

A Type System for Flexible Role Assignment in Multipartly Communicating Systems

Pedro Baltazar¹, Luís Caires³, Vasco T. Vasconcelos², and Hugo T. Vieira³

¹ Instituto de Telecomunicações, IST, Universidade Técnica de Lisboa

² LaSIGE, Faculdade de Ciências, Universidade de Lisboa

³ CITI-DI, Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa

Abstract. Communication protocols in distributed systems often specify the roles of the parties involved in the communications, namely for enforcing security policies or task assignment purposes. Ensuring that implementations follow role-based protocol specifications is challenging, especially in scenarios found, e.g., in business processes and web applications, where multiple peers are involved, single peers impersonate several roles, or single roles are carried out by several peers. We present a type-based analysis for statically verifying role-based multi-party interactions, based on a simple π -calculus model and prior work on conversation types. Our main result ensures that well-typed systems follow the role-based protocols prescribed by the types, including systems where roles are flexibly assigned to processes.

1 Introduction

Communication is a central feature of nowadays software systems, as more and more often systems are built using computational resources that are concurrently available and distributed in the web. Examples range from operating systems where functionality is distributed between distinct threads in the system, to services available on the Internet, which rely on third-party (remote) service providers to carry out subsidiary tasks, following the emerging model of SaaS (software as a service) and cloud computing. Building software from the composition of communicating interacting pieces is very flexible, at least in principle, since resources can be dynamically discovered and chosen according to criteria such as declared functionality, availability and work load. In such a setting, all interacting parties must agree on communication protocols without relying on centralized control. Verification mechanisms that automatically check whether the code meets some common protocol specification become then of crucial importance.

A protocol specification describes a set of message exchanges, recording when these should occur as well as the parties involved in the interaction. A party involved in a protocol may have a spatial meaning, for instance denoting a distinguished site or process, or, more generally, a party may have a behavioral meaning, a *role* in the interaction that may be realized by one or more processes or sites. Conversely, a process may impersonate different roles throughout its

execution. Such flexibility is essential to address systems, e.g., where a leader role is impersonated by different sites at different stages of the protocol, and the role of each site changes accordingly.

A challenge that arises is then to devise techniques to verify whether a system complies to a protocol specification, given such dynamic and distributed implementation of roles, just by inspecting the source code. A particular situation where roles must be traced is when checking conformance against security policies like, for example, those involving separation of duties.

In this paper we present a type-based analysis for verifying whether systems defined in a model programming language follow the role-based protocol descriptions as prescribed by types. Our development is based on conversation type theory [4], extending it with the ability to specify and analyze the roles involved in the interactions. The underlying model of our analysis is an extremely parsimonious extension of the π -calculus [13, 15], where communication actions specify a message label and the role performing the action, inspired by TyCO [16]. Conversations generalize sessions [10, 12] with support to multiparty interaction, addressing dynamically established collaborations between an unanticipated number of partners. A distinguishing feature of the conversation types approach is that multiple parties interact using labeled messages in a single medium of communication, while other works support multiparty communication via message queues [11] and indexed communication channels [2]. We choose to adopt the simplest possible setting where session-like multiparty interaction may be studied, and extend it in a minimal way so as to support general reasoning about roles. So, apart from retaining the simplicity of conversation types, our theory addresses systems where a single role may be realized by several parties and where processes may dynamically change the role on behalf of which they are interacting, as needed to model communicating workflows as present in actual business processes. This contrasts with related approaches (see, e.g., [7, 11]) where roles have a “spatial” meaning, as they are mapped into the structure of systems or sites in a static way.

In the remainder of this section we informally describe our type analysis by going through some examples. Consider the protocol specification given by type

$$\text{Sender} \rightarrow \text{Receiver } \mathit{hello}(). \text{Sender} \rightarrow \text{Receiver } \mathit{bye}()$$

which captures a binary interaction where messages *hello* and *bye* are sequentially exchanged, and the communicating partners are identified by **Sender** and **Receiver**, which send and receive the messages, respectively (read \rightarrow as “sends to”). A non surprising implementation of this interaction is given by process

$$\mathit{chat} \triangleleft_{\text{Sender}} \mathit{hello}(). \mathit{chat} \triangleleft_{\text{Sender}} \mathit{bye}() \mid \mathit{chat} \triangleright_{\text{Receiver}} \mathit{hello}(). \mathit{chat} \triangleright_{\text{Receiver}} \mathit{bye}()$$

where two concurrent processes interact on channel *chat* following the protocol above. The process on the left sends the two messages under role **Sender** ($\triangleleft_{\text{Sender}}$), as described by type $!\text{Sender } \mathit{hello}(). !\text{Sender } \mathit{bye}()$, while the process on the right receives the two messages under role **Receiver** ($\triangleright_{\text{Receiver}}$), described by type $?\text{Receiver } \mathit{hello}(). ?\text{Receiver } \mathit{bye}()$.

In this first example there is a perfect match between processes and the roles under which the processes interact. However, this does not need to be the case. Consider a different implementation of the same protocol

$$chat \triangleleft_{\text{Sender}} \text{hello}().chat \triangleright_{\text{Receiver}} \text{bye}() \mid chat \triangleright_{\text{Receiver}} \text{hello}().chat \triangleleft_{\text{Sender}} \text{bye}()$$

where the process on the left sends message *hello* as **Sender** and then receives message *bye* as **Receiver**, described by type $! \text{Sender } \text{hello}(). ? \text{Receiver } \text{bye}()$, and the process on the right first acts as **Receiver** and then as **Sender**, described by type $? \text{Receiver } \text{hello}(). ! \text{Sender } \text{bye}()$. Notice each role is carried out by two distinct processes and each process implements two distinct roles.

Our type analysis ensures that both implementations follow the prescribed protocol, since the protocol $\text{Sender} \rightarrow \text{Receiver } \text{hello}(). \text{Sender} \rightarrow \text{Receiver } \text{bye}()$ is decomposed in “complementary” types that describe the behavior of the individual processes (for instance, in type $! \text{Sender } \text{hello}(). ? \text{Receiver } \text{bye}()$ and type $? \text{Receiver } \text{hello}(). ! \text{Sender } \text{bye}()$). Although very simple, this example already distinguishes our approach from previous works, since the ability to specify roles is absent in [4] while [7, 11] do not support such role distribution. Conceivably channel delegation (channel-passing) supported by previous works may be used to represent a similar notion but, to model this example in particular, two channel delegations would be necessary, which implies it would not be possible to directly observe that the two interactions take place in a related medium (in our case the *chat* channel) and the ability to audit role participation locally would be lost (as the personification of a different role would be a consequence of channel-passing).

