mathcal{I}'_1}$$

with

$$\bullet_{\mathcal{S}_1}; \Gamma; \Theta \vdash C \triangleright \Delta'_1, \kappa^p : !(\tau).\alpha, \kappa^p : [!(\tau).\alpha]; \mathcal{I}'_1; \quad (18)$$

and $\mathcal{S}_1 = \Gamma; \Theta \vdash \Delta_1, \kappa^p : !(\tau).\alpha, \kappa^p : [!(\tau).\alpha]; \mathcal{I}_1$. Similarly, for subtree (19) we obtain (we show only the last step of the derivation):

$$(20) \quad \frac{\Gamma; \Theta \vdash \kappa^p(x).P_2 \mid \kappa^p[?(\tau).\beta] \triangleright \Delta_2, \kappa^p : ?(\tau).\beta, \kappa^p : [?(\tau).\beta]; \mathcal{I}_2}{\Gamma; \Theta \vdash D\{\overline{\kappa}^p(x).P_2 \mid \kappa^p[?(\tau).\beta]\} \triangleright \Delta'_2, \kappa^p : ?(\tau).\beta, \kappa^p : [?(\tau).\beta]; \mathcal{I}'_2}$$

with

$$\bullet_{\mathcal{S}_2}; \Gamma; \Theta \vdash D \triangleright \Delta'_2, \kappa^p : ?(\tau).\beta, \kappa^p : [?(\tau).\beta]; \mathcal{I}'_2 \quad (20)$$

and

$$\mathcal{S}_2 = \Gamma; \Theta \vdash \Delta_2, \kappa^p : ?(\tau).\beta, \kappa^p : [?(\tau).\beta]; \mathcal{I}_2$$

where Lemma B.3 ensures $\Delta_1 \subseteq \Delta'_1$, $\Delta_2 \subseteq \Delta'_2$.

Now, by Lemma B.1(2) we know $\Gamma; \Theta \vdash P_2[v/x] \triangleright \Delta_2, \kappa^p : \beta; \mathcal{I}_2$. Moreover by Lemma B.4(3) and Rules $\langle \text{T:PAR} \rangle$ and $\langle \text{T:FILL} \rangle$ we obtain the following type derivations:

$$\bullet_{\mathcal{S}_3}; \Gamma; \Theta \vdash C \triangleright \Delta'_1, \kappa^p : \alpha, \kappa^p : [\alpha]; \mathcal{I}'_1 \quad \Gamma; \Theta \vdash P_1 \mid \kappa^p[\alpha] \triangleright \Delta_1, \kappa^p : \alpha, \kappa^p : [\alpha]; \mathcal{I}_1$$

$$\bullet_{\mathcal{S}_4}; \Gamma; \Theta \vdash D \triangleright \Delta'_2, \kappa^p : \beta, \kappa^p : [\beta]; \mathcal{I}'_2 \quad \Gamma; \Theta \vdash P_2[v/x] \mid \kappa^p[\beta] \triangleright \Delta_2, \kappa^p : \beta, \kappa^p : [\beta]; \mathcal{I}_2$$

$$(21) \quad (22)$$

$$\Gamma; \Theta \vdash C\{P_1\} \mid D\{P_2[v/x]\} \triangleright \Delta'_1 \cup \Delta'_2, \kappa^p : \alpha, \kappa^p : \beta, \kappa^p : [\alpha], \kappa^p : [\beta]; \mathcal{I}'_1 \uplus \mathcal{I}'_2$$

with

$$\begin{aligned} \mathcal{S}_3 &= \Gamma; \Theta \vdash \Delta_1, \kappa^p : \alpha, \kappa^p : [\alpha]; \mathcal{I}_1 \\ \mathcal{S}_4 &= \Gamma; \Theta \vdash \Delta_2, \kappa^p : \beta, \kappa^p : [\beta]; \mathcal{I}_2 \\ \mathcal{S}_5 &= \Gamma; \Theta \vdash \Delta'_1 \cup \Delta'_2, \kappa^p : \alpha, \kappa^p : \beta, \kappa^p : [\alpha], \kappa^p : [\beta]; \mathcal{I}'_1 \uplus \mathcal{I}'_2 \end{aligned}$$

Since by inductive hypothesis Δ'_1 and Δ'_2 are balanced, we infer that $\Delta'_1 \cup \Delta'_2, \kappa^p : \alpha, \kappa^p : \beta$ is balanced as well and $\Delta \mapsto \Delta'_1 \cup \Delta'_2, \kappa^p : \alpha, \kappa^p : \beta$ (cf. Definition 4.5(1)). This concludes the proof for this case.