Now consider a more realistic scenario (adapted from [4]) described by type

$$\begin{aligned} \text{Buyer} \rightarrow \text{Seller } \text{buy}(). \text{Seller} \rightarrow \text{Buyer } \text{price}(). \\ \text{Seller} \rightarrow \text{Shipper } \text{product}(). \text{Shipper} \rightarrow \text{Buyer } \text{details}() \quad (1) \end{aligned}$$

which captures the interactions in a purchase system involving three parties. Messages *buy*, *price*, *product* and *details* are exchanged between a **Buyer**, a **Seller**, and a **Shipper**. First, the buyer sends the seller a buy request, then the seller replies the price back to the buyer. After that, the seller informs the shipper of the chosen product and the shipper sends the buyer the delivery details.

Fig. 1 shows a possible implementation of the purchase interaction system. Using the **new** construct, process **Buyer** creates a fresh channel *chat* that will host the purchase interaction described by (1). This newly created name is passed to a shop, via message *buyService*. Code $shop \triangleleft_{\text{Buyer}} \text{buyService}(chat)$ represents the output of message *buyService* on channel *shop*, passing name *chat* under role **Buyer**. The **Buyer** process then sends message *buy*, after which it is simultaneously active to receive *price* and to send name *chat* on *mailBox storeService*.

The **Shop** process starts by receiving a channel name (that instantiates variable *x*) in message *buyService*. Then, in this received channel the **Shop** impersonates the **Seller** role and receives message *buy*, after which it sends message *price*. At this point, process **Shop** simultaneously impersonates **Seller** and **Shipper**,

$$\begin{aligned}
\mathbf{Buyer} &\triangleq (\mathbf{new\ } chat) \\
&\quad shop \triangleleft_{\mathbf{Buyer}} buyService(chat). \\
&\quad\quad chat \triangleleft_{\mathbf{Buyer}} buy(). \\
&\quad\quad\quad (chat \triangleright_{\mathbf{Buyer}} price() \mid mailbox \triangleleft_{\mathbf{Buyer}} storeService(chat)) \\
\mathbf{Shop} &\triangleq shop \triangleright_{\mathbf{Shop}} buyService(x). \\
&\quad\quad x \triangleright_{\mathbf{Seller}} buy(). \\
&\quad\quad\quad x \triangleleft_{\mathbf{Seller}} price(). \\
&\quad\quad\quad\quad (x \triangleleft_{\mathbf{Seller}} product() \mid x \triangleright_{\mathbf{Shipper}} product().x \triangleleft_{\mathbf{Shipper}} details()) \\
\mathbf{Mail} &\triangleq mailbox \triangleright_{\mathbf{Mail}} storeService(x). \\
&\quad\quad x \triangleright_{\mathbf{Buyer}} details() \\
\mathbf{System} &\triangleq (*\mathbf{Buyer} \mid *\mathbf{Mail} \mid *\mathbf{Shop})
\end{aligned}$$

Fig. 1. Code for the Purchase System.

which exchanges message *product*, after which message *details* is sent. Notice that this particular **Shop** carries out both the role of the **Seller** and the role of the **Shipper**, allowing to represent a shop equipped with its own shipping service.

The **Mail** process defines a message storage service that impersonates the buyer in receiving the shipping delivery details. Notice that the buyer passes name *chat* to the mailbox, allowing in this way a third party to dynamically join the ongoing interaction, while still interacting on the delegated channel (via message *price*). Hence, in this system the **Buyer** role is actually carried out by two distinct processes (**Buyer** and **Mail**), which can be simultaneously active.

The implementation shown in Fig. 1 involves three distinguished processes that carry out the three roles identified in the protocol, albeit not in a one-to-one-correspondence. The type given in (1) captures the interaction in channel *chat*, which is passed from the buyer to the shop and to the mailbox in messages *buyService* and *storeService*, respectively. In order to analyze the protocol distribution between the three parties, we must consider the “slices” of protocol that are delegated in messages. Namely, the overall protocol is *split* in the type that captures the behavior that is sent to the shop (via message *buyService*)

$$?Seller\ buy(). !Seller\ price(). Seller \rightarrow Shipper\ product(). !Shipper\ details()$$

and in the type that captures the behavior retained by the buyer

$$!Buyer\ buy(). ?Buyer\ price(). \diamond ?Buyer\ details()$$

The \diamond type expresses the fact that the input of message *details* occurs “some-time in the future”, i.e., it does not necessarily occur exactly after the input of message *price*. In fact the **Buyer** process illustrated in Fig. 1 does not guarantee that the input is active only after the reception of message *price*. However, the sequentiality of the message exchanges is ensured by the **Shop** process, since the output of message *details* only occurs after the output of message *price*. A type $\diamond B$ denotes a behavior that must occur sometime, but not necessarily “now” — $\diamond B$ types obey the basic laws of the eventually temporal logic operator.

When typing the buyer process there is a further type decomposition, at the level of messages *price* and *details*, resulting in types $?Buyer\ price()$ and $\diamond?Buyer\ details()$, the former being retained by the buyer process and the latter delegated to the mailbox. When typing the shop process there is another type decomposition, at the level of message *product*, resulting in types $!Seller\ product()$ and $?Shipper\ product().!Shipper\ details()$, which explain the behaviors of the parallel processes in the shop code. All decompositions sketched above are captured by a *type split*, \circ , relation that explains how protocols may be split in two complementary slices, along with subtyping. $\diamond B$ types are crucial to the definition of type split, as they provide algebraic support to the flexibility required to sequentially order message exchanges among multiple parties.

In the previous example, the fact that message *details* is exchanged after message *price* is not observable just by looking at the source code of the buyer and mail. However, such ordering is guaranteed by the shop. If we specify that the buyer, in general, exhibits such behaviors concurrently (for example, $?Buyer\ price() | ?Buyer\ details()$) we would require (order preserving) decompositions of protocols into multiple threads of behavior. The flexibility introduced by $\diamond B$ types solves this problem as they support the specification of orderings that are guaranteed via synchronization. For example, the type $?Buyer\ price().\diamond?Buyer\ details()$ says that the reception of message *details* takes place (sometime) after the reception of message *price*. On the other hand, the type $!Seller\ price().Seller \rightarrow Shipper\ product().!Shipper\ details()$ says that the output of *details* necessarily occurs immediately after the output of message *price*. The combination of the two typed guarantees the overall ordering: first message *price*, then *product* and finally *details*.

The purchase interaction of the system shown in Fig. 1 follows the protocol specification given in (1). Notice that the Buyer role is distributed between two processes (Buyer and Mail), and that roles Seller and Shipper are carried out by a single process (Shop). From the point of view of our type analysis the system follows the prescribed protocols, regardless of the spatial configuration of the processes that implement the roles.

2 Process Language

In this section we present the process model, first by introducing the syntax and second by defining the operational semantics. Our process language is the π -calculus [13, 15] extended with labeled communication and role-based annotations. The syntax, inspired in TyCO [16], is illustrated in Fig. 2, where we consider given an infinite sets of labels \mathcal{L} , of channel names \mathcal{N} and of roles \mathcal{R} . Labels, used to index communication, are identifiers that may neither be created nor communicated (e.g., XML tags). Names are used to identify mediums of communication. For typing purposes, we distinguish two distinct usages of channels: public (*shared*) communication mediums (e.g., gateways to service providers, like the *shop* and *mailBox* channels in the example) and private (*linear*) mediums, where a set of related interactions among several parties may take place (cap-

$$P ::= \mathbf{0} \mid (\mathbf{new} \ x)P \mid P_1 \mid P_2 \mid *P \mid x \triangleright_r \{l_i(x_i).P_i\}_{i \in I} \mid x \triangleleft_r l(y).P$$

$$l \in \mathcal{L}(\text{abels}) \quad x, y \in \mathcal{N}(\text{ames}) \quad r, s \in \mathcal{R}(\text{oles})$$

Fig. 2. Process Syntax.

$$P \mid \mathbf{0} \equiv P \quad P_1 \mid P_2 \equiv P_2 \mid P_1 \quad (P_1 \mid P_2) \mid P_3 \equiv P_1 \mid (P_2 \mid P_3)$$

$$(\mathbf{new} \ x)(\mathbf{new} \ y)P \equiv (\mathbf{new} \ y)(\mathbf{new} \ x)P$$

$$P_1 \mid (\mathbf{new} \ x)P_2 \equiv (\mathbf{new} \ x)(P_1 \mid P_2) \quad (\text{if } x \notin \text{fn}(P_1))$$

$$(\mathbf{new} \ x)\mathbf{0} \equiv \mathbf{0} \quad *P \equiv *P \mid P \quad P_1 \equiv P_2 \quad (\text{if } P_1 \equiv_\alpha P_2)$$

Fig. 3. Structural Congruence.

turing, e.g., service instance interactions, like the *chat* channel in the example). Roles are used to identify the parties involved in communications.

A process is either an inactive process $\mathbf{0}$, a name restriction $(\mathbf{new} \ x)P$ where name x is known only to process P , a parallel composition $P_1 \mid P_2$ where P_1 and P_2 are simultaneously active, or a replication $*P$ where unlimited copies of P are simultaneously active. Process constructs described up to here (the static fragment) correspond exactly to the ones found in π -calculus. As for communication primitives, we extend the (monadic) π -calculus output and input primitives with labeled communication and role annotations. Process $x \triangleleft_r l(y).P$ is able to send a message on channel x , under role r , labeled by l . Upon synchronization the name y is sent and the continuation P activated. Notice that the r annotation identifies the role in which the emission is performed. The input summation process $x \triangleright_r \{l_i(x_i).P_i\}_{i \in I}$ is able to receive one message in name x , under role r , labeled by any of the l_i labels, where i ranges over index set I (we assume that all labels l_i in an input prefix are distinct). Upon synchronization with an l_j labeled message, the respective parameter x_j is instantiated and the respective continuation P_j activated. In $(\mathbf{new} \ x)P$ all occurrences of x are bound in P , and in $x \triangleright_r \{l_i(x_i).P_i\}_{i \in I}$ all occurrences of x_i are bound in P_i , for each $i \in I$.

We introduce some auxiliary notions: we use $\text{fn}(P)$ to denote the set of free names of process P , defined as expected, and $P[x \leftarrow y]$ to denote the process obtained by replacing all free occurrences of x by y in P . As usual, we omit inactive continuations (e.g., $x \triangleleft_r l(y)$ stands for $x \triangleleft_r l(y).\mathbf{0}$).

The operational semantics is given by a reduction relation and by a structural congruence. We consider the standard definition of structural congruence, denoted by \equiv , defined as the least congruence that satisfies the rules in Fig. 3. Structural congruence is used in the definition of the reduction relation to syntactically rearrange the process, in order to allow reduction to be defined, as usual, by capturing the basic case for synchronization and identifying the active contexts in which a synchronization may take place.

$$\begin{array}{c}
\frac{k \in I}{x \triangleright_r \{l_i(x_i).P_i\}_{i \in I} \mid x \triangleleft_s l_k(y).P \xrightarrow{x:s \rightarrow r l_k} P_k[x_k \leftarrow y] \mid P} \quad (\text{Red-Comm}) \\
\frac{P \xrightarrow{\lambda} P' \quad \lambda \in \{\tau, x : s \rightarrow r l\}}{(\mathbf{new} \ x)P \xrightarrow{\tau} (\mathbf{new} \ x)P'} \quad \frac{P \xrightarrow{x:s \rightarrow r l} P' \quad y \neq x}{(\mathbf{new} \ y)P \xrightarrow{x:s \rightarrow r l} (\mathbf{new} \ y)P'} \quad (\text{Red-New1,Red-New2}) \\
\frac{P_1 \xrightarrow{\lambda} P'_1}{P_1 \mid P_2 \xrightarrow{\lambda} P'_1 \mid P_2} \quad \frac{P_1 \equiv P'_1 \quad P'_1 \xrightarrow{\lambda} P'_2 \quad P'_2 \equiv P_2}{P_1 \xrightarrow{\lambda} P_2} \quad (\text{Red-Par,Red-Struct})
\end{array}$$

Fig. 4. Reduction Relation.