□

B.2. Multiparty Case

As in the binary case, the proof for multiparty sessions relies on an Inversion Lemma:

Lemma B.7 (Inversion Lemma for Networks). Let P and N be a process and a network, respectively. We have:

1. If $\Gamma \vdash_{\Sigma} \mathbf{new}(G) \triangleright \Delta$ then $\Sigma = \Delta = \emptyset$.
2. If $\Gamma \vdash_{\Sigma} \mathbf{end}[P] \triangleright \Delta$ then $\Sigma = \Delta = \emptyset$.
3. If $\Gamma \vdash_{\Sigma} \mathcal{M}[P] \triangleright \Delta$ and $\mathcal{M} \neq \mathbf{end}$ then $\Sigma = \emptyset$, $\Delta = \kappa[\mathbf{p}] : \mathcal{M}$, $\Gamma \vdash P \triangleright \kappa[\mathbf{p}] : T$ and $T \propto \mathcal{M}$.
4. If $\Gamma \vdash_{\Sigma} \kappa^e[h] \triangleright \Delta$ then $\Sigma = \Delta = \emptyset$.
5. If $\Gamma \vdash_{\Sigma} \kappa : \emptyset \triangleright \Delta$ then $\Sigma = \{\kappa\}$ and $\Delta = \emptyset$.
6. If $\Gamma \vdash_{\Sigma} \kappa : (\mathbf{p}, \mathbf{q}, \lambda(v)) \cdot h \triangleright \Delta$ then $\Sigma = \{\kappa\}$, $\Delta \approx \Delta' \# \{\kappa[\mathbf{p}] : \mathbf{q}! \lambda(S)\}$, $\Gamma \vdash_{\{\kappa\}} \kappa : h \triangleright \Delta'$ and $\Gamma \vdash v : S$.
7. If $\Gamma \vdash_{\Sigma} N_1 \mid N_2 \triangleright \Delta$ then $\Sigma = \Sigma_1 \cup \Sigma_2$, $\Delta \approx \Delta_1 \star \Delta_2$, $\Gamma \vdash_{\Sigma_1} N_1 \triangleright \Delta_1$ and $\Gamma \vdash_{\Sigma_2} N_2 \triangleright \Delta_2$.
8. If $\Gamma \vdash_{\Sigma} (\nu \kappa)N \triangleright \Delta$ then $\Sigma = \Sigma' \setminus \{\kappa\}$, $\Delta \approx \Delta' \setminus \{\kappa\}$, and $\Gamma \vdash_{\Sigma'} N \triangleright \Delta'$ and $\mathbf{bal}(\Delta', \kappa)$.
9. If $\Gamma \vdash_{\Sigma} \mathbf{loc}[\mathcal{P}; \mathbf{loc}[r]] \triangleright \Delta$ then $\Sigma = \Delta = \emptyset$, and $\forall (P, T_P) \in \mathcal{P}$, $\Gamma \vdash P \triangleright T_P$.
10. If $\Gamma \vdash_{\Sigma} \mathbf{loc}[\mathcal{M}[P]; \mathcal{P}; \mathbf{loc}[r]] \triangleright \Delta$ then $\Sigma = \emptyset$, $\Delta = \kappa[\mathbf{p}] : \mathcal{M}$, $\forall (P, T_P) \in \mathcal{P}$, $\Gamma \vdash P \triangleright T_P$ and $\Gamma \vdash_{\emptyset} \mathcal{M}[P] \triangleright \kappa[\mathbf{p}] : \mathcal{M}$.

Proof. Follows by the typing system definition (cf. Table 9). \square

We also need to prove the Substitution Lemma that allow us to type processes with variable substitutions.

Lemma B.8 (Substitution Lemma). Let P be a process.

1. If $\Gamma \vdash P \triangleright T$ then $\Gamma \vdash P[\kappa[\mathbf{q}/y]] \triangleright T$
2. If $\Gamma, x : S \vdash P \triangleright T$ and $\Gamma \vdash v \triangleright S$ then $\Gamma, x : S \vdash P[v/x] \triangleright T$.

Proof. Easily shown by induction on the structure of P . \square

Finally, the following lemma takes care of typability as preserved by Rule $\langle \text{M:EQUIV} \rangle$:

Lemma B.9 (Subject Congruence). If $\Gamma \vdash_{\Sigma} N \triangleright \Delta$ and $N \equiv N'$ then $\Gamma \vdash_{\Sigma} N' \triangleright \Delta$.

Proof. By induction on the derivation of $N \equiv N'$, with a case analysis on the last applied rule. \square

We can finally state our main theorem:

Theorem B.10 (Subject Reduction - 5.15). If $\Gamma \vdash_{\Sigma} N \triangleright \Delta$ with Δ balanced and $N \longrightarrow N'$ then $\Gamma \vdash_{\Sigma} N' \triangleright \Delta'$, for some balanced Δ' such that either $\Delta = \Delta'$ or $\Delta \mapsto \Delta'$.