For typing purposes, and since we intend to match process behaviors against type specifications, our reduction relation records (public) synchronization information in labels. Reduction labels (ranged over by λ) are of two forms: a τ label captures a private internal interaction, whereas an $x : s \rightarrow r l$ label captures an l -labeled message exchange on channel x , between roles s (ender) and r (eceiver).

We may now present the reduction relation, defined by the rules given in Fig. 4, where we use $P_1 \xrightarrow{\lambda} P_2$ to represent that process P_1 reduces to P_2 with label λ . Rule (RED-COMM) says that two parallel input and output processes may exchange message l_k on channel x , the interaction being captured by label $x : s \rightarrow r l_k$, where also the roles involved in the interaction are recorded. As the result of the synchronization, name y activates the continuation (respective to l_k) instantiating parameter x_k . The continuation of the output process is also activated as a consequence of the synchronization. Rule (RED-PAR) closes reduction under parallel contexts, while rules (RED-NEW1) and (RED-NEW2) close reduction under name restriction. (RED-NEW1) captures synchronization in private names in the scope of the name restriction, either by “hiding” a public synchronization in the restricted name or by allowing private synchronizations. (RED-NEW2) captures public synchronizations in the scope of the name restriction, not involving the restricted name. (RED-STRUCT) closes reduction under structural congruence.

3 Type System

In this section we present our type system. The type language is given in Fig. 5, where we distinguish between behavioral types that describe linear interactions (B) from types that describe shared interactions (T) (cf. conversation [4] or session [12] initiation primitives). We also use message (argument) types (M) that specify either a linear protocol or a shared message type, and communication prefixes (ρ) that describe role-based communication actions.

A behavioral type B specifies the inactive behavior **end**, the parallel composition $B_1 \mid B_2$ of two independent behaviors B_1 and B_2 , the sometime $\diamond B$, which says that behavior B may occur at any point in time, or a menu of labeled actions $\rho\{l_i(M_i).B_i\}_{i \in I}$, each one specifying the type of the name communicated

$$\begin{aligned}
B &::= \mathbf{end} \mid B \mid B \mid \diamond B \mid \rho\{l_i(M_i).B_i\}_{i \in I} \\
T &::= l(B) \quad M ::= B \mid T \quad \rho ::= !s \mid ?r \mid s \rightarrow r
\end{aligned}$$

Fig. 5. Conversation Types Syntax.

$$\begin{array}{c}
\vdash \mathbf{end} \quad \frac{B_1 \# B_2 \quad \vdash B_1 \quad \vdash B_2}{\vdash B_1 \mid B_2} \quad \frac{\forall i \in I \quad \vdash B_i \quad \rho\{l_i(M_i).\mathbf{end}\} \# B_i}{\vdash \rho\{l_i(M_i).B_i\}_{i \in I}} \\
\vdash \diamond \mathbf{end} \quad \frac{\vdash B_1 \mid B_2 \quad \vdash \diamond B_1 \quad \vdash \diamond B_2}{\vdash \diamond(B_1 \mid B_2)} \quad \frac{\forall i \in I \quad \vdash \rho\{l_i(M_i).B_i\} \quad \rho \in \{!s, ?r\}}{\vdash \diamond \rho\{l_i(M_i).B_i\}_{i \in I}}
\end{array}$$

Fig. 6. Well-Formed Type Predicate.

in the message M_i , and the respective continuation behavior B_i . Depending on the communication prefix ρ , an action menu represents either an input branching (when ρ is $?r$), an output choice (when ρ is $!s$)—cf. branch and choice session types [12]—or an internal choice $s \rightarrow r$, i.e., a matched communication between an output and an input. Notice that the communication roles are identified in the communication prefixes: the sender role in $!s$, the receiver role in $?r$, and the two roles involved in the interaction in $s \rightarrow r$ (s sends to r). Notice also that input and output actions (interface types that capture interactions with the environment) are mixed with matched actions (capturing internal interactions) at the same level in the type language.

The conversation type language is extended with role-based annotations and sometime types ($\diamond B$). Although a specification is not expected to use $\diamond B$ types, these are crucial to allow the decomposition of protocols into slices, some of which related to interactions that occur later in the protocol.

A message argument type M either specifies a behavioral linear type B , in case a linear name is communicated in the message, or a shared type T , in case a shared name is communicated in the message. A shared type T abbreviates $l(B)$, identifying the label of the message exchanged and the (linear) type of the name sent in the message — to simplify the presentation we consider that only linear names can be communicated in shared messages (communicating shared names can be easily encoded).

We now introduce some auxiliary notions, namely the type apartness, well-formed types, and matched types, all defined as predicates. Type apartness is used to identify non-interfering concurrent behaviors that may be safely composed in a linear interaction. To define type-apartness we use $lab(B)$ to denote the set of labels occurring in type B , defined as expected. We say that two types B_1 and B_2 are apart, and we write $B_1 \# B_2$, if the set of labels of B_1 is disjoint from the set of labels of B_2 ($lab(B_1) \cap lab(B_2) = \emptyset$). Building on apartness, we introduce well-formed type predicate, noted $\vdash B$, given by the rules in Fig. 6. Informally, in a well-formed type labels do not appear repeatedly in parallel (to

$$\begin{array}{c}
\frac{B_1 <: B'_1}{B_1 | B_2 <: B'_1 | B_2} \quad \frac{\forall i \in I \ B_i <: B'_i}{\rho\{l_i(M_i).B_i\}_{i \in I} <: \rho\{l_i(M_i).B'_i\}_{i \in I}} \quad \frac{\vdash \diamond B}{B <: \diamond B} \\
(B_1 | B_2) | B_3 \equiv B_1 | (B_2 | B_3) \quad B_1 | B_2 \equiv B_2 | B_1 \quad B | \mathbf{end} \equiv B \\
\diamond(B_1 | B_2) \equiv \diamond B_1 | \diamond B_2 \quad \diamond \mathbf{end} \equiv \mathbf{end}
\end{array}$$

Fig. 7. Subtyping Relation.

$$\begin{array}{c}
\frac{\vdash B}{B = \mathbf{end} \circ B} \quad \frac{B_1 = B'_1 \circ B''_1 \quad B_2 = B'_2 \circ B''_2 \quad \vdash B_1 | B_2}{B_1 | B_2 = B'_1 | B'_2 \circ B''_1 | B''_2} \quad (\text{S-END,S-PAR}) \\
\frac{\forall i \in I \ B_i = B'_i \circ B''_i \quad \{\rho_1, \rho_2\} = \{!r_1, ?r_2\} \quad \vdash r_1 \rightarrow r_2\{l_i(M_i).B_i\}_{i \in I}}{r_1 \rightarrow r_2\{l_i(M_i).B_i\}_{i \in I} = \rho_1\{l_i(M_i).B'_i\}_{i \in I} \circ \diamond \rho_2\{l_i(M_i).B''_i\}_{i \in I}} \quad (\text{S-TAU}) \\
\frac{\forall i \in I \ B_i = B'_i \circ \diamond B \quad \vdash \rho\{l_i(M_i).B_i\}_{i \in I}}{\rho\{l_i(M_i).B_i\}_{i \in I} = \rho\{l_i(M_i).B'_i\}_{i \in I} \circ \diamond B} \quad (\text{S-BRK}) \\
\frac{\forall i \in I \ B_i = B'_i \circ \diamond B \quad \vdash \diamond \rho\{l_i(M_i).B_i\}_{i \in I}}{\diamond \rho\{l_i(M_i).B_i\}_{i \in I} = \diamond \rho\{l_i(M_i).B'_i\}_{i \in I} \circ \diamond B} \quad (\text{S-BRKS}) \\
\frac{B = B_2 \circ B_1 \quad B'_1 = B'_2 \circ B'_3 \quad B_1 \equiv B'_1 \quad B_2 \equiv B'_2 \quad B_3 \equiv B'_3}{B = B_1 \circ B_2 \quad B_1 = B_2 \circ B_3} \quad (\text{S-SYM,S-EQU})
\end{array}$$

Fig. 8. Type Split Relation.

ensure race-free behavior) or in sequence (useful to simplify presentation). Also well-formed \diamond types are not applied directly to message exchanges ($s \rightarrow r$), since we are interested in specifying message exchanges that happen exactly at some point in the protocol. Also used by our typing is the notion of matched types, which captures systems where all input actions have a matching output. We say that type B is matched, noted $\mathit{matched}(B)$, if all communication prefixes in B are of the form $s \rightarrow r$.

The subtyping relation between behavioral types, noted $B_1 <: B_2$, is the least reflexive and transitive relation satisfying the rules in Fig. 7, where we write $B_1 \equiv B_2$ when $B_1 <: B_2$ and $B_2 <: B_1$. We remark on the use of subtyping to introduce flexibility at the level of \diamond types: type B is a subtype of $\diamond B$, which, intuitively, means that carrying out behavior B immediately is a safe implementation of eventually carrying out behavior B .

We may now introduce type split, a ternary relation that explains how a behavioral type may be safely decomposed in two slices of behavior, capturing, in a compositional way, the behavioral contribution of distinct processes to the overall interaction. The split relation is defined by the rules given in Fig. 8, where $B = B_1 \circ B_2$ denotes that type B may be decomposed in parts B_1 and B_2 .

We briefly discuss the splitting rules. Rule (S-END) specifies that a behavioral type may be decomposed in itself and the inactive behavior, typing

processes that contribute “all or nothing” to the interaction. Rule (S-PAR) explains the decomposition of two independent behaviors in two slices of behaviors each, capturing the decomposition of a system in two processes that contribute both to independent interactions.

Rule (S-TAU) separates a matched communication, between roles r_1 and r_2 , in the respective output by role r_1 and input by role r_2 , given a splitting of the continuation behaviors. The rule captures the decomposition of a system in two processes that synchronize in a message, each with a given role in the interaction, where one of them carries out the behavior immediately, while the other may carry out the behavior at some point in time (\diamond). In such way, since one of the behaviors occurs immediately we ensure that also the message exchange takes place immediately. Notice that a rule to separate the message exchange in two immediate behaviors is not necessary since the sometime behavior may also take place immediately (via subtyping).

Rule (S-BRK) separates a \diamond (sometime) distinguished slice of behavior from a communication prefixed type, provided this behavior can be split from the continuations in all branches. The rule thus captures the decomposition of a system in two parts, where one retains the (entire) interaction capability specified by the communication prefixed type while the other contributes to ensuing interactions—singled out by the \diamond . Notice that (S-BRK) allows to split behaviors such that the same slice is shared between all branches, useful when addressing, e.g., a branching protocol where every branch terminates with an *ok* or *ack* message. Rule (S-BRKS) expresses the same principle as (S-BRK) but for \diamond prefixed types. Rule (S-SYM) closes the relation under symmetry and rule (S-EQU) closes the relation under type equivalence.

To simplify the presentation, we sometimes write $B_1 \circ B_2$ to represent a type B such that $B = B_1 \circ B_2$ (if any such B exists). Notice that $B_1 \circ B_2$ does not uniquely identify a type, as B_1 and B_2 may be the result of splitting distinct types. Notice also that a type may be split in several ways. In prior work on conversation types [4], we use “merge” instead of “split”, in the sense that if $B = B_1 \circ B_2$ then we may see B as the result of merging the behaviors B_1 and B_2 . The merge was originally inspired in the (non-algebraic) end-point projection introduced in [5]. We can show that split is an associative relation, which is a crucial property to our type system since we rely on the flexibility of the type decomposition to address the behavioral contributions of multiple parties.