Proof. By case analysis and induction on the reduction rules for networks (cf. Table 7). The cases can be divided into three families:

1. Session establishment: Rule $\langle \text{N:OPEN} \rangle$;
2. Intra-session communication: Rules $\langle \text{N:SEND} \rangle$, $\langle \text{N:RECV} \rangle$, $\langle \text{N:TAU} \rangle$, and $\langle \text{N:INMSG} \rangle$;
3. Adaptation: Rules $\langle \text{N:INUPD} \rangle$, $\langle \text{N:ADDLOC1} \rangle$, $\langle \text{N:ADDLOC2} \rangle$, $\langle \text{N:UPDLOC1} \rangle$ and $\langle \text{N:UPDLOC2} \rangle$.

We give the proof for each one of these families:

1. In this case $N \equiv \mathbf{new}(G) \mid \mathcal{L}oc \mid N$ with $\mathcal{L}oc = \prod_{\mathbf{p} \in G} \mathbf{loc}_{\mathbf{p}}[\mathcal{P}_{\mathbf{p}}; \mathbf{loc}_{\mathbf{p}}[\epsilon]]$ by hypothesis and by the Inversion Lemma we know that $\Gamma \vdash_{\emptyset} \mathbf{new}(G) \mid \mathcal{L}oc \triangleright \emptyset$ and $\Gamma \vdash_{\Sigma} N \triangleright \Delta$. After applying Rule $\langle \text{N:OPEN} \rangle$ we have

$$N' \equiv (\nu \kappa) \left(\prod_{\mathbf{p} \in \Pi} \mathbf{loc}_{\mathbf{p}}[\mathcal{M}_{\mathbf{p}}[P_{\mathbf{p}}[\kappa[\mathbf{p}]/y]]; \mathcal{P}_{\mathbf{p}}; \mathbf{loc}_{\mathbf{p}}[\epsilon] \mid \kappa : \emptyset \mid \kappa^e[\epsilon] \mid N \right).$$

By hypothesis we know that $\Gamma \vdash_{\Sigma} N \triangleright \Delta$. By Rules $\langle \text{M:QINIT} \rangle$ and $\langle \text{M:CQUEUE} \rangle$ we have $\Gamma \vdash_{\kappa} \kappa : \emptyset \triangleright \emptyset$ and $\Gamma \vdash_{\emptyset} \kappa^e[\epsilon] \triangleright \emptyset$. Now for each $\mathbf{p} \in \Pi$ by applying the Inversion Lemma (Lemma B.7) and the Substitution Lemma (Lemma B.8) we have:

$$\frac{\forall (P, T_P) \in \mathcal{P}_{\mathbf{p}} \quad \Gamma \vdash P \triangleright T_P \quad \frac{\Gamma \vdash_{\emptyset} P[\kappa[\mathbf{p}]/y] \triangleright \kappa[\mathbf{p}] : T_P \quad \mathcal{M}_{\mathbf{p}} \neq \mathbf{end} \quad T_P \propto \mathcal{M}_{\mathbf{p}}}{\Gamma \vdash_{\emptyset} \mathcal{M}_{\mathbf{p}}[P[\kappa[\mathbf{p}]/y]] \triangleright \kappa[\mathbf{p}] : \mathcal{M}_{\mathbf{p}}}}{\Gamma \vdash_{\emptyset} \mathbf{loc}_{\mathbf{p}}[\mathcal{M}_{\mathbf{p}}[P_{\mathbf{p}}[\kappa[\mathbf{p}]/y]]; \mathcal{P}_{\mathbf{p}}; \mathbf{loc}_{\mathbf{p}}[\epsilon]] \triangleright \kappa[\mathbf{p}] : \mathcal{M}_{\mathbf{p}}}$$

Finally, we apply Rules $\langle M:PAR \rangle$ and $\langle M:RES \rangle$, noticing that by Proposition 5.12 the type of

$$\prod_{p \in \Pi} \text{loc}_p[\mathcal{M}_p[P_p[\kappa[p]/y]]; \mathcal{P}_p; \text{loc}_p[\epsilon]]$$

is balanced for κ ; we then have that $\Gamma \vdash_{\emptyset} (\nu\kappa)(\prod_{p \in \Pi} \text{loc}_p[\mathcal{M}_p[P_p[\kappa[p]/y]]; \mathcal{P}_p; \text{loc}_p[\epsilon]] \mid \kappa : \emptyset \mid \kappa^e[\epsilon]) \triangleright \emptyset$. Thus we conclude $\Gamma \vdash_{\Sigma} N' \triangleright \Delta$ and Δ is unchanged thus balanced by hypothesis.