We may now present the type system. A typing judgment is of the form $\Delta; \Gamma \vdash P$ where Δ is the typing environment that describes the interactions on linear channels, and Γ is the typing environment that describes the interactions on shared channels. We write $\Delta; \Gamma$ only when the domains of Δ and Γ are disjoint. A typing environment Δ is an assignment of identifiers to behavioral types ($\Delta \triangleq x_1 : B_1, \dots, x_k : B_k$) and a typing environment Γ is an assignment of identifiers to shared types ($\Gamma \triangleq x_1 : T_1, \dots, x_k : T_k$). We introduce some auxiliary notation to simplify presentation: we use $(x_1 : B'_1, \dots, x_k : B'_k, \Delta_1) \circ (x_1 : B''_1, \dots, x_k : B''_k, \Delta_2)$ to denote $x_1 : B_1, \dots, x_k : B_k, \Delta_1, \Delta_2$ such that $B_i = B'_i \circ B''_i$, for all i in $1, \dots, k$ and the domains of Δ_1 and Δ_2 are disjoint.

$$\begin{array}{c}
\frac{\Delta_{\text{end}}; \Gamma \vdash \mathbf{0} \quad \frac{\Delta_1; \Gamma \vdash P_1 \quad \Delta_2; \Gamma \vdash P_2}{\Delta_1 \circ \Delta_2; \Gamma \vdash P_1 \mid P_2}}{\Delta_{\text{end}}; \Gamma \vdash \mathbf{0}} \quad (\text{T-END, T-PAR}) \\
\frac{\Delta, x : B; \Gamma \vdash P \quad \text{matched}(B)}{\Delta; \Gamma \vdash (\mathbf{new } x)P} \quad \frac{\Delta; \Gamma, x : l(B) \vdash P}{\Delta; \Gamma \vdash (\mathbf{new } x)P} \quad (\text{T-NEW, T-SNEW}) \\
\frac{\Delta, y : B; \Gamma, x : l(B) \vdash P}{\Delta; \Gamma, x : l(B) \vdash x \triangleright_r \{l(y).P\}} \quad \frac{\Delta; \Gamma, x : l(B) \vdash P}{\Delta \circ y : B; \Gamma, x : l(B) \vdash x \triangleleft_r l(y).P} \quad (\text{T-SIN, T-SOUT}) \\
\frac{\forall i \in I \quad \Delta \circ x : B_i, y_i : B'_i; \Gamma \vdash P_i \quad ?r\{l_i(B'_i).B_i\}_{i \in I} <: B}{\Delta \circ x : B; \Gamma \vdash x \triangleright_r \{l_i(y_i).P_i\}_{i \in I}} \quad (\text{T-IN}) \\
\frac{k \in I \quad \Delta \circ x : B_k; \Gamma \vdash P \quad !r\{l_i(B'_i).B_i\}_{i \in I} <: B}{\Delta \circ x : B \circ y : B'_k; \Gamma \vdash x \triangleleft_r l_k(y).P} \quad (\text{T-OUT}) \\
\frac{\forall i \in I \quad \Delta \circ x : B'_i; \Gamma, y_i : T_i \vdash P_i \quad ?r\{l_i(T_i).B'_i\}_{i \in I} <: B}{\Delta \circ x : B; \Gamma \vdash x \triangleright_r \{l_i(y_i).P_i\}_{i \in I}} \quad (\text{T-LSIN}) \\
\frac{\Delta \circ x : B'_k; \Gamma, y : T_k \vdash P \quad !r\{l_i(T_i).B'_i\}_{i \in I} <: B}{\Delta \circ x : B; \Gamma, y : T_k \vdash x \triangleleft_r l_k(y).P} \quad (\text{T-LSOUT}) \\
\frac{\Delta_1; \Gamma \vdash P \quad \Delta_1 <: \Delta_2}{\Delta_2; \Gamma \vdash P} \quad \frac{\Delta_{\text{end}}; \Gamma \vdash P}{\Delta_{\text{end}}; \Gamma \vdash *P} \quad (\text{T-SUB, T-REP})
\end{array}$$

Fig. 9. Typing Rules.

Also, we use $x_1 : B_1, \dots, x_k : B_k <: x_1 : B'_1, \dots, x_k : B'_k$ when $B_i <: B'_i$, for all i in $1, \dots, k$, and Δ_{end} to denote $x_1 : \mathbf{end}, \dots, x_k : \mathbf{end}$.

We say process P is well-typed if $\Delta; \Gamma \vdash P$ may be derived using the rules given in Fig. 9. We discuss the key features of the typing rules. Rule (T-END) says the inactive process has no linear behavior (but complies to any shared behavior specification). Rule (T-PAR) types the parallel composition process with the linear types that are split in the behaviors of the two parallel branches, while ensuring both branches comply to the same usage of shared types. Rule (T-NEW) types a restricted linear name provided its usage is matched, i.e., it has no outstanding unmatched (? or !) communications. Rule (T-SNEW) types a restricted shared name, if it is used according to a shared type.

Rules for communication prefixes are divided in three groups, depending on the shared or linear usage of both communication subject and object. Rules (T-SIN) and (T-SOUT) address the case when the communication subject has shared usage while the object has linear usage. Notice that the behavioral type B , specified in the argument type of the shared type $l(B)$ of x , captures the slice of behavior that is delegated in the communication. Type B describes the linear usage of the input parameter in the premise of (T-SIN). Argument type B is also used in the conclusion of (T-SOUT), singled out via splitting so as to take into account the usage of y (the sent name) by the continuation (crucial to type processes that delegate a name and continue to interact in it).

Rules (T-IN) and (T-OUT) address the cases when both the communication subject and object have linear usage, and follow the lines above. Both rules record the prefixed type $\rho\{l_i(B'_i).B_i\}_{i \in I}$ in the conclusions, where ρ is either $?r$ or $!r$ for

$$\frac{k \in I}{s \rightarrow r\{l_i(M_i).B_i\}_{i \in I} \xrightarrow{s \rightarrow rl_k} B_k} \quad \frac{B_1 \xrightarrow{s \rightarrow rl} B_2}{B_1 \mid B \xrightarrow{s \rightarrow rl} B_2 \mid B} \quad \frac{B_1 \xrightarrow{s \rightarrow rl} B_2}{B \mid B_1 \xrightarrow{s \rightarrow rl} B \mid B_2}$$

Fig. 10. Type Reduction.

$$\Delta; \Gamma \xrightarrow{\tau} \Delta; \Gamma \quad \Delta; \Gamma, x : l(B) \xrightarrow{x : s \rightarrow rl} \Delta; \Gamma, x : l(B) \quad \frac{B_1 \xrightarrow{s \rightarrow rl} B_2}{\Delta, x : B_1; \Gamma \xrightarrow{x : s \rightarrow rl} \Delta, x : B_2; \Gamma}$$

Fig. 11. Typing Environment Reduction.

input and output, respectively. A single output is typed with a communication menu (containing the label of the emitted message) so as to directly match input menus. Notice that the prefixed type is taken up to subtyping, so as to allow to introduce \diamond types that may be necessary for the split in the conclusion. Notice also that the prefixed type is singled out via splitting, so as to take into account behaviors of x originally assigned to other threads (due to name delegation). Rules (T-LSIN) and (T-LSOUT) follow similar lines, addressing the case when the communication subject/object have linear/shared use. The last two rules are (T-REP), which types the replicated process, provided it uses no linear names, and the subsumption rule (T-SUB).

We can show that typing is preserved by substitution and by structural congruence. Given that our main result involves relating process actions and type specifications, we introduce type reduction, defined by the rules given in Fig. 10. In this way, we are able to precisely describe process reductions via the corresponding type reductions. Type reduction specifies how matched types reduce, explaining a message exchange that activates the respective continuation. Type reduction relies on reduction labels of the form $s \rightarrow rl$, identifying the roles involved in the communication and the label of the exchanged message.

Type reduction provides the expected semantics of behavioral types. Building on type reduction and in order to simplify the presentation of the results we introduce typing environment reduction, given by the rules in Fig. 11. Typing environment reduction specifies that environments seamlessly mimic internal τ (non public) reductions as well as synchronizations on shared channels. Also, typing environments exhibit linear reductions provided the reduction is observable at the level of the type of the respective channel. We may now state our main result that explains process reduction via typing environment reduction.

Theorem 1 (Type Preservation).

If $\Delta; \Gamma \vdash P$ and $P \xrightarrow{\lambda} P'$ then $\Delta; \Gamma \xrightarrow{\lambda} \Delta'; \Gamma$ and $\Delta'; \Gamma \vdash P'$.

The proof follows by induction on the length of the derivation of $P \xrightarrow{\lambda} P'$ (full details can be found in the supporting technical report [1]). Theorem 1 states

that any reduction of a well-typed process is explained by the corresponding type reduction, thus ensuring that processes interact according to the protocols prescribed by the types. Notice that this compliance entails that the protocols are actually carried out by the roles accordingly to the type specifications. We provide a precise characterization of this property as follows.

We now define role-based protocol fidelity. Let P be a process and Δ, Γ typing environments. We say P follows the role-based protocols prescribed by Δ, Γ if for any reduction sequence of the process $P \xrightarrow{\lambda_1} P_1 \xrightarrow{\lambda_2} \dots \xrightarrow{\lambda_k} P_k$ there is a matching reduction sequence of the typing environments $\Delta; \Gamma \xrightarrow{\lambda_1} \Delta_1; \Gamma \xrightarrow{\lambda_2} \dots \xrightarrow{\lambda_k} \Delta_k; \Gamma$. We have that well-typed processes satisfy role-based protocol fidelity as a direct consequence of Theorem 1.

Corollary 1 (Role-Based Protocol Fidelity).

If $\Delta; \Gamma \vdash P$ then P follows the role-based protocols prescribed by Δ, Γ .

In order to provide further intuition on our type system we show part of the typing of the running example (see Fig. 1). The type of name *chat*, as described in (1), is checked by successively splitting and matching resulting types with subprocesses. For example, in the typing of the **Buyer** process, after the first delegation, the type of *chat* can be decomposed by using rules (S-END), (S-BRK) and (S-BRK) (*b* abbreviates *Buyer*).

$$\frac{\frac{\frac{\diamond ?\mathbf{b}\{details().\mathbf{end}\} = \mathbf{end} \circ \diamond ?\mathbf{b}\{details().\mathbf{end}\}}{\frac{?\mathbf{b}\{price().\diamond ?\mathbf{b}\{details().\mathbf{end}\}} = ?\mathbf{b}\{price().\mathbf{end}\} \circ \diamond ?\mathbf{b}\{details().\mathbf{end}\}}{\mathbf{!}\mathbf{b}\{buy().?\mathbf{b}\{price().\diamond ?\mathbf{b}\{details().\mathbf{end}\}}\} = \mathbf{!}\mathbf{b}\{buy().?\mathbf{b}\{price().\mathbf{end}\} \circ \diamond ?\mathbf{b}\{details().\mathbf{end}\}}}}$$

Now, the split given above appears when typing the subprocess

$$chat \triangleleft_{\mathbf{Buyer}} buy().(chat \triangleright_{\mathbf{Buyer}} price() \mid mailBox \triangleleft_{\mathbf{Buyer}} storeService(chat))$$

Here, the delegation of name *chat*, on message *storeService*, requires that the behavior of *chat* is split between the two processes. Using (T-SUB), (T-SOUT) and (T-END) we have the following derivation.

$$\frac{\frac{\frac{chat : \mathbf{end}; mailBox : storeService(?\mathbf{b}\{details().\mathbf{end}\}) \vdash \mathbf{0}}{chat : ?\mathbf{b}\{details().\mathbf{end}\}; mailBox : storeService(\cdot) \vdash mailBox \triangleleft_{\mathbf{Buyer}} storeService(chat)}}{chat : \diamond ?\mathbf{b}\{details().\mathbf{end}\}; mailBox : storeService(\cdot) \vdash mailBox \triangleleft_{\mathbf{Buyer}} storeService(chat)}}$$

The example shows that the *sometime* operator behaves as a delayed choice between a *dot*, which expresses the sequentiality of behaviors, and a *parallel* composition, which types concurrent actions. These alternatives are introduced by rules (S-TAU), in only one of the branches of the split types, and in order to preserve, globally, the specified order of labels. Conceivably, the same flexibility would be achieved by a different (S-TAU) rule, which would immediately select between *dot* and *parallel*. Nevertheless, such rule would need to “look inside” the types and pull *parallels* to the top level. Therefore, this extension of session types with a new modality for breaking sequentiality, enriches the languages of types with an operator that enables us to perform choices locally and as needed. Such innovation supports a simple algebraic definition of the split operation.