2. We consider only the case of Rule $\langle N:SEND \rangle$; the other cases are similar or simpler.

We have that $N \equiv \mathcal{E}[\text{loc}[\mathcal{M}[P]; \mathcal{P}; \text{loc}[\epsilon]] \mid \kappa : (p, q, \lambda(v)) \cdot h \mid \kappa^e[\epsilon]]$. By Rules $\langle M:QSEND \rangle$ and $\langle M:PAR \rangle$ we have that

$$\Gamma \vdash_{\kappa} \text{loc}[\mathcal{M}[P]; \mathcal{P}; \text{loc}[\epsilon]] \mid \kappa : (p, q, \lambda(v)) \cdot h \mid \kappa^e[\epsilon] \triangleright \Delta_1$$

with $\Delta_1 = \{\kappa[q] : \mathcal{M}\} \star (\Delta_2 \# \{\kappa[p] : q! \lambda(S)\})$. By hypothesis we have that there exists Δ_0 such that $\Delta = \Delta_0 \star \Delta_1$ and Δ balanced.

By applying Lemmas B.7 and B.8 and Rule $\langle <:PAR \rangle$ we have:

$$\Gamma \vdash_{\kappa} \text{loc}[\mathcal{M}[P']; \mathcal{P}; \text{loc}[\epsilon]] \mid \kappa : h \mid \kappa^e[\epsilon] \triangleright \Delta'_1$$

with $\Delta'_1 = \{\kappa[q] : \mathcal{M}'\} \star \Delta_2$, it is easy to see that $\Delta' = \Delta_0 \star \Delta'_1$ is still balanced and that $\Delta \mapsto \Delta'$ (cf. Definition 5.14(1)).

3. We give details of the proof for Rules $\langle N:INUPD \rangle$ and $\langle N:UPDLOC1 \rangle$ the other cases are similar or simpler.

- The case of Rule $\langle N:INUPD \rangle$ is similar to the one of Rule $\langle N:OPEN \rangle$ above. Let

$$N \equiv (\nu\kappa) \left(\prod_{p \in \Pi} \text{loc}_p[\mathcal{M}_p[P_p]; \mathcal{P}_p; \text{loc}_p[\epsilon]] \mid \kappa : h \mid \kappa^e[G'] \mid \prod_{q \in \Pi' \setminus \Pi} \text{loc}_q[\mathcal{P}_q; \text{loc}_q[\epsilon]] \right)$$

where $\Pi' = \text{part}(G')$. We have $\Gamma \vdash_{\emptyset} N \triangleright \emptyset$ and after the application of Rule $\langle N:INUPD \rangle$ we also have

$$N' \equiv (\nu\kappa) \left(\prod_{p \in \Pi'} \text{loc}_p[\mathcal{M}'_p[P'_p]; \mathcal{P}_p; \text{loc}_p[\epsilon]] \mid \kappa : \epsilon \mid \kappa^e[\epsilon] \mid \prod_{q \in \Pi \setminus \Pi'} \text{loc}_q[\mathcal{P}_q; \text{loc}_q[\epsilon]] \right)$$

By Rules $\langle M:LOC1 \rangle$ and $\langle M:PAR \rangle$ we know that

$$\Gamma \vdash_{\epsilon} \prod_{q \in \Pi \setminus \Pi'} \text{loc}_q[\mathcal{P}_q; \text{loc}_q[\epsilon]] \triangleright \emptyset.$$

By Proposition 5.12 and as queue κ has been emptied, the type Δ of $\prod_{p \in \Pi'} \text{loc}_p[\mathcal{M}'_p[P'_p]; \mathcal{P}_p; \text{loc}_p[\epsilon]]$ is balanced; therefore, by applying typing rules $\langle M:PAR \rangle$ and $\langle M:RES \rangle$ we can conclude $\Gamma \vdash_{\emptyset} N' \triangleright \emptyset$.

- For Rule $\langle N:UPDLOC1 \rangle$, let

$$N \equiv \mathcal{E}[\text{loc}_p[(\mathcal{M}[P]; \mathcal{P}; \text{loc}_p[(upd : \text{case } x \text{ of } \{(T_i) : Q_i\}_{i \in I}]])]]$$

and $\Gamma \vdash_{\Sigma} N \triangleright \Delta$ with Δ balanced. After the application of Rule $\langle N:UPDLOC1 \rangle$ we have

$$N' \equiv \text{loc}_p[\mathcal{M}[R]; \mathcal{P}; \text{loc}_p[\epsilon]].$$

Now notice that the substitution of process P with R has no influence on the type of the system N' as the reduction rule guarantees that the type is left unchanged (cf. premise $\text{match}_I^{\infty}(\{\mathcal{M}\}, \{T_R\})$). Thus $\Gamma \vdash_{\Sigma} N' \triangleright \Delta$, which concludes the proof.

□