4 Concluding Remarks

Our development is based on previous work on conversation types [4], extended so as to address assignment of dynamic roles to the several parties involved in a concurrent system. Technically, we identified a minimal set of ingredients to add to a core process specification language (the π -calculus [13, 15], TyCO [16] more precisely) so as to address role-based protocol verification (labeled channels and role annotations) and extended the type analysis accordingly. Noticeably, the split relation defined in this paper is much more readable and also more expressive than the merge relation in [4] — in particular, it allows for splitting (the same) behavior out of the continuations of a branching behavior. Crucial to our development is the introduction of the \diamond type to control behavior interleaving.

We discuss some possible extensions to our work. A necessary further development is the extension of the model to consider infinite behavior. An essential feature of any type analysis is a verification procedure. We are yet to implement such a procedure, but we may already assert there exists such a procedure in a setting where all bound names are type annotated. Another crucial property left out of this paper is progress. However, we expect that the progress analysis introduced in [4] for a labeled π -calculus, combined with our typing analysis, may be used to single-out systems that enjoy progress. An interesting further development to be addressed is the dynamic delegation of roles. In our setting roles are statically annotated in processes. Extending the language with role delegation would allow parties to dynamically assume unanticipated roles.

Several works address role-based type specifications to enforce security concerns (for example [8] introduces a type analysis to discipline role-based access control to data). We focus on communication protocol assignment and leave security to be handled orthogonally. Our approach builds on conversation type theory, introduced as a generalization of session types [10, 12] to discipline multiparty interaction, including dynamically established conversations with an unanticipated number of participants. Other works share the goal to address multiparty interaction [2, 3, 6, 11, 14]. In particular with respect to the works more closely related to ours [2, 11], we single out the approach of conversation types since it addresses multiparty interaction where the number of participants is not fixed a priori, while considering a simpler underlying model. We remark that in [2, 6, 11] a notion of role assignment is explicit, unlike in [4] where types do not mention identities of communicating partners. However, such role assignment is achieved via a structural projection, forcing single roles to be carried out by single threads. A different notion of dynamic roles is also considered in the approaches described in [7, 9], allowing for several processes, much like a thread pool, to simultaneously carry out a single role.

In this work we have presented a type-based analysis that ensures that systems follow the prescribed role-based protocol specifications. Novel to our approach is the flexibility of role assignment, allowing us to address dynamic distributed implementations of role specifications, where a single role can be distributed among several processes and a single process can dynamically switch between roles. To the best of our knowledge, ours is the only (session-type like)

approach that addresses such configurations, that are actually found in, e.g., real world business protocols. Our development extends conversation types with role-based protocol specifications, retaining the simplicity of the approach, simplifying and generalizing the underlying technical framework, and contrasting with related approaches in the dynamic and flexible nature of roles.

Acknowledgments Work partially supported by FCT Pest UI527 2011 and projects PTDC/EIA-CCO/113033/2009 ComFormCrypt, PTDC/EIA-CCO/104583/2008 StreamLine, PTDC/EIA-CCO/117513/2010 Liveness, Statically, PTDC/EIA-CCO/122547/2010 Multicore, and CMU-PT NGN-44A Interfaces.

References

1. Baltazar, P., Caires, L., Vasconcelos, V.T., Vieira, H.: Dynamic Roles in Multiparty Communicating Systems. UNL-DI-1-2012, Universidade Nova de Lisboa (2012)
2. Bettini, L., Coppo, M., D'Antoni, L., De Luca, M., Dezani-Ciancaglini, M., Yoshida, N.: Global Progress in Dynamically Interleaved Multiparty Sessions. In: CONCUR 2008. LNCS, vol. 5201, pp. 418–433. Springer (2008)
3. Bravetti, M., Zavattaro, G.: A Foundational Theory of Contracts for Multi-party Service Composition. *Fundamenta Informaticae* 89(4), 451–478 (2008)
4. Caires, L., Vieira, H.: Conversation Types. *Theoretical Computer Science* 411(51-52), 4399–4440 (2010)
5. Carbone, M., Honda, K., Yoshida, N.: Structured Communication-Centred Programming for Web Services. In: ESOP 2007. LNCS, vol. 4421, pp. 2–17. Springer (2007)
6. Castagna, G., Dezani-Ciancaglini, M., Padovani, L.: On Global Types and Multiparty Sessions. In: FMOODS/FORTE 2011. LNCS, vol. 6722, pp. 1–28. Springer (2011)
7. Deniérou, P.M., Yoshida, N.: Dynamic Multirole Session Types. In: POPL 2011. pp. 435–446. ACM (2011)
8. Ghilezan, S., Jaksic, S., Pantovic, J., Dezani-Ciancaglini, M.: Types and Roles for Web Security. *Transactions on Advanced Research* 8(2), 16–21 (2012)
9. Giachino, E., Sackman, M., Drossopoulou, S., Eisenbach, S.: Softly Safely Spoken: Role Playing for Session Types. In: PLACES 2009 (2009)
10. Honda, K.: Types for Dyadic Interaction. In: CONCUR 1993. LNCS, vol. 715, pp. 509–523. Springer (1993)
11. Honda, K., Yoshida, N., Carbone, M.: Multiparty Asynchronous Session Types. In: POPL 2008. pp. 273–284. ACM (2008)
12. Honda, K., Vasconcelos, V.T., Kubo, M.: Language Primitives and Type Discipline for Structured Communication-Based Programming. In: ESOP 1998. LNCS, vol. 1381, pp. 122–138. Springer (1998)
13. Milner, R., Parrow, J., Walker, D.: A Calculus of Mobile Processes, Part I and II. *Information and Computation* 100(1), 1–77 (1992)
14. Padovani, L.: On Projecting Processes into Session Types. *Mathematical Structures Computer Science* 22, 237–289 (2012)
15. Sangiorgi, D., Walker, D.: *The π -Calculus: A Theory of Mobile Processes*. Cambridge University Press (2001)
16. Vasconcelos, V.T., Tokoro, M.: A typing system for a calculus of objects. In: ISO-TAS 1993. LNCS, vol. 472, pp. 460–474. Springer (1993)